

TITLE:

VARIATIONAL METHODS FOR ELLIPTIC PROBLEMS IN FLUID MODELS

AUTHOR(S): Piotr K Smolarkiewicz, National Center for Atmospheric Research
Len G Margolin, CNLS/T-Division

SUBMITTED TO

Conference at European Centre for Medium Range Weather Forecasting on "Developments in Numerical Methods for Very High resolution Global Models" -- Reading, UK - June 5-7, 2000

By acceptance of this article, the publisher recognized that the U S Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution or to allow others to do so for U S Government purposes.

The Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U S Department of Energy.



LOS ALAMOS

Los Alamos National Laboratory
Los Alamos, New Mexico 87545

VARIATIONAL METHODS FOR ELLIPTIC PROBLEMS IN FLUID MODELS

Piotr K Smolarkiewicz* and Len G Margolin†

*National Center for Atmospheric Research, Boulder, Colorado, U.S.A.

† Los Alamos National Laboratory, Los Alamos, New Mexico, U.S.A.

1. INTRODUCTION

Numerical models used for simulating fluid flows often require solutions to general, second-order elliptic partial differential equations, which are implied by the governing model equations themselves (e.g., the incompressible or anelastic approximations) or their implicit discretizations. In meteorological applications, this generality may arise from the effects of rotation, from the use of general curvilinear coordinates in the governing equations, or from free-slip conditions imposed along an irregular lower boundary. There are several difficulties associated with solving such a general problem on discrete meshes. First, modern prognostic atmospheric models employ large grids ($\sim 10^7$ grid points) so that solving the elliptic equation at each of numerous time steps represents a significant computational task. Second, in atmospheric applications the discretized elliptic operator frequently does not possess certain regularity properties that renders some otherwise attractive methods inadequate (cf. Leslie and McAvaney, 1973). Third, in spite of a number of solvers proposed in the mathematical literature that in principle appear suitable for prognostic atmospheric models, only limited numerical experience exists at present with their performance in complex practical applications (cf. Navon and Cai 1993). Such difficulties are recognized in the atmospheric literature as a serious drawback of those fundamental formulations of the equations of motion that lead to complex elliptic problems (cf. Durran 1989, Skamarock and Klemp 1992).

Here, we discuss iterative, conjugate-gradient (alias: Krylov, variational, CG) methods for solving the general linear elliptic equation

$$\sum_{I=1}^M \frac{\partial}{\partial x^I} \left(\sum_{J=1}^M C^{IJ} \frac{\partial \phi}{\partial x^J} + D^I \phi \right) - A\phi = R , \quad (1)$$

with variable coefficients A , C^{IJ} , D^I , R , and periodic, Dirichlet, or Neumann boundary conditions. We emphasize a particular class of conjugate-gradient methods, the conjugate-residual (CR) schemes, which are attractive for atmospheric applications because of their robustness, relative simplicity and computational efficiency. Another important feature is that the convergence of CR schemes does not depend on the self-adjointness of the elliptic operator $\mathcal{L}(\phi)$ on the lhs of (1). Self-adjointness—i.e., the symmetry of the matrix representation of \mathcal{L} on the grid—is usually assumed in textbooks on numerical methods, but cannot be generally assured when, e.g., Coriolis forces or curvilinear boundaries are present, as in atmospheric flows. Finally, CR schemes do not require any explicit knowledge of the matrix resulting from the discretization of $\mathcal{L}(\phi)$, and so are particularly well-suited for models where the discrete representation of \mathcal{L} follows from the general definition of spatial derivative operators such as gradient and divergence on the grid.¹

This paper is organized as follows. In section 2, we outline the essential concepts on which Krylov solvers are based while emphasizing the key assumptions underlying various schemes. In section 3, we discuss in some detail how to enforce proper boundary conditions in the context of the Krylov methods. In section 4, we outline the idea of preconditioning. In order to appeal to the meteorologist’s experience with integrating PDEs of weather and climate, we exploit the unorthodox exposition of Smolarkiewicz and Margolin (1994; hereafter SM94) where the elliptic problem is augmented with a pseudo-time dependence (Richardson 1910). The coefficients of the resulting transient problem are determined by minimizing solution errors in the course of the pseudo-time integration toward the steady state. In essence, this allows for circumventing the idea of quadratic forms (alias “energy functionals”) and the associated interpretation of CG schemes as searching algorithms (for the functionals’ extrema)—standards in the culture of the optimization theory; cf. Shewczuk 1994. While working through the basics, we describe academic algorithms such as the method of steepest descent, the method of minimum residual, and the conjugate gradient method itself, because they aid in the understanding of CG methods especially well. These schemes, however, are about as useless for solving elliptic problems in atmospheric fluid models, as the straightforward upwind or leap-frog schemes are for evaluating advective transport of interest to NWP.

¹Note that values of $\mathcal{L}_i(\phi)$ at the grid points $\mathbf{x}_i = \Delta X \cdot \mathbf{i}$, where $\forall_{I=1,\dots,M} i^I = 1, \dots, n^I$, may be easily computed following (1); whereas, coefficients of the matrix representation $\mathcal{L}_j(\phi) = \sum a_{jk} \phi_k$ [where $k, j = 1, \dots, n^1 \cdots n^M$] number grid points in the lexicographic order, e.g., $j(\mathbf{i}) = n^1 n^2 \cdots n^{M-1} (i^M - 1) + \dots + n^1 (i^2 - 1) + i^1$ may be cumbersome to evaluate and may require storing up to 27 matrices in a three-dimensional problem with second-order-accurate approximations to the partial derivatives of ϕ .

In section 5, we focus on practical applications. We present a preconditioned generalized conjugate-residual scheme (GCR(k); Eisenstat et al. 1983)—substantially more complex than its elementary predecessors yet truly useful—and outline one particular preconditioner that we have found to be effective in meteorological applications. Remarks in section 6 conclude the paper.

