

High Resolution Mesh Convergence Properties and Parallel Efficiency of a Spectral Element Atmospheric Dynamical Core

John Dennis* Aimé Fournier[†] William F. Spotz[‡]
Amik St.-Cyr* Mark A. Taylor^{‡§} Stephen J. Thomas*
Henry Tufo*

Abstract

We first demonstrate the parallel performance of the dynamical core of a spectral element atmospheric model. The model uses continuous Galerkin spectral elements to discretize the surface of the Earth, coupled with finite differences in the radial direction. Results are presented from two distributed memory, mesh interconnect supercomputers (ASCI Red and BlueGene/L), using a two-dimensional space filling curve domain decomposition. Better than 80% parallel efficiency is obtained for fixed grids on up to 8938 processors. These runs represent the largest processor counts ever achieved for a geophysical application. They show that the upcoming Red Storm and BlueGene/L supercomputers are well suited for performing global atmospheric simulations with a 10km average grid spacing. We then demonstrate the accuracy of the method by performing a full 3D mesh refinement convergence study, using the primitive equations to model breaking Rossby waves on the polar vortex. Due to the excellent parallel performance, the model is run at several resolutions up to 36km

*Scientific Computing Division, National Center for Atmospheric Research, Boulder Colorado

[†]Department of Meteorology, College of Computer, Mathematical, Physical Sciences, University of Maryland at College Park

[‡]Computation, Computers, Information and Mathematics, Sandia National Laboratories, Albuquerque NM

[§]Corresponding author. email: mataylo@sandia.gov

with 200 levels using only modest computing resources. Isosurfaces of scaled potential vorticity (PV) exhibit complex dynamical features, e.g., a primary PV tongue, and a secondary instability causing roll-up into a ring of five smaller sub-vortices. As the resolution is increased, these features are shown to converge while PV gradients steepen.

1 Introduction

The spectral element method is a finite element method in which a high degree spectral method is used within each element. The method provides spectral accuracy while retaining both parallel efficiency and the geometric flexibility of unstructured finite elements grids. The method has proven accurate and efficient for a wide variety of geophysical problems, including global atmospheric circulation modeling (???????) ocean modeling (????) and planetary scale seismology (??). The method has unsurpassed parallel performance; it was used for earthquake modeling by the 2003 Gordon Bell Best Performance winner, running on 1944 CPUs of the Earth Simulator (?) and for climate modeling by a 2002 Gordon Bell Award honorable mention, running on 2048 processors of an IBM SP (?).

In this work we focus on a spectral element atmospheric climate model. The model is a prototype dynamical core for the Community Atmospheric Model (CAM) component of the Community Climate System Model (CCSM). The CCSM simulates the total earth system, coupling together atmosphere, ocean, sea ice and land surface models. The CAM consists of a dynamical core based on the hydrostatic primitive equations, coupled to sub-grid scale models of physical processes such as boundary layer turbulence, moist convection and the effects of clouds on the radiative forcing (??). We describe our model in Section 2 and then in Section 3 we document its parallel performance on very large processor counts. We use these results to demonstrate the ability of two upcoming parallel computers to run the spectral element model at a global resolution of 10km. This is an important long term goal of DOE's climate modeling program (?). In Section 4, we use the polar vortex problem to perform a high resolution mesh convergence study of the model. The ability to run at the resolutions necessary to start to achieve convergence were made possible by the parallel performance of the model.

2 Spectral Element Dynamical Cores

Our spectral element dynamical core solves the primitive equations using a hybrid η pressure vertical coordinate system (?), using the continuous Galerkin spectral element discretization for the horizontal directions (the surface of the sphere), and second order finite differences in the vertical direction (??). The spectral element method relies on quadrilateral elements. We use a subdivided inscribed cube to generate quasi-isotropic tilings of the sphere with such elements. An example mesh is shown in Fig. 1. To characterize the horizontal resolution of these meshes, let M be the total number of elements and N be the number of polynomials in each direction used within each element. In the figure, $M = 384$. The spectral transforms performed within each element rely on an $N \times N$ grid and thus the total horizontal resolution is specified by $N \times N \times M$. The number of nodes along the equator is given by $4(M/6)^{1/2}N$, and the average equatorial grid spacing in kilometers is given by $2.45 \times 10^4 M^{-1/2} N^{-1}$.

Using a spectral element discretization on the sphere has several advantages for global climate modeling. First of all, handling the spherical geometry presents no problem since the sphere can be tiled with quadrilateral elements of approximately the same size, thus avoiding clustering points at the poles. Secondly, by using a local coordinate system within each element, the singularities associated with spherical coordinates can also be avoided. Additionally, for climate applications, the method can obtain the accuracy of traditional spherical harmonics based models with only a slight increase in the number of degrees of freedom (??).

In what follows we present results from two models, SEAM and HOMME. SEAM is a spectral element atmospheric model research code (?). Most of the algorithms in SEAM have been reimplemented in NCAR's High Order Multi-scale Modeling Environment (HOMME). HOMME adds modern software engineering practices in addition to many new features and algorithms, including advanced time stepping, discontinuous Galerkin, adaptive mesh refinement and several domain decomposition strategies (????). Here we only use the simplest form of HOMME: explicit time stepping, a continuous Galerkin treatment of the prognostic variables and quasi-isotropic conforming element meshes. In this form, the numerical results of HOMME and SEAM are indistinguishable.

3 Parallel Scalability

Spectral element methods are well suited to modern cache-based parallel computers for several reasons. First, the basic data structure in the method, the spectral element, is naturally cache blocked. Secondly, due to the $O(N^3)$ cost of the spectral transforms, the method has a very low ratio of communication to computation. Finally, the spectral element discretization allows for efficient two-dimensional domain decomposition strategies.

We demonstrate this performance by presenting results for HOMME running on the IBM BlueGene/L (up to 7776 processors) and ASCI Red (up to 8938 processors). We use vertical resolutions from 20 to 100 levels, and horizontal resolutions from 156km down to 10km. All cases use $N = 8$. For the benchmark problem, we use timings from HOMME configured to run the Held-Suarez test (?). The Held-Suarez tests were designed to allow for the intercomparison of the climate produced by different atmospheric dynamical cores. Results from SEAM and HOMME have been previously reported (??).

The Held-Suarez tests are also useful for benchmarking since they represent the entire dynamical core of an atmospheric model such as CAM. Only the physics components (sub-grid scale models of physical processes) in CAM are not represented. CAM physics is column based, meaning the calculations are performed using only data from the vertical column associated with each element. Using an element based domain decomposition, all the column physics will be performed on processor and require no additional communication. Thus the addition of physics can only improve the parallel scalability of the dynamical core; however, the single processor performance will be affected by the single processor performance of the column physics models. One other difference between CAM and the Held-Suarez tests which can effect parallel performance is that CAM contains additional constituents which must be advected by the dynamical core. These prognostic variables will have an identical parallel behavior as the prognostic variable for temperature, and thus they are represented in the Held-Suarez test but going from one such variable to many such variables will result in a proportional increase in the required CPU time.

These factors show that the time-to-solution from a Held-Suarez test problem can provide a reasonable estimate of the time-to-solution for a full climate simulation. For a given resolution, adding column physics and additional constituents may actually increase the parallel scalability of the model and the number of floating point operations per second. The time-to-solution

will increase, but in the worst case only proportional to the complexity of the column physics and number of additional constituents. In typical climate simulations, this increase is a factor in the range of 1.2 to 2.

