Spectral Preconditioners for Nonhydrostatic Atmospheric Models

STEPHEN J. THOMAS, JOSHUA P. HACKER, AND PIOTR K. SMOLARKIEWICZ

National Center for Atmospheric Research, Boulder, Colorado

ROLAND B. STULL

Department of Earth and Ocean Sciences, University of British Columbia, Vancouver, British Columbia, Canada

(Manuscript received 23 October 2002, in final form 20 March 2003)

ABSTRACT

The elliptic problems in semi-implicit nonhydrostatic atmospheric models are difficult. Typically, they are poorly conditioned, nonseparable, contain cross-derivative terms, and are often nonsymmetric. Here, the resulting linear system is solved using a preconditioned Krylov subspace method—the generalized conjugate residual (GCR) algorithm. A horizontal spectral preconditioner is developed as an alternative to a more standard and much simpler line Jacobi relaxation scheme. However, the efficacy of the spectral preconditioner requires neglecting the cross-derivative terms and homogenization (e.g., averaging) metric coefficients over the computational domain. Because such a compromise causes a substantial departure of the preconditioner from the original elliptic operator, it is not obvious a priori whether it leads to a competitive solver. The robustness of the proposed approach over a broad range of representative meteorological applications is evaluated, in the context of a three-time-level semi-implicit semi-Lagrangian all-scale weather-prediction/research model.

1. Introduction

A semi-implicit time-stepping scheme (Robert 1969) is commonly applied to the terms responsible for fast waves in large-scale weather prediction and general circulation models. This removes the associated time step restrictions and makes computation practical. More recently, semi-implicit integrators have also been applied at mesoscale resolutions (Tapp and White 1976; Cullen 1990; Tanguay et al. 1990; Golding 1992; Yeh et al. 2002; Grabowski and Smolarkiewicz 2002) and convective scales (Robert 1993; Grabowski and Smolarkiewicz 2002). The resulting elliptic partial differential equations are nontrivial. They are poorly conditioned, nonseparable, contain cross-derivative terms, and are typically nonsymmetric. This is due to domain anisotropy, effects of planetary rotation, ambient stratification, the use of general curvilinear coordinates in the governing equations [e.g., the Gal-Chen and Somerville (1975) terrain-following transformation], or the imposition of partial-slip conditions along an irregular lower boundary. Among the most effective methods reported for solving such problems are the preconditioned nonsymmetric conjugate-gradient-type (alias Krylov subspace, hereafter Krylov) iterative schemes (e.g., Kapitza and Eppel 1992; Skamarock et al. 1997; Thomas et al.

1998, 2000). A number of alternative nonsymmetric Krylov solvers are used in computational research and engineering (Axelsson 1994; Greenbaum 1997). Although this paper focuses on a particular scheme, the issues discussed are relevant to other iterative methods as well. Our method of choice is the restarted generalized conjugate residual [GCR(k)] algorithm¹ akin to the popular generalized minimum residual [GMRES(k)] solver (Eisenstat et al. 1983; Saad and Schultz 1986; Smolarkiewicz and Margolin 1994, 1997; Smolarkiewicz et al. 1997).² For the reader's convenience, GCR(k) is summarized in appendix A, where we also introduce some terminology, notation, and a notion of preconditioning (left, as opposed to right) used throughout this paper. A brief discussion of line-relaxation preconditioners is included in appendix B.

Designing a suitable preconditioner is important, as it can dramatically accelerate solver convergence, thereby reducing the overall computational expense of a model. In principle, the preconditioner $\mathcal{P}$ can be any linear operator such that $\mathcal{L}\mathcal{P}^{-1}$ is definite, where $\mathcal{L}$ symbolizes the original elliptic operator implied by the physical problem at hand. The role of the preconditioner is

¹ Throughout this paper GCR (4) is employed unless stated otherwise.

² Other popular alternatives include the biconjugate gradients stabilized (BiCGStab; van der Vorst 1992) and transpose-free quasi-minimal residual (TFQMR; Freund 1993) algorithms.

Corresponding author address: Dr. Stephen J. Thomas, NCAR/SCD, P.O. Box 3000, Boulder, CO 80307-3000.
E-mail: thomas@ucar.edu

to substitute the governing elliptic problem $\mathcal{L}(\Psi) - Q = 0$ with an auxiliary problem $\mathcal{P}^{-1}[\mathcal{L}(\Psi) - Q] = 0$ that converges faster (than the original problem) because of a closer clustering of the eigenvalues of the auxiliary elliptic operator $\mathcal{P}^{-1}\mathcal{L}$, where Ψ and Q symbolize the dependent variable and the rhs, respectively. For the preconditioner to be useful, the convergence of the auxiliary problem must be sufficiently rapid to overcome the effort associated with “inverting” the preconditioner itself [i.e., computing $\mathcal{P}^{-1}(\cdot)$]. In general, the closer the preconditioner approximates the original operator, the faster the solver converges but the more difficult it is to compute $\mathcal{P}^{-1}(\cdot)$.

There is no general method for designing an optimal preconditioner (Axelsson 1994, section 7). Kadoglu and Mudrick (1992) argued that basic point-Jacobi relaxation provides an effective preconditioner for GMRES(k) applied to atmospheric flow problems. However, their conclusion is derived from solving the diagnostic quasigeostrophic omega equation in a periodic channel—a fairly simple elliptic problem that does not reflect the full complexity of a multiscale atmosphere. In particular, it does not extrapolate to prognostic problems when the shallowness of the earth’s atmosphere dictates condition numbers $\kappa(\mathcal{L}) \sim O(10^{10})$; recall that $\kappa(\mathcal{L})$ can be interpreted as the squared ratio of the longest to the shortest wave present in the system. Because the asymptotic convergence rate of conjugate-gradient-type methods is proportional to $\kappa(\mathcal{L})^{-1/2}$, a direct preconditioner in the vertical is the “categorical imperative” of an effective iterative solver for all-scale atmospheric models (e.g., Marshall et al. 1997; Skamarock et al. 1997; Thomas et al. 1998; Smolarkiewicz et al. 2001). In this paper we take this preconditioning strategy in the vertical for granted and focus on the optimization of the horizontal component of $\mathcal{P}$.

Guided by the development of elliptic solvers for early anelastic models, we extend a fast direct preconditioning strategy to three spatial dimensions via horizontal spectral decomposition. This offers an alternative to alternating-direction-implicit (ADI) iterative methods (Skamarock et al. 1997), of which the line Jacobi preconditioner is a special case (appendix B). The use of the FFT approach has a long history in the design of elliptic solvers. Because FFTs are effective for separable problems, early elliptic solvers for meteorological models often employed operator decomposition (into a separable part and the residual) and a *stationary* block-iteration with the nonseparable residual lagging behind (cf. Schumann and Volkert 1987 and Schumann and Sweet 1988, and references therein). Bernardet (1995) investigated an alternative approach (conceptually the opposite) where the full operator is retained in a *non-stationary* Krylov iteration, but FFTs are used to invert the “flat” preconditioner obtained by neglecting orography. He used the classical, symmetric, conjugate-gradient solver of Hestenes and Stiefel (1952), thereby necessitating careful modifications to the boundary con-

ditions (in order to assure a self-adjoint elliptic operator). His development addressed 2D mesoscale problems in the x - z plane. Because of the inherent difficulties with extending this approach to 3D, Lafore et al. (1998) adopted an FFT-preconditioned Richardson iteration. Recently, Elman and O’Leary (1998) analyzed the performance of nonsymmetric Krylov solvers for indefinite 3D Helmholtz equations with Sommerfeld radiation boundary conditions and found that FFT methods can result in effective preconditioners even on parallel architectures. Because of the indefiniteness, they addressed a more difficult problem than those typically encountered in meteorological applications. On the other hand, their Helmholtz operator contained the standard Laplacian on isotropic grids and had constant coefficients. Both the coefficients and solutions were smooth.

Although there exists a substantial body of evidence indicating the potential of FFT preconditioners, a systematic study in the context of all-scale meteorological applications is needed. This is because modeling natural flows leads to elliptic problems characterized by highly inhomogeneous anisotropic coefficients and full spectra forced on the rhs. We complement earlier work with a comprehensive study based on a fully compressible non-hydrostatic all-scale atmospheric model (Thomas et al. 1998). A semi-implicit time discretization leads to a nonsymmetric, but definite, Helmholtz problem. In order to design an effective FFT preconditioner, we derive a separable constant-coefficient approximation to the semi-implicit Helmholtz equation by dropping cross-derivative terms and averaging metric coefficients across the domain. Neumann boundary conditions are specified to close the problem, and thus a discrete cosine transform (DCT) is required to diagonalize the discrete horizontal Laplacian component of the preconditioner.