2. VARIATIONAL ITERATIVE SCHEMES

2.1 Fundamentals

2.1.1 Pseudo-time augmentation, solution errors, and negative definiteness

The model elliptic equation (1), can be written symbolically as

$$\mathcal{L}(\phi) - R = 0 . \quad (2)$$

Iterative solvers for (2) have been derived sometimes by augmenting (2) with a pseudo-time dependence (cf. Richardson 1910, Frankel 1952; Birkhoff and Lynch 1984), e.g.,

$$\frac{\partial \phi}{\partial \tau} = \mathcal{L}(\phi) - R . \quad (3)$$

Denoting the exact solution in (2) by $\bar{\phi}$, the augmented equation (3) implies

$$\frac{\partial e}{\partial \tau} = \mathcal{L}(e) , \quad (4)$$

where $e \equiv \phi - \bar{\phi}$ is the solution error. In order to derive (4) from (3), we have employed $R = \mathcal{L}(\bar{\phi})$, assumed that \mathcal{L} is *linear*, and used $\partial \bar{\phi} / \partial \tau = 0$. Multiplying both sides of (4) by e and integrating the results over the entire domain gives

$$\frac{\partial \langle e^2 \rangle}{\partial \tau} = 2 \langle e \mathcal{L}(e) \rangle , \quad (5)$$

where $\langle \dots \rangle$ denotes the domain integral. The resulting equation (5) implies that the augmented problem (3) will yield the exact solution as $\tau \rightarrow \infty$, given $\forall e \neq 0 \quad \langle e \mathcal{L}(e) \rangle < 0$. The latter is the definition of the *negative definiteness* of the operator \mathcal{L} —a property sometimes referred to as *dissipativity*.² The Laplacian operators that form the core of elliptic problems in atmospheric models tend to possess this property, but not exactly as stated. In discrete

²In order to associate \mathcal{L} with the diffusion-type operator, we refer to negative definiteness rather than, as is more traditional in the mathematical literature, positive definiteness, and adjust all signs accordingly.

form, \mathcal{L} may have a nontrivial *null space*, i.e., a set of such $\phi \neq 0$ that $\mathcal{L}(\phi) = 0$.³ This means that $\langle e\mathcal{L}(e) \rangle \leq 0$ for all $e \neq 0$ —i.e., \mathcal{L} is only *negative semi-definite*—and that for the elements of the operator's null space the exact solution may be never achieved. However, this is not particularly bothersome in atmospheric applications. In Helmholtz-type elliptic problems (resulting typically from semi-implicit discretizations of compressible systems), the Helmholtz term will assure negative definiteness; alternately, in Poisson-type problems (resulting typically from the discretization of incompressible or anelastic systems) only pressure gradients are of interest, and the exact solution is not required.

Negative-definiteness appears as a natural property of elliptic problems in atmospheric flows, and we shall assume this property (or at least negative semi-definiteness) throughout the rest of this paper. In general, however, this important property must not be taken for granted. It is possible to conceive a situation (e.g., a semi-implicit approximations to chemical reactions in atmospheric flows) where \mathcal{L} is dissipative for some ϕ and “energy increasing” for others. In such a situation, the iterative schemes founded on this assumption may not work at all.

2.1.2 *Self-adjointness*

Conjugate-gradient methods are often interpreted as algorithms that search for the extrema of certain functionals. The quadratic form

$$\mathcal{J}(\phi) \equiv -\frac{1}{2}\langle\phi\mathcal{L}(\phi)\rangle + \langle\phi R\rangle, \quad (6)$$

often referred to as the “energy” functional, is perhaps the most familiar (Birkhoff and Lynch, 1984) of these.

Exploiting the definition of the error e and replacing R with $\mathcal{L}(\bar{\phi})$, we are led straightforwardly to the equation

$$\mathcal{J}(\phi) = -\frac{1}{2}\langle e\mathcal{L}(e)\rangle + \mathcal{J}(\bar{\phi}) + \frac{1}{2}[\langle\phi\mathcal{L}(\bar{\phi})\rangle - \langle\bar{\phi}\mathcal{L}(\phi)\rangle]. \quad (7)$$

Note that if the third term on the rhs of (7) can be assumed to vanish, then the negative definiteness of \mathcal{L} implies positivity of the first term on the rhs of (7), and so $\mathcal{J}(\phi) > \mathcal{J}(\bar{\phi}) \quad \forall \phi$, i.e., the exact solution $\bar{\phi}$ minimizes the energy functional (6). The property of the linear operator that $\langle \xi \mathcal{L}(\phi) \rangle = \langle \phi \mathcal{L}(\xi) \rangle \quad \forall \phi, \xi$ is referred to as *self-adjointness* (or *symmetry* in

³Familiar examples are $2\Delta X$ -wave on the Arakawa A grid, the hourglass pattern on the Arakawa B grid, and a constant on the staggered grid C.

the context of the matrix representing \mathcal{L} on the grid). Since $\xi\Delta\phi \equiv \nabla \cdot (\xi\nabla\phi - \phi\nabla\xi) + \phi\Delta\xi$, self-adjointness is a common property of continuous Laplacian or Helmholtz operators, given suitable boundary conditions.

Although self-adjointness may appear natural, in practical applications it is difficult to achieve. Curvilinear boundaries are notorious for destroying the symmetry of finite-difference operators, and a great care may be required to assure this property in atmospheric models (cf. Bernardet 1995). We shall return to this issue throughout the paper.

2.1.3 Residual error

Equation (3) implies another useful entity. Note that the rhs of (3) vanishes only for the exact solution $\bar{\phi}$, and otherwise it defines a *residual error*

$$r \equiv \mathcal{L}(\phi) - R \quad (\Rightarrow -r = \mathcal{L}(e)) . \quad (8)$$

Thus, (3) can be rewritten as

$$\frac{\partial\phi}{\partial\tau} = r . \quad (9)$$

Acting on both sides of (9) with \mathcal{L} , and subtracting R under the partial τ -derivative on the lhs gives

$$\frac{\partial r}{\partial\tau} = \mathcal{L}(r) . \quad (10)$$

This equation describes the evolution of the residual, just as (4) describes the evolution of the solution error. Further, in analogy with (5), we multiply (10) by r , and integrate the result over the entire domain, deriving

$$\frac{\partial\langle r^2 \rangle}{\partial\tau} = 2\langle r\mathcal{L}(r) \rangle , \quad (11)$$

which shows that $r \rightarrow 0$ as $\tau \rightarrow \infty$, for negative definite \mathcal{L} .