3.1 Cache Blocking

The elements in the spectral element method provide a natural cache blocking. By way of example, consider the 64 bit data storage requirement for a typical 8×8 element with 20 vertical levels requires only 10 Kbytes of cache per variable. One hundred such 3-D variables will fit into a typical 1 Mbyte data cache. If the data is stored in memory so as to avoid cache conflicts, then all the computations performed within an element can be done entirely in cache. This blocking is independent of resolution since we can increase the resolution by simply using more elements. This situation is unlike codes optimized for vector architecture, where the natural block size and data access patterns grow with resolution. Thus a spectral element model maintains good performance even at high resolutions.

3.2 Domain decomposition

The most natural way to parallelize the spectral element method is to simply assign several elements to each processor. Each element only needs information from adjacent elements, so the domain decomposition reduces to a standard graph partitioning problem. To solve this problem, we use an algorithm based on space filling curves (?).

The resulting communication patterns are similar to finite difference/finite element methods which parallelize with the same type of domain decomposition. Thus the parallel efficiency of this method would be similar to these other methods, except for the fact that the spectral element method involves computing high order derivatives using spectral transforms instead of low order stencils. These computationally intensive transforms are performed within each element and are localized to each processor thus requiring no communication. The result is that for a wide range of problem sizes, number of processors, and computer architectures, the parallel efficiency remains above 80%, for parallel decompositions as fine as two elements per processor, and reasonable efficiency is obtained at the finest decomposition of one element per processor.

Resolution	Processors	GFLOPS	Simulated days per day
156km/26L	384	13	1600
40km/50L	6144	181	154
20km/70L	8192	265	18
9.8km/100L	8936	294	1.5

Table 1: HOMME dynamical core (dry dynamics) benchmark runs on ASCI-Red

3.3 HOMME on ASCI Red

We first present results from the ASCI-Red computer at Sandia National Laboratories. ASCI-Red was the first of DOE’s Advanced Strategic Computing Initiative (ASCI) machines. It has 4510 nodes in a mesh interconnect, each with two 333 MHz Pentium II processors. In Fig. 2, we present the total MFLOPS obtained per processor at several resolutions and processor counts. Each curve represents a fixed resolution, so that the the total amount of work was kept constant while the processor count was increased. Data for two resolutions includes parallel decompositions as fine as one element per processor. As can be seen in the figure, all resolutions achieve good scalability - better than 80% in all cases. In this range the performance fluctuates between 30 and 40 MFLOPS per processor. The best performance obtained for each resolution, along with the integration rate is given in Table 1.

ASCI Red is of interest because of Cray’s upcoming Red Storm architecture. As of this writing, several Red Storm computers are being built with up to 10,000 Opteron processors running at 2 GHz. Red Storm was designed to have the same balance between processor performance and inter-processor communications as ASCI Red, and thus we expect to obtain very similar scalability results on Red Storm. Initial results show the single processor performance of HOMME on a 2 GHz Opteron is better than 600 MFLOPS. Thus with an 80% parallel efficiency, we expect to sustain at least 4.3 TFLOPS on 8936 processors of Red Storm. This is an integration rate of 22 simulated days per day for a 9.8km/100L resolution.

3.4 HOMME on IBM BlueGene/L

Our benchmark runs for the IBM BlueGene/L system were run on systems at IBM Watson and Rochester, on up to 7776 processors using one processor

Resolution	Processors	GFLOPS	Simulated days per day
80km/20L	512	117	7200
71km/20L	1944	372	6500 [†]
35km/40L	7776	1320	1100

Table 2: HOMME dynamical core benchmark runs on BlueGene/L. The run denoted by a † included one prognostic moisture variable.

per node and only using one of the floating-point pipelines. BlueGene/L uses a toroidal interconnect and PowerPC processors. Sustained MFLOPS per processor are shown in Fig. 3, for several resolutions and processor counts. Each curve represents a fixed resolution, so that the total amount of work was kept constant while the processor count was increased. Data for two resolutions includes parallel decompositions as fine as one element per processor. The best performance obtained for each resolution, along with the integration rate is given in Table 2.

BlueGene/L is of interest since IBM intends to build some very large systems, with over 64,000 processors. Noting that the performance on the dry dynamics problem is around 250 MFLOPS per processor even when running with as few as 2 elements per processor, we can estimate a performance at 9.8km/100L on 49152 processors (2 elements per processor) to be 9.8 TFLOPS, thus allowing an integration rate of 50 simulated days per day.

4 The Polar Vortex

The stratospheric polar vortex is bounded by strong potential-vorticity gradients which isolate polar air from lower latitudes. A phenomenon of major atmospheric-research importance is the eventual mixing of these air masses, when the vortex is distorted by breaking planetary-scale Rossby waves. This was the focus of ?. In addition to that study and those referenced therein, the importance of this phenomenon is underscored by a number of recent studies of its various aspects, in various journals. ? studied the positive feedbacks involving ozone destruction and other phenomena that may delay polar-vortex breakdown and enhance spring ozone loss. ? provided objective diagnostics of characteristic winter polar-vortex displacements towards Europe or Canada, linked to transient dynamically induced ozone mini-holes. Planetary Rossby-wave refraction was found by ? to link stratospheric polar-vortex

strength to tropospheric annular-mode variability. ? evaluated the crucial choice of boundaries for computing material transport across the polar-vortex edge. Finally, ? found that an “unusually strong upper stratospheric vortex . . . caused enhancements” of NO_x that led to “reductions in O_3 of more than 60% in some cases”.

All these authors also cite numerous other examples of this topic’s importance. Here we document a minimally complex model for the numerical simulation of polar-vortex breakdown based on that presented in ?, and then present results from SEAM for several different resolutions up to 36km with 200 levels.

4.1 Initial Condition

Following ?, our initial model state consists of a gradient-flow balanced axisymmetric-vortex: at time $t = 0$, as a function of longitude λ , co-latitude φ and pressure p , let the horizontal wind vector

$$\mathbf{u}(\lambda, \varphi, p) = \mathbf{u}^0(\varphi, p) \equiv a_0 \Omega_0 u^0(\varphi, p) \mathbf{i},$$

the isobaric velocity $\omega(\lambda, \varphi, p) = 0$, the surface pressure $p_{\text{sfc}}(\lambda, \varphi) = 10^4 \text{Pa}$, the temperature

$$T(\lambda, \varphi, p) = T^0(\varphi, p) \equiv T_0 - \frac{a_0^2 \Omega_0^2}{R_0} \frac{\partial \Phi^0}{\partial \ln p}$$

and the geopotential

$$\Phi(\lambda, \varphi, p) = g_0 z_p + a_0^2 \Omega_0^2 \Phi^0(\varphi, p),$$

where a_0 , Ω_0 and g_0 are the earth’s radius, angular frequency and gravity, \mathbf{i} is the longitudinal unit vector, $T_0 = g_0 H_0 / R_0 = 239.14 \text{K}$ is the isotherm for a scale height $H_0 = 7 \text{km}$, R_0 is the dry-air gas constant and $z_p \equiv H_0 \ln p_{\text{sfc}} / p$ is log-pressure height. The nonlinear gradient-flow balance is

$$\frac{\partial \Phi^0}{\partial \varphi} = - (2\mu + u^0 \tan \varphi) u^0,$$

where $\mu = \sin \varphi$, so the complete initial state only depends on the dimensionless velocity component u^0 , assuming $\int_{-1}^1 \Phi^0(\varphi, p) d\mu = 0$.

