A Schwarz Preconditioner for the Cubed-Sphere

Stephen J. Thomas, John M. Dennis
National Center for Atmospheric Research

Henry M. Tufo and Paul F. Fischer
Argonne National Laboratory

Abstract

A spectral element formulation of the atmospheric 2-D shallow-water equations on the cubed-sphere is described. The equations are written in tensor form using the contravariant and covariant velocity components. A semi-implicit time discretization results in a Helmholtz problem for the pressure. The Laplacian operator is approximated by the L^2 pseudo-Laplacian arising in the P_N/P_{N-2} spectral element formulation of the incompressible Stoke's problem. The overlapping Schwarz preconditioner of Fischer et al. (1998), based on the fast diagonalization method (FDM), is extended to generalized curvilinear coordinates. To obtain a separable operator for the linear finite-element tensor-product approximation within each spectral element, extrema of the inverse metric tensor and its determinant are employed. Convergence rates and parallel CPU timings are compared against a block-Jacobi preconditioner.

1. Introduction

A semi-implicit time-stepping scheme is commonly applied to the terms responsible for fast gravity waves in atmospheric general circulation models (Robert 1969). The scheme removes the associated time step restrictions but requires the solution of a Helmholtz problem. Climate models have traditionally employed the spectral transform method. The resulting Helmholtz problem is trivial to solve in spectral space because the spherical harmonics are eigenfunctions of the Laplacian on the sphere. The spectral element method offers several potential advantages over the spectral transform approach and has been proven effective in geophysical fluid dynamics by Taylor et al. (1997) and Iskandarani et al. (1995). In particular, spectral elements maintain the accuracy and exponential convergence rate exhibited by the spectral transform method. The spherical harmonic basis functions are global and require inefficient non-local operations on parallel computers, whereas spectral elements use local basis functions and nearest neighbor communication. A semi-implicit scheme for a spectral element atmospheric model on the cubed-sphere

was developed by Thomas and Loft (2002), requiring a weak formulation of the governing equations. Initially, a block-Jacobi preconditioned conjugate gradient solver was implemented, following Fischer and Ronquist (1994). Both explicit and semi-implicit implementations of this model were shown to be highly scalable, even at relatively low climate resolutions with $\mathcal{O}(10^6)$ degrees of freedom (Loft et al. 2001).

A suitable preconditioner can significantly accelerate solver convergence, thereby reducing the overall computational expense of the semi-implicit relative to the explicit time-step. It is imperative that the longer time-step overcome the higher cost of the Helmholtz solver, resulting in an acceleration of the integration rate (measured in simulated days per day). A block-Jacobi preconditioner performs well at low climate resolutions, but leads to a growth in the iteration count at higher resolutions as both the number of elements K and the polynomial order N are increased. In this paper we investigate whether a generalization of the additive overlapping Schwarz approach of Fischer et al. (1998) leads to a more robust preconditioner. The 2-D block-Jacobi preconditioner has a computational cost of $\mathcal{O}(KN^4)$. The subdomain problems in the Schwarz preconditioner are based upon linear finite element approximations of the local (separable) elliptic operators. Tensor products of linear 1-D finite elements within the overlap region admit solution via the fast diagonalization method (FDM) of Lynch et al. (1964). The computational cost in this case drops to $\mathcal{O}(KN^3)$. The generalized Laplacian-like elliptic operator appearing in the Helmholtz problem is analogous to the L^2 pseudo-Laplacian E (viz. the Schur complement system) appearing in the P_N/P_{N-2} spectral element formulation of the Stoke's problem. In generalized curvilinear coordinates on the cubed-sphere, this operator includes cross-derivative terms and is nonseparable due to the presence of metric factors. To obtain a separable operator, cross-terms are dropped and extrema of the metric factors are specified within each spectral element.

A standard test suite for evaluating numerical approximations of the shallow-water equations in spherical geometry was proposed by Williamson et al (1992). The test suite was designed to identify the merits and trade-offs that might be encountered by various numerical schemes. Test Cases 1 through 4 have analytic solutions. The remaining cases, zonal flow over an isolated mountain (Test Case 5), Rossby-Haurwitz wave (Test Case 6), and analyzed 500 mb initial conditions (Test Case 7), do not have closed form solutions. In this paper, the integration rate of the spectral element model on the cubed sphere is evaluated by using a zonal geostrophic flow (Test Case 2). The overlapping Schwarz and block-Jacobi preconditioners are compared on the basis of the achieved integration rate acceleration over explicit time-stepping. The polynomial order N is also varied in order to examine the affect on solver performance and demonstrate exponential convergence.

