

## 1 Introduction

Climate simulation is a grand challenge problem requiring multiple, century-long integrations of the equations governing the Earth's atmosphere. Consequently, grid resolutions in atmospheric climate models are coarser than in numerical weather models, where accurate predictions are limited to about ten days. Current climate models are typically run with a 300 km equatorial grid spacing (T42 spectral truncation), whereas global weather model resolutions are approaching 10 km. Recently, it has been recognized that localized flow structures may play an important role in obtaining the correct climate signal. Higher-resolution climate simulations may be required in the near future at NCAR and other climate modeling centers. To date, achieving high integration rates for climate models on highly parallel clusters of microprocessors has been problematic. Operating at low resolutions effectively limits the scalability of climate models on such computer architectures. This limitation is further compounded by the widespread choice of dynamical algorithms based on the spectral transform method. The spectral transform method is mathematically attractive because the global spherical harmonic basis functions provide an isotropic representation on the sphere. In addition, it is trivial to implement semi-implicit time-stepping schemes, because the spherical harmonics are eigenfunctions of the Laplacian on the sphere and the resulting Helmholtz problem is embarrassingly parallel in spectral space. Nevertheless, although spectral models have exhibited good performance on modestly parallel vector architectures, they require nonlocal operations, such as array transpositions, which inhibit scaling to thousands of processors. A promising alternative is the spectral element method (SEM). Spectral elements maintain the accuracy and exponential convergence rate exhibited by the spectral transform method. They also offer several computational advantages on massively parallel systems composed of cache-based microprocessors. The computations are naturally cache efficient and derivatives can be computed using *only* nearest neighbor communication. Loft et al. [39], for example, implemented and studied the performance of multilayer shallow water and 3-D primitive equation models based on scalable spectral element numerics. The shallow water test suite of Williamson et al. [72] and the Held-Suarez [32] idealized climate tests are used to evaluate these for climate modeling. These authors were awarded second place in the 2001 IEEE Gordon Bell award for demonstrating the scalability of a dynamical core for an atmospheric general circulation model. An integration rate of over 100 years per day was achieved with sustained performance of 370 gigaflops for a T170 equivalent resolution spectral element model [39]. This represents a major advance in geophysical fluid dynamics simulations.

The geometric flexibility of spectral element methods are ideally suited to tackling several outstanding problems related to multiscale phenomena in geophysical fluid dynamics. This flexibility has not yet been fully exploited in the 3-D spectral element dynamical core developed at NCAR. The Intergovernmental Panel on Climate Change (IPCC) has noted several outstanding issues related to resolution, physical parameterizations, and the convergence of climate models. McAvaney et al. [47], Chapter 8.9, state: "The importance of numerical aspects of climate models continues to be well recognized and new numerical techniques are beginning to be tested for use in climate simulation. However, there has been very little systematic investigation of the impact of improved numerics for climate simulation and many important questions remain unanswered. The degree of interaction between horizontal and vertical resolution in climate models and the interaction of physical parameterizations at differing resolutions has made it extremely difficult to make general statements about the convergence of model solutions and hence the optimum resolution that should be used. An important question regarding the adequacy of resolution is deciding whether the information produced at finer scales at higher resolution feeds back on the larger scales or do the finer scales simply add to local effects (Williamson, [73]). Insufficient systematic work has been done with coupled models to answer this question. As well as improving model accuracy in advection, improved horizontal resolution can also improve the representation of the lower boundary of a model (the mountains) and the land-sea mask; this may improve the regional climate of a model but little systematic work has been carried out to assess this aspect."

The need for multiscale physical parameterizations was explicitly identified in Chapter 8.9.1 of the IPCC report: “A series of experiments that explores convergence characteristics has been conducted with the NCAR Community Climate Model (CCM) by Williamson [73]. In these experiments the grid and scale of the physical parameterizations was held fixed while the horizontal resolution of the dynamical core was increased. As the dynamical resolution was increased, but the parameterization resolution held fixed, the local Hadley circulation in the dual-resolution model simulations converged to a state close to that produced by a standard model at the fixed parameterization resolution. The mid-latitude transient aspects did not converge with increasing resolution when the scale of the physics was held fixed. Williamson (1999) concludes that the physical parameterizations used in climate models should explicitly take into account the scale of the grid on which it is applied. That does not seem to be common in parameterizations for global models today.”

We propose to study the effects of enhanced local resolution using simplified experiments such as the aqua-planet of Williamson. Working with collaborators at NCAR and universities we will explore how physical parameterizations in climate models can be made scale dependent and possibly adaptive. An adaptive mesh refinement capability will allow us to explicitly resolve small-scale flow features or combine fine- and coarse-scale parameterizations in simulations.

## 2 Adaptive Mesh Refinement for Climate and Weather Simulation

Limited-area atmospheric models with boundaries have been employed to provide high-resolution weather forecasts at continental, regional, and higher resolutions. The boundary conditions for these models are typically provided by a large-scale global forecast model. There are well-known problems associated with this approach related to the well-posedness of the governing equations and numerical approximations at boundaries. The two approaches adopted in atmospheric models are one-way and two-way nesting. In a one-way nesting strategy, the boundary conditions for a fine-grid model are supplied by a coarse-grid driving model; however, information is not passed from the fine-grid to the coarse-grid model. Two-way interactive nesting is analogous to adaptive mesh refinement (AMR), with a different time step in the refined region. Information leaving the fine-grid domain is passed to the coarse-grid model at the boundary. The hydrostatic primitive equations commonly used for large-scale models are ill-posed for the initial boundary value problem (IBVP) according to the seminal work of Olinger and Sundstrom [50]. There also exists the potential for dynamical imbalances caused by driving a nonhydrostatic limited-area model with boundary conditions supplied by a hydrostatic global model. For one-way nesting, well-posed inflow/outflow boundary conditions can be specified (e.g., Browning and Kreiss [10]). To minimize the effects of spurious gravity wave propagation back into the problem domain, Rayleigh damping ‘sponge’ layers have been applied in finite difference models. Alternative schemes include the Davies [19] flow relaxation or radiation boundary conditions.

To avoid the boundary problem altogether, variable resolution meshes have been implemented (e.g., Cote et al. [17], Fox-Rabinowitz et al. [28]). In these models a uniform mesh is placed over a region of interest on the sphere and the grid spacing is gradually increased away from this region. These models are based on low-order finite-difference or finite-volume numerics. The Meteo-France Arpege model [48] combines a high-order pseudo-spectral method and global spherical harmonic basis functions with a coordinate transformation on the sphere to achieve variable resolution. However, the mesh remains fixed in space and time for each of these models. For both the variable resolution and nested grid approaches, there remain difficulties related to implementing subgrid-scale parameterizations when disparate space and time scales are present in the models. Convective parameterizations are particularly difficult in this context (see Weisman et al. [71]). Nevertheless, grid refinement strategies are being pursued because of the potential for computational savings and much higher resolution in regions where small-scale flow features are present. The challenge is to detect when these small-scale structures are present and relevant, then to provide a feedback mechanism to the large-scale flow that avoids introducing nonphysical numerical artifacts. Unresolved subgrid-scale processes are still present and must be properly handled on a variable resolution adaptive mesh that is changing as a function of the model dynamics.

Both variable resolution and limited-area models are currently being evaluated for regional climate modeling as part of the international AMIP program [1]. Prof. M. Fox-Rabinowitz at the University of Maryland, along with Environment Canada and Prof. R. Laprise at the University of Quebec at Montreal, is comparing the variable resolution approach with a nonhydrostatic compressible limited-area regional climate model (RCM) [15]. These models employ low-order space discretization schemes. Prof. F. Baer at the University of Maryland and Dr. J. Tribbia at NCAR are investigating static grid refinement with high-order spectral element methods on a quasi-uniform cubed-sphere grid. In particular, they have applied conforming spectral elements to greatly enhance resolution near the poles in simulations of polar vortices. The University of Michigan in collaboration with NASA is investigating AMR for climate modeling using the semi-Lagrangian finite-volume dynamical core of Lin and Rood. The flux-form of the governing equations employed by Lin-Rood has the potential to allow for AMR in combination with variable time-stepping schemes. This approach is similar to the Berger-Colella algorithms for AMR applied to compressible hydrodynamics. The Lin-Rood finite-volume formulation is low order in space and implements an explicit time stepping scheme. An adaptive implementation of the Lin-Rood dynamical core would also be limited by the ill-posed nature of the hydrostatic primitive equations. There is a need for adaptive high-order methods in space and time combined with more efficient implicit time-stepping schemes. The governing equations should also be well-posed for adaptive schemes to avoid the generation of spurious gravity waves at interelement boundaries.