Let us now describe the initial profile as a function of (φ, p) . Observing equatorward from $\varphi = \pi/2$, we assume the absolute vorticity

$$\zeta \equiv 2\mu - \frac{\partial u \cos \varphi}{\partial \mu} \quad (1)$$

is initially positive, almost constant until the vortex-edge latitude $\varphi_v \equiv 60^\circ$, where it decreases rapidly over a zone of width $\Delta\varphi \equiv 6^\circ$, then again nearly constant until the *surf-zone* edge $\varphi_s \equiv 35^\circ$, followed by cubic-sine decrease to zero at the equator, and solid-body rotation for $\varphi \leq 0$. We impose $u_z^0 = 0$, where subscript z denotes evaluation at $\varphi = \varphi_z \equiv 37^\circ$ and any p . The equatorward progression just described is controlled by four vertical-structure coefficients $c_n(p)$ to be determined:

$$\zeta^0(\varphi, p) \equiv \begin{cases} r'(\mu)c_1 + c_2, & \mu_s \leq \mu, \\ 2\mu c_3 - 4\mu^3 c_4, & 0 \leq \mu \leq \mu_s, \\ 2M_e^0 \mu, & \mu \leq 0, \end{cases} \quad (2)$$

where the step-like function

$$r'(\mu) \equiv \frac{1}{2} + \frac{1}{2} \tanh \frac{\mu - \mu_v}{\Delta\mu}, \quad (3)$$

the specific absolute angular momentum from (1)

$$M(\mu, p) \equiv 1 - \mu^2 + u \cos \varphi = \int_{\mu}^1 \zeta(\mu') d\mu', \quad (4)$$

subscript e denotes evaluation at the equator $\varphi = 0$, subscripts s and v denote evaluation at φ_s and φ_v , and $\Delta\mu \equiv \Delta\varphi \cos \varphi_v$. The coefficients $c_n(p)$ are determined as follows. From (2,4) one has

$$M^0(\mu, p) = \begin{cases} -r(\mu)c_1 + (1 - \mu)c_2, & \mu_s \leq \mu, \\ M_e^0 - \mu^2 c_3 + \mu^4 c_4, & 0 \leq \mu \leq \mu_s, \\ (1 - \mu^2)M_e^0, & \mu \leq 0, \end{cases} \quad (5)$$

where by (3) the hyperbolic ramp-like function

$$r(\mu) \equiv \int_1^{\mu} r'(\mu') d\mu' = \frac{\mu - 1}{2} + \frac{\Delta\mu}{2} \left(\ln \cosh \frac{\mu - \mu_v}{\Delta\mu} - \ln \cosh \frac{1 - \mu_v}{\Delta\mu} \right).$$

Since $\varphi_v > \varphi_z > \varphi_s$ we may immediately solve a linear system to obtain

$$c_2 = \frac{M_v^0 r_z - M_z^0 r_v}{(1 - \mu_v) r_z - (1 - \mu_z) r_v},$$

$$c_1 = [(1 - \mu_v) c_2 - M_v^0] r_v^{-1}.$$

Ensuring μ -continuity of (2,5) leads to another linear system with solution

$$c_4 = \mu_s^{-4} (M_e^0 - M_s^0) - 2^{-1} \mu_s^{-3} \zeta_s^0,$$

$$c_3 = (2\mu_s)^{-1} \zeta_s^0 + 2\mu_s^2 c_4.$$

Observing that

$$M_z^0 = 1 - \mu_z^2$$

and

$$M_s^0 = -r_s c_1 + (1 - \mu_s) c_2,$$

it only remains to define $a_0 \Omega_0 u_e^0 \equiv -20 \text{m s}^{-1}$ in order to get M_e^0 from (4), and then

$$u_v^0 \equiv (1 + z_p/z_{\text{top}}) u_e^0$$

in order to get M_v^0 from (4), where z_{top} is the model height. Finally, (4,5) yields $u^0(\varphi, p)$, as shown in Fig. 4a.

The initial scaled-potential-vorticity (Π) profile is shown in Fig. 4b. $\Pi \equiv P(T, \zeta)/P(T_0, 2\Omega_0)$, where

$$P(T, \zeta) \equiv \left(\sec \varphi \frac{\partial v}{\partial p} \frac{\partial}{\partial \lambda} - \frac{\partial u}{\partial p} \frac{\partial}{\partial \varphi} - \zeta \frac{\partial}{\partial p} \right) \theta(p, T)$$

approximates the Ertel potential vorticity in isobaric coordinates, $\theta(p, T) \equiv (p_0/p)^{R_0/c_p} T$ is the potential temperature and $p_0 = 10^5 \text{Pa}$. In fact Π thus defined is identically the *SPV* of ?, section 3.

4.2 Time-dependent forcing at the bottom

The model is forced by setting Φ at the surface for all time equal to

$$\Phi_{\text{stc}}(\lambda, \varphi, t) = \begin{cases} \Phi_{\text{max}} \cos \lambda \left(\sin \frac{9}{2} \varphi \right)^2 (1 - e^{-t/\tau}), & 40^\circ < \varphi < 80^\circ, \\ 0, & \text{otherwise,} \end{cases}$$

with decay time $\tau = 3 \text{d}$ and amplitude $\Phi_{\text{max}} = g_0 \times 800 \text{m}$. The forcing peaks at $\varphi = 60^\circ = \varphi_v$, and increases monotonically in time. The effect is to instigate upward-propagating Rossby waves.

4.3 Sponge layer at the top

Depending on z_{top} , spurious reflection may be prevented by adding the following sponge-layer forcing near the model top, which amounts to Rayleigh damping of \mathbf{u} and Newton damping of T :

$$\left. \frac{\partial(\mathbf{u}, T)}{\partial t} \right|_{\text{sl}} = \nu \cdot (\mathbf{u}^0 - \mathbf{u}, T^0 - T), \quad (6)$$

where the damping coefficient

$$\nu \equiv \begin{cases} \frac{1}{2} + \frac{1}{2} \tanh \frac{z_p - \bar{z}_{\text{sl}}}{\Delta z}, & z_{\text{sl}} < z_{\text{top}}, \\ 0, & z_{\text{sl}} \geq z_{\text{top}}, \end{cases} \quad (7)$$

$\bar{z}_{\text{sl}} \equiv 2^{-1}(z_{\text{sl}} + z_{\text{top}})$, $\Delta z \equiv \frac{3}{10}(z_{\text{top}} - z_{\text{sl}})$ and $z_{\text{sl}} = 45\text{km}$.

4.4 Results

Using SEAM, we were able to carry out the polar-vortex simulation at several resolutions, up to 36km with 200 levels. The highest resolution required $\approx 10^5$ time steps on $\approx 10^8$ collocation points in ≈ 2 wall-clock days on 256 IBM SP RS/6000 processors. The Π isosurfaces from that run are shown in in Fig. 5. They indicate the complex dynamical features of the polar vortex: a primary Π ‘‘tongue,’’ succumbing to a secondary instability, leading to a roll-up into a ring of five or six smaller sub-vortices.