2. Spectral Element Formulation

The shallow water equations have been used for many years by the atmospheric modeling community to test 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 ϕ. For curvilinear coordinates, the shallow water equations in perturbation form can be written as follows (see Sadourny 1972)

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

where $\phi = \phi' + \phi_0$, f is the Coriolis force and ζ is the relative vorticity. ϵ_{jk} 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 (1996) showed that an equal angular projection results in a more uniformly spaced grid. For equal angular coordinates $(\,x_1,\,x_2\,)$, $-\pi/4 \le x_1, x_2 \le \pi/4$, and the metric tensor for all six faces of the cube is

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

$r = \{1+\tan^2 x_1+\tan^2 x_2\}^{1/2}$, 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, 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, the relationship is $A_i\,\mathbf{u}_i = A_j\,\mathbf{u}_j$, $\mathbf{u}_i = A_i^{-1}A_j\,\mathbf{u}_j$.

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 $\delta\, \phi = \phi^{n+1} - \phi^{n-1}$

$$\delta\, u_i + \Delta t\, \frac{\partial}{\partial x^i}\, (\, \delta\, \phi\,) \;=\; 2\Delta t \left[-\frac{\partial}{\partial x^i}\, (\, \phi\,)^{n-1} + f_u^{(n)} \right] \tag{2.1}$$

$$\delta\, \phi + \Delta t\, \frac{\phi_0}{g}\, \frac{\partial}{\partial x^j}\, (\, g\, \delta u^j\,) \;=\; 2\Delta t \left[-\frac{\phi_0}{g}\, \frac{\partial}{\partial x^j}\, (\, g\, u^j\,)^{n-1} + f_\phi^n \right] \tag{2.2}$$

where the tendencies f_u and f_ϕ contain nonlinear terms.

In the spectral element discretization, the computational domain Ω is partitioned into K elements Ω^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 Ronquist (1988),

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

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 $(\, \xi_i,\, \rho_i\,)$, $i = 0, \ldots N$ are the Gauss-Lobatto nodes and weights and $(\, \zeta_i,\, \sigma_i\,)$, $i = 1, \ldots N-1$ 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 inter-element boundaries which share Gauss-Lobatto points and direct stiffness summation is applied to assemble the global matrices. Derivative matrices $\tilde{\mathbf{D}} = (\tilde{D}_1, \tilde{D}_2)$ are

rectangular. $\mathbf{B} = (B_1, B_2)$ and $\tilde{B}$ are diagonal velocity and geopotential mass matrices, with $\tilde{D}_j = \tilde{B}\, D_j$. The assembled discrete shallow water equations are then

$$B_i\, \delta\, u^i - g^{ij}\, \Delta t\, \tilde{D}_j^T\, \delta\, \phi \;=\; R_u^i \tag{2.3}$$

$$\tilde{B}\, \delta\, \phi + \Delta t\, \frac{\phi_0}{g}\, \tilde{D}_i\, g\, \delta\, u^i \;=\; R_\phi \tag{2.4}$$

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 (2.3) 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 $\delta\, u^i$

$$\delta\, u^i = B_i^{-1}\, \left(R_u^i + \Delta t\, g^{ij}\, \tilde{D}_j^T\, \delta\, \phi \right) \tag{2.5}$$

and then applying back-substitution to obtain

$$g\, \tilde{B}\, \delta\, \phi + \Delta t^2\, \phi_0\, \tilde{D}_i\, g\, B_i^{-1}\, g^{ij}\, \tilde{D}_j^T\, \delta\, \phi = R'_\phi \tag{2.6}$$

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 δu^i is obtained from (2.5). 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 thus a preconditioned conjugate gradient solver is applied. The generalized Laplacian-like operator $\tilde{D}_i\, g\, B_i^{-1}\, g^{ij}\, \tilde{D}_j^T$ is analogous to the L^2 pseudo-Laplacian $E = \tilde{D}_i\, B_i^{-1}\, \tilde{D}_i^T$ for the Stoke's problem. The metric factors g and g^{ij} are both functions of (x_1, x_2) and thus the operator is nonseparable. In order to construct an approximate separable Laplacian for the FDM based preconditioner (section 3), cross-derivative terms are dropped and $\max g \min g^{ij} E$ is employed within each element.

Rather than reducing the relative residual to a specified error tolerance, we observe that the rhs R'_ϕ depends on the flow divergence. The CG solver need only be iterated until the norm of the divergence remains constant, Smolarkiewicz et al. (1997). A stopping criterion for the solver is therefore

$$\| H\, \delta\, \phi - R'_\phi \|_\infty \sim \Delta t\, \|\nabla \cdot \mathbf{v}\|_\infty < \varepsilon_c\, C\ .$$

where $C \sim \mathcal{O}(1)$ is the advective Courant number for spectral elements.

3. Overlapping Schwarz Method

In this paper we generalize the overlapping Schwarz preconditioner of Fischer et al. (1998) to the cubed-sphere. Their work is based on the overlapping additive Schwarz procedure developed by Dryja and Widlund (1987) and Widlund (1992). The subdomain solver relies on the fast diagonalization method of Lynch et al. (1964), successfully used in a number of spectral element preconditioning applications by Couzy and Deville (1995) and by Couzy (1995). The additive variant was adopted because it is intrinsically parallel and symmetric. An overlapping method is appropriate, as opposed to a Schur complement or substructuring approach, because it does not require interface data structures. This is an important consideration because the Gauss points do not lie on the natural subdomain interfaces.