2.1.4 Richardson iteration

Discretizing (3) in pseudo-time with an increment $\Delta\tau = \beta$, while using one-sided differencing for the $\partial\phi/\partial\tau$, leads to a two-term recurrence formula

$$\phi^{n+1} = \phi^n + \beta(\mathcal{L}(\phi^n) - R) , \quad (12)$$

known as the Richardson iteration, or the Richardson diffusion scheme. When applied to the integration of diffusion equations, β must be properly limited to assure numerical stability.

This limiting must take into account the “diffusion coefficient(s)” embedded within \mathcal{L} and local increments of the spatial discretization. This consideration also applies when integrating (3) in pseudo-time; however the value of β determined by this limiting is not necessarily optimal for the convergence of the Richardson iteration to the steady state. We describe two choices of an optimal β in the next section.

2.2 Steepest descent and minimum residual

By the same arguments that convert (3) into (4) and (10) for the continuous evolution of error, (12) implies

$$e^{n+1} = e^n + \beta \mathcal{L}(e^n) , \quad (13)$$

and

$$r^{n+1} = r^n + \beta \mathcal{L}(r^n) , \quad (14)$$

equations that predict the evolution of the error and the residual in the discretized case.

For self-adjoint \mathcal{L} , (7) implies that

$$\mathcal{J}(e) = -\frac{1}{2} \langle e \mathcal{L}(e) \rangle = -\frac{1}{2} \langle e r \rangle . \quad (15)$$

Since the exact solution minimizes \mathcal{J} , one way of assuring an optimal convergence in (12) is to choose β (the independent variable in pseudo-time) to minimize $-\langle e^{n+1} r^{n+1} \rangle$. Taking the product of (13) and (14), integrating the result over the entire domain, differencing over β , and demanding that the resulting derivative vanishes, results in

$$\langle r^n r^n \rangle + \langle e^n \mathcal{L}(r^n) \rangle + 2\beta \langle r^n \mathcal{L}(r^n) \rangle = 0 \quad (16)$$

In general, (16) is useless since e is not computable within the iteration. However for self-adjoint \mathcal{L} , $\langle e^n \mathcal{L}(r^n) \rangle = \langle \mathcal{L}(e^n) r^n \rangle = \langle r^n r^n \rangle$, whereupon (16) implies

$$\beta = -\frac{\langle r^n r^n \rangle}{\langle r^n \mathcal{L}(r^n) \rangle} . \quad (17)$$

The elementary variational iterative scheme built on the Richardson iteration (12), and (17)—known as the *steepest descent* scheme—can be written as follows:

For any initial guess ϕ^0 , set $r^0 = \mathcal{L}(\phi^0) - R$; then iterate :

For $n = 0, 1, 2, \dots$ until convergence do

$$\begin{aligned}\beta &= -\frac{\langle r^n r^n \rangle}{\langle r^n \mathcal{L}(r^n) \rangle}, \\ \phi^{n+1} &= \phi^n + \beta r^n, \\ r^{n+1} &= r^n + \beta \mathcal{L}(r^n), \\ \text{exit if } &\| r^{n+1} \| \leq \varepsilon.\end{aligned}$$

DIGRESSION 1: In the literature, conjugate gradient methods are often derived by using certain orthogonality relationships. The steepest descent scheme introduced above can be derived requiring the orthogonality of the residual errors from two subsequent iterations, i.e., $\langle r^n r^{n+1} \rangle = 0$. In order to show this, note that

$$\frac{\partial}{\partial \beta} \langle e^{n+1} r^{n+1} \rangle = 0 \implies \langle \frac{\partial e^{n+1}}{\partial \beta} r^{n+1} \rangle + \langle e^{n+1} \frac{\partial r^{n+1}}{\partial \beta} \rangle = 0.$$

With (13), (14), and self-adjointness of \mathcal{L} the latter implies

$$\langle r^n r^{n+1} \rangle + \langle e^{n+1} \mathcal{L}(r^n) \rangle = 2 \langle r^n r^{n+1} \rangle = 0.$$

We have emphasized in section 2.1.2 that self-adjointness is difficult to achieve in practical applications. If \mathcal{L} does not possess this property, then there is no reason to expect the convergence of the method of steepest descent. Fortunately, there is an alternative way to optimize the convergence of the Richardson iteration (12) that does not require use of either self-adjointness or an energy functional. Since (14) predicts the residual error at the next iteration, it can be used to choose β that minimizes $\langle r^{n+1} r^{n+1} \rangle$ norm of the residual error. So, taking the product of (14) with itself, integrating the result over the entire domain, differencing over β , and demanding the resulting derivative vanish, we are led straightforwardly to

$$\langle r^n \mathcal{L}(r^n) \rangle + \beta \langle \mathcal{L}(r^n) \mathcal{L}(r^n) \rangle = 0, \quad (18)$$

which implies

$$\beta = -\frac{\langle r^n \mathcal{L}(r^n) \rangle}{\langle \mathcal{L}(r^n) \mathcal{L}(r^n) \rangle}. \quad (19)$$

The variational scheme built on the Richardson iteration (12) with (19) is known as the *minimum residual*. It takes precisely the same algorithmic form as the steepest descent except it uses a different choice of β . The convergence of this scheme requires no reference

to the self-adjointness, symmetry, or quadratic forms. The scheme can be derived from first principles assuming merely a linear negative-definite \mathcal{L} .