We look at the convergence question systematically, as shown in Fig. 6. It appears that at moderate horizontal resolution of 156km, some additional Π structure is resolved in increasing from 50 to 100 levels (Fig. 6a-b), but not as much, from 100 to 200 levels (Fig. 6b-c). At the higher horizontal resolution of 70km, the sub-vortices are much better resolved for 100 levels, with little change at 200 levels (Fig. 6d-f). Increasing the horizontal resolution to 36km at 200 levels (Fig. 6g) produces mainly small and qualitative changes, and only in the details, not in the overall simulation. The problem appears to have converged w.r.t. horizontal and vertical resolution.

The precise dissipation and filtering were different for all runs, but for each run were empirically minimized while sufficient to stabilize the simulation. That is, we attempt to investigate the inviscid limit. This differs from other convergence studies in which dissipation and filtering are fixed, and only resolution is varied.

To investigate convergence under mesh refinement more systematically, we consider an the isentropic Π -tongue-tip position diagnostic. This may be quantified as the point $\mathbf{x}_{\max}(\theta, t)$ of maximum angular change $\beta(\mathbf{x})$ of the $\Pi = 1$ contour in stereographic projection coordinates

$$\mathbf{x} \equiv \tan(\pi/4 - \varphi/2) \begin{bmatrix} \cos \lambda \\ \sin \lambda \end{bmatrix}. \quad (8)$$

The diagnostic involves computing the L tangent vectors \mathbf{t}_j along the L points of each contour:

$$\mathbf{t}_j \equiv \Delta \begin{cases} \mathbf{x}_j, & j = 1, \\ (\mathbf{x}_j + \mathbf{x}_{j-1})/2, & j = 2, \dots, L \\ \mathbf{x}_{j-1}, & j = L + 1, \end{cases}$$

where $\Delta s_j \equiv s_{j+1} - s_j$. From the \mathbf{t}_j one computes the angular change

$$\beta(\mathbf{x}_j) \equiv \arccos \hat{\mathbf{t}}_j \cdot \hat{\mathbf{t}}_{j+1} \leq \beta(\mathbf{x}_{\max}(\theta, t)), \quad (9)$$

where $\hat{\mathbf{a}} \equiv \mathbf{a}/|\mathbf{a}|$. This is illustrated in Fig. 7.

The area $A(\theta, t)$ enclosed by values $\Pi \geq 1$ evolves with a qualitatively similar pattern as resolution is increased, as shown in Fig. 8. This is a good measure of vortex erosion, as discussed in further detail by ?. We estimated $A(\theta, t)$ by summing over northern-hemisphere points with $\Pi(\lambda, \varphi, \theta, t) > 1$:

$$A(\theta, t) = \frac{2\pi}{N_\lambda} \sum_{m=1}^{N_\varphi} \Delta \sin \varphi_m^+ \sum_{l=0}^{N_\lambda-1} \begin{cases} 1, & \Pi(\lambda_l, \varphi_m, \theta, t) > 1, \\ 0, & \Pi(\lambda_l, \varphi_m, \theta, t) \leq 1, \end{cases} \quad (10)$$

where $\lambda_l \equiv (2N_\lambda^{-1}l - 1)\pi$, $\varphi_m \equiv (2N_\varphi)^{-1}m\pi$ and $\varphi_m^+ \equiv \varphi_m$ ($m = 1, \dots, N_\varphi$), $2\varphi_m - \varphi_{m-1}$ ($m = N_\varphi + 1$).

As an indication of convergence as resolution is increased, the ratio

$$\frac{s_{36\text{km}}/s_{70\text{km}} - 1}{s_{70\text{km}}/s_{156\text{km}} - 1}$$

for $s = A(\theta, t)$ stayed below 0.35 for all θ in Fig. 8, averaged over $t \in [0, 12]\text{d}$, and was usually much smaller.

5 Discussion

The spectral element method is known to perform well on parallel computers. This performance is maintained for a spectral element atmospheric model on very large, previously unattained, processor counts. The parallel efficiency remains between 70–80%, for parallel decompositions as fine as one element per processor and for all problem sizes and processor counts. These results give evidence that the massively parallel systems being built now (IBM's BlueGene/L and Cray's Red Storm) will be able to sustain between 4-10 TFLOPS of performance, allowing one to run 10km global atmospheric simulations at a rate of 22–50 simulated days per day. This establishes a performance level for atmospheric modeling competitive with that obtained using the specialized vector supercomputer architectures of the Japanese Earth Simulator. On that machine, the AFES (Atmospheric model for the Earth Simulator) obtained 27 TFLOPS of performance and an integration rate of 57 simulated days per day (?). AFES uses a global spectral model for its dynamical core, and thus requires more flops to achieve similar integration rates.

The parallel performance of the spectral element method allowed us to conduct a mesh convergence study using a polar vortex model problem. This problem is the focus of much recent research. It has a complex unstable evolution dominated by strong potential-vorticity gradients. Only at high resolution (36km and 200 vertical levels) does evidence for mesh convergence of large scale features start to become apparent.

6 Acknowledgments

We thank Lorenzo M. Polvani and R. Saravanan for several informative discussions. A.F. thanks the Office of Biological and Environmental Research, Office of Science, US Department of Energy, for support. M.T. and W.S. thank the Sandia National Laboratories' Lab Directed Research and Development program for support.

List of Figures

- 1 An example spectral element grid for the Earth. Continental outlines are shown for reference. The grid is generated by projecting an inscribed cube onto the surface, and then subdividing each of the 6 faces of the cube into an 8×8 grid of elements. Typically each element uses a degree 7 polynomial expansion. Finite differences are used in the vertical direction (not shown). 16
- 2 Performance of the HOMME dynamical core (dry dynamics) benchmark runs on ASCI Red. Each curve shows the performance for a fixed resolution as the number of processors is increased. 17
- 3 Performance of the HOMME dynamical core (dry and moist dynamics) benchmark runs on BlueGene/L. Each curve shows the performance for a fixed resolution as the number of processors is increased. 18
- 4 Initial state profiles vs latitude φ ($^\circ$, abscissa). (a) Zonal-wind isotachs (m s^{-1}) vs log-pressure height z_p (km, ordinate). Contour interval is 10 m s^{-1} and $u(\varphi = 37^\circ, z_p) = 0$. (b) Isopleths of scaled potential vorticity vs potential temperature θ (K, ordinate). Contour interval is 0.2 and $\Pi(\varphi = 0, \theta) = 0$ 19
- 5 Isosurfaces of scaled potential vorticity $\Pi = 1$ for the 36km run. Horizontal coordinates have been stereographically projected (8), and there are 22 vertical levels $\theta_l = 10^{2/7} T_0 \exp(2.5 \frac{l+3/2}{49})$. View is towards stereographic origin and $\theta_1 = 524\text{K}$, from $\lambda = 110^\circ\text{W}$, elevation 40° . Top plane shows contours of $\Pi > 1$ on $\theta_{22} = 1530\text{K}$. Times (a) $t = 13\text{d}$, (b-f) $t \in \{16, \dots, 20\}\text{d}$ 20
- 6 Stereographic projection in the the $\theta = 1500\text{K}$ surface of $\Pi(\lambda, \varphi)$ contours from -0.3 to 1.3 by 0.1 , at $t = 20\text{d}$. Values below 0.2 are dark gray and above 0.8 are white. Resolutions (a) $156\text{km}/50\text{L}$, (b) $156\text{km}/100\text{L}$, (c) $156\text{km}/200\text{L}$; (d-f) as in (a-c) but for 70km ; (g) as in (f) but for 36km 21
- 7 $\Pi = 1$ contours, for $t \in \{5, 7, 9, 11\}\text{d}$ (light gray to black). Diamonds indicate $\mathbf{x}_{\max}(\theta, t)$ (Eq. 9). (a-d) Descending levels $\theta = 3000, 2500, 2000, 1500\text{K}$ for 156km resolution. (e-h) As (a-d) but for 70km . (i-l) As (a-d) but for 36km 22