Robustness is vital for the elliptic solver when a numerical model admits a broad range of scales of motion. We assess the relative efficiency of line relaxation versus DCT preconditioners in both linear and nonlinear flow regimes (prototypes for, respectively, laminar and turbulent flows) throughout the range $O(10^1)$ – $O(10^7)$ m of spatial scales. Our study encompasses the simulation of shallow thermal convection, small and large mesoscale mountain flows, and synoptic-scale winter storm prediction. We measure the performance of the two preconditioners by comparing the amount of work (the iteration count and total CPU time of the model) required to achieve a specified convergence threshold. Because the residual error of the elliptic problem has a sense of Δt -normalized first-order spatial partial derivatives of velocity, the convergence criterion is based on the magnitude of the local Courant number and its variations. This specifies both the absolute admissible error and its magnitude compared to the flow. In effect, both numerical and physical error metrics are employed (cf. Smolarkiewicz et al. 1997).

The paper is organized as follows: Section 2 summarizes the governing equations of the model along with

their discretization and formulation of the elliptic problem. Section 3 presents the derivation of a spectral preconditioner. Section 4 discusses the results of our simulations. Remarks in section 5 conclude the paper.

2. Model description

a. Analytic formulation

To focus the discussion, consider the standard compressible Euler model for a dry, adiabatic, stratified, rotating fluid. Cast in a noninertial frame rotating with angular velocity $\mathbf{\Omega}$, it consists of mass, momentum, and entropy evolution laws, and the equation of state:

$$\begin{aligned} \frac{D\rho}{Dt} &= -\rho \nabla \cdot \mathbf{v}, \\ \frac{D\mathbf{v}}{Dt} &= -\frac{1}{\rho} \nabla p - \mathbf{g} - 2\mathbf{\Omega} \times \mathbf{v}, \\ \frac{D\Theta}{Dt} &= 0, \quad p = \rho RT, \end{aligned} \quad (1)$$

where $D/Dt = \partial/\partial t + \mathbf{v} \cdot \nabla$ is the material-derivative operator, $\mathbf{v}$ the velocity vector, p thermodynamic pressure, ρ density, and $\mathbf{g}$ combines gravitational and centrifugal accelerations. Also, $\Theta = T \exp(-\kappa q)$ is the potential temperature, where T denotes the temperature, $\kappa = R/c_p$ (with R and c_p being the gas constant and specific heat at constant pressure), and $q = \ln(p/p_{00})$, with $p_{00} = 100$ kPa.³

In order to isolate the fast acoustic-gravity modes (for the sake of implicit numerical treatment), the prognostic equations in (1) are decomposed by introducing a horizontally homogeneous, static reference atmosphere (hereafter denoted with subscript 0). Thermodynamic variables are decomposed into a sum of the base state and perturbations

$$T = T_0 + T', \quad q = q_0(z) + q', \quad (2)$$

where

$$\frac{\partial q_0}{\partial z} = -\frac{g}{RT_0}. \quad (3)$$

After grouping terms associated with “fast” linear waves on the lhs, while retaining the “slow” residual on the rhs and eliminating ρ using the gas law, the governing evolution equations take the form

$$\frac{1}{\gamma} \frac{Dq'}{Dt} + \nabla \cdot \mathbf{v} - \frac{g}{c_0^2} w = 0,$$

³ Using the auxiliary “pressure” variable q , rather than the pressure p itself, facilitates isolating the $1/\rho \nabla p$ nonlinearity in the numerical-model equations, analogous to $q = \ln p_s$ (surface pressure p_s) in hydrostatic primitive equation models (cf. Machenhauer and Daley 1972; Bourke 1974; Robert et al. (1985; Tanguay et al. 1990). The Exner function $\Pi = (p/p_{00})^{\kappa}$ is an alternative (Skamarock et al. 1997).

$$\frac{D\mathbf{v}}{Dt} + RT_0 \nabla q' - B\mathbf{k} = -RT' \nabla q' - f\mathbf{k} \times \mathbf{v},$$

$$\frac{DT'}{Dt} - \frac{RT_0}{c_p} \frac{Dq'}{Dt} + T_0 \frac{N_0^2}{g} w = -\frac{RT'}{c_v} \nabla \cdot \mathbf{v}, \quad (4)$$

where $B = gT'/T_0$ denotes the buoyancy, and $f = 2\Omega \sin\phi$ is the Coriolis parameter at latitude ϕ . Also, $\gamma = c_p/c_v$ (with $c_v = c_p - R$ denoting the specific heat at constant volume), and N_0 and c_0 are, respectively, the Brunt–Väisälä frequency and speed of sound:

$$N_0^2 = \frac{g^2}{c_p T_0}, \quad c_0^2 = \gamma RT_0. \quad (5)$$

Equations (1) and (4), as written, are valid in a rotating Cartesian reference frame. To account for the curvature of the earth when simulating atmospheric flows and produce results consistent with standard weather maps, Eqs. (4) are written in terms of a Cartesian coordinate system projected onto the plane intersecting the sphere at reference latitude ϕ_0 (cf. Tanguay et al. 1990). The resulting equations take the form

$$\begin{aligned} \frac{1}{\gamma} \frac{Dq'}{Dt} + D_H(U, V) + \frac{\partial q}{\partial z} - \frac{g}{c_0^2} w &= 0, \\ \frac{DU}{Dt} + RT_0 \frac{\partial q'}{\partial X} &= fV - K \frac{\partial S}{\partial X} - RT' \frac{\partial q'}{\partial X}, \\ \frac{DV}{Dt} + RT_0 \frac{\partial q'}{\partial Y} &= -fU - K \frac{\partial S}{\partial Y} - RT' \frac{\partial q'}{\partial Y}, \\ \frac{Dw}{Dt} + RT_0 \frac{\partial q'}{\partial z} - B &= -RT' \frac{\partial q'}{\partial z}, \\ \frac{DT'}{Dt} - \frac{RT_0}{c_p} \frac{Dq'}{Dt} + T_0 \frac{N_0^2}{g} w &= -\frac{RT'}{c_v} D, \end{aligned} \quad (6)$$

where the (X, Y) conformal coordinate-based form of D/Dt , horizontal divergence $D_H(U, V)$ of the wind image (U, V) , pseudokinetic energy K , and metric coefficient S are all detailed in appendix C for the sake of completeness. To account for the irregularity of the lower boundary, the system in (6) is transformed to the terrain-following curvilinear framework of Gal-Chen and Somerville (1975). Optionally, variable spacing of vertical levels is allowed by composing the Gal-Chen transformation with a smooth stretching (Thomas et al. 1998). Rather than complicating (6) with details of the composition of these mappings, all metric terms associated with the vertical transformations are absorbed into the definitions of the partial derivatives (appendix C).

b. Numerical approximations

The semi-Lagrangian approach employed for approximating (6) on a discrete mesh (Robert 1993) aims

at second-order accuracy in space and time.⁴ A staggered Arakawa C grid (Arakawa and Lamb 1977) is employed in the horizontal, and a staggered Tokioka B grid (Tokioka 1978) in the vertical. The associated spatial partial-derivative discrete operators are detailed in appendix D.

In general, the governing equations—here, either (1), (4), or (6)—can be viewed symbolically as

$$\frac{D\psi}{Dt} = F_\psi, \quad (7)$$

where ψ stands for any of the kinematic or thermodynamic dependent variables of the model, and F_ψ for the associated rhs. Then, semi-Lagrangian schemes can be interpreted as approximations to the trajectory integrals of the governing equations:

$$\psi = \psi_0 + \int_{\Gamma} F_\psi d\tau, \quad (8)$$

where ψ and ψ_0 denote, respectively, the values of $\psi(\mathbf{X}, t)$ at the two endpoints of the trajectory Γ connecting $[\mathbf{X}_0(\mathbf{X}, t_1), t_0]$ and $(\mathbf{X}, t_1)$. For our study (of spectral preconditioners), we adopt the Mesoscale Compressible Community (MC2) model (Tanguay et al. 1990; Benoit et al. 1997; Laprise et al. 1997; Thomas et al. 1998, and references therein), which implements the semi-implicit semi-Lagrangian algorithm of Robert (1981). In the context of the integral representation, this scheme blends midpoint three-time-level and trapezoidal two-time-level quadrature rules. The midpoint and trapezoidal approximations result, respectively, in explicit (leapfrog) and implicit (Crank–Nicholson) time discretizations of the various forcing terms. In order to concisely present the resulting finite-difference equations, we define the discrete operators

$$D_t\psi = \frac{\psi_i^{n+1} - \psi_0^{n-1}}{2\Delta t}, \quad \mu_t\psi = \frac{\psi_i^{n+1} + \psi_0^{n-1}}{2}, \quad (9)$$

where ψ_i^{n+1} denotes the unknown (sought) value of ψ at the grid point $(\mathbf{X}_i, t^{n+1})$, and ψ_0^{n-1} is the value of ψ at the foot $(\mathbf{X}_0, t^{n-1})$ of the trajectory arriving at $(\mathbf{X}_i, t^{n+1})$, evaluated using cubic interpolation from the neighboring grid points. With the definitions in (9), the compact discretized form of the equations (6) becomes (after some regrouping)

$$\begin{aligned} D_t q' + \gamma \mu_t D_H(U, V) + \gamma \left[D_Z - \frac{g}{c_0^2} \mu_Z \right] \mu_t w &= F_q^n, \\ D_t U + RT_0 \mu_t D_X q' &= F_U^n, \\ D_t V + RT_0 \mu_t D_Y q' &= F_V^n, \\ D_t w + RT_0 \mu_t D_Z q' - \mu_t B &= F_w^n, \end{aligned}$$

$$D_t T' - \frac{RT_0}{c_p} D_t q' + T_0 \frac{N_0^2}{g} \mu_t w = F_T^n, \quad (10)$$

where spatial partial-derivative and average operators (D_X , D_Y , D_Z , and μ_Z) are defined in appendix D, and the F_ψ^n forcings combine the Coriolis and slow-residual terms that are interpolated at the trajectory midpoint, $(\mathbf{X}, t^n)$, with $\mathbf{X} \equiv (\mathbf{X}_i + \mathbf{X}_0)/2$ [cf. Thomas et al. (1998) for a discussion of the trajectory scheme)].