DIGRESSION 2: It is easy to show that the minimum residual scheme yields the \mathcal{L} -orthogonality of the residual errors from two subsequent iterations, i.e., $\langle r^{n+1} \mathcal{L}(r^n) \rangle = 0$.

2.3 Conjugate gradient and conjugate residual

The elementary schemes discussed in section 2.2 are merely a convenient vehicle to introduce the fundamental ideas underlying Krylov solvers. Since the convergence rate of these schemes is relatively slow, their utility is insignificant. Much better performance can be obtained by beginning with the damped wave equation in lieu of (3)

$$\frac{\partial^2 \phi}{\partial \tau^2} + \frac{1}{T} \frac{\partial \phi}{\partial \tau} = \mathcal{L}(\phi) - R. \quad (20)$$

Discretizing (20) in pseudo-time with an increment $\Delta\tau$ and a yet to be determined damping scale $T \equiv \eta^{-1}\Delta\tau$ and using, respectively, centered- and one-sided differencing for the first and second term on the lhs of (20), leads to a three-term recurrence formula (also known as the second-order Richardson scheme, due to Frankel 1950)

$$\phi^{n+1} = \gamma\phi^n + (1 - \gamma)\phi^{n-1} + \beta(\mathcal{L}(\phi^n) - R), \quad (21)$$

where $\gamma = (2 + \eta)/(1 + \eta)$, and $\beta = (\Delta\tau)^2/(1 + \eta)$. The recurrence formula in (21) implies the corresponding recurrence relations for the solution and residual errors:

$$e^{n+1} = \gamma e^n + (1 - \gamma)e^{n-1} + \beta\mathcal{L}(e^n), \quad (22)$$

$$r^{n+1} = \gamma r^n + (1 - \gamma)r^{n-1} + \beta\mathcal{L}(r^n), \quad (23)$$

As with the two-term recurrence in (12), one can minimize directly either $\langle er \rangle$ or $\langle rr \rangle$ by setting their β and γ derivatives to zero and solving the resulting linear system for β and γ . The first scheme becomes a special form of the classical conjugate gradient, whereas the second becomes a special form of the classical conjugate residual (both due to Hestenes and Stiefel, 1952); see SM94 for a discussion. In order to obtain the standard forms of these algorithms, as they appear in the literature, we rewrite (21) as

$$\phi^{n+1} = \phi^n + \beta^n \left(\frac{(\gamma^n - 1)\beta^{n-1}}{\beta^n} \cdot \frac{(\phi^n - \phi^{n-1})}{\beta^{n-1}} + r^n \right) \equiv \phi^n + \beta^n(\alpha^n p^{n-1} + r^n); \quad (24)$$

where $\alpha^n \equiv (\gamma^n - 1)\beta^{n-1}/\beta^n$, $p^n \equiv (\phi^{n+1} - \phi^n)/\beta^n$, and the superscripts appearing on γ , β , and α refer to values of the coefficients at different iterations. This leads to the algorithm in the form

For any initial guess ϕ^0 , set $p^0 = r^0 = \mathcal{L}(\phi^0) - R$;

For $n = 0, 1, 2, \dots$ until convergence do

$$\beta^n = \dots ,$$

$$\phi^{n+1} = \phi^n + \beta^n p^n ,$$

$$r^{n+1} = r^n + \beta^n \mathcal{L}(p^n) ,$$

$$\alpha^{n+1} = \dots ,$$

$$p^{n+1} = \alpha^{n+1} p^n + r^{n+1} .$$

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

where the coefficients β and α can be derived by either a direct minimization of the proper error norms or by employing orthogonality relationship derivable from the minimization (see SM94, for further discussion). For the conjugate gradient scheme,

$$\beta^n = -\frac{\langle r^n r^n \rangle}{\langle p^n \mathcal{L}(p^n) \rangle} , \quad \alpha^{n+1} = \frac{\langle r^{n+1} r^{n+1} \rangle}{\langle r^n r^n \rangle} ; \quad (25)$$

whereas for the conjugate residual scheme,

$$\beta^n = -\frac{\langle r^n \mathcal{L}(p^n) \rangle}{\langle \mathcal{L}(p^n) \mathcal{L}(p^n) \rangle} , \quad \alpha^{n+1} = -\frac{\langle \mathcal{L}(r^{n+1}) \mathcal{L}(p^n) \rangle}{\langle \mathcal{L}(p^n) \mathcal{L}(p^n) \rangle} . \quad (26)$$

Although the operator \mathcal{L} appears in the algorithm as acting both on p and r , it needs to be evaluated only once per iteration as the recurrence of p implies

$$\mathcal{L}(p^{n+1}) = \mathcal{L}(r^{n+1}) + \alpha^{n+1} \mathcal{L}(p^n) . \quad (27)$$

As is the case for the method of steepest descent and of minimum residual, the conjugate gradient requires both the self-adjointness and definiteness of \mathcal{L} , whereas only the definiteness of \mathcal{L} suffices for the convergence of the conjugate residual. For examples of relative performance of the two algorithms, see SM94.

3. BOUNDARY CONDITIONS

Elliptic problems are boundary value problems, and so are extremely sensitive to the imposed boundary conditions. This is a trivial statement in the context of the analytic equations, yet its consequence for discrete solvers is often underappreciated. Careful design of the discretized boundary conditions, especially along curvilinear boundaries, may have a dramatic impact on the rate of convergence of Krylov solvers (for example, see Bernardet, 1995) and the overall accuracy of the fluid model. In order to illustrate a principle for imposing boundary conditions in Krylov solvers, consider the conjugate residual scheme from the preceding section. For either Dirichlet or Neumann boundaries, the recurrence relation for the solution ϕ implies, respectively,

$$\phi_B^{n+1} = \phi_B^n + \beta^n p_B^n , \quad (28)$$

$$\mathbf{n} \cdot \nabla \phi^{n+1}|_B = \mathbf{n} \cdot \nabla \phi_B^n + \beta^n \mathbf{n} \cdot \nabla p^n|_B , \quad (29)$$

where \mathbf{n} is the unit vector normal to the boundary, and the subscript B refers to the boundary values. Note that if the boundary conditions were satisfied at the preceding iteration, they will be satisfied at the subsequent iteration, given that the boundary conditions on p are homogeneous. This latter implies that the residual also satisfies homogeneous boundary conditions, a result of the recurrence relation for p . Thus, to ensure the correct boundary conditions throughout the iteration process, it is important to satisfy them from the outset—i.e., at the initialization of the iteration loop, and to maintain the equivalent homogeneous boundary conditions while computing $\mathcal{L}(r^{n+1})$ right before evaluating α^{n+1} .