8	Horizontal vortex-area $A(\theta, t)$ (10) vs t (d, abscissa) for (a) $\theta = 3000$, (b) 2500, (c) 2000, (d) 1500K. Resolutions are 156km/50L (light gray), 70km/100L (gray) and 36km/200L (dark gray).	23
---	--	----

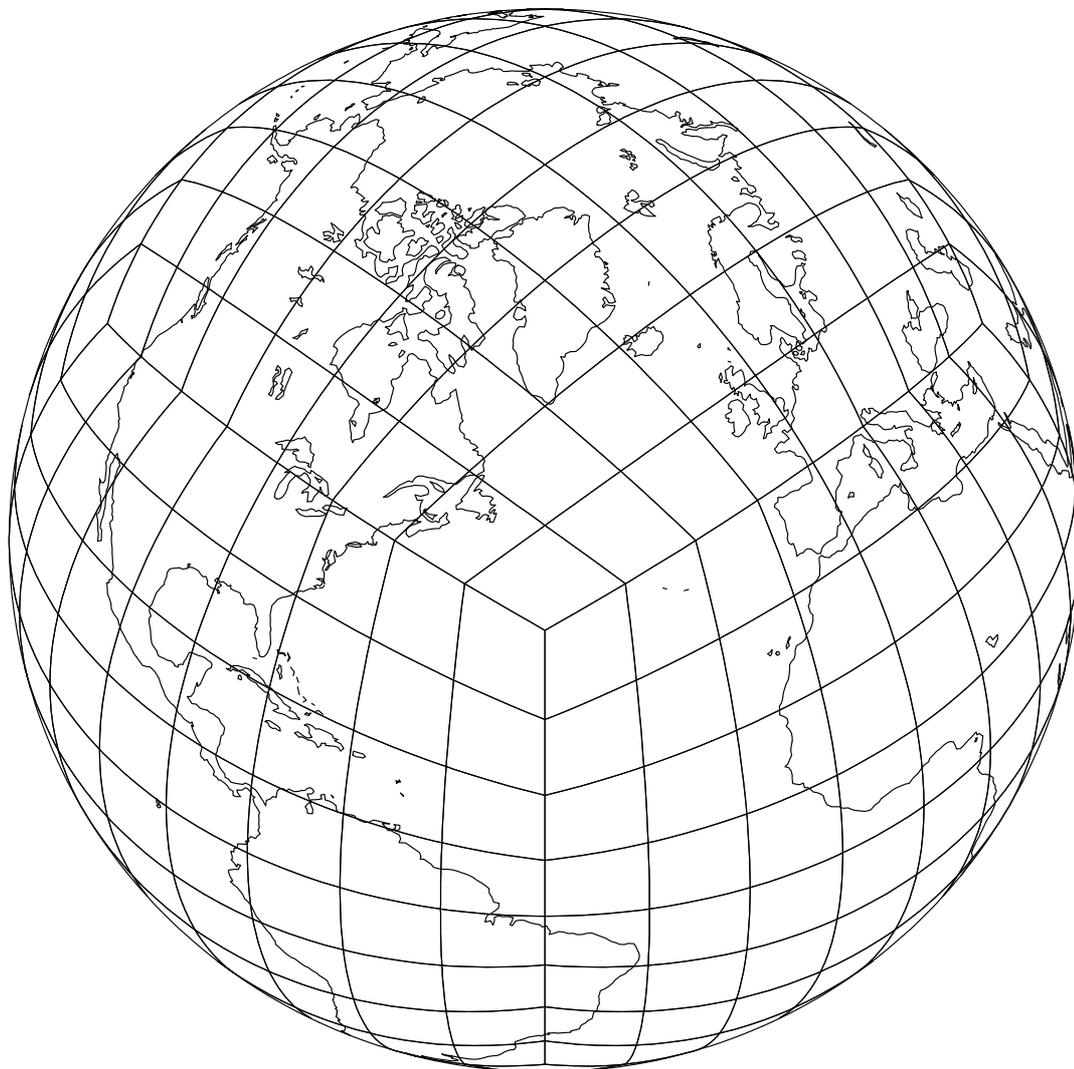


Figure 1: An example spectral element grid for the Earth. Continental outlines are shown for reference. The grid is generated by projecting an inscribed cube onto the surface, and then subdividing each of the 6 faces of the cube into an 8×8 grid of elements. Typically each element uses a degree 7 polynomial expansion. Finite differences are used in the vertical direction (not shown).

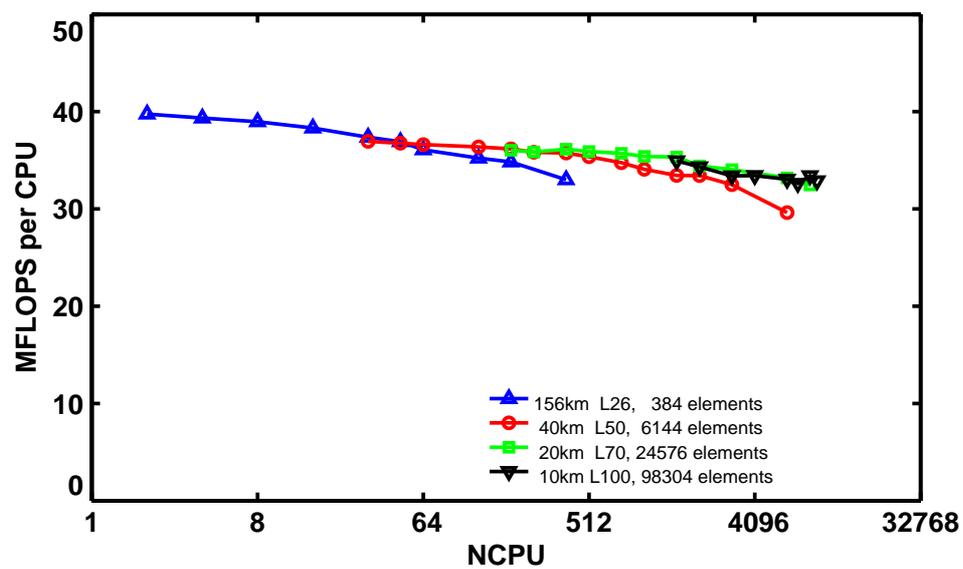


Figure 2: Performance of the HOMME dynamical core (dry dynamics) benchmark runs on ASCI Red. Each curve shows the performance for a fixed resolution as the number of processors is increased.