Formally, the additive Schwarz preconditioner is expressed as the sum of outputs from several subproblems:

$$M_o^{-1} = R_0^T A_0^{-1} R_0 + \sum_{k=1}^{K} R_k^T A_k^{-1} R_k \quad .$$

The subproblems for $k \geq 1$ correspond to the solution of local Helmholtz problems on overlapping subdomains, $\tilde{\Omega}^k$. The restriction and prolongation operators, R_k and R_k^T, $k \geq 1$, are Boolean matrices which transfer data to and from the subdomain problems. The product $\underline{p}_k = R_k \underline{p}$ extracts the components of a vector $\underline{p}$ which belong to $\tilde{\Omega}^k$, while $\underline{p} = R_k^T \underline{p}_k$ copies the components of a subdomain solution, $\underline{p}_k$, to a global vector, $\underline{p}$, and sets components outside of $\tilde{\Omega}^k$ to zero. In addition to the local problems, the Schwarz preconditioner has a coarse grid component, denoted here by subscript 0, which serves to efficiently eliminate low-wave number components of the residual. The coarse grid problem corresponds to a Helmholtz problem discretized on a mesh defined by a triangulation of the subdomain vertices. The prolongation operator, R_0^T, is simply an interpolant from the coarse grid to the Gauss points. For this study, a coarse grid solver was not implemented.

In this section we consider the development of solvers for the local problems which are particularly well suited to the spectral element method in $\mathbb{R}^2$. For ease of exposition, we restrict our attention to the Laplacian. However, the preconditioner is developed for the full Helmholtz operator. The local stiffness matrices, A_k, $k \geq 1$, are derived from a tensor-product of one-dimensional finite element bases and this is reflected in Figure 1, where the overlapping domains $\tilde{\Omega}^k$ are obtained from the Ω^k by extending the boundary by an amount δ in each direction. This allows the use of FDM based solvers which require

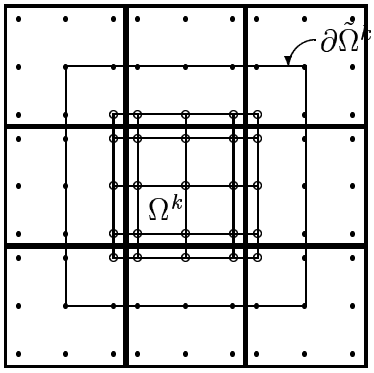

Figure 1: Degrees-of-freedom for tensor-product based discretizations of local problems. Zero Dirichlet boundary conditions are applied on $\partial\tilde{\Omega}^k$.

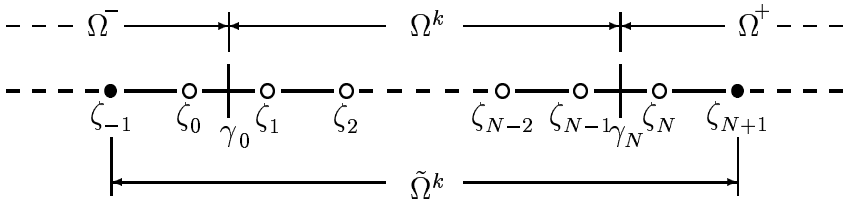

Figure 2: Depiction of overlapping subdomain $\tilde{\Omega}^k$ in one dimension, minimal overlap case.

only $O(N^2)$ storage and $O(N^3)$ work per solve.

We begin with the definition of the overlapping subdomains by considering the one-dimensional example shown in Figure 2. Degrees-of-freedom are associated with the nodes (open circles) in $\tilde{\Omega}^k$. The points ζ_i, $i \in \{1, \ldots, N-1\}$ are the images of the Gauss points in $]-1, 1[$ mapped onto Ω^k. Similarly, ζ_i, $i \leq 0$ and $i \geq N$ are the images of the corresponding Gauss points mapped onto the left and right subdomains, respectively. The overlapping region, $\tilde{\Omega}^k \in [\zeta_{-1}, \zeta_{N+1}]$, is obtained by extending Ω^k by two nodal points in each direction. Homogeneous Dirichlet boundary conditions are applied at ζ_{-1} and ζ_{N+1} when Ω^k is in the interior of Ω so the extension only adds two degrees-of-freedom to the local problem. We refer to this as the minimal overlap case. If the left (right) side of $\partial\Omega^k$ is coincident with the boundary, $\partial\Omega$, then the domain is not extended beyond γ_0 (γ_N) and homogeneous Dirichlet or Neumann boundary conditions are imposed at that point in accordance with the boundary conditions on $\partial\Omega$.