We propose a three-year research program with two parallel streams: (1) extending the geometrically nonconforming spectral element formulations and algorithms of Fischer, Tufo, Kruse to the 3-D hydrostatic primitive equations and (2) investigating the use of discontinuous Galerkin (DG) methods in adaptive methods for the well-posed primitive equations of Teman and Tribbia (2001). We will implement AMR within the context of a geometrically nonconforming spectral element code. The semi-implicit time stepping scheme of Thomas and Loft [65] will be retained, with the time step determined by the advective CFL at the highest resolution in the mesh hierarchy. Prof. J. Hesthaven and collaborators at Brown University have made significant progress in the development of discontinuous Galerkin (DG) methods, and these appear well suited to adaptive implementations of geophysical fluid flows. This group of researchers have applied DG to the 2-D shallow-water equations, which admit the same gravity and Rossby wave solutions as the full 3-D primitive equations. However, extension to a 3-D adaptive scheme could lead to potential problems related to the ill-posed nature of the primitive equations for the IBVP. Prof. J. Levin and the ocean modeling group at Rutgers are also investigating the use of DG methods in the context of 3-D hydrostatic flows. These methods are similar to spectral element methods because they are based on high-order orthogonal polynomial basis functions and Gaussian quadrature within each element. However, DG methods do not require  $C^0$  continuity at interelement boundaries and instead rely on numerical fluxes. We propose to investigate the use of discontinuous Galerkin methods for constructing a 3-D climate dynamical core with AMR including adaptive time stepping. Given that the hydrostatic primitive equations for atmospheric flow are ill-posed for the IBVP, the well-posed primitive equations of Teman and Tribbia [62] will instead be employed in our studies.

### 3 Primitive versus Anelastic or Compressible Equations

The primitive equations for the IBVP on a domain with open boundaries are ill-posed according to Olinger and Sundstrom [50]. Teman and Tribbia [62] have proposed a regularization procedure for rendering the primitive equations well-posed. This modification consists of a mild vertical viscosity added to the hydrostatic equation. We propose to evaluate the Teman and Tribbia [62] equations in the context of AMR, using a cubed-sphere grid, for climate and weather modeling applications. Initially our schemes will be adaptive in space but not in time. In particular, we propose first to apply horizontal AMR using the high-order nonconforming spectral element techniques of Fischer and Tufo to investigate whether resolving/tracking eddies and small-scale structures can lead to the same or improved accuracy compared with a uniform resolution model. Another important criterion is to achieve this accuracy at reduced computational cost. Therefore, we will also use the grid refinement criterion for vortex tracking developed by Jeong and Hussain [36] and employed by Fischer and Tufo in their isosurface rendering and immersive visualization system [67, 68]. As a first step in this direction we will track dry vortex structures, which provide a good initial test before

moving to vortex tracking with moisture (both transport and subgrid-scale parameterized convection).

We will also use variations of the Held-Suarez idealized climate test to evaluate our adaptive schemes. The resulting zonal jets have a localized structure that may be amenable to local grid refinement in the mid-latitude regions. The correct position and shape of the zonally averaged wind and temperature profiles depends directly on the resolution. Rather than introducing a turbulence closure scheme such as Smagorinsky [59], we will employ a spectral element filtering procedure such as that proposed by Boyd [8] and Fischer and Mullen [25]. The advantage of such an approach is that the filter automatically adapts the dissipation to the scales of motion present in the model. These filters are equally applicable in the case of discontinuous Galerkin formulations. Semi-implicit time stepping schemes both damp and retard fast moving gravity waves. The semi-implicit time-stepping scheme of Thomas and Loft [65] for the 3-D primitive equations will be extended to the adaptive nonconforming spectral element algorithms of Fischer and Kruse [26, 38]. This scheme relies on a conjugate gradient iterative Krylov solver preconditioned using overlapping Schwarz techniques. It is more difficult to develop implicit time-stepping schemes for the flux-form of the governing equations because of the need to linearize and solve the equations responsible for fast gravity waves. Another component in the development of an adaptive discontinuous Galerkin formulation, therefore, is the implementation of an efficient implicit solver to avoid excessive time-step restrictions on the finest mesh.

A major contribution resulting from this proposal will be to provide feedback to the geophysical modeling community on whether the Teman and Tribbia [62] well-posed primitive equations are appropriate for multiscale modeling. The ill-posed primitive equations would result in the generation of spurious gravity waves. For spectral elements, the element size versus the local Rossby radius of deformation (length scale) would govern the amplitude, phase speed, and location of these spurious waves. We propose to investigate whether the Teman and Tribbia [62] modification will eliminate these spurious waves in the case of both static and adaptive grid refinement. There may be potential unresolved difficulties with using these equations, and it may be necessary to consider the nonhydrostatic, anelastic, or fully compressible equations. Fischer and Tufo have implemented a spectral element code for the Boussinesq equations applied to thermal convection studies in spherical geometry and are currently working on an anelastic spectral element formulation. This model would serve as starting point for comparing nonconforming spectral element and discontinuous Galerkin formulations of the 3-D governing equations.

## 4 Spectral Element Formulation

The shallow water equations have been used for many years by the atmospheric modeling community to test promising numerical methods. These equations contain the same horizontal wave propagation mechanisms found in more complete models. In particular, they admit the Rossby and gravity wave solutions found in 3-D primitive equations models. The governing equations for the inviscid flow of a thin layer of fluid in 2-D are the horizontal momentum and continuity equations for the velocity  $\mathbf{v}$  and geopotential height  $\phi$ . In curvilinear coordinates, the shallow water equations can be written as follows (see Sadourny [58]):

$$\begin{aligned}\frac{\partial u^i}{\partial t} &= -g^{ij} \left[ \epsilon_{jk} u^k g (f + \zeta) + \frac{\partial}{\partial x^j} \left( \frac{1}{2} u_k u^k \right) + \frac{\partial \phi}{\partial x^j} \right] \\ \frac{\partial \phi'}{\partial t} &= -u^j \frac{\partial \phi}{\partial x^j} - \frac{\phi}{g} \frac{\partial}{\partial x^j} (g u^j),\end{aligned}$$

where  $\phi = \phi' + \phi_0$ ,  $f$  is the Coriolis force,  $\zeta$  is the relative vorticity, and  $\epsilon_{j k}$  is a permutation matrix. Covariant and contravariant vectors are related through the metric tensor  $g_{ij}$ ,  $u^i = g^{ij} u_j$ ,  $g_{ij}^{-1} = g^{ij}$  and  $g = \{ \det(g_{ij}) \}^{1/2}$ . Divergence and vorticity are given by

$$g \nabla \cdot \mathbf{v} = \frac{\partial}{\partial x^j} (g u^j), \quad g \zeta = \epsilon_{ij} \frac{\partial u_j}{\partial x^i}.$$

The sphere is tiled with rectangular elements by subdividing the six faces of the cube, which circumscribes the sphere, and then using a gnomonic projection to map these elements onto the surface of the sphere.

Rancic et al. [55] showed that an equal angular projection results in a more uniformly spaced grid. For equal angular coordinates  $(x_1, x_2)$ ,  $-\pi/4 \leq x_1, x_2 \leq \pi/4$ , and the metric tensor for all six faces of the cube is

$$g_{ij} = \frac{1}{r^4 \cos^2 x_1 \cos^2 x_2} \begin{bmatrix} 1 + \tan^2 x_1 & -\tan x_1 \tan x_2 \\ -\tan x_1 \tan x_2 & 1 + \tan^2 x_2 \end{bmatrix},$$

where  $r = \{1 + \tan^2 x_1 + \tan^2 x_2\}^{1/2}$  and  $g = 1/r^3 \cos^2 x_1 \cos^2 x_2$ . A vector  $\mathbf{v} = (v_1, v_2)$  in spherical coordinates is entirely defined by its covariant and contravariant components  $u_i$  and  $u^i$  on the cube. For the vector  $(u_1, u_2)$  on the cube, we define the mapping

$$A = \begin{bmatrix} \cos \theta \partial \lambda / \partial x_1 & \cos \theta \partial \lambda / \partial x_2 \\ \partial \theta / \partial x_1 & \partial \theta / \partial x_2 \end{bmatrix}, \quad A^T \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} = \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}, \quad A \begin{bmatrix} u^1 \\ u^2 \end{bmatrix} = \begin{bmatrix} v_1 \\ v_2 \end{bmatrix},$$

where  $A^T A = g_{ij}$ . Requiring that the velocities in spherical coordinates match along the cube edge shared by face  $i$  and face  $j$ , we obtain the relationship  $A_i \mathbf{u}_i = A_j \mathbf{u}_j$ ,  $\mathbf{u}_i = A_i^{-1} A_j \mathbf{u}_j$ . Our semi-implicit scheme applied to the shallow water equations combines an explicit leapfrog scheme for the advection terms with a Crank-Nicholson scheme for the gradient and divergence terms. For the shallow water equations the scheme can be written in terms of the differences  $\delta u^i = u^{i(n+1)} - u^{i(n-1)}$  and  $\delta \phi = \phi^{n+1} - \phi^{n-1}$ :

$$\delta u^i + \Delta t g^{ij} \frac{\partial}{\partial x^j} (\delta \phi) = 2\Delta t \left[ -g^{ij} \frac{\partial}{\partial x^j} (\phi)^{n-1} + f_u^{i(n)} \right] \quad (1)$$

$$\delta \phi + \Delta t \frac{\phi_0}{g} \frac{\partial}{\partial x^j} (g \delta u^j) = 2\Delta t \left[ -\frac{\phi_0}{g} \frac{\partial}{\partial x^j} (g u^j)^{n-1} + f_\phi^n \right], \quad (2)$$

where the tendencies  $f_u$  and  $f_\phi$  contain nonlinear terms.