c. Elliptic problem

In order to derive the discrete Helmholtz problem that underlies the semi-implicit scheme in (10), all *unknown* t^{n+1} terms are grouped on the lhs, and all *known* t^n and t^{n-1} terms are separated on the rhs. Denoting the rhs by Q and dropping all temporal superscripts (as there is no ambiguity at this point), the resulting linear system of algebraic equations (valid at each grid point) is written as follows:

$$q' + \Delta t \gamma D_H(U, V) + \Delta t \gamma \left[D_Z - \frac{g}{c_0^2} \mu_Z \right] w = Q_{q'}, \quad (11)$$

$$U + \Delta t RT_0 D_X q' = Q_U, \quad (12)$$

$$V + \Delta t RT_0 D_Y q' = Q_V, \quad (13)$$

$$w + \Delta t RT_0 D_Z q' - \Delta t B = Q_w, \quad (14)$$

$$T' - \frac{RT_0}{c_p} q' + \Delta t T_0 \frac{N_0^2}{g} w = Q_T. \quad (15)$$

To arrive at the implicit equation for q' , all remaining dependent variables are eliminated from (11)–(15). Combining (14) and (15) eliminates the buoyancy B and results in

$$[1 + \Delta t^2 N_0^2] w + \Delta t RT_0 \left[D_Z - \frac{g}{c_p T_0} \mu_Z \right] q' = Q'_w, \quad (16)$$

where

$$Q'_w = Q_w + \Delta t \frac{g}{T_0} Q_T. \quad (17)$$

Next, taking the divergence of (12) and (13) and substituting for $D_H(U, V)$ in (11),

$$[1 - \Delta t^2 c_0^2 D_H(D_X, D_Y)] q' + \Delta t \gamma \left[D_Z - \frac{g}{c_0^2} \mu_Z \right] w = Q'_{q'}, \quad (18)$$

where

$$Q'_{q'} = Q_{q'} - \Delta t \gamma D_H(Q_U, Q_V). \quad (19)$$

Following the original notation of Tanguay et al. (1990), we define the operators,

$$D_Z^{(1)} = \left[D_Z - \frac{g}{c_p T_0} \mu_Z \right] \quad \text{and}$$

⁴ Numerical dissipative terms absent from the analytic model (1)—absorbers, filters, subgrid-scale turbulence parameterizations, etc.—are split and integrated to first order.

$$D_Z^{(2)} = \left[D_Z - \frac{g}{c_0^2} \mu_Z \right], \quad (20)$$

and the coefficient $\mathcal{N} = (1 + \Delta t^2 N_0^2)$. Then, solving directly for w in (16) and substituting the result in (18), we arrive at the discrete Helmholtz equation,

$$\{[1 - \Delta t^2 c_0^2 D_H(D_X, D_Y)] - \Delta t^2 c_0^2 D_Z^{(2)} \mathcal{N}^{-1} D_Z^{(1)}\} q' = Q_q^*, \quad (21)$$

where the modified forcing on the rhs takes the form

$$Q_q^* = Q'_{q'} - \Delta t D_Z^{(2)} \mathcal{N}^{-1} Q'_w. \quad (22)$$

d. Boundary conditions

Lateral boundaries are open with inflow/outflow determined by the normal components of the velocity. From the perspective of the Helmholtz problem in (21), normal velocities specified at the lateral boundaries imply Neumann boundary conditions on the pressure via (12) and (13). The model top and bottom boundaries are impermeable. At the model top z_{top} , the terrain-following coordinate Z becomes flat [cf. (C7) in appendix C], whereupon the contravariant “vertical” velocity becomes identical to the vertical component of the physical velocity and they both vanish:

$$\dot{Z}|_{z_{\text{top}}} = w|_{z_{\text{top}}} = 0. \quad (23)$$

At the lower boundary, the normal contravariant velocity vanishes ($\dot{Z}|_0 = 0$), implying that the physical vertical velocity w follows coordinate surfaces according to

$$w|_0 = GS(G^{13}U + G^{23}V)|_0. \quad (24)$$

The boundary conditions on q' associated with (23) and (24) are derived by substituting the momentum equations (12) and (13) into, respectively, (23) and (24) and requiring that the resulting relations be consistent with (16). Then Q'_w and $D_Z^{(1)}$ can be specified at $Z = 0$ and $Z = z_{\text{top}}$.⁵

3. Spectral preconditioner

The discrete Helmholtz equation (21) takes the form

$$\mathcal{L}q' = Q_q^*, \quad (25)$$

referred to early in the introduction. The symbol $\mathcal{L}$ on the lhs of (25) represents the entire operator appearing between the brackets $\{\cdot\}$ on the lhs of (21). For the terrain-following transformation and conformal coordinates, the generalized Laplacian-like operator $D_H(D_X, D_Y)$ embedded in $\mathcal{L}$ includes cross-derivative terms and is nonseparable. The discrete horizontal divergence and 3D gradient, (D7) and (D4) respectively, couple the hor-

izontal to the vertical operators through a complex blend of X - Z and Y - Z derivatives and averages. The first step in constructing an effective preconditioner $\mathcal{P}$ (for atmospheric models) is to neglect these cross terms—to facilitate a direct inversion of $\mathcal{P}$ in the vertical (recall the discussion in the introduction). The resulting simplification leads to the preconditioned linear system of equations:

$$\mathcal{P}e = r \quad (26)$$

or, equivalently,

$$\{[1 - \Delta t^2 c_0^2 S \nabla_H^2] - \Delta t^2 c_0^2 D_Z^{(2)} \mathcal{N}^{-1} D_Z^{(1)}\} e = r, \quad (27)$$

where $\nabla_H^2 = \delta_{xx} + \delta_{yy}$ denotes the standard discrete horizontal Laplacian; e is the current iterate (of the Krylov solver) estimation of the solution error $e = q' - \bar{q}'$, with the overbar denoting the sought exact solution; and r is the current iterate value of the residual error $r = \mathcal{L}q' - Q_q^*$ (cf. appendix A). Because the operator on the lhs of (27) is a simple sum of the vertical and horizontal differential operators, iterative line-relaxation preconditioners (outlined in appendix B) can be easily applied at this point. Here, we advocate an alternate strategy, where a separable (in the horizontal) constant-coefficient approximation of the operator on the lhs of (27) is specified—thereby departing even further with $\mathcal{P}$ from the original operator $\mathcal{L}$ in (25)—enabling a fast direct inversion in the horizontal.

The horizontal component of $\mathcal{P}$ in (27) is nonseparable because the metric coefficient $S = S(X, Y)$ is a function of both X and Y .⁶ In general, a separable operator is obtained by a cross-directional homogenization of metric coefficients—for example, by averaging S that multiplies the δ_{xx} and δ_{yy} parts of ∇_H^2 , respectively, in the Y and X directions—whereupon the tensor-product method (Lynch et al. 1964) can be applied to construct a direct solver. This requires computing separately the eigenvalues and eigenvectors of the matrices representing the discrete operators $\bar{S}^Y \delta_{xx}$ and $\bar{S}^X \delta_{yy}$, and diagonalization follows by combining two unitary similarity transformations. Insofar as stand-alone solvers for separable problems are concerned, this approach requires $O(n_z n^3)$ operations and is not competitive with ADI schemes [an $O(n_z n^2 \log^2 n)$ operation count] except when a fast transform [e.g., FFT with $O(n_z n^2 \log n)$ operations] is available (Lynch et al. 1964). Here, $n = n_x = n_y$ refers to the number of mesh points in either horizontal direction, and n_z is the number of points in the vertical. Because (27) is merely a preconditioner, S can be globally averaged or extremized (Bernardet 1995; Smolarkiewicz and Margolin 2000), thereby enabling application of the FFT. Here, we replace S in (27) with the global mean; flat topography is specified to assure

⁵ The reader interested in implicit formulations of boundary conditions in the context of Krylov solvers is referred to Smolarkiewicz and Margolin (2000) for a discussion.