To construct the finite element operators for the standard (interior) one-dimensional case,

we consider the space of piecewise linear functions, $\phi_i(\zeta)$, $\zeta \in [\zeta_{-1}, \zeta_{N+1}]$, $i = 0, \ldots, N$:

$$\phi_i(\zeta) \;=\; \begin{cases} \frac{\zeta-\zeta_{i-1}}{\zeta_i-\zeta_{i-1}} & \zeta_{i-1} \le \zeta < \zeta_i \\ \frac{\zeta-\zeta_{i+1}}{\zeta_i-\zeta_{i+1}} & \zeta_i \le \zeta < \zeta_{i+1} \quad i \in \{0, \ldots, N\} \\ 0 & \text{otherwise.} \end{cases} \tag{3.1}$$

The variational form for the homogeneous Dirichlet problem, $-u''(x) = f(x)$ in $\tilde{\Omega}^k$, $u = 0$ on $\partial\tilde{\Omega}^k$, gives rise to the tridiagonal stiffness matrix:

$$\tilde{A}_{ij} \;=\; \int_{\zeta_{-1}}^{\zeta_{N+1}} \frac{d\phi_i}{d\zeta}\frac{d\phi_j}{d\zeta} d\zeta \quad i,j \;\in\; \{0, \ldots, N\}^2 \,,$$

and associated diagonal (lumped) mass matrix:

$$\tilde{B}_{ij} \;=\; \delta_{ij} \int_{\zeta_{-1}}^{\zeta_{N+1}} \phi_j(\zeta) d\zeta \quad i,j \;\in\; \{0, \ldots, N\}^2 \,.$$

The matrices are modified in the usual way if either end of Ω^k coincides with $\partial\Omega$.

The construction of the one-dimensional problem is extended to $\mathbb{R}^2$ by taking the tensor product of the bases and operators just described. A typical overlapping domain in $\mathbb{R}^2$ is shown in Figure 1 (right). The degrees-of-freedom correspond to Lagrangian basis coefficients associated with the nodes (open circles) in the interior of $\tilde{\Omega}^k$. Corner points are included in the overlap region on the cubed-sphere, except at the eight cube corner points where values at the missing Gauss points are set to zero. If the nodes are numbered lexicographically, then the stiffness matrix for the two-dimensional Laplacian on $\tilde{\Omega}^k$ can be written as the Kronecker product:

$$A_k \;=\; \tilde{B}_2 \otimes \tilde{A}_1 \;+\; \tilde{A}_2 \otimes \tilde{B}_1 \,. \tag{3.2}$$

Here, the subscript on the one-dimensional matrices, $\tilde{A}$ and $\tilde{B}$, indicates the associated coordinate direction in the reference element.

Matrices which satisfy (3.2) have a particularly simple inverse based upon the FDM. If $\tilde{A}$ is symmetric and $\tilde{B}$ is symmetric positive definite, then the following similarity transformation holds:

$$S^T \tilde{A} S = \Lambda, \qquad S^T \tilde{B} S = I,$$

where $\Lambda = diag(\lambda_1, \ldots, \lambda_n)$ the matrix of eigenvalues, and $S = (\underline{s}_1, \ldots, \underline{s}_n)$ is the matrix of eigenvectors associated with the generalized eigenvalue problem $\tilde{A}\underline{s} = \lambda\tilde{B}\underline{s}$. As a result, A_k is readily diagonalized, and its inverse is given by:

$$A_k^{-1} \;=\; (S_2 \otimes S_1)\,(I \otimes \Lambda_1 \;+\; \Lambda_2 \otimes I)^{-1}\,(S_2^T \otimes S_1^T)\,.$$

4. Numerical Results

A standard test set for evaluating numerical approximations to the shallow water equations in spherical geometry has been proposed by Williamson et al. (1992). Global l_1, l_2 and l_∞ relative error metrics on the sphere are also defined in this paper. Test case 2 is a zonal geostrophic flow which is a steady-state solution of the nonlinear equations. The velocity field on the sphere is specified initially (and for all time) as

$$\begin{aligned} u &= u_0 \left(\cos\theta \cos\alpha + cos\lambda \sin\theta \sin\alpha \right) \\ v &= -u_0 \sin\lambda \sin\alpha \ . \end{aligned}$$

where (λ, θ) are spherical longitude and latitude coordinates. α is the angle between the axis of solid body rotation and the polar axis. The corresponding stream function and velocity potential are given by

$$\begin{aligned} \psi &= -au_0 \left(\sin\theta \cos\alpha + cos\lambda \sin\theta \sin\alpha \right) \\ \chi &= 0 \end{aligned}$$

where a is the radius of the earth. The analytic geopotential field $\phi = gh$ is specified as

$$\phi = \phi_0 - \left(a\Omega u_0 + \frac{u_0^2}{2} \right) \times \left(-\cos\lambda \cos\theta \sin\alpha + \sin\theta \cos\alpha \right)^2 \ .$$

Ω is the rotation rate of the earth. Parameter values are specified to be $u_0 = 2\pi a/(12\text{days})$ and $\phi_0 = 2.94 \times 10^4$ m^2/s^2. The Coriolis parameter associated with this solution is

$$f = 2\Omega \left(-\cos\lambda \cos\theta \sin\alpha + \sin\theta \cos\alpha \right)$$