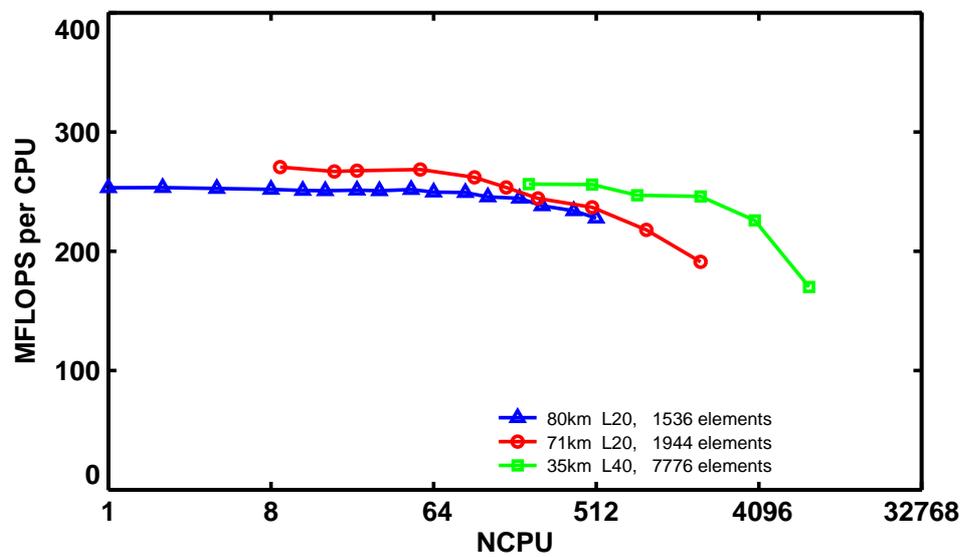


Figure 3: Performance of the HOMME dynamical core (dry and moist dynamics) benchmark runs on BlueGene/L. Each curve shows the performance for a fixed resolution as the number of processors is increased.

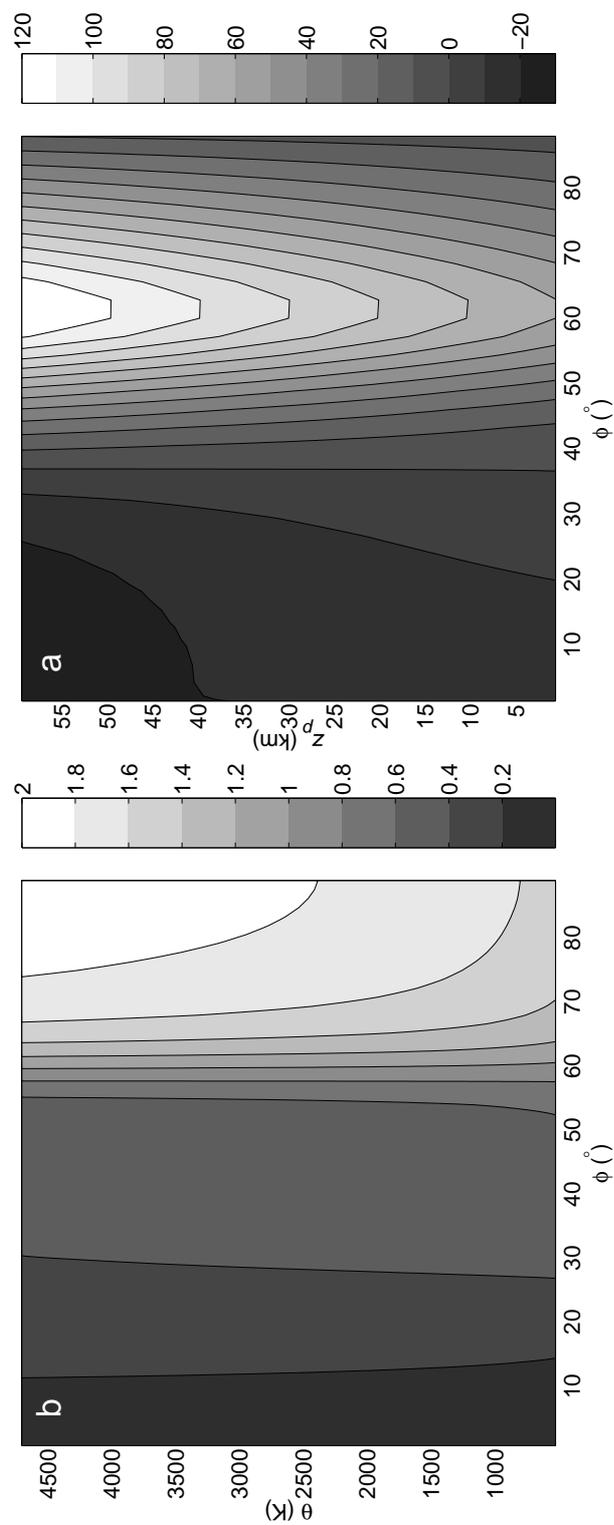


Figure 4: Initial state profiles vs latitude φ ($^\circ$, abscissa). (a) Zonal-wind isotachs (m s^{-1}) vs log-pressure height z_p (km, ordinate). Contour interval is 10 m s^{-1} and $u(\varphi = 37^\circ, z_p) = 0$. (b) Isopleths of scaled potential vorticity vs potential temperature θ (K, ordinate). Contour interval is 0.2 and $\Pi(\varphi = 0, \theta) = 0$.