In the spectral element discretization, the computational domain  $\Omega$  is partitioned into  $K$  elements  $\Omega_k$  in which the dependent and independent variables are approximated by  $N$ th order tensor-product polynomial expansions. The velocity is expanded in terms of the  $N$ th degree Lagrangian interpolants  $h_i$  defined in Rønquist [56], as follow:

$$\mathbf{u}_h^k(r_1, r_2) = \sum_{i=0}^N \sum_{j=0}^N \mathbf{u}_{ij} h_i(r_1) h_j(r_2),$$

and the geopotential is expanded using the  $(N-2)$ th degree interpolants  $\tilde{h}_i$

$$\phi_h^k(r_1, r_2) = \sum_{i=1}^{N-1} \sum_{j=1}^{N-1} \phi_{ij} \tilde{h}_i(r_1) \tilde{h}_j(r_2),$$

A weak variational problem is obtained by integrating the equations with respect to test functions and directly evaluating inner products using Gaussian quadrature. Two integration rules are defined for a staggered mesh by taking the tensor-product of Gauss and Gauss-Lobatto quadrature rules on each element:

$$(f, g)_{GL} = \sum_{k=1}^K \sum_{i=0}^N \sum_{j=0}^N f^k(\xi_i, \xi_j) g^k(\xi_i, \xi_j) \rho_i \rho_j$$

$$(f, g)_G = \sum_{k=1}^K \sum_{i=1}^{N-1} \sum_{j=1}^{N-1} f^k(\zeta_i, \zeta_j) g^k(\zeta_i, \zeta_j) \sigma_i \sigma_j,$$

where  $(\xi_i, \rho_i)$ ,  $i = 0, \dots, N$  are the Gauss-Lobatto nodes and weights and  $(\zeta_i, \sigma_i)$ ,  $i = 1, \dots, N-1$  are the Gauss nodes and weights on  $\Lambda = [-1, 1]$ . Physical coordinates are mapped according to  $\mathbf{x} \in \Omega_k \Rightarrow \mathbf{r} \in \Lambda \times \Lambda$ .  $C^0$  continuity of the velocity is enforced at interelement boundaries which share Gauss-Lobatto points and direct stiffness summation is applied to assemble the global matrices. The discrete divergence  $\tilde{\mathbf{D}} = (\tilde{D}_1, \tilde{D}_2)$  and discrete gradient (weak form)  $\tilde{\mathbf{D}}^T$  operators are rectangular matrices.  $\mathbf{B} = (B_1, B_2)$  and  $\tilde{B}$  are the

diagonal velocity and geopotential mass matrices, with  $\tilde{D}_j = \tilde{B} D_j$ .  $D_j$  is the strong form of the derivative operator. The assembled discrete shallow water equations are then

$$B_i \delta u^i - g^{ij} \Delta t \tilde{D}_j^T \delta \phi = R_u^i \quad (3)$$

$$\tilde{B} \delta \phi + \Delta t \frac{\phi_0}{g} \tilde{D}_i g \delta u^i = R_\phi. \quad (4)$$

The pressure is defined on the interior of an element and is ‘communicated’ between elements through the divergence in the continuity equation. An averaging procedure is required at element boundaries to enforce continuity, where velocity mass matrix elements in equation (3) are summed. This averaging procedure is related to the fact that basis functions interpolating boundary nodes are not local to a specific element.

A Helmholtz problem for the geopotential perturbation is obtained by solving for the velocity difference  $\delta u^i$

$$\delta u^i = B_i^{-1} \left( R_u^i + \Delta t g^{ij} \tilde{D}_j^T \delta \phi \right) \quad (5)$$

and then applying back-substitution to obtain

$$g \tilde{B} \delta \phi + \Delta t^2 \phi_0 \tilde{D}_i g B_i^{-1} g^{ij} \tilde{D}_j^T \delta \phi = R'_\phi, \quad (6)$$

where

$$R'_\phi \equiv g R_\phi - \Delta t \phi_0 \tilde{D}_i g B_i^{-1} R_u^i$$

Once the change in the geopotential  $\delta \phi$  is computed, the velocity difference  $\delta u^i$  is obtained from (5). Given that the metric tensor  $g_{ij}$  is symmetric,  $g$ ,  $\tilde{B}$  and  $B_i$  are diagonal, it can be easily shown that the Helmholtz operator

$$H = g \tilde{B} + \Delta t^2 \phi_0 \tilde{D}_i g B_i^{-1} g^{ij} \tilde{D}_j^T$$

is symmetric positive definite and a preconditioned conjugate gradient solver is applied. An effective preconditioner can be constructed by using local element direct solvers for the Helmholtz problem with zero Neumann pressure gradient boundary conditions. The inverse of these local matrices is computed once and applied as a matrix-vector product during time-stepping. This block-Jacobi preconditioner is strictly local to an element and requires no communication. Application of the Helmholtz operator as a matrix-vector multiply requires element edge communication for each CG iteration. A more detailed description of the solver and preconditioner is given in Thomas and Loft [63].

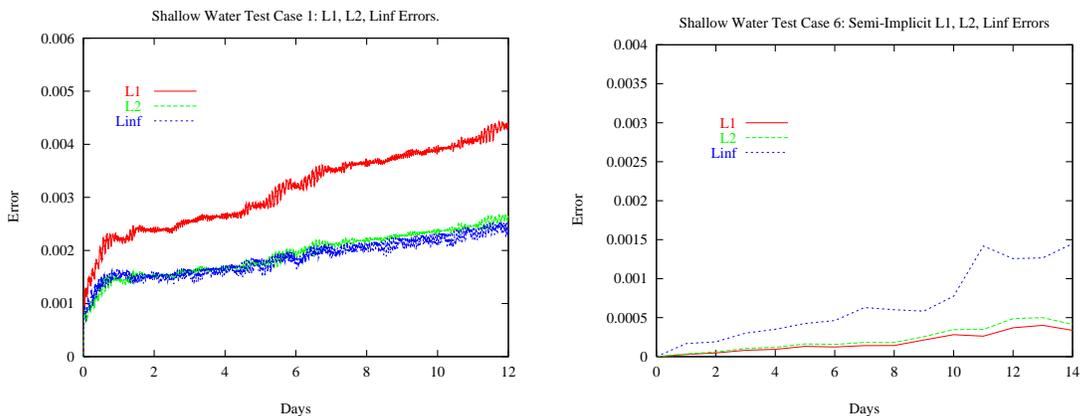


Figure 1: SWE Test Case 1. Cosine Bell. Explicit model.  $l_1$ ,  $l_2$  and  $l_\infty$  errors for 12 day integration. Grid size:  $96 \times 16 \times 16$  pressure points (left). SWE Test Case 6. Rossby-Haurwitz Wave. Semi-implicit model.  $l_1$ ,  $l_2$  and  $l_\infty$  errors for 14 day integration (right).

## 5 Numerical Results

A standard test set for evaluating numerical approximations to the shallow water equations (SWE) in spherical geometry has been proposed by Williamson et al. [72]. Test case 1 is the advection of a cosine bell with compact support. This test case is designed to evaluate the advective component of the numerical scheme in isolation and is the only test that does not involve the full equation set. Since the first derivative of the initial field is discontinuous, the cosine bell is a challenge for spectral methods because of the resulting Gibbs phenomena or so-called spectral ringing. Several values for the angle  $\alpha$  between the axis of solid body rotation and the spherical coordinate pole are specified, and results are reported here for  $\alpha = \pi/4$  which advects the cosine bell over the cubed-sphere corner points. The length of the integration is 12 days, corresponding to one complete rotation around the sphere. Figure 1 (left) is a plot of the error metrics for a grid with  $96 \times 16 \times 16$  pressure points, where each cube face contains 16 spectral elements. An explicit time step of length  $\Delta t = 30$  s was used for this run. A Boyd-Vandeven filter was applied every 20 min in order to suppress the growth of aliasing errors. Figure 1 (left) should be compared with Figure 5 of Taylor et al. [61], where it can be seen that the  $l_1$  error is reduced by a factor of two. The oscillations visible in the error curves are similar to those observed in Jakob-Chien et al. [35] and are characteristic of spectral methods. They are due to sampling errors as the cosine bell moves through different parts of the grid.