We use test case 2 with $\alpha = 0$ to compare the integration rates of the semi-implicit and explicit time-stepping schemes. Initially, three resolutions of the multi-layer shallow-water equations were examined with 30 vertical layers. The number of elements on a cube face edge was set to $ne = 8$, 12 and 16 with 6×6 pressure points ($np = 6$) per element (these resolutions are denoted C56N6, C84N6 and C112N6, referring to the number of degrees of freedom along a cube face edge). The integration rate acceleration factor (ratio of integration times) of the semi-implicit over the explicit scheme for each resolution, as a function of processor count on an IBM SP parallel computer, is plotted in Figures 4 and 5 for the block-Jacobi and Schwarz preconditioners, respectively. The l_2 error after 15 days is reported along with the time step Δt, CG solver convergence tolerance ε_c, average number of CG iterations during the integration $iter$. A Boyd-Vandeven filter (Boyd 1996) was applied every $freq$ time steps with weight factor μ to stabilize the model (see Taylor

et al. 1997). The l_1, l_2 and l_∞ errors for the 15 day semi-implicit integration are plotted in Figure 3. In order to evaluate the robustness of the preconditioners with h-refinement, simulation results are reported for $np = 6$ with increasing ne in Table 1. To demonstrate exponential convergence, results for $np = 12$ pressure points are presented in Table 2.

The global relative error metrics plotted in Figure 3 indicate that the semi-implicit spectral element model at C56N6 resolution is producing results which are comparable in accuracy to the explicit integration. From Table 2, a C26N12 spectral element simulation ($ne = 2$, $np = 12$) produces the same order of magnitude l_2 errors as a spectral transform model at T42 resolution (Jacob-Chien et al. 1995). The gradual error growth observed in these plots is due to the accumulation of round-off errors rather than truncation errors. The fact that the semi-implicit and explicit error metrics are similar suggests that the Helmholtz solver tolerance is sufficient to ensure a converged solution. A comparison of the integration rate acceleration factors in Figures 4 and 5 indicates that the overlapping Schwarz preconditioner with an overlap of one (FDM-1), without the coarse grid solve, performs better than the block-Jacobi preconditioner. In particular, the FDM-1 scheme shows a marked improvement over block-Jacobi at low processor counts below 64. We have implemented a stacked multi-layer shallow water code to assess the performance of both preconditioners in the 3-D case. In the low processor count regime, the block-Jacobi scheme is hampered by cache-effects due to a larger memory footprint (corresponding to a different matrix for each of the 30 vertical layers).

Even though the Schwarz preconditioner requires extra communication associated with the overlap, it exhibits good parallel scalability at relatively low climate resolutions. Despite the use of the latency tolerant CG algorithm of d'Azevedo et al. (1992), communication tends to dominate at high processor counts when eight or fewer elements are assigned to a processor. For example, communication requires 30% of the total simulation time using 48 processors for the semi-implicit model run with the FDM-1 preconditioner at C56N6 resolution, whereas communication in the block-Jacobi run represents 20% of the time. This is the cross-over point where block-Jacobi achieves a higher acceleration factor. We also note that the acceleration rate for overlapping Schwarz does not decrease as rapidly as block-Jacobi at higher resolutions. Tables 1 and 2 also show a growth in the iteration count as the number of elements is increased under h-refinement. However, the FDM-1 preconditioner requires fewer iterations than block-Jacobi. We also found that an FDM-0 scheme with zero overlap typically required twice the number of CG iterations compared to FDM-1, which is in agreement with theory. A coarse solver is required in order to keep the iteration count from growing as the number of spectral elements increases.

5. Conclusions

A semi-implicit time stepping scheme for a shallow water model on the cubed-sphere has been implemented using spectral element methods. In this paper we have compared block-Jacobi and overlapping Schwarz preconditioners for the solution of the symmetric Helmholtz problem via the conjugate gradient algorithm. Our comparison is based on the acceleration of the model integration rate (number of simulated days per day) of the semi-implicit over the explicit schemes. In addition, the acceleration factor was computed for parallel runs of a multi-layer model on an IBM SP at NCAR. Both schemes exhibit good parallel scalability, however, the extra communication required by overlapping Schwarz is a negative factor at large processor counts and low resolutions. The block-Jacobi scheme performs poorly at low processor counts due to a larger memory footprint that adversely impacts cache utilization. The advantage of the semi-implicit scheme with a block-Jacobi preconditioner is diminished at higher resolutions because the low-wave number components of the residual are not effectively eliminated. The overlapping Schwarz preconditioner requires slightly fewer iterations than block-Jacobi. However, the iteration count also increases under h refinement. A coarse grid solver is therefore required to maintain a constant iteration count and this will be studied in future work. A larger overlap region may also improve the robustness of the solver. We also plan to extend the 2-D semi-implicit algorithms to the full 3-D primitive equations with a hybrid pressure vertical coordinate. The resulting atmospheric general circulation model can then be tested using the the Held and Suarez (1994) idealized forcings.