For illustration, consider the Euler equations for an ideal incompressible fluid in Cartesian geometry, integrated with a standard projection of the preliminary velocity \mathbf{v}^* onto a solenoidal flow (cf. Smolarkiewicz et al. 1997)

$$\nabla \cdot \mathbf{V}^{t+dt} = 0 \implies \nabla \cdot (\mathbf{v}^* - \nabla \phi) = 0 \quad (30)$$

with the velocity boundary conditions

$$\mathbf{n} \cdot \mathbf{V}_B^{t+dt} = V_B \implies \mathbf{n} \cdot (\mathbf{v}^* - \nabla \phi)|_B = V_B \implies \mathbf{n} \cdot \nabla \phi|_B = \mathbf{n} \cdot \mathbf{v}^* - V_B \quad (31)$$

The normal component of the pressure gradient in the last equation of (31) is equal to the appropriate spatial partial derivative. Thus, one can express the boundary pressure gradient term in (30) with the last equation in (31), thereby assuring correct boundary condition

at the initialization of the iteration loop. In the iterations that follow, the corresponding gradient term of the residual error must be set to zero.

At Neumann boundaries in general, standard centered partial derivative operators in a discrete representation of \mathcal{L} may be undefined wherever they require data from outside the computational domain. A common procedure is to replace the centered approximations with one-sided difference formulae at the boundaries where required (cf. Chorin 1968, Glowinski 1992). This seemingly minor aspect of fluid model design has important consequences. Because of local anisotropies of the difference formulae, the numerical operator \mathcal{L} may not be symmetric, and spurious vorticity may be generated at the free-slip boundaries. In the general curvilinear three-dimensional case, this is a nontrivial issue that so far does not seem to have a rigorous yet practical solution (cf. Bernard and Kapitza 1992).⁴ In our models (Smolarkiewicz and Margolin 1997, Smolarkiewicz et al. 1999), we adopt the following approach. At Neumann boundaries, the normal component of the pressure gradient in (29) combines all spatial derivatives

$$\mathbf{n} \cdot \nabla \phi|_B = CX \frac{\partial \phi}{\partial x} + CY \frac{\partial \phi}{\partial y} + CZ \frac{\partial \phi}{\partial y} \quad (32)$$

where the coefficients CX , CY , and CZ depend in general on all the spatial coordinates. At the initialization of the conjugate residual solver, we evaluate all but the normal partial derivatives explicitly from pressure field values available on the grid (e.g., from the previous time step of the model). The normal derivatives are then computed from the velocity boundary conditions (in the spirit of the example above) and are substituted for the other-than-normal components of the pressure gradient. Within the iteration loop, we do the same for the residual error while computing the normal derivatives from the homogeneous boundary conditions. More specifically, we evaluate all but the normal partial derivatives explicitly from the residual error field available from the preceding iteration. The normal derivatives are then computed from the homogeneous boundary conditions and substituted for other-than-normal components of the residual-error gradient. This procedure has several virtues important for applications: a) it assures that the velocity boundary conditions are satisfied exactly at each iteration of the solver; b) it assures that the correct pressure boundary conditions (viz. free of splitting errors) are employed in the limit of convergence and c) in our experience, it minimizes the production of spurious vorticity at the curvilinear

⁴In three spatial dimensions, a possible formal solution may require incorporation of six two-dimensional, and twelve one-dimensional additional elliptic equations.

ear boundaries. Since the symmetry of the discrete elliptic operator is not assured, such a procedure cannot be employed within solvers that rely on self-adjointness.

4. OPERATOR PRECONDITIONING

The convergence of variational schemes may be further accelerated by operator preconditioning. In essence, preconditioning procedures replace (2) (e.g., by means of operator splitting and/or composition) with a modified governing equation $\mathcal{L}'(\phi) - R' = 0$ that can be more easily inverted on the grid. There is no general theory of how to design an optimal preconditioner (Axelsson 1994, section 7). Here, we consider the so-called left preconditioning that we have found particularly useful in atmospheric models.

To illustrate the concept of the preconditioning, let us return to the augmented parabolic equation (3) and consider $\mathcal{L} = \partial/\partial z K(z) \partial/\partial z$; i.e., a one-dimensional diffusion operator with variable coefficient K . Let us assume that K is large near the ground and decreases rapidly with height. Numerical integration of such a diffusion problem with, say, the Richardson scheme (12) requires limiting the time step $\beta \leq 0.5\Delta z^2 K^{-1}$ for stability. In effect, the convergence toward steady state at higher altitudes will be much slower than near the ground. If only the steady state is of interest, $\partial\phi/\partial\tau$ in (3) may be replaced with $\partial K\phi/\partial\tau$, thereby reducing the stiffness of the problem and accelerating the convergence at the higher elevations. So, effectively the left preconditioning replaces (3) with

$$\frac{\partial \mathcal{P}(\phi)}{\partial \tau} = \mathcal{L}(\phi) - R . \quad (33)$$

where \mathcal{P} is the preconditioner. The preconditioner \mathcal{P} can be (in principle) any linear operator such that $\mathcal{L}\mathcal{P}^{-1}$ is negative definite. Its goal, however, is to augment the original recurrence (12) with