⁶ In models with X - Y cross differentiation [e.g., due to the coordinate mapping for the sake of mesh adaptivity (Smolarkiewicz and Prusa 2002)] either ADI or spectral preconditioners would require neglecting these cross terms.

constant coefficients in the vertical derivatives. The discrete operator ∇_H^2 , with *homogeneous* Neumann boundary conditions,⁷ is then readily diagonalized by the real part of the Fourier transform (via the discrete cosine transform; Sweet 1973). This results in $n_x \times n_y$ independent tridiagonal vertical problems:

$$\{[1 - \Delta t^2 c_0^2 \bar{S}^{xy}(\lambda_l + \lambda_m)] - \Delta t^2 c_0^2 D_Z^{(2)} \mathcal{K}^{-1} D_Z^{(1)}\} \hat{e} = \hat{f}, \quad (28)$$

where

$$\lambda_l = -\frac{4}{\Delta X^2} \sin^2 \left[\frac{\pi(l-1)}{2n_x} \right], \quad l = 1, \dots, n_x,$$

$$\lambda_m = -\frac{4}{\Delta Y^2} \sin^2 \left[\frac{\pi(m-1)}{2n_y} \right], \quad m = 1, \dots, n_y,$$

are the eigenvalues of ∇_H^2 in the X and Y directions, and $\hat{\cdot}$ denotes Fourier transform.

4. Numerical experiments

A suite of model runs⁸ is designed to compare the performance of the DCT preconditioner against the line Jacobi preconditioner, proven effective in meteorological applications (Thomas et al. 1998, 2000). The relative efficiency of the two preconditioners is measured with two Jacobi/DCT ratios [hereafter the performance ratios (PR)]: the ratio of total GCR solver iterations and the ratio of the total model CPU times. Both ratios are informative. For a given convergence criterion, the first one estimates the computational cost of the solver, independent of the implementation and machine, whereas the second assesses the total work required for the application at hand. Values of $PR > 1$ show efficiency gains by the DCT preconditioner. Because the residual error of the Helmholtz problem in (21) combines Δt -normalized first-order spatial partial derivatives of the velocity components, it motivates the stopping criterion:

$$\|Lq' - Q_q^*\| \leq \varepsilon \min(C, L). \quad (29)$$

In (29), C and L represent the Courant $\|\Delta t \mathbf{v} / \Delta \mathbf{X}\|$ and Lipschitz $\|\Delta t (\partial \mathbf{v} / \partial \mathbf{x})\|$ numbers—correspondingly, the measures of the flow magnitude and its variations. Typically, in order to control the stability/accuracy of the Eulerian and semi-Lagrangian computations, $C \sim L \sim O(1)$, whereupon ε specifies both the absolute admissible error and its magnitude compared to the flow.

The test cases considered here are designed to evaluate the robustness and efficiency of the DCT preconditioner over a wide range of meteorological applica-

tions. The first test is very special; it probes the *reflexivity* of our preconditioned Krylov approach. In all problems where the preconditioner is an exact solver, GCR should converge in one iteration to machine precision, regardless of the value of ε . In the MC2 model, all small-to-mesoscale problems with map scale factor $m = 1$ in (C1) and flat lower boundary with $h(X, Y) \equiv \text{const}$ in (C7) belong to this class, which constitutes a broad range of research applications including studies of fully developed turbulence and deep moist convection. For illustration, we consider a full-spectrum turbulent thermal convection simulation.

The physical scenario is a buoyant thermal rising in a neutrally stratified environment—a 3D extension of the earlier experiment in Smolarkiewicz and Pudykiewicz (1992). The grid employed consists of $100 \times 100 \times 150$ points with constant interval $\Delta X = \Delta Y = \Delta Z = 10$ m; $\Delta t = 2.5$ s. Initially, a spherical temperature anomaly of radius $r_0 = 250$ m and $\Theta' = 0.5$ K is centered in the horizontal, $r_0 + \Delta Z$ above the lower boundary. Figure 1 shows potential temperature (Θ) xz and xy cross sections through the center of the convective bubble and at $z = 800$ m, respectively; at the end of the integration, $t = 180\Delta t$, by which time the flow becomes turbulent and exhibits a full range of scales (see Fig. 2).

The solution in Figs. 1 and 2 is for the DCT preconditioner. It is independent of the stopping criterion in (29). This is documented in Fig. 3, which compares the rms divergence of the solution using the DCT and line Jacobi preconditioners for a range of ε . Following Skamarock et al. (1997), this diagnostic is used to measure the solution convergence in elastic models. For the Jacobi preconditioner, as $\varepsilon \rightarrow 0$ the curve flattens, indicating the critical value $\varepsilon = 10^{-3}$, below which the solution can be considered physically converged; that is, further solver iterations do not contribute to the overall accuracy of the results. Because the DCT preconditioner is a direct solver of the elliptic problem at hand, convergence is reached in a single GCR iteration regardless of the specified value of ε . The latter is not the case for the Jacobi preconditioner, which requires an increasing number of GCR iterations as $\varepsilon \rightarrow 0$. Figure 4 demonstrates this by displaying the PR for iteration count and total CPU time as a function of ε . Here, the DCT-preconditioned solver converges in one iteration, and the PR iteration count is equivalent to the average number of GCR iterations (per model time step) using the Jacobi preconditioner. The PR CPU time documents that, even with the overhead associated with computing FFTs, the DCT preconditioner is at least an order of magnitude more efficient for all values of ε representing the convergent solutions.

The second test simulates a smooth, vertically propagating 3D mountain wave, identical to those in Saito et al. (1998) and Thomas et al. (2000). The resulting flow is characterized by small-amplitude gravity wave response at Froude number $F \equiv U/Nh = 4.4$, where the

⁷ Nonperiodic boundary conditions are always homogeneous in the preconditioner (Smolarkiewicz and Margolin 2000, section 3).

⁸ All calculations reported were performed with 64-bit real arithmetic.

FIG. 1. Potential temperature θ (a) x - z and (b) x - y cross sections through the center of the convective bubble at $t = 450$ s. The xy cross sections are at $z = 800$ m, shown by the horizontal dashed line in (a). Contour interval is 0.1 K.

linear theory of Smith (1980) is qualitatively valid. To contrast with the buoyant thermal case, the linear flow regime is chosen because it emphasizes a low signal-to-noise ratio. Here, unlike in the previous case, the energy spectrum of a fully developed solution is compact, and the grid aspect ratio significantly departs from unity, reflecting the anisotropy of the problem. To quantify the effect of domain anisotropy (shallow versus deep flow scenarios) on solver performance, both hydrostatic and nonhydrostatic variants are addressed [cf.

Skamarock et al. (1997) for a discussion]. The parameters of the experiment are as follows.

The bell-shaped mountain,

$$h(X, Y) = \frac{h_0}{[1 + (X/a)^2 + (Y/a)^2]^{3/2}},$$

with the half-width $a = 3\Delta X = 3\Delta Y$ and the height $h_0 = 100$ m is centered at the (30, 30) grid point of the 61×61 horizontal mesh. In the vertical, 31 exponentially stretched levels extend to 19 km, resulting in ver-

FIG. 2. Spanwise (a) x -average and (b) y -average 1D power spectra at three height levels of the solution in Fig. 1. The thin solid line shows the “ $-5/3$ ” slope for reference.

FIG. 3. The rms divergence as a function of ε in Eq. (29) for the convective bubble case shown in Fig. 1. Results with both the Jacobi (solid line) and the DCT (dashed line) preconditioners are shown.

tical grid increments ranging from 40 to 1180 m, respectively, for the bottom and top thermodynamic layers. The model is integrated over 480 time steps, when the solution is considered to have reached an approximate steady state. A constant inflow velocity of $U = 8 \text{ m s}^{-1}$ and an isothermal atmosphere with buoyancy frequency $N = 0.018 \text{ s}^{-1}$ are assumed. The Coriolis force and the curvature of the earth are neglected. No subgrid-scale parameterizations are employed. The Davies (1976) relaxation scheme is used at the lateral boundaries to minimize wave reflections.

In the hydrostatic case, the grid length $\Delta X = 2000 \text{ m}$ and time step $\Delta t = 60 \text{ s}$ result in the advective Courant number $C = 0.24$. The horizontal scale of the mountain is $a = 6000 \text{ m}$, so $aN/U = 13.5$. In the nonhydrostatic case, $\Delta X = 400 \text{ m}$ and $\Delta t = 15 \text{ s}$, giving $C = 0.3$. The horizontal scale decreases to $a = 1200 \text{ m}$, so that $aN/U = 2.7$, whereupon nonhydrostatic effects are significant. For illustration, Fig. 5 displays the converged solution. Vertical velocity cross sections are shown through the vertical center plane and on the horizontal plane at $z = 2100 \text{ m}$ ($3/4$ of the vertical wavelength) for both cases. The maximum vertical velocities at this level are 0.06 and 0.2 m s^{-1} in the hydrostatic and nonhydrostatic cases, respectively. Note the downstream tilt of the wave packet (manifested by the downstream shift of the maxima) in the nonhydrostatic case. The corresponding energy spectra in the mean flow direction at $z = 1800 \text{ m}$ are shown in Fig. 6.