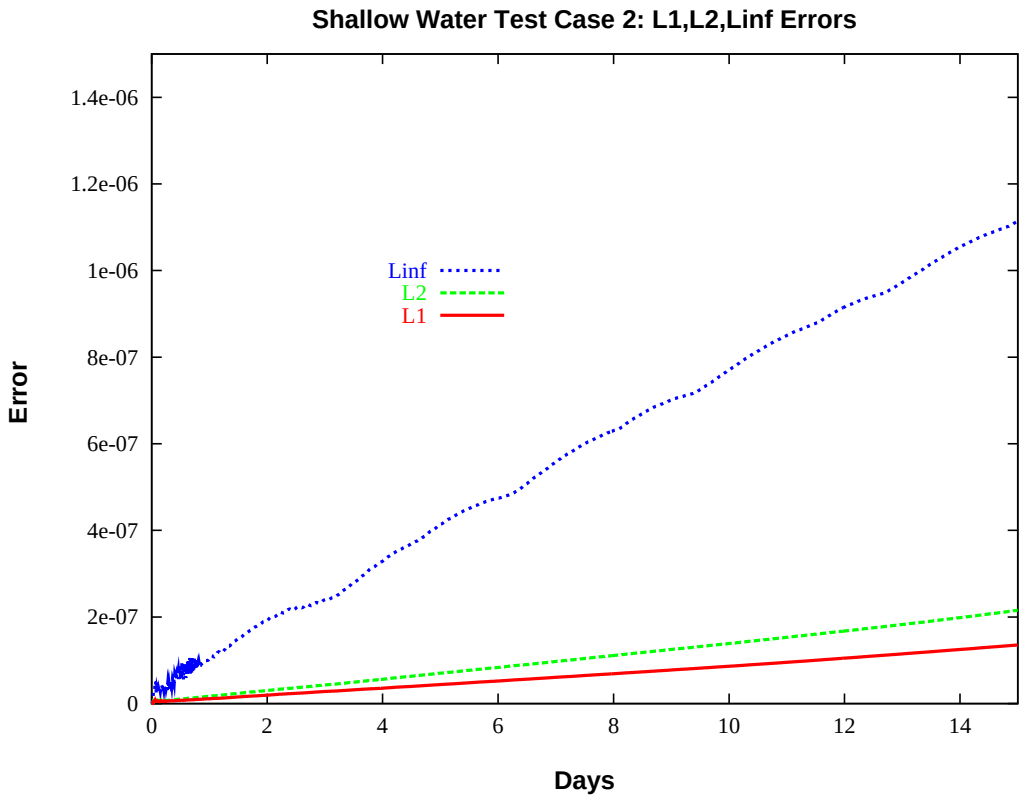

Figure 3: SWE Test Case 2. Zonal geostrophic flow. Semi-implicit model. C56 resolution. l_1, l_2 and l_∞ errors for 15 day integration.