$$\phi^{n+1} = \phi^n + \beta\mathcal{P}^{-1}(\mathcal{L}(\phi^n) - R) = \phi^n + \beta\mathcal{P}^{-1}(r) , \quad (34)$$

which converges faster (than the original problem) due to the smaller condition number (i.e., a closer clustering of the eigenvalues of the auxiliary elliptic operator $\mathcal{P}^{-1}\mathcal{L}$). For the preconditioner to be useful, the convergence of the auxiliary problem must be sufficiently rapid to overcome the additional effort associated with inverting \mathcal{P} in (34). In general, the closer \mathcal{P} approximates \mathcal{L} , the faster the scheme converges, but the more difficult it is to invert the operator in (34). For example: in the limit $\mathcal{P} \equiv \mathcal{L}$, (34) converges in one iteration but the entire effort of solving (2) is placed in inverting \mathcal{P} (bringing us back to the starting

point); whereas in the $\mathcal{P} \equiv \mathcal{I}$ limit, inverting \mathcal{P} is effortless but there is no acceleration of the convergence. In between, there is great flexibility in designing preconditioners that exploit either direct or relaxation methods. This flexibility adds another degree of freedom to the study of Krylov methods, which in itself constitutes an established research area (see Axelsson 1994, for a review). The choice of an effective preconditioner is both problem and computer dependent (cf. Shadid and Tuminaro, 1994) and a detailed discussion of the associated issues is beyond the scope of this paper. In the discussions that follow, we focus on aspects particularly important for meteorological models.

In order to derive the preconditioned conjugate-residual algorithm from section 2.3, one can start with

$$\frac{\partial^2 \mathcal{P}(\phi)}{\partial \tau^2} + \frac{1}{T} \frac{\partial \mathcal{P}(\phi)}{\partial \tau} = \mathcal{L}(\phi) - R , \quad (35)$$

and repeat the entire derivation while acting with \mathcal{P}^{-1} on both sides of the equations. The recurrence relations for ϕ , r , and p then take, respectively, the form

$$\phi^{n+1} = \phi^n + \beta^n \mathcal{P}^{-1}(p^n) , \quad (36)$$

$$r^{n+1} = r^n + \beta^n \mathcal{L} \mathcal{P}^{-1}(p^n) , \quad (37)$$

$$\mathcal{P}^{-1}(p^{n+1}) = \alpha^{n+1} \mathcal{P}^{-1}(p^n) + \mathcal{P}^{-1}(r^{n+1}) . \quad (38)$$

Redefining $p_{new} = \mathcal{P}^{-1}(p_{old})$ the complete preconditioned conjugate residual scheme can be written as follows

For any initial guess ϕ^0 , set $r^0 = \mathcal{L}(\phi^0) - R$, $p^0 = \mathcal{P}^{-1}(r^0)$

For $n = 0, 1, 2, \dots$ until convergence do

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

$$\phi^{n+1} = \phi^n + \beta^n p^n ,$$

$$r^{n+1} = r^n + \beta^n \mathcal{L}(p^n) ,$$

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

$$q^{n+1} = \mathcal{P}^{-1}(r^{n+1}) ,$$

$$\alpha^{n+1} = -\frac{\langle \mathcal{L}(q^{n+1}) \mathcal{L}(p^n) \rangle}{\langle \mathcal{L}(p^n) \mathcal{L}(p^n) \rangle} ,$$

$$p^{n+1} = \alpha^{n+1} p^n + q^{n+1} ,$$

$$\mathcal{L}(p^{n+1}) = \mathcal{L}(q^{n+1}) + \alpha^{n+1} \mathcal{L}(p^n) .$$

Since the recurrence relation for ϕ retains its original form, our discussion of the boundary conditions from section 3 holds, except that the recurrence relation for p implies now homogeneous boundary conditions on q .

The technical difference between the preconditioned and unpreconditioned scheme is that there is an auxiliary elliptic problem to be solved, respectively, at the initialization $p^0 = \mathcal{P}^{-1}(r^0)$, and at each iteration $q^{n+1} = \mathcal{P}^{-1}(r^{n+1})$. This auxiliary problem has two degrees of freedom: a) the definition of the operator \mathcal{P} itself; and b) the solution method. In principle, \mathcal{P} could be identical with \mathcal{L} and the solution method could be the unpreconditioned conjugate residual scheme. Such a preconditioner would reduce the number of iterations in the outer, preconditioned conjugate residual solver, but its overall cost would be about the same (slightly larger, in fact) as that of the unpreconditioned outer solver with an accordingly larger number of iterations. The goal is not to reduce the number of iterations in the main solver, but to reduce overall CPU time of the entire solver! The possibilities for exercising ingenuity are endless and “goals justify means”. In particular, understanding the physics of the modeled flows may be helpful for designing effective preconditioners. In the following section we shall return to this topic and outline the preconditioner we have found useful for practical applications in meteorology.

5. PRACTICAL APPLICATIONS

5.1 Generalized Conjugate Residual

Given a suitable preconditioner, the conjugate residual solver is quite effective for a broad class of applications. Nonetheless, as the symmetry of the discrete operator deteriorates, the convergence rate decreases. The generalized conjugate residual, GCR(k), scheme of Eisenstat et al. 1983 is capable of maintaining the optimal CG convergence for nonsymmetric problems. It is mathematically equivalent to the truncated ORTHOMIN and GMRES schemes (Kapitza and Eppel 1992, Saad and Schultz 1986), yet it has a clear interpretation in terms of the pseudo-time augmentation of (2).

We generalize our derivation of the conjugate residual method by starting with a k th-order damped oscillation equation

$$\frac{\partial^k \mathcal{P}(\phi)}{\partial \tau^k} + \frac{1}{T_{k-1}(\tau)} \frac{\partial^{k-1} \mathcal{P}(\phi)}{\partial \tau^{k-1}} + \dots + \frac{1}{T_1(\tau)} \frac{\partial \mathcal{P}(\phi)}{\partial \tau} = \mathcal{L}(\phi) - R \quad (39)$$

in lieu of (20). We proceed with the formalism of section 2.3—viz. discretize (39) in a pseudo-time τ , 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 $\langle rr \rangle$ —leading to the following algorithm.