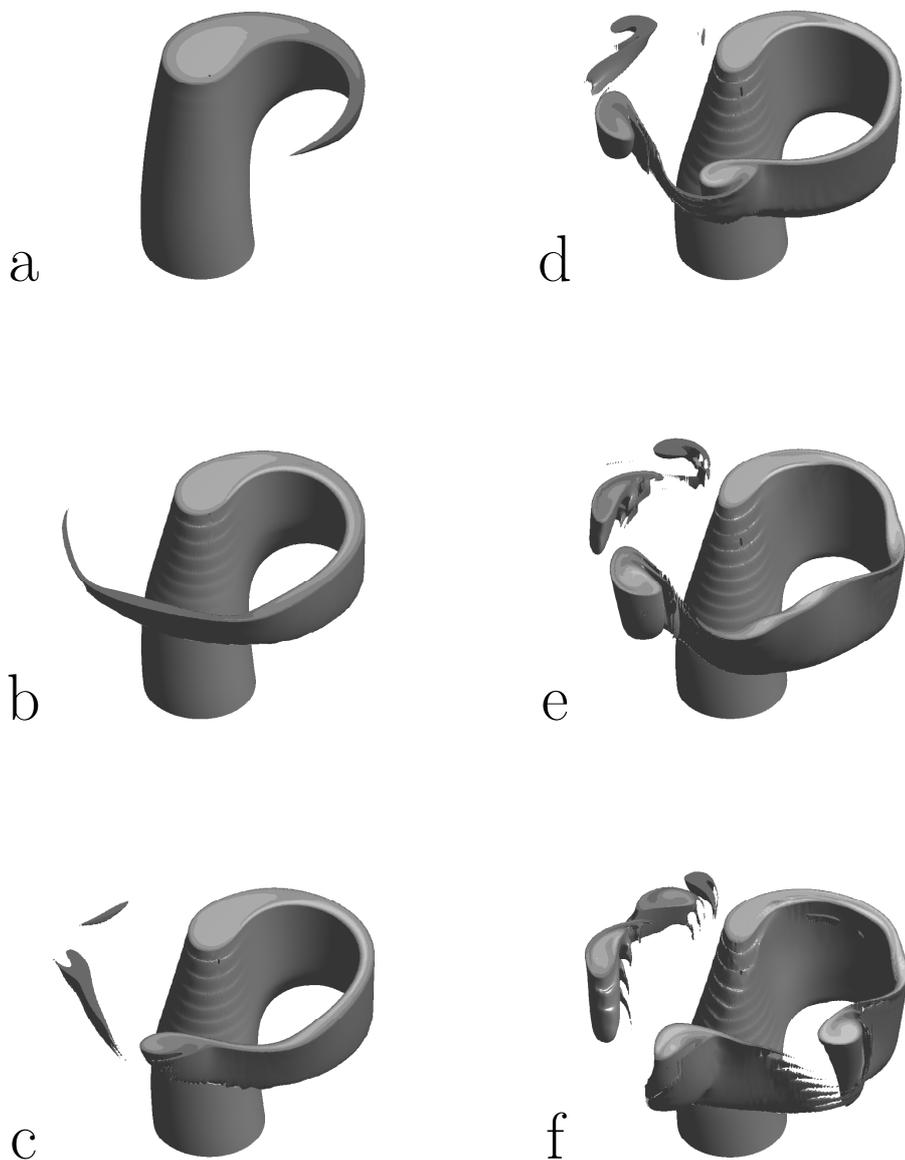


Figure 5: Isosurfaces of scaled potential vorticity $\Pi=1$ for the 36km run. Horizontal coordinates have been stereographically projected (8), and there are 22 vertical levels $\theta_l = 10^{2/7} T_0 \exp(2.5 \frac{l+3/2}{49})$. View is towards stereographic origin and $\theta_1 = 524\text{K}$, from $\lambda = 110^\circ\text{W}$, elevation 40° . Top plane shows contours of $\Pi > 1$ on $\theta_{22} = 1530\text{K}$. Times (a) $t = 13\text{d}$, (b-f) $t \in \{16, \dots, 20\}\text{d}$.

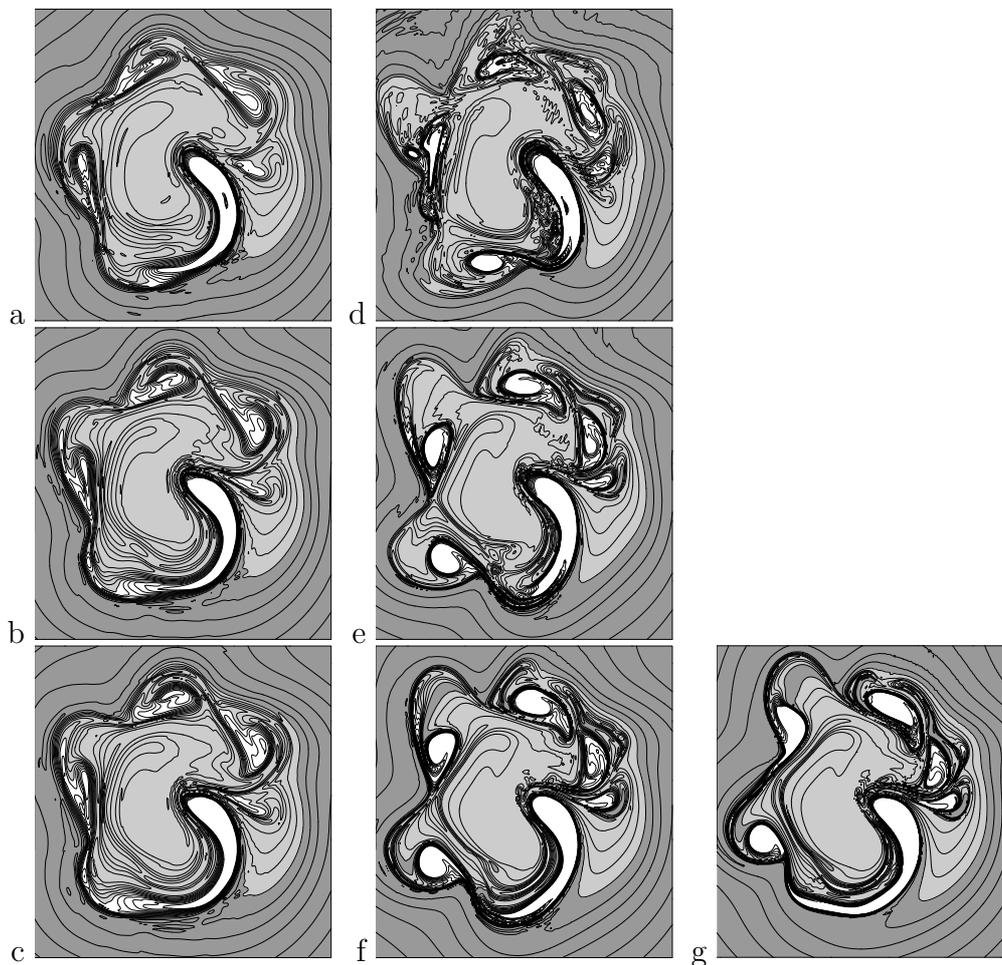


Figure 6: Stereographic projection in the the $\theta = 1500\text{K}$ surface of $\Pi(\lambda, \varphi)$ contours from -0.3 to 1.3 by 0.1 , at $t = 20\text{d}$. Values below 0.2 are dark gray and above 0.8 are white. Resolutions (a) $156\text{km}/50\text{L}$, (b) $156\text{km}/100\text{L}$, (c) $156\text{km}/200\text{L}$; (d-f) as in (a-c) but for 70km ; (g) as in (f) but for 36km .