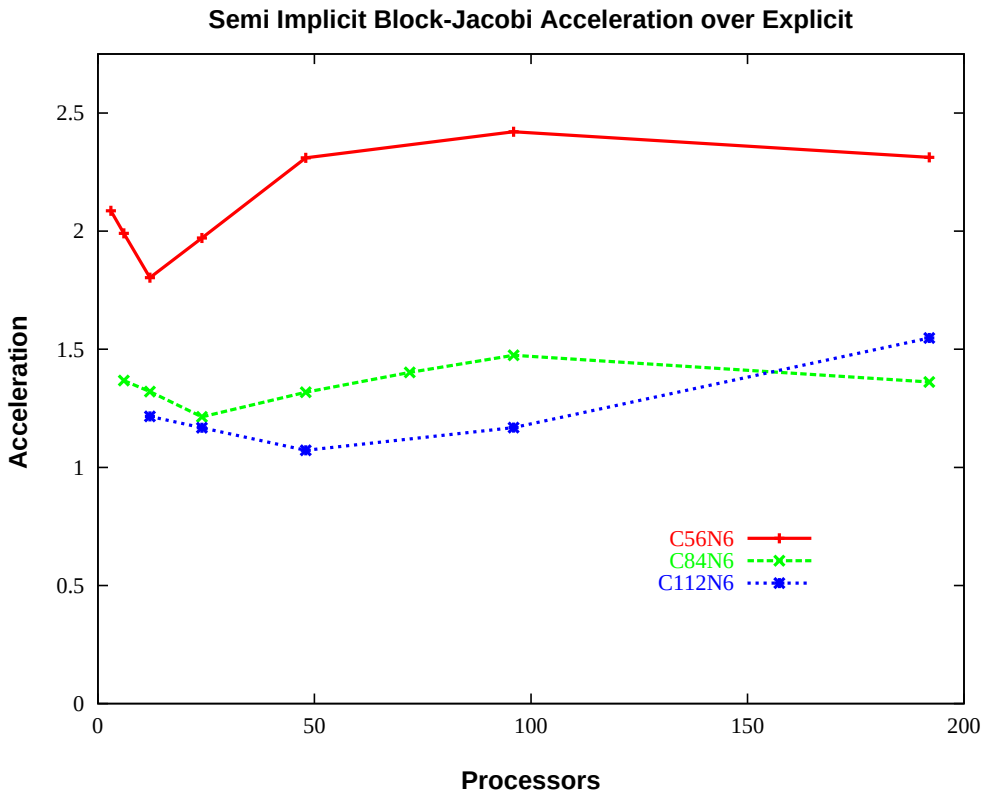

Figure 4: Semi-implicit integration rate acceleration factor over explicit. Block-Jacobi preconditioner. $np = 6$ Gauss pressure points.

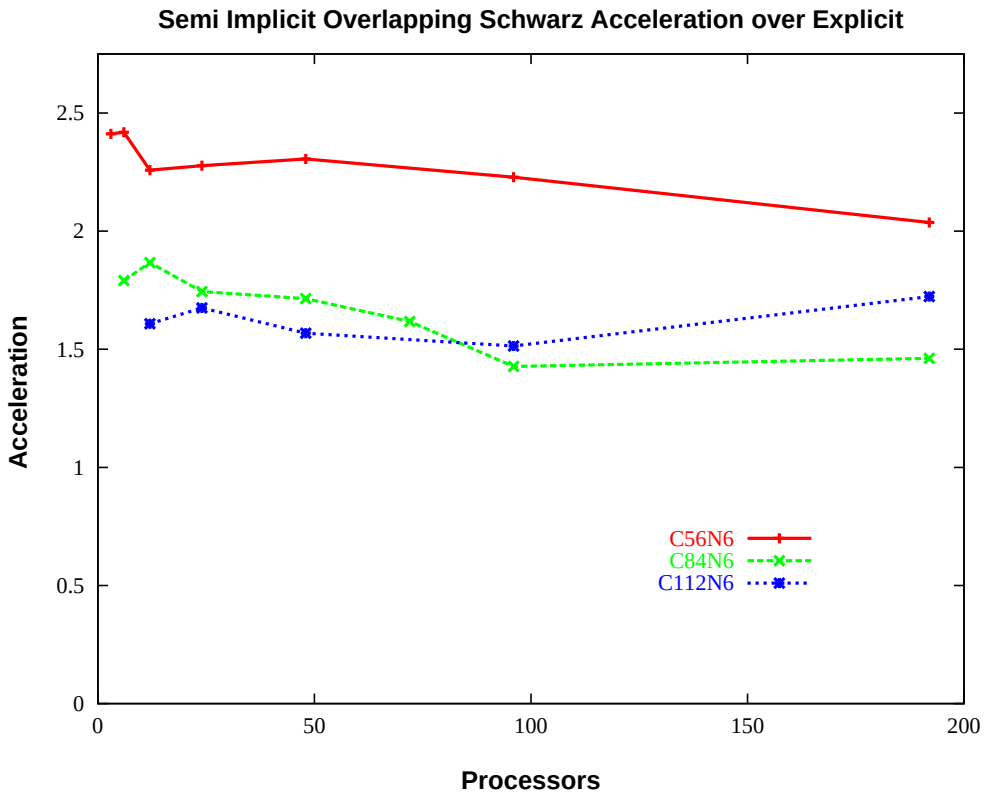

Figure 5: Semi-implicit integration rate acceleration factor over explicit. Overlapping Schwarz FDM-1 preconditioner. $np = 6$ Gauss pressure points.

	l_2	Δt	ε_c	$iters$	$freq$	μ
$ne = 8$						
exp	0.44e-6	100	-	-	1	0.002
BJ	0.10e-6	864	1.0e-12	2.03	1	0.020
FDM-1	0.21e-6	864	1.0e-12	2.00	1	0.020
$ne = 12$						
exp	0.17e-7	72	-	-	1	0.002
BJ	0.57e-8	640	1.0e-13	4.99	1	0.010
FDM-1	0.68e-8	640	1.0e-13	3.99	1	0.010
$ne = 16$						
exp	0.96e-9	50	-	-	1	0.002
BJ	0.10e-8	400	1.0e-13	4.99	1	0.020
FDM-1	0.36e-8	400	1.0e-13	3.99	1	0.020
$ne = 32$						
exp	0.17e-10	25	-	-	1	0.001
BJ	0.17e-10	200	1.0e-13	5.99	1	0.01
FDM-1	0.69e-10	200	1.0e-13	4.99	1	0.01

Table 1: h-refinement for $np = 6$. 15 day integration. l_2 error after 15 days. Time step Δt. CG solver tolerance ε_c. Average CG iterations $iter$. Filter frequency $freq$. Filter weight μ.