The iteration count of the Krylov solver using the DCT (solid line) and line Jacobi (dashed line) preconditioners is compared in Fig. 7 as a function of the convergence threshold ε in (29). The rms divergence (thick solid line) shows the critical value of $\varepsilon = 10^{-5}$, below which the solution is considered physically converged.⁹ It is apparent that the number of iterations per

FIG. 4. Performance ratios for CPU and iteration count, as a function of ε , for the convective bubble case shown in Fig. 1.

time step required for the line Jacobi preconditioner is about 5(10) times that required for the DCT scheme in the hydrostatic (nonhydrostatic) case. However, this does not imply that the efficiency of the DCT preconditioner follows the same pattern. The computational cost of a single GCR iteration with the DCT preconditioner is much higher than with line Jacobi, whereupon the CPU time performance ratio is about 1.5 for the hydrostatic result and somewhat over 2 for the nonhydrostatic solution (see Fig. 8). The trend of decreasing relative efficacy of the DCT preconditioner with increasing hydrostaticity of the problem is consistent with that observed for ADI preconditioners (Skamarock et al. 1997). It is likely related to the propagation characteristics of gravity waves present in the system.

Next, the complex Scandinavian orography is used for the lower boundary to probe the solver performance for natural mesoscale flow scenarios containing many scales of motion. For this case, $\Delta X = \Delta Y = 10 \text{ km}$, and $\Delta Z = 300 \text{ m}$, giving a grid aspect ratio of 0.01. The total horizontal domain size is $2000 \text{ km} \times 2000 \text{ km}$ (200×200 horizontal grid points), with the model lid placed at 15 km . A uniform northwesterly ambient flow with speed of 10 m s^{-1} and buoyancy frequency $N = 0.01 \text{ s}^{-1}$ are assumed. Unlike in the previous cases, a β -plane approximation is employed, with the map scale factor m varying accordingly in (C1). The model is integrated over 24 h with 288 5-min time steps. The converged wind solution on the lowest momentum level (approximately 150 m above the ground) is shown in Fig. 9, for illustration. Based on experience gained from the earlier test cases, a safe value $\varepsilon = 10^{-6}$ is assumed to assure physical convergence. Performance-ratio values are 4.04 and 5.46 for CPU time and iteration count, respectively, documenting much higher gains for the DCT preconditioner compared to the isolated mountain case discussed above. This is not surprising because compact spectra limit the number of significant eigenmodes present in the solution, thereby accelerating the convergence of Krylov solvers. For the sake of com-

⁹ For the exact steady-state solution, the rms divergence should vanish.

FIG. 5. Vertical velocity (left) x - z and (right) x - y cross sections through the mountain wave simulations at time step 480: (top) hydrostatic case and (bottom) nonhydrostatic case. The x - y cross sections are at $z = 2100$ m. Contour intervals are 0.01 m s^{-1} for the hydrostatic case and 0.04 m s^{-1} for the nonhydrostatic case.

pleteness, a 3D ADI preconditioner was also tested in this case. The number of GCR iterations required (for the given convergence threshold) is the same as with the DCT preconditioner, but the computational overhead is substantial, owing to multiple ADI iterations within the preconditioner itself.

Finally, we test the performance of the DCT preconditioner on a synoptic-scale, weather-prediction problem. An arbitrarily chosen forecast from 0000 UTC 16 September 2002 is run using initial and boundary conditions supplied by the National Centers for Environmental Prediction (NCEP) Eta Model. The MC2 model is integrated over 48 h with 480 6-min time steps. The computational domain is $200 \times 160 \times 19$, with $\Delta X = \Delta Y = 50$ km on a polar stereographic projection true at 60°N , and the vertically stretched grid ranges from $\Delta Z = 72$ m near the surface to 3800 m near the top at 23 km. The resulting condition number of the Helmholtz operator is extreme, $O(10^{12})$, and the energy spectrum full. As in the Scandinavian case, $\varepsilon = 10^{-6}$ is assumed. The resulting performance ratios are 1.93 and 4.90 for

CPU time and iteration count, respectively. The DCT preconditioner again requires fewer iterations, and the model integration rate is about twice as fast. This demonstrates that the DCT is far superior to line Jacobi, despite the large departure of the preconditioner from the full Helmholtz operator emphasized by a synoptic-scale forecast.

To assess the relative merits of alternative Krylov methods, we also implemented the right-preconditioned flexible GMRES (FGMRES) algorithm of Saad (1993).¹⁰ For comparison, we ran the MC2 model using FGMRES(4), with DCT preconditioner, for the Scandinavian topography simulation and winter storm synoptic forecast. In both cases, we found that FGMRES requires additional iterations because it tends to converge more slowly than GCR after a restart. The model run times were nearly identical, with FGMRES giving a slight advantage in the synoptic case. To compare

¹⁰ Computational complexity estimates for the various algorithms are provided by Saad (1996).

FIG. 6. One-dimensional power spectra through the centers of the right-hand panels in Fig. 5. Solid and dashed lines denote the hydrostatic and nonhydrostatic case, respectively.

solutions, we examined the maximum Courant numbers (equivalent to a maximum norm) and residual L_2 norm (least squares error) of the elliptic problem over the computational domain for the GCR and FGMRES solvers. In both cases, these agree to six digits, and thus we conclude that the solutions are the same.

5. Concluding remarks

We have developed a spectral preconditioner for a generalized conjugate residual (GCR) Krylov solver applied to the Helmholtz problem arising in the semi-implicit time discretization of a compressible nonhydrostatic atmospheric model. Our approach offers an alternative to the more standard and much simpler line-

FIG. 8. The CPU time performance ratio as a function of convergence criterion ϵ for both the hydrostatic and nonhydrostatic mountain wave cases.

relaxation, stationary block-iteration-type preconditioners used in meteorology. Despite substantial departures of the spectral preconditioner from the governing Helmholtz operator, we observed dramatic performance gains over these simpler schemes. The relative improvement in the model integration rate ranges from a factor of 2 to one order of magnitude, depending on the problem at hand.

The relative performance gains are, of course, greatest when the DCT preconditioner happens to be a direct solver for the governing Helmholtz problem, that is, zero orography and constant mapscale factor characteristic of small-scale dynamics applications. In a meso-scale flow scenario over Scandinavian topography, with a broad range of scales of motion present, a factor of 4 improvement was achieved. For smooth idealized mountain wave flows, frequently used to benchmark numerical model performance, the relative gains are more modest. But extrapolating these results to natural problems underestimates the potential gains because idealized flows often have compact spectra that favor simpler schemes. For synoptic weather prediction,

FIG. 7. The rms divergence and average number of GCR iterations per time step for the (a) hydrostatic and (b) nonhydrostatic mountain wave cases. Results with both the Jacobi (JAC) and DCT preconditioner are shown.

FIG. 9. Wind solution for the first model level ($Z = 150$ m) over the Scandinavian topography at 24 h with the DCT preconditioner. Topography is also shown at 500-m intervals.

where the shallowness of the atmosphere implies an extremely poorly conditioned problem, the performance improvement is still a factor of 2.

Spectral methods rely on global basis functions, and computing their expansion coefficients requires multiple evaluations of global sums. Consequently, spectral transform methods are criticized as being inappropriate for modern, distributed-memory parallel architectures. Nevertheless, in the context of FFT-based preconditioners for iterative Krylov methods, Elman and O'Leary (1998) documented overall parallel speedups despite the large data communication overheads. In meteorology, the vertical direction is physically distinct and typically not partitioned across the processors in numerical models. Thus, a massively parallel implementation of spectral preconditioners can be based on a data-transposition strategy (Thomas et al. 2002) similar to that used in the European Centre for Medium-Range Weather Forecasts (ECMWF) weather prediction model Integrated Forecast System (IFS; Skålin 1997; Barros and Kauranne 1994).

The performance of the next-generation (all-scale, nonhydrostatic semi-implicit, weather and climate prediction/research) numerical models in meteorology depends critically on the performance of elliptic solvers employed. It is widely recognized that the state-of-the-art methods rely on preconditioned nonsymmetric Krylov algorithms. Given the extensive body of mathematical literature, still relatively little numerical experience exists with such solvers in advanced meteorological

applications (cf. Navon and Cai 1993). We have advocated a spectrally preconditioned GCR solver in the context of a particular model, representative of current trends in NWP. Further studies with a focus on alternative numerics, mesh refinement, complex geometries, and extreme meteorological scenarios will pose/define extreme challenges/benchmarks to the advocated approach.

Acknowledgments. The authors are grateful to Claude Girard and Michel Desgagné of RPN and Environment Canada, for maintaining and improving the MC2 model, and to Xingxiu Deng and Henryk Modzelewski from UBC, for preparing the synoptic case and adept computer administration, respectively. This work was supported by the National Science Foundation, the Department of Energy Climate Change Prediction Program (CCPP), the Canadian Natural Science and Engineering Council, RPN and Environment Canada, the Canadian Foundation for Innovation, BC Knowledge Development Fund, the Geophysical Disaster Computational Fluid Dynamics Center, and UBC Research Endowments. The comments of two anonymous referees helped to improve the presentation.