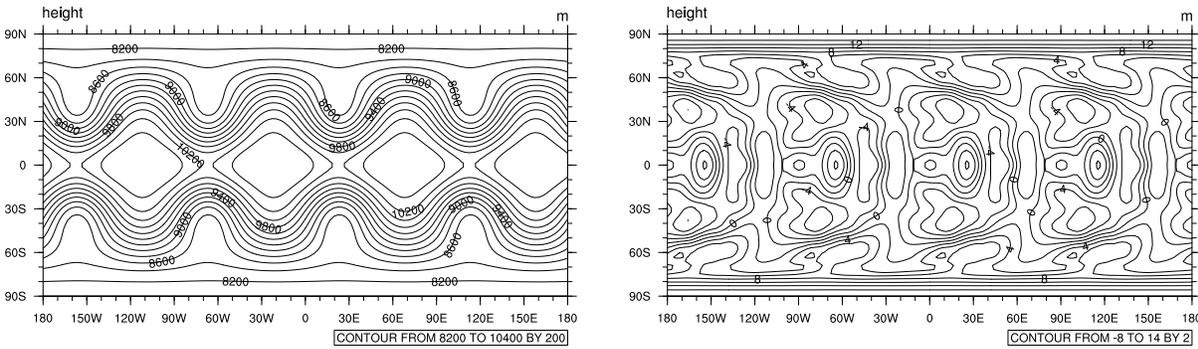


Figure 2: SWE Test Case 6. Rossby-Haurwitz Wave. Semi-implicit model. Geopotential height field after 14 day integration. Grid size:  $1734 \times 8 \times 8$  pressure points. 200 m contours. SWE Test Case 6. Rossby-Haurwitz Wave. Semi-implicit model (left). Geopotential height error field (vs. T213 reference) after 14 day integration. Grid size:  $1734 \times 8 \times 8$  pressure points. 2 m contours (right).

The initial condition for test case 6 is a  $R = 4$  wavenumber Rossby-Haurwitz wave. These waves are an ideal test because they represent exact analytic solutions to the nonlinear nondivergent barotropic vorticity equation. Rossby-Haurwitz waves are not closed-form solutions of the barotropic shallow water equations. However, a high-resolution integration of a spectral model can be used to generate a reference solution. A grid resolution of  $1734 \times 8 \times 8$  pressure points was chosen in order to compare against the simulation results presented in Figure 11 of Taylor et al. [61]. A stable time step for this simulation was  $\Delta t = 120$  s which is eight times the explicit step of  $\Delta t = 15$  s. Test case 6 is particularly difficult for a semi-implicit solver since the mean geopotential height is set at 8000 m and the maximum gravity wave phase velocity approaches  $300 \text{ ms}^{-1}$ . In this case, the CG solver requires 4 to 5 iterations to converge. The advective Courant number is also limited by wind speeds approaching  $100 \text{ ms}^{-1}$ . The  $l_1$ ,  $l_2$ , and  $l_\infty$  errors are plotted in Figure 1 (right), where it can be seen that all three error measures at 14 days are reduced by a factor of two when compared with the earlier results of Taylor et al [61]. These improvements can be attributed to the use of more velocity points in the staggered grid and exponential convergence. A contour plot of the Rossby-Haurwitz solution after 14 days of integration by the semi-implicit spectral element model is presented in Figure 2 (left). The error field comparing a T213 reference and spectral element solutions at 14 days is plotted in Figure 2 (right). The reference solution has been interpolated to a spherical latitude-longitude projection of the cubed-sphere spectral element grid using the spherical harmonic basis functions and T213 spectral coefficients, truncated to T106 and stored as 32-bit reals.

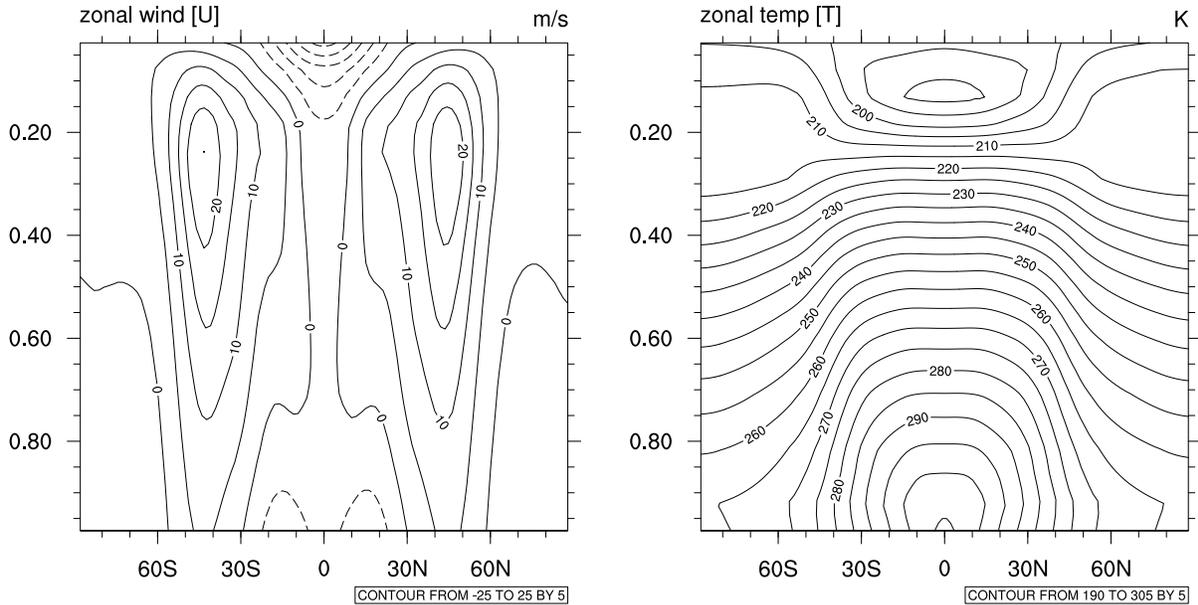


Figure 3: Held-Suarez idealized climate. 1200 day integration. 150 elements. Left, zonally averaged wind  $[U]$ . Right, zonally averaged temperature  $[T]$ .

The shallow water equations encapsulate many of the dynamical aspects of climate modeling. Any numerical method being considered as a basis for a climate dynamical core should perform well on the above test cases. However, the main purpose of these tests is to compare the accuracy of numerical methods and their ability to handle spherical geometry. Because the longest integration is 15 days, these tests do not assess the ability of a dynamical core to generate accurate long term climate statistics. The Held-Suarez [32] idealized climate forcing is designed to test the dry dynamical core of a GCM for longer integrations. It assumes an ideal gas atmosphere over a rotating sphere with no topography. The flow is not specified as hydrostatic, however, the hydrostatic primitive equations may be employed. The prescribed forcing consists of a simple Newtonian relaxation of the temperature field to a zonally symmetric state and Rayleigh damping of the lower-level wind field to approximate friction or drag caused by the atmospheric boundary layer near the surface. The initial state of the atmosphere is hydrostatic and isothermal  $T = 300$  K. The model is integrated for 1200 days. Zonally averaged wind and temperature fields are then reported. The Held-Suarez forcings take the following form,

$$\frac{\partial \mathbf{v}}{\partial t} = \dots - k_{\mathbf{v}}(\phi, \sigma)$$

$$\frac{\partial T}{\partial t} = \dots - k_T(\phi, \sigma) [T - T_{\text{eq}}(\phi, \sigma)]$$

where  $\phi$  is the latitude and  $\sigma = p/p_s$  is the vertical sigma coordinate level. The temperature is relaxed to the equilibrium temperature  $T_{\text{eq}}$  and relaxation rate  $k_T$ . The linear damping rate of the wind is given by  $k_{\mathbf{v}}$ . A simulation was performed using 150 spectral elements and 19 equally spaced sigma levels in the vertical direction. The results are plotted in figure 3. These plots closely match those given in Held-Suarez [32] and exhibit the characteristic formation of jets in the upper atmosphere.

Finally, we present performance of the spectral element atmospheric model (SEAM) code developed by Thomas and Loft [39]. Figure 4 shows the scalability of SEAM for the multi-layer shallow water equations and the explicit 3-D primitive equations for the Held-Suarez idealized climate test. An integration rate of over 100 years per day was achieved with sustained performance of 370 gigaflops for T170 equivalent resolution. This represents a major advance in geophysical fluid dynamics simulations.