For any initial guess ϕ^0 , set $r^0 = \mathcal{L}(\phi^0) - R$, $p^0 = \mathcal{P}^{-1}(r^0)$; then iterate:

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

```

    for  $\nu = 0, \dots, k - 1$  do
         $\beta = -\frac{\langle r^\nu \mathcal{L}(p^\nu) \rangle}{\langle \mathcal{L}(p^\nu) \mathcal{L}(p^\nu) \rangle}$  ,
         $\phi^{\nu+1} = \phi^\nu + \beta p^\nu$  ,
         $r^{\nu+1} = r^\nu + \beta \mathcal{L}(p^\nu)$  ,
        exit if  $\| r^{\nu+1} \| \leq \varepsilon$  ,
         $q = \mathcal{P}^{-1}(r^{\nu+1})$  ,
         $\mathcal{L}(q) = \left[ \sum_{I=1}^M \frac{\partial}{\partial x^I} \left( \sum_{J=1}^M C^{IJ} \frac{\partial q}{\partial x^J} + D^I q \right) - Aq \right]$  ,
         $\forall_{l=0, \nu} \alpha_l = -\frac{\langle \mathcal{L}(q) \mathcal{L}(p^l) \rangle}{\langle \mathcal{L}(p^l) \mathcal{L}(p^l) \rangle}$  ,
         $p^{\nu+1} = q + \sum_{l=0}^{\nu} \alpha_l p^l$  ,
         $\mathcal{L}(p^{\nu+1}) = \mathcal{L}(q) + \sum_{l=0}^{\nu} \alpha_l \mathcal{L}(p^l)$  ,
    end do ,
    reset  $[\phi, r, p, \mathcal{L}(p)]^k$  to  $[\phi, r, p, \mathcal{L}(p)]^0$  ,
    end do .

```

Direct evaluation of the elliptic operator on the grid takes place only once per iteration following the preconditioning $q = \mathcal{P}^{-1}(r^{\nu+1})$.

The common wisdom for nonsymmetric solvers such as *GMRES* is to use a quite long expansion in (39) of $k \approx 10$. The price to be paid for this is the necessity of storing many matrices of p and $\mathcal{L}(p)$ from preceding iterations. In our experience with atmospheric problems, $k = 4$ appears sufficient. It handles satisfactorily flows on all scales from micro to global, with possibly steep orography and large planetary rotations.

5.2 Implicit Richardson-iteration preconditioner

A distinctive feature of meteorological flows is their anisotropy in the vertical direction. The larger the ratio of the horizontal scale of the problem to the fluid depth, the stiffer is the elliptic problem. In modeling small scale motions (e.g., large-eddy-simulations of planetary boundary layer) even an unpreconditioned conjugate residual scheme performs reasonably well, as the solver takes advantage of having an accurate first guess, namely the model solution from the preceding time step. On the other end of the spectrum, there are nonhydrostatic global models with full three-dimensional elliptic problems and outrageous condition numbers. For the purpose of practical applications, the condition number may be thought as the squared ratio of the longest to the shortest wavelength present in the system. In global model simulations, it can easily be $\mathcal{O}(10^{10})$. Since the asymptotic convergence rate of CR schemes is inversely proportional to the square root of the condition number, in global flows it can be so slow⁵ that iterative updates may be smaller than the machine precision; i.e., the solver will stall and there will be no convergence at all. Thus, designing a preconditioner capable of mitigating this aspect of meteorological models is a necessary prerequisite for any solver suitable for simulating all-scale flows.

In our nonhydrostatic anelastic models for meteorological flows on scales from micro to planetary (Smolarkiewicz and Margolin 1997, Smolarkiewicz et al. 1999), we solve a general Poisson equation

$$-\frac{\Delta t}{\tilde{\rho}} \sum_{I=1}^3 \frac{\partial}{\partial x^I} \left[\tilde{\rho} \mathcal{J}^I \mathcal{E} \left(\mathcal{V}^I - \sum_{J=1}^3 \mathcal{C}^{IJ} \frac{\partial \pi}{\partial x^J} \right) \right] = 0, \quad (40)$$

where all coefficients (due to coordinate transformations and implicit integrals of the rhs of the momentum equations) are significant functions of coordinates, $\tilde{\rho}$ is the reference density premultiplied by the Jacobian of the coordinate transformation, and π is a normalized pressure perturbation. Equation (40) has the form of the generalized elliptic equation (1) where the coefficients $A \equiv D \equiv 0$, the rhs R is included under the divergence operator,⁶ and where the entire equation has been premultiplied by $-\Delta t/\tilde{\rho}$. The factor (-1) assures the formal negative-definiteness of the elliptic operator on the lhs of (40); further normalization by $\Delta t/\tilde{\rho}$ gives the residual errors of (40) the sense of being the divergence of a dimensionless velocity on the grid. The latter compares directly to the magnitudes of the Courant

⁵For example, it may take 10^5 iterations to reduce the residual error by the factor e^{-1} .

⁶ $\mathcal{E}(\vec{\mathcal{V}} - \dots)$ in (40) is the updated transformed velocity, so the impermeability condition at the model surface and lid, translates into the implicit Neumann boundary conditions for pressure $\mathcal{V}^3 = \sum_{J=1,3} \mathcal{C}^{3J} (\partial \pi / \partial x^J)$.

and Lipschitz numbers and facilitates the design of a physically meaningful stopping criteria (Smolarkiewicz et al. 1997). The coefficients \mathcal{E} , \mathcal{V}^I and C^{IJ} appearing in (40) are given explicitly in Smolarkiewicz et al. 1999.