APPENDIX A

Generalized Conjugate-Residual Approach

Here we describe the preconditioned GCR(k) algorithm used in this study. In general, we consider a linear elliptic problem,

$$\mathcal{L}(\Psi) \equiv \sum_{j=1}^3 \frac{\partial}{\partial x^j} \left(\sum_{j=1}^3 C^{jj} \frac{\partial \Psi}{\partial x^j} + D^j \Psi \right) - A \Psi = Q, \quad (\text{A1})$$

with variable coefficients A , C^{jj} , D^j , Q , and either periodic, Dirichlet, or Neumann boundary conditions and adopt the following notation: The discrete representation of a field on the grid is denoted by the subscript i ; the discrete representation of the elliptic operator on the lhs of (A1) is denoted by $\mathcal{L}_i(\Psi)$; and the inner product $\langle \xi \xi \rangle \equiv \sum_i \xi_i \xi_i$. The preconditioner $\mathcal{P}$ is a linear operator that approximates $\mathcal{L}$ to a greater or lesser degree but is easier to invert. Its role is to substitute (A1) with an auxiliary problem that converges faster (than the original problem) because of a closer clustering of the eigenvalues of the auxiliary operator resulting from the superposition of $\mathcal{L}$ and $\mathcal{P}^{-1}$. In this paper, we are primarily concerned with “left” preconditioning that substitutes (A1) with an auxiliary problem $\mathcal{P}^{-1}[\mathcal{L}(\Psi) - Q] = 0$. Its accelerated convergence exploits spectral properties of $\mathcal{P}^{-1}\mathcal{L}$. In general, “right” preconditioning is also possible. It augments (A1) with $\mathcal{L}\mathcal{P}^{-1}[\mathcal{P}(\Psi)] = Q$, and its convergence relies on reducing the spectral radius of $\mathcal{L}\mathcal{P}^{-1}$. Left preconditioning assumes $\mathcal{P}$ constant during the iteration process, whereas right preconditioning allows for variable $\mathcal{P}$ as, for example, in the FGMRES solver of Saad (1993).

The GCR(k) method of Eisenstat et al. (1983) may be derived via variational arguments (cf. Smolarkiewicz and Margolin 1994, 2000). In essence, we augment (A1) with a k th-order damped oscillation equation:

$$\frac{\partial^k \mathcal{P}(\Psi)}{\partial \tau^k} + \frac{1}{T_{k-1}(\tau)} \frac{\partial^{k-1} \mathcal{P}(\Psi)}{\partial \tau^{k-1}} + \dots + \frac{1}{T_1(\tau)} \frac{\partial \mathcal{P}(\Psi)}{\partial \tau} = \mathcal{L}(\Psi) - Q. \quad (\text{A2})$$

Then we discretize (A2) in a pseudotime τ to form the affine discrete equation for the progression of the residual errors r and determine the optimal parameters $T_1, \dots, T_{k-1}$ and integration increment $\Delta \tau$ (variable in τ) that assure minimization of the residual errors in the norm defined by the inner product $\langle rr \rangle$. This leads to the following algorithm.

For any initial guess Ψ_i^0 , set $r_i^0 = \mathcal{L}_i(\Psi^0) - Q_i$, $p_i^0 = \mathcal{P}_i^{-1}(r^0)$; then iterate:

For $n = 1, 2, \dots$ until convergence do

for $\nu = 0, \dots, k-1$ do

$$\beta = -\frac{\langle r^\nu \mathcal{L}(p^\nu) \rangle}{\langle \mathcal{L}(p^\nu) \mathcal{L}(p^\nu) \rangle},$$

$$\Psi_i^{\nu+1} = \Psi_i^\nu + \beta p_i^\nu,$$

$$r_i^{\nu+1} = r_i^\nu + \beta \mathcal{L}_i(p^\nu),$$

exit if $\|r^{\nu+1}\| \leq \varepsilon$,

$$e_i = \mathcal{P}_i^{-1}(r^{\nu+1}),$$

$$\mathcal{L}_i(e) = \left[\sum_{l=1}^3 \frac{\partial}{\partial x^l} \left(\sum_{j=1}^3 C^{lj} \frac{\partial e}{\partial x^j} + D^l e \right) - A e \right]_i,$$

$$\forall_{l=0,\nu} \alpha_l = -\frac{\langle \mathcal{L}(e) \mathcal{L}(p^l) \rangle}{\langle \mathcal{L}(p^l) \mathcal{L}(p^l) \rangle},$$

$$p_i^{\nu+1} = e_i + \sum_{l=0}^{\nu} \alpha_l p_i^l,$$

$$\mathcal{L}_i(p^{\nu+1}) = \mathcal{L}_i(e) + \sum_{l=0}^{\nu} \alpha_l \mathcal{L}_i(p^l),$$

end do,

reset $[\Psi, r, p, \mathcal{L}(p)]_i^k$ to $[\Psi, r, p, \mathcal{L}(p)]_i^0$,

end do.

For the convergence, the GCR(k) algorithm above requires $\mathcal{L}\mathcal{P}^{-1}$ to be negative definite^{A1} but not necessarily self-adjoint.^{A2} Direct evaluation of the elliptic op-

erator on the grid takes place only once per iteration following the preconditioning $e = \mathcal{P}^{-1}(r^{\nu+1})$, which provides an estimate of the solution error $e^{\nu+1} = \Psi^{\nu+1} - \Psi_{\text{exact}}$.

APPENDIX B

Line-Relaxation Preconditioners

A distinctive feature of meteorological flows is their anisotropy in the vertical direction—the larger the ratio of the horizontal scale of the flow to the fluid depth, the stiffer the elliptic problem. A class of simple yet effective preconditioners that mitigate this aspect (of elliptic problems in meteorology) derives from an implicit Richardson iteration (Smolarkiewicz and Margolin 2000),

$$\frac{e^{\mu+1} - e^\mu}{\Delta \tilde{\tau}} = \mathcal{P}^h(e^\mu) + \mathcal{P}^z(e^{\mu+1}) - r^{\nu+1}, \quad (\text{B1})$$

a realization of the preconditioning step $e = \mathcal{P}^{-1}(r^{\nu+1})$ of the GCR(k) solver summarized in appendix A. Here, $\mathcal{P}^h$ and $\mathcal{P}^z$ are the horizontal and the vertical counterparts of the operator $\mathcal{P}$, respectively; $\Delta \tilde{\tau}$ is a parameter of the iteration (a pseudo-time step) based on spectral properties of $\mathcal{P}^h$ [viz., linear stability analysis of (B1)]; μ numbers successive Richardson iterations; and ν numbers the outer iterations of the Krylov solver. Equation (B1) leads to a linear problem,

$$(I - \Delta \tilde{\tau} \mathcal{P}^z) e^{\mu+1} = \tilde{R}^\mu, \quad (\text{B2})$$

where $\tilde{R}^\mu \equiv e^\mu + \Delta \tilde{\tau} [\mathcal{P}^h(e^\mu) - r^{\nu+1}]$, that can be solved readily using the well-known tridiagonal algorithm (cf. appendix A in Roache 1972). Alternating implicit discretization between the vertical and horizontal counterparts of $\mathcal{P}$ in fractional steps of $\tilde{\tau}$ leads to the ADI preconditioners of Skamarock et al. (1997).

The implicit Richardson preconditioner in (B1) can be further improved. Consider extending the Richardson diffusion scheme with respect to $\mathcal{P}^h$ to the diagonally preconditioned Duffort–Frankel type implicit algorithm,

$$\frac{\mathcal{D}e^{\mu+1} - \mathcal{D}e^\mu}{\Delta \tilde{\tau}} = \mathcal{P}^h(e^\mu) - \mathcal{D}(e^{\mu+1} - e^\mu) + \mathcal{P}^z(e^{\mu+1}) - r^{\nu+1}, \quad (\text{B3})$$

where $(-1)\mathcal{D}$ stands for the diagonal coefficient embedded within the matrix representing $\mathcal{P}^h$ on the grid. Note that adding the relaxation term on the rhs of (B1) has the effect of replacing the $\mathcal{D}e^\mu$ term with $\mathcal{D}(e^{\mu+1})$ in $\mathcal{P}^h(e^\mu)$ without complicating flux boundary conditions imposed in constructing $\mathcal{P}^h(e)$ (cf. Smolarkiewicz and Margolin 2000 for discussions). In the limit $\Delta \tilde{\tau} \rightarrow \infty$, (B3) is equivalent to the block Jacobi preconditioner proven effective in meteorological applications (Thomas et al. 1998, 2000).

^{A1} An operator $\mathcal{A}$ is said to be definite if $\langle \xi \mathcal{A}(\xi) \rangle$ is either strictly positive (positive definite) or strictly negative (negative definite) for all ξ .