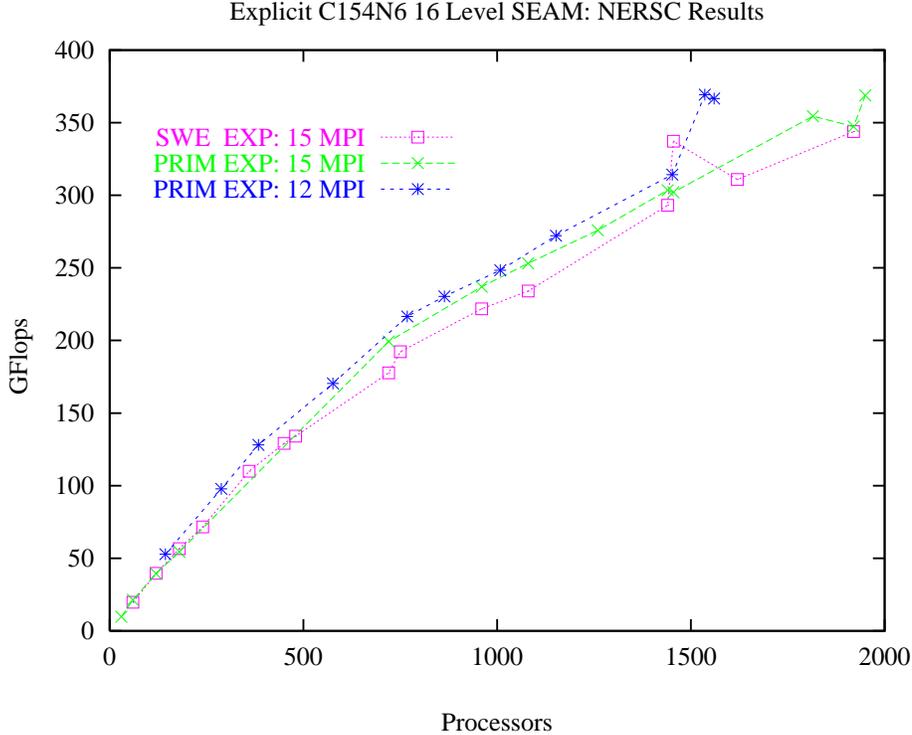


Figure 4: Total gigaflops for C154N6 16 level SEAM simulation on the IBM SP system at NERSC.

## 6 Nonconforming Spectral Elements

We follow closely the work of Kruse [38] and Fischer, Kruse and Loth [26]. Element-by-element operator evaluation is central to the efficiency of the SEM because it allows the use of tensor-product forms, which reduce the work and storage complexity from  $O(KN^{2d})$  to  $O(KN^{d+1})$  and  $O(KN^d)$ , respectively [51]. The extension to nonconforming spaces preserves this feature and essentially involves redefining the interface operators that impose the matching conditions across element interfaces. Here, we consider the development of interpolation-based interface conditions that leave the approximation spaces  $X^N$  and  $Y^N$  unchanged but allow for nonconforming meshes of the type illustrated in Figure 5a. On the nonconforming interface  $\Gamma$  we refer to the large element as the parent element, the two (or more) smaller elements as children, and the interface between them as a parent-child (PC) interface. We do not restrict the number of child elements per PC edge. However, we insist that the union of the closure of the child faces constitutes the closure of the parent face. While this restriction rules out configurations such as shown in Figure 5b, it allows us to preserve local (*element-to-element*) interactions. For similar reasons, we also exclude configurations in which

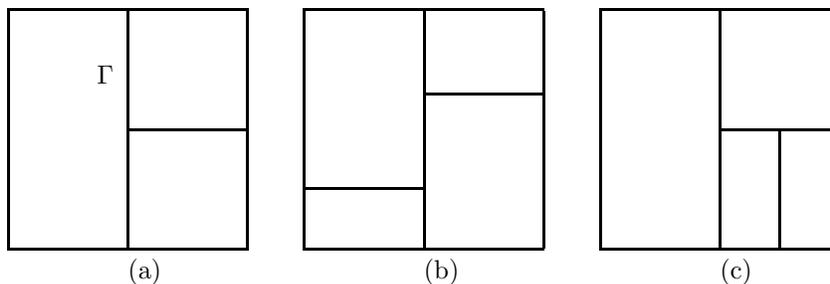


Figure 5: (a) Valid and (b,c) invalid nonconforming meshes in  $\mathbb{R}^2$ .

the endpoint of one PC interface connects to the interior of another, as shown in Figure 5c.

Much work has been done on nonconforming spectral element methods, starting with the early work of Mavriplis [46], Anagnostou et al. [2], and others [3, 14, 33]. Most of these have employed “mortar” elements which increase flexibility through the use of  $L^2$ -projection operators to enforce weak continuity at the nonconforming interface. In particular, the “vertex-free” mortar spaces of Ben Belgacem and Maday [3] alleviate the restriction of Figure 5c. The conforming-space/nonconforming-mesh approach used here was motivated by the results of Rønquist [57], who reported spurious eigenvalues in the convection operator for certain combinations of convection and nonconforming formulations. For brevity, *conforming* shall refer to conforming meshes (no hanging vertices) and *nonconforming* to conforming spaces with nonconforming meshes. We further assume that the polynomial degree  $N$  is the same in each spectral element.

To introduce the interface matching conditions, we begin by considering enforcement of continuity of a function  $u(\mathbf{x})$ ,  $\mathbf{x} \in \Omega \subset \mathbb{R}^2$  for the conforming case. For isoparametrically mapped elements, the geometry within each element is represented as,

$$\mathbf{x}^k(r, s)|_{\Omega^k} = \sum_{i=0}^N \sum_{j=0}^N \mathbf{x}_{ij}^k h_i^N(r) h_j^N(s). \quad (7)$$

Because the basis functions are Lagrangian, function continuity for  $u(\mathbf{x})$  is enforced by simply equating coincident nodal values, that is,

$$\mathbf{x}_{ij}^k = \hat{\mathbf{x}}_{ij}^k \implies u_{ij}^k = \hat{u}_{ij}^k. \quad (8)$$

If  $n$  is the number of distinct nodes on  $\Omega$ , then (8) represents  $K(N+1)^d - n$  constraints on the set of local nodal values  $\{u_{ij}^k\}$ .

For convenience we cast the constraint (8) in matrix form. Let  $\underline{u} \in \mathbb{R}^n$  denote the vector of nodal values associated with a global numbering of the distinct nodes in all of  $\Omega$ , as illustrated in Figure 6a. Let  $\underline{u}^k \in \mathbb{R}^{(N+1)^d}$  denote the vector of local basis coefficients associated with  $\Omega^k$ :

$$\underline{u}^k := (u_{00}^k, u_{10}^k, \dots, u_{NN}^k)^T, \quad k = 1, \dots, K,$$

as illustrated in Figure 6b, and let  $\underline{u}_L$  be the  $K(N+1)^d \times 1$  collection of these local vectors. If  $u$  is to be continuous, then there exists a Boolean connectivity matrix,  $Q$ , that maps the global form  $\underline{u}$  to its local counterpart  $\underline{u}_L$  such that (8) is satisfied. The operation

$$\underline{u}_L = Q\underline{u} \quad (9)$$

is referred to as a *scatter* from the global ( $\underline{u}$ ) to local ( $\underline{u}_L$ ) vector. For example, in Figure 6 the global value  $u_3$  is copied to local coefficients  $u_{2,0}^1$  and  $u_{0,0}^2$ . Note that for every global vector,  $\underline{u}$ , there is a corresponding local vector,  $\underline{u}_L$ , given by (9). The converse is not true because  $Q$  is not invertible. However, we will frequently employ the closely related *gather* operation

$$\underline{v} = Q^T \underline{u}_L \quad (10)$$

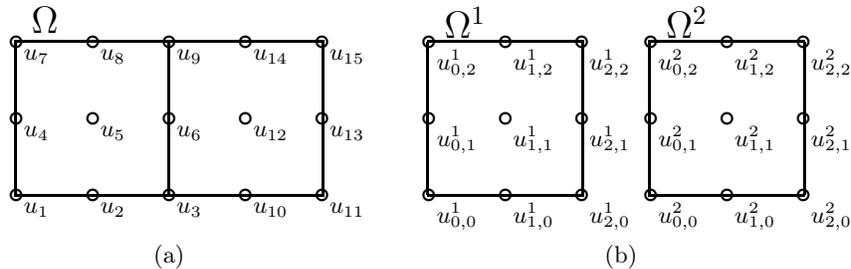


Figure 6: (a) Global and (b) local numbering for spectral element mesh.

and denote the output ( $\underline{v}$ ) with a different notation from the input ( $\underline{u}_L$ ). Whereas the action of  $Q$  is to *copy* entries of  $\underline{u}$  to  $\underline{u}_L$ , the action of  $Q^T$  is to *sum* entries from corresponding nodes. In practice, the matrix  $Q$  is never constructed. Rather, the *actions* of  $Q$  and  $Q^T$  are implemented via indirect addressing (and message passing, in parallel implementations). The combined gather-scatter operation  $\Sigma' := QQ^T$  is often referred to as *direct stiffness summation* in the spectral element literature.