	l_2	Δt	ε_c	$iters$	$freq$	μ
$ne = 2$						
exp	0.25e-9	150		-	1	0.005
BJ	0.29e-9	1600	1.0e-13	3.50	1	0.001
FDM-1	0.29e-9	1600	1.0e-13	2.10	1	0.001
$ne = 4$						
exp	0.78e-12	75	-	-	1	0.005
BJ	0.29e-12	800	1.0e-13	4.33	1	0.02
FDM-1	0.21e-11	800	1.0e-13	3.99	1	0.02
$ne = 8$						
exp	0.97e-12	36	-	-	1	0.2
BJ	0.10e-12	400	1.0e-13	5.75	1	0.1
FDM-1	0.14e-12	400	1.0e-13	4.99	1	0.05

Table 2: h-refinement for $np = 12$. 15 day integration.

References

D'Azevedo, E., V. Eijkhout, and C. Romine, 1992: Reducing communication costs in the conjugate gradient algorithm on distributed memory multiprocessors. LAPACK working note 56, University of Tennessee.

Boyd, J. P., 1996: The erfc-log filter and the asymptotics of the Euler and Vandeven sequence accelerations. *Proceeding of the Third International Conference on Spectral and High Order Methods*, ed. A. V. Ilin and L. Ridgway Scott, Houston Journal of Mathematics, 267–276.

Couzy, W., 1995: Spectral element discretization of the unsteady Navier-Stokes equations and its iterative solution on parallel computers. Thèse No. 1380, Ecole Polytechnique Fédérale de Lausanne.

Couzy, W., and M. O. Deville, 1995: A fast Schur complement method for the spectral element discretization of the incompressible Navier-Stokes equations *J. Comp. Phys.*, **116**, 135-142.

Dryja, M., and O. B. Widlund, 1987: An additive variant of the Schwarz alternating method for the case of many subregions. Technical Report 339, Dept. of Comp. Sci., Courant Institute.

Fischer, P. F., N. I. Miller, and H. M. Tufo, 1998: An overlapping Schwarz method for spectral element simulation of three-dimensional incompressible flows. 1997 IMA Domain Decomposition workshop proceedings.

Held, I. H., and M. J. Suarez, 1994: A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bull. Amer. Met. Soc.*, **75**, 1825–1830.

Iskandarani, M., D. B. Haidvogel, and J. P. Boyd, 1995: A staggered spectral element model with application to the oceanic shallow water equations. *Int. J. Numer. Meth. Fluids*, **20**, 394–414.

Jakob-Chien, R., J. J. Hack, and D. L. Williamson, 1995: Spectral transform solutions to the shallow-water test set. *J. Comput. Phys.*, **119**, 164–187.

Loft, R. D., S. J. Thomas, and J. M. Dennis, 2001: Terascale spectral element dynamical core for atmospheric general circulation models. Proceedings of ACM/IEEE Supercomputing 2001.

Lynch, R. E., J. R. Rice, and D. H. Thomas, 1964: Direct solution of partial difference

equations by tensor product methods. *Numer. Math.*, **6**, 185–199.

Rancic, M., R. J. Purser, and F. Mesinger, 1996: A global shallow-water model using an expanded spherical cube: Gnomic versus conformal coordinates. *Q. J. R. Meteorol. Soc.*, **122**, 959–982.

Robert, A. J., 1969: The integration of a spectral model of the atmosphere by the implicit method. WMO/IUGG symposium on NWP. Japan Meteorological Agency, vol. VII, 19–24.

Ronquist, E. M., 1988: *Optimal Spectral Element Methods for the Unsteady Three Dimensional Navier Stokes Equations*, Ph.D Thesis, Massachusetts Institute of Technology, 176p.

Sadourny, R., 1972: Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids. *Mon. Wea. Rev.*, **100**, 136–144.

Smolarkiewicz, P. K., V. Grubišić, and L. G. Margolin, 1997: On forward-in-time differencing for fluids: Stopping criteria for iterative solutions of anelastic pressure equations. *Mon. Wea. Rev.*, **125**, 647–654.

Taylor, M., J. Tribbia, and M. Iskandarani, 1997: The spectral element method for the shallow water equations on the sphere. *J. Comp. Phys.*, **130**, 92–108.

Thomas, S. J. and R. D. Loft, 2002: Semi-implicit spectral element atmospheric model. *J. Sci. Comput.*, in press.

Widlund, O. B., 1992: Some Schwarz methods for symmetric and nonsymmetric elliptic problems. *Fifth Conf. on Domain Decomposition Methods for Partial Differential Equations*, eds. T. .F. Chan, D. E. Keyes, G. A. Meurent, J. S. Scroggs and R. G. Voigt, SIAM Philadelphia PA, 19–36.

Williamson, D. L., J. B. Drake, J. J. Hack, R. Jakob, and P. N. Swarztrauber, 1992: A standard test set for numerical approximations to the shallow water equations in spherical geometry. *J. Comp. Phys.*, **102**, 211–224.