^{A2} An operator $\mathcal{A}$ is said to be self-adjoint if $\langle \xi \mathcal{A}(\xi) \rangle = \langle \xi \mathcal{A}(\xi) \rangle$ for all ξ and ζ .

APPENDIX C

Coordinate Systems

For a polar stereographic projection, true at longitude ϕ_0 , m is the map scale factor and $S = m^2$:

$$m = \frac{1 + \sin\phi_0}{1 + \sin\phi}. \quad (\text{C1})$$

The components (U , V) of the wind “image” are defined in terms of the scaled and rotated components of the true horizontal velocity (u , v) in local (x , y , z) coordinates,

$$\begin{bmatrix} U \\ V \end{bmatrix} = \frac{1}{m} \begin{bmatrix} -\sin\lambda & -\cos\lambda \\ \cos\lambda & -\sin\lambda \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}, \quad (\text{C2})$$

where λ represents longitude. The (X , Y) components of the conformal coordinates are scaled and rotated according to

$$\begin{bmatrix} dX \\ dY \end{bmatrix} = m \begin{bmatrix} -\sin\lambda & -\cos\lambda \\ \cos\lambda & -\sin\lambda \end{bmatrix} \begin{bmatrix} dx \\ dy \end{bmatrix}. \quad (\text{C3})$$

In conformal coordinates (X , Y , z), the Euler equations take the following form:

$$\begin{aligned} \frac{1}{\gamma} \frac{Dq}{Dt} &= -S \left(\frac{\partial U}{\partial X} + \frac{\partial V}{\partial Y} \right) - \frac{\partial w}{\partial z}, \\ \frac{DU}{Dt} &= fV - K \frac{\partial S}{\partial X} - RT \frac{\partial q}{\partial X}, \\ \frac{DV}{Dt} &= -fU - K \frac{\partial S}{\partial Y} - RT \frac{\partial q}{\partial Y}, \\ \frac{Dw}{Dt} &= -g - RT \frac{\partial q}{\partial z}, \\ \frac{DT}{Dt} - \frac{RT}{c_p} \frac{Dq}{Dt} &= 0, \end{aligned} \quad (\text{C4})$$

where $K = (U^2 + V^2)/2$ is the pseudohorizontal kinetic energy per unit mass. The material derivative is written as

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + S \left(U \frac{\partial}{\partial X} + V \frac{\partial}{\partial Y} \right) + w \frac{\partial}{\partial z}. \quad (\text{C5})$$

The divergence of (U , V , w) is given by

$$D = S \left(\frac{\partial U}{\partial X} + \frac{\partial V}{\partial Y} \right) + \frac{\partial w}{\partial z}, \quad (\text{C6})$$

where the horizontal divergence of (U , V) on the rhs is denoted by $D_H(U, V)$ throughout the body of the paper.

To account for the orography one can introduce a terrain-following coordinate system, where any monotonic function of geometrical height can be used as a vertical coordinate, denoted by Z . Throughout this paper

we employ the classical transformation of Gal-Chen and Sommerville (1975), where

$$Z = z_{\text{top}} \frac{z - h(X, Y)}{z_{\text{top}} - h(X, Y)}. \quad (\text{C7})$$

Because Z in (C7) is assumed to be time independent [cf. Prusa et al. (1996) for time-dependent extensions], the resulting Z -coordinate system represents a nondeformable system with surfaces of constant Z fixed in physical space—in contrast to various pressure-based vertical coordinates used in hydrostatic weather-prediction models. The mapping from conformal (X , Y , z) coordinates to curvilinear (X , Y , Z) coordinates is represented by the metric coefficients G^U of the transformation:

$$G^{13} = \frac{1}{G} \left(\frac{\partial z}{\partial X} \right)_z, \quad G^{23} = \frac{1}{G} \left(\frac{\partial z}{\partial Y} \right)_z, \quad G = \frac{\partial z}{\partial Z}. \quad (\text{C8})$$

Admitting an optional, variable spacing of vertical levels modifies the metric coefficients in (C8) via multiplicative factors (cf. Thomas et al. 1998). The terrain-following system is right-handed, where Z increases from the bottom of the model domain to the top, and the Z surface corresponding to the lower boundary is conformal to the terrain. Thus, by definition, G is always positive. Partial derivatives of any variable ψ are to be interpreted for terrain-following coordinates as follows:

$$\begin{aligned} \left(\frac{\partial \psi}{\partial X} \right)_z &= \left(\frac{\partial \psi}{\partial X} \right)_Z + G^{13} \frac{\partial \psi}{\partial Z}, \\ \left(\frac{\partial \psi}{\partial Y} \right)_z &= \left(\frac{\partial \psi}{\partial Y} \right)_Z + G^{23} \frac{\partial \psi}{\partial Z}, \\ \frac{\partial \psi}{\partial z} &= \frac{1}{G} \frac{\partial \psi}{\partial Z}. \end{aligned} \quad (\text{C9})$$

The contravariant “vertical” velocity is related to the physical velocity components by the equation

$$W \equiv \dot{Z} = G^{-1}w + S(G^{13}U + G^{23}V). \quad (\text{C10})$$

APPENDIX D

Numerical Operators

The centered averaging μ_ζ and differencing δ_ζ operators are defined as

$$\mu_\zeta \psi = \frac{1}{2} \left\{ \psi \left(\zeta + \frac{\Delta \zeta}{2} \right) + \psi \left(\zeta - \frac{\Delta \zeta}{2} \right) \right\}, \quad (\text{D1})$$

$$\delta_\zeta(\psi) = \frac{1}{\Delta \zeta} \left\{ \psi \left(\zeta + \frac{\Delta \zeta}{2} \right) - \psi \left(\zeta - \frac{\Delta \zeta}{2} \right) \right\}, \quad (\text{D2})$$

where ψ is a dependent model variable that may be

defined at the center of a grid cell or on the grid cell faces. Multiple averages are denoted by

$$\mu_{xy}\psi = \overline{\overline{\psi}}^{xy} = \overline{(\overline{\psi}^x)^y} = \overline{(\overline{\psi}^y)^x}. \quad (\text{D3})$$

Discrete horizontal and vertical derivative operators are defined as follows:

$$\begin{aligned} D_x &= \delta_x + G^{13}\mu_{xz}\delta_z, \\ D_y &= \delta_y + G^{23}\mu_{yz}\delta_z, \\ D_z &= \frac{1}{\overline{G}^{xyz}}\delta_z. \end{aligned} \quad (\text{D4})$$

The discrete divergence operator D , acting on the contravariant wind components, can be written as follows:

$$\begin{aligned} D(U, V, W) &= \frac{1}{\overline{G}^{xyz}}\{S[\delta_x\overline{G}^{yz}U + \delta_y\overline{G}^{xz}V] \\ &+ \delta_z\overline{G}^{xy}W\}. \end{aligned} \quad (\text{D5})$$

The operator in (D5) satisfies the Gaussian divergence theorem in discrete form (Thomas et al. 1998). The Jacobian G of the coordinate transformation is located at grid cell vertices and $D(U, V, W)$ at cell centers. The contravariant velocity component W , defined in (C10), is computed from

$$\begin{aligned} \overline{G}^{xy}W &= w + S[G_1\overline{U}^x + G_2\overline{V}^y], \\ G_1 &= \overline{GG}^{13x}, \quad G_2 = \overline{GG}^{23y}, \end{aligned} \quad (\text{D6})$$

and the formal discrete flow divergence (D5) is then implemented as $D(U, V, W) = D_H(U, V) + \delta_z w$, where

$$\begin{aligned} D_H(U, V) &= S\delta_x U + \frac{S}{\overline{G}^{xyz}}[\delta_x\overline{G}^{yz}\mu_x + \mu_z\delta_z G_1\mu_x]U \\ &+ S\delta_y V \\ &+ \frac{S}{\overline{G}^{xyz}}[\delta_y\overline{G}^{xz}\mu_y + \mu_z\delta_z G_2\mu_y]V. \end{aligned} \quad (\text{D7})$$