We illustrate the use of the above notation by considering an integral that arises in the weak formulation of the Poisson equation. Assuming  $u, v \in H^1$ , we have

$$\int_{\Omega} \nabla u \cdot \nabla v dV = \sum_{k=1}^K \int_{\Omega^k} \nabla u \cdot \nabla v dV. \quad (11)$$

Restricting  $u$  and  $v$  to  $X^N$ , inserting the SEM basis, and substituting GL quadrature for integration, we obtain

$$\int_{\Omega^k} \nabla u \cdot \nabla v dV \simeq (\underline{v}^k)^T A^k \underline{u}^k, \quad (12)$$

where  $A^k$  is the local elemental stiffness matrix and the approximation ( $\simeq$ ) results from substitution of quadrature for integration. An example of  $A^k$  is given by the tensor-product form

$$A^k = \frac{L_s^k}{L_r^k} (\widehat{B} \otimes \widehat{A}) + \frac{L_r^k}{L_s^k} (\widehat{A} \otimes \widehat{B}),$$

for the case where  $\Omega^k$  is an  $L_r^k \times L_s^k$  rectangle. Here,  $\widehat{A}$ , and  $\widehat{B}$  are the respective stiffness and mass matrices on  $[-1, 1]$ , with entries

$$\widehat{A}_{ij} = \sum_{l=0}^N \widehat{D}_{li} \rho_l \widehat{D}_{lj}, = (\widehat{D}^T \widehat{B} \widehat{D})_{ij}, \quad \widehat{B}_{ij} = \rho_i \delta_{ij} = \delta_{ij} \int_{-1}^1 h_i^N(r) dr,$$

where  $\rho_i$  is the GL quadrature weight,  $\widehat{D}_{ij} = h_j^{N'}(\xi_i^N)$  is the one-dimensional derivative matrix, and  $\delta_{ij}$  is the Kronecker delta. Substituting (12) into (11) yields

$$\int_{\Omega} \nabla u \cdot \nabla v dV = \sum_{k=1}^K (\underline{v}^k)^T A^k \underline{u}^k = \underline{v}_L^T A_L \underline{u}_L = \underline{v}^T Q^T A_L Q \underline{u}, \quad (13)$$

where  $A_L := \text{block-diag}(A^k)$  comprises the unassembled local stiffness matrices. Note that the final equality is a result of the interface matching conditions,  $u, v \in H^1$ .

Equation (13) illustrates how the matrix assembly process ( $Q, Q^T$ ) is decoupled from the local spectral element operators contained in  $A_L$ . In the nonconforming case,  $Q$  must be modified at the PC interfaces, where global nodal values are stored along the parent edge. Application of  $Q$  involves interpolation of the associated Lagrange polynomial to nodal points distributed along the corresponding child faces. To ease parallelism, this is implemented by the two-step process illustrated in Figure 7a. Data is first copied from the parent data structure to the corresponding child edges. This step may involve communication if the parent and child elements are not on the same processor. After the copy, an interpolation operator,  $J^{cp}$ , is locally applied to produce the desired nodal values on the child face. This two-step procedure can be represented in matrix form as

$$Q = J_L \widetilde{Q},$$

where  $\widetilde{Q}$  is a Boolean matrix similar to the  $Q$  operator used in the conforming case, and  $J_L$  is block-diagonal and comprises local matrices  $J^{cp}$  that interpolate from  $\partial\Omega^p$  to  $\partial\Omega^p \cap \partial\Omega^c$ . The entries of  $J^{cp}$  are

$$(J^{cp})_{ij} = h_j^N(c_i^{cp}),$$

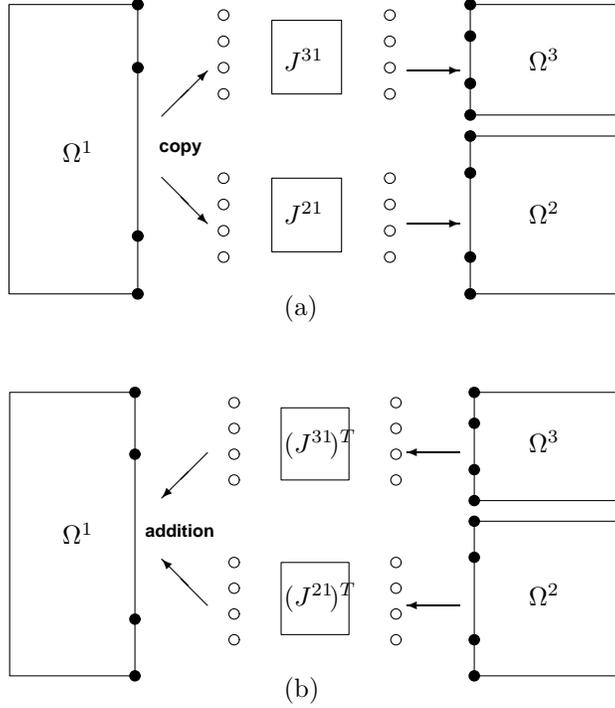


Figure 7: Schematic of (a)  $Q$  and (b)  $Q^T$  implementation.

where  $\zeta_i^{cp}$  represents the mapping of the GL points from the child edge to its parent. In  $\mathbb{R}^3$ , the local interpolation operators mapping from the parent to child face have the tensor-product form  $J_s^{cp} \otimes J_r^{cp}$ . Application of  $Q^T$  follows in the reverse order, with summation replacing the copy, as illustrated in Figure 7b.

For time advancement of the incompressible Navier-Stokes equations, it is desirable to have a diagonal mass matrix [33]. If  $\psi_i, \psi_j$  are two elements of the Lagrangian basis set spanning  $X^N$ , the entries of the mass matrix for the standard spectral element formulation are  $B_{ij} := (\psi_i, \psi_j)_{GL}$ . Equivalently, we have  $B = Q^T B_L Q$ , where  $B_L := \text{block-diag}(B^k)$  comprises the local mass matrices. For the two dimensional case, an entry of  $B^k$  for a nodal point  $\mathbf{x}_{pq}^k$  is simply  $\rho_p \rho_q \mathcal{J}_{pq}^k$ , where  $\mathcal{J}_{pq}^k$  is the Jacobian associated with the mapping  $\hat{\Omega} \rightarrow \Omega^k$ . Diagonality of the mass matrix in the conforming case is assured because of the coincidence of the quadrature points and the Lagrangian nodal points. In the nonconforming case, this property does not hold because the nodal basis functions along the parent edge do not coincide with the quadrature points along the child edge. However, a diagonal (lumped) mass matrix  $\tilde{B}$  can be recovered by setting

$$\tilde{b} := B \hat{e} = Q^T B_L \hat{e}_L \quad (14)$$

and then setting  $\tilde{B}_{ij} = \delta_{ij} \tilde{b}_i$ . Here,  $\hat{e}$  and  $\hat{e}_L$  are the respective global and local vectors containing all ones. Note that, because  $B_L$  is diagonal, (14) amounts to applying  $Q^T$  to the local vector  $\hat{e}_L$  containing the entries of  $B_L$ . In [38] it was shown that this mass lumping procedure simply amounts to replacing the more accurate quadrature in the child elements by quadrature at the nodal points along the parent edge.

*Conditioning.* The nonconforming discretization is particularly effective for external flow problems. In addition to reducing the number of gridpoints in the far field, it allows one to avoid the creation of high-aspect ratio elements that can lead to illconditioning [22]. This point is illustrated by the two-dimensional meshes in Figure 8, which are used to solve the problem of start-up flow past a cylinder at  $Re = 5000$ , following [22, 24]. The conforming mesh (left) exhibits a few high-aspect ratio elements in the far-field that have been eliminated in the nonconforming mesh (right). Table 1 shows the number of Schwarz PCG

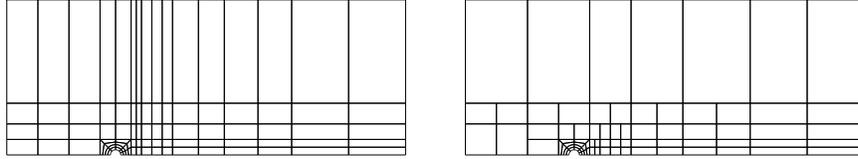


Figure 8:  $K=93$  conforming (left) and  $K = 77$  nonconforming (right) spectral element meshes for flow past a cylinder.

iterations to reduce the pressure residual on the first timestep by  $10^{-5}$  for the case  $N = 7$  and two successive quad-refinements of the meshes in Figure 8. The conforming case shows a marked increase in iteration count with refinement. In contrast, the nonconforming case exhibits a nearly bounded iteration count that is lower in all cases than that achieved by even the coarsest conforming mesh. We note that the extension of the Schwarz method to the nonconforming case required the development of a nonconforming coarse-grid operator, which was done by allowing  $N$  to vary in the  $(Q, Q^T)$  routines and calling them with  $N = 1$  during assembly of the linear finite elements that define the coarse problem.