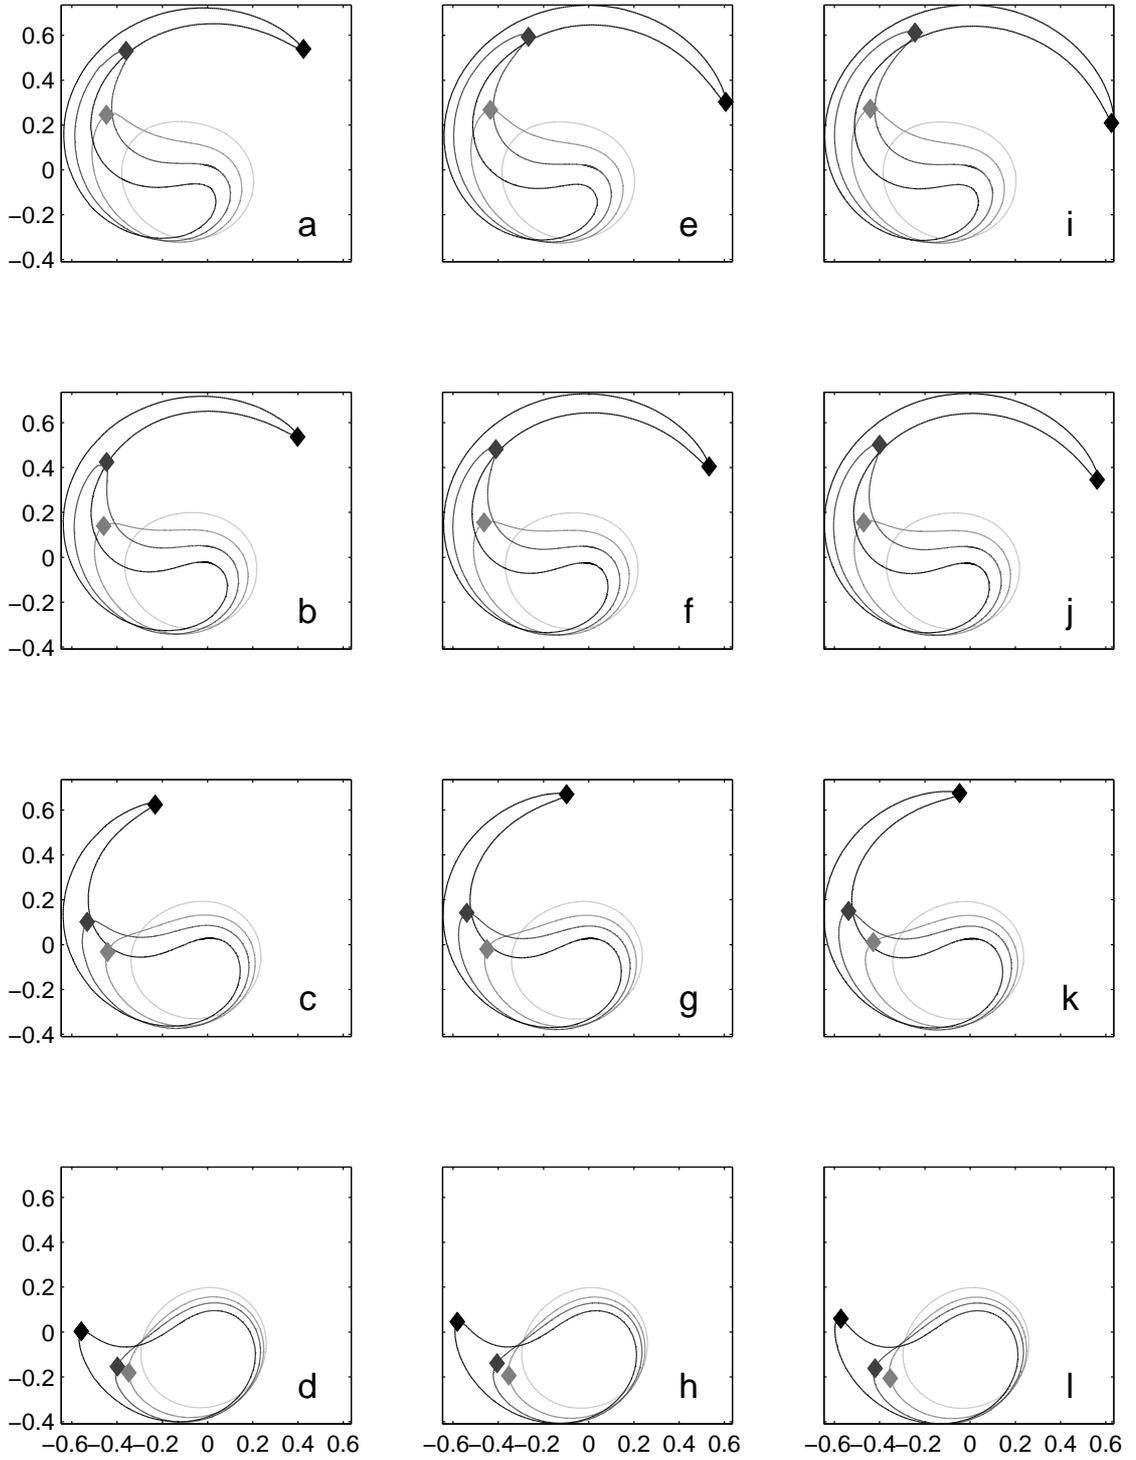


Figure 7: $\Pi=1$ contours, for $t \in \{5, 7, 9, 11\}d$ (light gray to black). Diamonds indicate $\mathbf{x}_{\max}(\theta, t)$ (Eq. 9). (a-d) Descending levels $\theta = 3000, 2500, 2000, 1500K$ for 156km resolution. (e-h) As (a-d) but for 70km. (i-l) As (a-d) but for 36km.

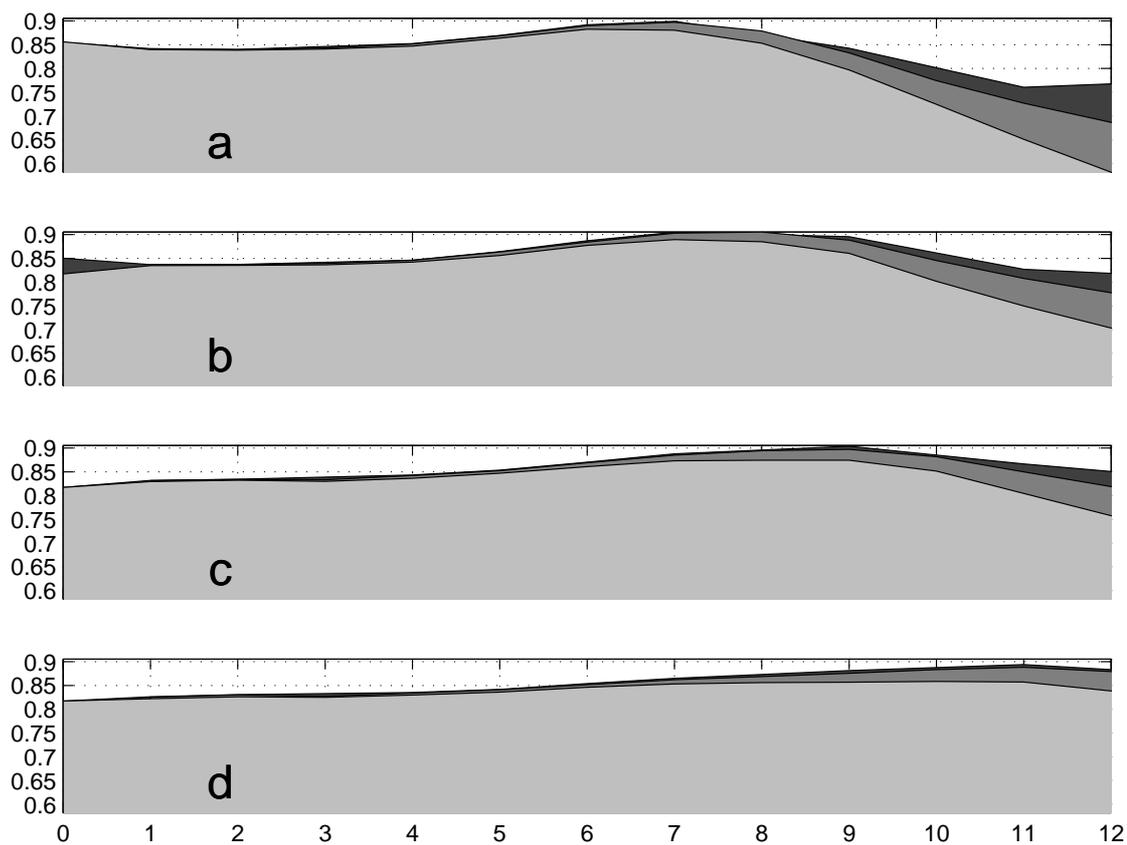


Figure 8: Horizontal vortex-area $A(\theta, t)$ (10) vs t (d, abscissa) for (a) $\theta = 3000$, (b) 2500, (c) 2000, (d) 1500K. Resolutions are 156km/50L (light gray), 70km/100L (gray) and 36km/200L (dark gray).