REFERENCES

- Arakawa, A., and V. Lamb, 1977: Computational design of the basic dynamical processes of the UCLA general circulation model. *Methods in Computational Physics*, Vol. 17, J. Chang, Ed., Academic Press, 174–265.
- Axelsson, O., 1994: *Iterative Solution Methods*. Cambridge University Press, 654 pp.
- Barros, S., and T. Kauranne, 1994: On the parallelization of global spectral weather models. *Parallel Comput.*, **20**, 1335–1356.
- Benoit, R., M. Desgagné, P. Pellerin, S. Pellerin, Y. Chartier, and S. Desjardins, 1997: The Canadian MC2: A semi-Lagrangian, semi-implicit wide-band atmospheric model suited for fine-scale process studies and simulations. *Mon. Wea. Rev.*, **125**, 2382–2415.
- Bernardet, P., 1995: The pressure term in the anelastic model: A symmetric elliptic solver for an Arakawa C grid in generalized coordinates. *Mon. Wea. Rev.*, **123**, 2474–2490.
- Bourke, W., 1974: A multi-level spectral model. I. Formulation and hemispheric integrations. *Mon. Wea. Rev.*, **102**, 687–701.
- Cullen, M., 1990: A test of a semi-implicit integration technique for a fully compressible nonhydrostatic model. *Quart. J. Roy. Meteor. Soc.*, **116**, 1253–1258.
- Davies, H., 1976: A lateral boundary formulation for multi-level prediction models. *Quart. J. Roy. Meteor. Soc.*, **102**, 405–418.
- Eisenstat, S., H. Elman, and M. Schultz, 1983: Variational iterative methods for nonsymmetric systems of linear equations. *SIAM J. Numer. Anal.*, **2**, 345–357.
- Elman, H., and D. O’Leary, 1998: Efficient iterative solution of the three-dimensional Helmholtz equation. *J. Comput. Phys.*, **142**, 163–181.
- Freund, R., 1993: A transpose-free quasi-minimum residual algorithm for non-Hermitian linear systems. *SIAM J. Sci. Stat. Comput.*, **14**, 470–482.
- Gal-Chen, T., and R. Somerville, 1975: On the use of a coordinate transformation for the solution of the Navier–Stokes equations. *J. Comput. Phys.*, **17**, 209–228.
- Golding, B., 1992: An efficient nonhydrostatic forecast model. *Meteor. Atmos. Phys.*, **50**, 89–103.
- Grabowski, W., and P. Smolarkiewicz, 2002: A multiscale model for meteorological research. *Mon. Wea. Rev.*, **130**, 939–956.
- Greenbaum, A., 1997: *Iterative Methods for Solving Linear Systems*. Society of Industrial and Applied Mathematics, 220 pp.
- Hestenes, M., and E. Stiefel, 1952: Method of conjugate gradients for solving linear systems. *J. Res. Natl. Bur. Stand. (U.S.)*, **49**, 409–436.
- Kadogliu, H., and S. Mudrick, 1992: On the implementation of the GMRES(m) method to elliptic equations in meteorology. *J. Comput. Phys.*, **102**, 348–359.
- Kapitza, H., and D. Eppel, 1992: The nonhydrostatic mesoscale model GESIMA. Part I: Dynamical equations and tests. *Beitr. Phys. Atmos.*, **65**, 129–146.
- Lafore, J., and Coauthors, 1998: The Meso–NH atmospheric modeling system. Part I: Adiabatic formulation and control simulations. *Ann. Geophys.*, **16**, 90–109.
- Laprise, R., D. Caya, G. Bergeron, and M. Giguère, 1997: The formulation of the André Robert MC2 (Mesoscale Compressible Community) model. *Atmos.–Ocean*, **35**, 127–152.
- Lynch, R., J. Rice, and D. Thomas, 1964: Direct solution of partial difference equations by tensor product methods. *Numer. Math.*, **6**, 185–199.
- Machenhauer, B., and R. Daley, 1972: A baroclinic primitive equation model with a spectral representation in three dimensions. Tech. Rep. 4, Institute of Theoretical Meteorology, University of Copenhagen, 63 pp.
- Marshall, J., C. Hill, L. Perelman, and A. Adcroft, 1997: Hydrostatic, quasi-hydrostatic and non-hydrostatic ocean modeling. *J. Geophys. Res.*, **102**, 5733–5752.
- Navon, I., and Y. Cai, 1993: Domain decomposition and parallel processing of a finite element model of the shallow water equations. *Comput. Methods Appl. Mech. Eng.*, **106**, 179–212.
- Prusa, J., P. Smolarkiewicz, and R. Garcia, 1996: On the propagation and breaking at high altitudes of gravity waves excited by tropospheric forcing. *J. Atmos. Sci.*, **53**, 2186–2216.
- Roache, P., 1972: *Computational Fluid Dynamics*. Hermosa, 446 pp.
- Robert, A., 1969: Integration of a spectral model of the atmosphere by the implicit method. *Proc. WMO/IUGG Symp. on NWP*, Vol. 7, Tokyo, Japan, Japan Meteorological Agency, 19–24.
- , 1981: A stable numerical integration scheme for the primitive meteorological equations. *Atmos.–Ocean*, **19**, 35–46.
- , 1993: Bubble convection experiments with a semi-implicit formulation of the Euler equations. *J. Atmos. Sci.*, **50**, 1865–1873.
- , T. Yee, and H. Ritchie, 1985: A semi-Lagrangian semi-implicit numerical integration scheme for multilevel atmospheric models. *Mon. Wea. Rev.*, **113**, 388–394.
- Saad, Y., 1993: A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Stat. Comput.*, **14**, 461–469.
- , 1996: *Iterative Methods for Sparse Linear Systems*. PWS, 447 pp.
- , and M. Schultz, 1986: GMRES: A generalized minimal residual

- algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, **7**, 856–869.
- Saito, K., U. Schättler, and U. Steppeler, 1998: 3D mountain waves by the Lokal-Modell of DWD and the MRI mesoscale nonhydrostatic model. *Pap. Meteor. Geophys.*, **49**, 7–19.
- Schumann, U., and H. Volkert, 1987: Three-dimensional mass- and momentum-consistent Helmholtz-equation in terrain-following coordinates. *Notes on Numerical Fluid Mechanics*, W. Hackbusch, Ed., Vol. 10, Springer-Verlag, 109–131.
- , and R. Sweet, 1988: Fast Fourier transforms for direct solution of Poisson's equation with staggered boundary conditions. *J. Comput. Phys.*, **75**, 123–137.
- Skålin, R., 1997: Scalability of parallel gridpoint limited-area atmospheric models. Part II: Semi-implicit time-integration schemes. *J. Atmos. Oceanic Technol.*, **14**, 442–455.
- Skamarock, W., P. Smolarkiewicz, and J. Klemp, 1997: Preconditioned conjugate-residual solvers for Helmholtz equations in nonhydrostatic models. *Mon. Wea. Rev.*, **125**, 587–599.
- Smith, R., 1980: Linear theory of stratified hydrostatic flow past an isolated mountain. *Tellus*, **32**, 348–364.
- Smolarkiewicz, P., and J. Pudykiewicz, 1992: A class of semi-Lagrangian approximations for fluids. *J. Atmos. Sci.*, **49**, 2082–2096.
- , and L. Margolin, 1994: Variational solver for elliptic problems in atmospheric flows. *Appl. Math. Comput. Sci.*, **4**, 527–551.
- , and —, 1997: On forward-in-time differencing for fluids: An Eulerian/semi-Lagrangian nonhydrostatic model for stratified flows. *Atmos.–Ocean*, **35**, 127–152.
- , and —, 2000: Variational methods for elliptic problems in fluid models. *Proc. ECMWF Workshop on Developments in Numerical Methods for Very High Resolution Global Models*, Reading, United Kingdom, ECMWF, 137–159.
- , and J. Prusa, 2002: VLES modelling of geophysical fluids with nonoscillatory forward-in-time schemes. *Int. J. Numer. Methods Fluids*, **39**, 779–819.
- , V. Grubišić, and L. Margolin, 1997: On forward-in-time differencing for fluids: Stopping criteria for iterative solutions of anelastic pressure equations. *Mon. Wea. Rev.*, **125**, 647–654.
- Sweet, R., 1973: Direct methods for the solution of Poisson's equation on a staggered grid. *J. Comput. Phys.*, **12**, 422–428.
- Tanguay, M., A. Robert, and R. Laprise, 1990: A semi-implicit semi-Lagrangian fully compressible regional forecast model. *Mon. Wea. Rev.*, **118**, 1970–1980.
- Tapp, M., and P. White, 1976: A nonhydrostatic mesoscale model. *Quart. J. Roy. Meteor. Soc.*, **102**, 277–296.
- Thomas, S., C. Girard, R. Benoit, M. Desgagné, and P. Pellerin, 1998: A new adiabatic kernel for the MC2 model. *Atmos.–Ocean*, **36**, 241–270.
- , —, G. Doms, and U. Schättler, 2000: Semi-implicit scheme for the DWD Lokal-Modell. *Meteor. Atmos. Phys.*, **73**, 105–125.
- , J. Hacker, M. Desgagné, and R. Stull, 2002: An ensemble analysis of forecast errors related to floating point performance. *Wea. Forecasting*, **17**, 898–906.
- Tokioka, T., 1978: Some considerations on vertical differencing. *J. Meteor. Soc. Japan*, **56**, 98–111.
- van der Vorst, H., 1992: Bi-CGStab: A fast and smoothly converging variant of Bi-CG for the solution of non-symmetric linear systems. *SIAM J. Sci. Stat. Comput.*, **12**, 631–644.
- Yeh, K.-S., J. Cote, S. Gravel, A. Methot, A. Patoiné, M. Roch, and A. Staniforth, 2002: The CMC-MRB Global Environmental Multiscale (GEM) model. Part III: Nonhydrostatic formulation. *Mon. Wea. Rev.*, **130**, 339–356.