Table 1: Iteration Count for Cylinder Problem

	Conforming			Nonconforming		
$K$	93	372	1488	77	308	1232
iter	68	107	161	50	58	60

## 7 AMR Development Plan

Kruse and co-workers have recently implemented the nonconforming formulation in the context of NEK5000, the incompressible flow code developed by Fischer and Tufo [49], and are using it to simulate transitional flows in three-dimensional domains. We will use this code as a vehicle to explore spectral element AMR techniques. The geometrically nonconforming formulation naturally lends itself to a block-structured AMR approach. Such techniques were developed in [4, 5, 6] and have been recently scaled to thousands of processors [11]. In particular, we will investigate block-structured AMR approaches which perform refinement by bisecting elements in *each* spatial direction (i.e., quad- and oct-refinement in 2- and 3-D respectively) and employ a tree data structure to maintain the hierarchy of subdomains.

We have extensive experience with PARAMESH [40, 41, 52], the parallel block-structured AMR package used in the FLASH code developed by the ASCI Center for Astrophysical Thermonuclear Flashes at the University of Chicago [27]. Though the package performs well on massively parallel architectures, it lacks the ability to take advantage of and store spectral element data, provides no mechanism for detecting invalid configurations (e.g., Figure 5c), provides no support for curved or deformed surfaces, and lacks robust non-Cartesian grid support. Though the latter shortcomings could be addressed, it is impractical to augment PARAMESH to support spectral element data due to the fact that the data structures are tightly bound to the implementation. Ideally, we would like to plug in a freely available AMR package. We have reviewed several packages (e.g., Chombo [16] and DAGH [18]) and, unfortunately, found them unsuitable for our purposes as they do not provide support for spectral element methods.

We propose to leverage our experience with PARAMESH and the nonconforming NEK5000 code to develop a lightweight, robust, and easy-to-use AMR toolkit for spectral element methods. There are four primary components: tree data management, load balancing, direct stiffness summation, and elliptic solvers. We will use the tree data structure techniques developed in PARAMESH to manage the hierarchy of subdomains and coordinate refinement/derefinement. Element to processor distribution capability will be provided through a suite of load balancing tools. For calculating initial or static distributions we will use the recursive spectral bisection module in NEK5000. For dynamic load balancing we will employ the space filling curve techniques of PARAMESH. Direct stiffness summation operations will be performed by our general-purpose,

stand-alone gather-scatter package that supports a number of associative/commutative operations, including user-defined functions. The recently developed bucket-sort initialization option has low setup overhead that is appropriate for dynamically adaptive meshing. Elliptic solver technology is discussed in section 9.

We will use the a posteriori error estimators for spectral element methods developed by Mavriplis [46] to determine when to refine/derefine elements. To identify vortices and other coherent structures we will employ the  $\lambda_2$  method of Jeong and Hussain [36]. Identification of a vortex in viscous flows is challenging because the classic rules governing vortex dynamics generally apply only in the inviscid limit. In boundary layer flows, viscosity is nonnegligible, and standard approaches such as integrating vortex lines or using pressure minima or vorticity maxima can lead to improper vortex identification. Jeong and Hussain have established a robust criterion for the identification of vortex (or coherent) structures in viscous flows based on the eigenvalues of the symmetric  $3 \times 3$  tensor

$$M_{ij} := \sum_{k=1}^3 (\Omega_{ik}\Omega_{kj} + S_{ik}S_{kj}), \quad (15)$$

where

$$\Omega_{ij} := \frac{1}{2} \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad S_{ij} := \frac{1}{2} \left( \frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) \quad (16)$$

represent the symmetric and antisymmetric components of the velocity gradient tensor,  $\nabla \mathbf{u}$ . To minimize noise, the gradients are computed using the original polynomial description of the data. Given the three (real) eigenvalues of  $M$  at each grid point, a vortex core is identified as any contiguous region having two negative eigenvalues. If the eigenvalues are sorted such that  $\lambda_1 \leq \lambda_2 \leq \lambda_3$ , then any region for which  $\lambda_2 < 0$  corresponds to a vortex core. Moreover, the criterion  $\lambda_2(\mathbf{x}) < 0$  is scale invariant, so there is in principle no ambiguity in selecting coherent structure cores within the extent of the domain given a  $\lambda_2$  threshold (in practice, one usually biases the isosurface to a value that is below zero by a small fraction of the full dynamic range in order to avoid noise in regions where the velocity is close to zero). We note that this procedure has been successfully employed by Tufo and colleagues to render vortex isosurfaces in transitional flows for display in an immersive environment (e.g., CAVE or ImmersaDesk) [67, 68].

## 8 Filtering

One of principal attractions of spectral element methods is that, for smooth solutions, the error decreases exponentially fast with increasing polynomial degree  $N$ . However, spectral element methods can also be highly effective in solving transport problems in which the solution is *not* smooth. This property is illustrated by the convected-cone example of Figure 9, which was introduced by Gottlieb and Orszag [30]. A unit-height cone with a base-radius of 0.1 centered at  $(x, y) = (0, .25)$  is subjected to plane rotation in the domain  $\Omega = [0, 1]^2$ . The solution is evolved according to  $u_t + \mathbf{c} \cdot \nabla u = 0$ , with periodic boundary conditions and convecting field  $\mathbf{c} = (y - .5, .5 - x)$ . Figure 9 shows the results after a single revolution for three spectral element discretizations,  $(K, N)$ , where  $K$  is the number of (square) elements, and  $N$  is the polynomial degree in each spatial direction. Each case corresponds to a  $32 \times 32$  grid. Time-stepping is based on third-order Adams-Bashforth (AB3) with  $\Delta t = \pi/1000$ . (Fourth-order Runge-Kutta results are identical.) The low-order

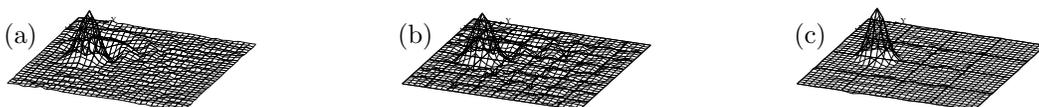


Figure 9: Spectral element results for convected cone problem [30] on  $32 \times 32$  grids: (a)  $(K, N) = (256, 2)$ , (b)  $(64, 4)$ , (c)  $(16, 8)$ .

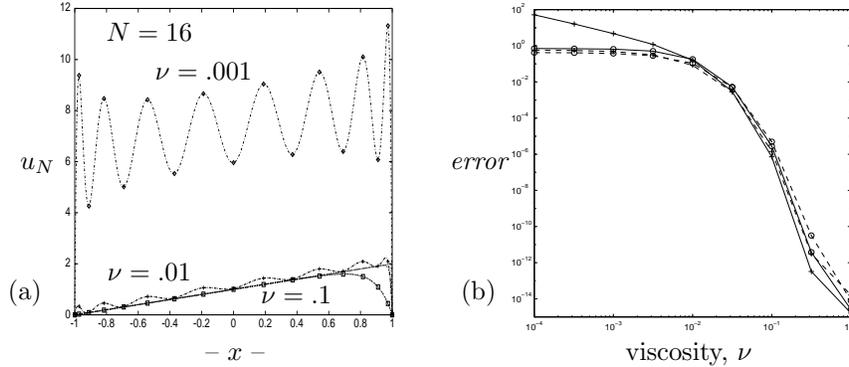


Figure 10: Spectral Galerkin results for steady advection-diffusion problem: (a)  $u_N(x)$  for  $N = 16$ ,  $\nu = .001$ – $.1$ , (b) maximum pointwise error with (---) and without (—) filtering for  $N = 15$  (o) and  $N = 16$  (+).

cases ( $N=2, 4$ ) show evidence of significant numerical dispersion. By contrast, the dispersion is diminished for the moderately high-order case ( $N = 8$ ), and the solution produces a reasonable representation of the original cone. The minima for the three respective cases are  $-.1419$ ,  $-.1127$ , and  $-.0371$ , while the maxima are  $.7693$ ,  $.7413$ , and  $.8652$ .

Unfortunately, Galerkin formulations suffer from well-known instabilities in convection-dominated flows when underresolved boundary layers are encountered. A classic example is the one-dimensional steady convection-diffusion problem  $u_x - \nu u_{xx} = f$ ,  $u(-1)=u(1)=0$ ,  $f=1$ . The spectral Galerkin formulation of this problem is given by: Find  $u \in \mathbb{P}_N^0$  such that

$$(u_x - \nu u_{xx} - f, v) = 0 \quad \forall v \in \mathbb{P}_N^0,$$

where  $\mathbb{P}_N^0$  is the space of all polynomials of degree  $\leq N$  vanishing at  $\pm 1$ , and  $(\cdot, \cdot)$  is the standard  $L^2$  inner-product on  $[-1, 1]$ . As shown by Canuto [12], and illustrated in Fig. 10, the spectral solution is unbounded as  $\nu \rightarrow 0$  when  $N$  is even. As shown in Fig. 10b, for large  $\nu$  (smooth solutions), the error is smaller for  $N = 16$  than for  $N = 15$ . However, as  $\nu \rightarrow 0$ , the error grows without bound for  $N=16$  but remains bounded for  $N=15$ .

Ideally, one would like to retain the good transport properties illustrated in Fig. 9, without the sensitivity to parameters exemplified in Fig. 10. Several proposed strategies for stabilizing convection-dominated problems involve a reformulation of the Galerkin procedure, for example, Petrov-Galerkin schemes [54], shifted grids [29], the addition of bubble functions [13], or the addition of higher-order derivative terms, such as in the spectrally vanishing viscosity method [60, 45]. Related to this latter approach are filtering schemes [9, 31]. A significant advantage of filtering is that it can be applied as a post-processing step and therefore does not require a change to the underlying discretization or solver. We propose to investigate and implement the filters of Fischer and Mullen [25] and Boyd [8, 9]. These filters are *local* in nature, and hence are ideally suited for parallel applications.

## 9 Solvers

Efficient solution of elliptic problems in complex domains depends on the availability of fast solvers for sparse linear systems. Our solution strategy involves two stages. First, we exploit the fact that we are solving similar problems from one step to the next, by projecting the current solution onto a subspace of previous solutions [23]. The remaining component is then computed using conjugate gradients, preconditioned by the additive overlapping Schwarz method introduced by Dryja and Widlund [20] and developed in the spectral element context in [22, 24]. The key components of our overlapping Schwarz implementation are fast local solves that exploit the tensor-product formulation, and a parallel coarse-grid solver that scales to thousands of processors[21, 66].

While the overlapping Schwarz method has provided a significant reduction in time to solution over previous multilevel solvers [22], we have nonetheless observed relatively high iteration counts for some transitional simulations. While we expect some performance recovery with improved mesh generation and smoothing algorithms, we feel that there is merit in revisiting the early spectral element multigrid (SEMG) work carried out by Maday et al. [44]. In particular, we will investigate the use of additive overlapping Schwarz methods as a smoother. This has the potential to alleviate degradation associated with high-aspect ratio elements in two ways. First, because the local solves are direct, ill-conditioning associated with local element stretching will not be manifest. Second, because the smoother on the coarser multigrid levels will have greater overlap, the aspect-ratio problems encountered in classical overlapping Schwarz methods should be mitigated [22]. As shown in [24], the local tensor-product forms can be exploited to implement the local Schwarz smoothers at a cost that is lower than that for operator evaluation. Initial trials in Matlab have shown that the overlapping Schwarz smoother with minimal overlap can provide  $> 2.5\times$  reduction in iteration count over earlier SEMG methods that employed Jacobi-based smoothers. We intend to exploit our existing local fast solvers, interelement communication utilities, and parallel coarse grid solvers to test these ideas in production settings.

## 10 Preliminary Results

We have begun to incorporate the scalable solver technology from Tufo and Fischer [24, 69] into SEAM, the 3-D climate dynamical code of Thomas and Loft. To date we have extended the local solution technique to the cubed-sphere spectral element formulation. Our results were documented in a paper, entitled “An overlapping Schwarz preconditioner for a spectral-element method on the cubed-sphere”, which was accepted to the 7th Copper Mountain Conference on Iterative Methods. In addition, the filtering techniques of Fischer and Mullen [25] have been added to suite of Boyd filters in SEAM, resulting in lower communication overhead because they are interpolation based and applied directly to the sphere-valued data.

## 11 Project Team Qualifications

We have been collaborating on the development of codes to study atmospheric flows for over six months. The integration of our team is reflected in the success of a recently held workshop (<http://www.asp.ucar.edu/gtp/methods2001.htm>), which brought together the community of researchers working in high-order methods with leading experts in adaptive mesh refinement algorithms and wavelets as well as those working on the several aspects of the physics of turbulence in the context of geophysical and astrophysical flows (participant list below). Thomas has over 10 years’ experience in climate and weather simulations and is one of the primary developers of SEAM, a freely available spectral element code for climate study. As a 2001 Gordon Bell finalist, Thomas and co-workers demonstrated scalability of a dynamical core for an atmospheric general circulation model, achieving an unprecedented integration rate of over 100 years per day for T170 resolution. Tufo has over 8 years’ experience with spectral element methods and implementations and is co-developer of NEK5000, a state-of-the-art code for simulation of unsteady incompressible flows in complex geometries. As a 1999 and 2000 recipient of the Gordon Bell award, Tufo has demonstrated experience in high-performance and large-scale parallel computing.

## 12 Workplan

A post-doctoral fellow will work under the supervision of Dr. Thomas at NCAR and will contribute directly to the ongoing development of the 3-D spectral element dynamical core (SEAM). This work will be coordinated with the introduction of the NCAR CCM physical parameterization package with the spectral element model. A graduate student will be located at the University of Chicago and work under the supervision of the co-PI, Dr. Henry Tufo. Travel funds will be used to support site visits by the participants. The post-doctoral fellow will be expected to also interact with scientific staff within the Climate and Global Dynamics (CGD) division at NCAR.

Our research time-line is given below

year 1:

1. Enhance SEAM filters/solvers.
2. Add geometric nonconforming spectral elements to SEAM.
3. Configure nonconforming version of NEK5000 for anelastic simulations.
4. Development of AMR package for SEAM.
5. Extension of SEAM code framework to include DG.

year 2:

1. Investigate functionality of nonconforming SEAM.
2. Integrate AMR package into SEAM.
3. Add vortex tracking module.
4. Add a posteriori error estimators.
5. Code verification and validation of SEAM.
6. Release error plots and results for standard test flows.
7. Compare “fully resolved” with new AMR code.
8. Compare both codes with idealized climate simulations.
9. Begin verification and validation of DG formulation.

year 3:

1. Complete verification and validation of DG formulation.
2. Enhance tracking to include a posteriori estimators.
3. Add physical parameterizations and explore multi-scale physics.
4. Leverage anelastic work in NEK5000 and compare against SEAM.
5. Compare time/space adaptive DG vs. space adaptive SE.

## 13 Modes of Collaboration and Training

We recently held the first NCAR Workshop on Adaptive and High-Order Methods with Applications in Turbulence. As part of this proposal we seek to continue this on a yearly basis. Among the participants were the following:

Paul Fischer, Argonne  
Cathy Mavriplis, George Washington, NSF  
Jan Hesthaven, Brown  
Tim Warburton, New Mexico  
Jean-Francois Remacle, RPI  
Charles Meneveau, John Hopkins  
Beth Wingate, LANL  
Peter Bartello, McGill  
Greg Beylkin, Colorado  
Marc Brachet, Nice, France  
Phil Colella, LBL

Timely distribution of test cases, results, and code will occur through the SEAM Web page (<http://www.scd.ucar.edu/css/seam.html>). In addition, results will be distributed via peer-reviewed journals and through the PDEs on the Sphere conference series organized by NCAR.

## 14 Management Plan

We envision a close collaboration on this project between the participants and scientists in the climate and global dynamics (CGD) division at NCAR, as well as university researchers interested in numerical modeling issues and physical parameterizations. There will also be interaction with the Geophysical Turbulence Program at NCAR through Annick Pouquet and Joe Tribbia. Dr. Thomas will coordinate the research activities of the group and also ensure that the results are relevant to the broader NCAR and university modeling community. Dr. Tufo will serve as a bridge to the CFD and spectral element modeling community through the University of Chicago and Argonne National Laboratory.

Our main focus is to provide feedback to the climate and weather modeling communities on the mechanisms by which small scales feedback and influence the large flow structures. The combination of workshops and frequent site visits by the graduate student and post-doctoral fellow should provide an excellent opportunity for interaction with the geophysical modeling community.

## Acknowledgments

Sections of text and figures from [26] used with permission of the authors.