For complex atmospheric models, we have found it beneficial to keep the preconditioner \mathcal{P} close to the main operator \mathcal{L} . Based on our experience, we choose \mathcal{P} identical to \mathcal{L} implied by (40), except for the cross-derivative terms which are set to zero. Such a \mathcal{P} can be written as

$$\frac{\Delta t}{\tilde{\rho}} \sum_{I=1}^3 \frac{\partial}{\partial x^I} \left[\tilde{\rho} \mathcal{J}^I \mathcal{E} \mathcal{C}^{II} \frac{\partial}{\partial x^I} \right], \quad (41)$$

Following the experience with ADI preconditioners summarized in Skamarock et al. 1997, we invert \mathcal{P} only approximately (in the horizontal) using a few iterations of the implicit Richardson iteration⁷

$$\frac{q^{\mu+1} - q^\mu}{\Delta \tilde{\tau}} = \mathcal{P}^h(q^\mu) + \mathcal{P}^z(q^{\mu+1}) - r^{\nu+1}, \quad (42)$$

where \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 (42)], μ numbers successive Richardson iterations, and ν numbers the outer iterations of the GCR solver. The equation (42) leads to a linear problem

$$(\mathcal{I} - \Delta \tilde{\tau} \mathcal{P}^z) q^{\mu+1} = \tilde{R}^\mu, \quad (43)$$

where $\tilde{R}^\mu \equiv q^\mu + \Delta \tilde{\tau}(\mathcal{P}^h(q^\mu) - r^{\nu+1})$, that can be solved readily using the well-known tridiagonal algorithm (cf. Appendix A in Roache 1972).⁸

For this preconditioner to work, it is essential that the details of differencing and boundary conditions within \mathcal{P}^z are consistent with those within \mathcal{L} . Otherwise unstable modes will leak to the solver and the entire preconditioning effort may be partially or totally lost. When implemented consistently, this fairly simple preconditioner works quite effectively and in large-scale models makes a dramatic difference. In fact, this preconditioner is what makes our nonhydrostatic global models run efficiently (cf. Smolarkiewicz et al. 1999). Without the preconditioner, the convergence of the GCR is so slow that the solver stalls. In small scale models, the impact of the preconditioner (42) is smaller, but is still substantial and can accelerate the convergence of entire pressure solver by a factor of about 2.

⁷As the horizontal scale of the model increases, it may be beneficial to use a few more (several) iterations within the preconditioner; however, in a massively parallel environment, the opposite may be true.

⁸On the nonstaggered mesh that we use, the resulting linear problem is formally pentadiagonal and requires a customized tridiagonal algorithm for its inversion.

The preconditioner in (42) can be further improved. The simplest thing to consider is to extend the Richardson diffusion scheme with respect to \mathcal{P}^h on the diagonally-preconditioned Duffort-Frankel type implicit algorithm

$$\frac{\mathcal{D}q^{\mu+1} - \mathcal{D}q^\mu}{\Delta\tilde{\tau}} = \mathcal{P}^h(q^\mu) - \mathcal{D}(q^{\mu+1} - q^\mu) + \mathcal{P}^z(q^{\mu+1}) - r^{\nu+1}, \quad (44)$$

where $-\mathcal{D}$ 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 (42) has the effect of replacing the $\mathcal{D}q^\mu$ term with $\mathcal{D}(q^{\mu+1})$ in $\mathcal{P}^h(q^\mu)$ without complicating flux boundary conditions imposed in constructing $\mathcal{P}^h(q)$. In the limit $\Delta\tilde{\tau} \rightarrow \infty$, (44) is equivalent to the block Jacobi preconditioner. In our small scale models, (44) improves the effectiveness of the elliptic solver by about 10 percent compared to (42), but it has no impact in the global model.

In general, the larger the horizontal scale of the nonhydrostatic fluid model, the higher the relative cost of the elliptic solver. In small-scale models, the time step is accordingly short, limited by the local advection and nonlinear dynamics. There, elliptic solvers can make effective use of the first guess provided by the previous time step(s) of the fluid model. Global models tend to have a much greater degree of implicitness in order to allow for long-term integrations with time steps as large as possible without sacrificing accuracy. The latter is particularly important in climate simulations. When a large time step is employed, most of the computational effort is already in the elliptic solver, so that the overall model efficiency strongly depends on effective preconditioning.

Although the preconditioner in (42) makes our nonhydrostatic global model perform reasonably well, there is no doubt that further improvements are important. The preconditioner (42) removes the worst aspect of the elliptic problem's stiffness, i.e., that due to the anisotropy in the vertical direction. However, there is another problem to be conquered for flows on the sphere. In spherical coordinates, the zonal grid elements become small near the poles, thereby slowing the overall convergence of the solver. In order to mitigate this aspect of the problem stiffness in (42), we have decomposed the horizontal operator \mathcal{P}^h into its zonal and meridional counterparts \mathcal{P}^Z and \mathcal{P}^M , respectively, and allowed for the implicit (in pseudo-time) discretization of the zonal part

$$\frac{q^{\mu+1} - q^\mu}{\Delta\tilde{\tau}} = \mathcal{P}^Z(q^{\mu+1}) + \mathcal{P}^M(q^\mu) + \mathcal{P}^z(q^{\mu+1}) - r^{\nu+1}, \quad (45)$$

To invert (45) with respect to $q^{\mu+1}$, we have used a rigorous spectral decomposition in the zonal direction. Based on the few results available in the CFD literature (cf. Elman and

O’Leary, 1998), one might expect that such a preconditioner would accelerate substantially the convergence of the Krylov solver, thereby leading to a much faster model for large ‘climatic’ time steps. Indeed, for a standard benchmark problem of the zonal flow past a smooth Gaussian hill, the new preconditioner reduces the total number of the Krylov solver’s iterations by the factor of about 30, thereby reducing the overall cost of the model by the factor of about 2—even with FFTs, the cost of the spectral preconditioner itself is relatively high. However, when the smooth hill is replaced by the real orography of Earth, the new approach works just the opposite: it degrades substantially the convergence of the default preconditioner in (42). This is because the spectral preconditioner zonally averages the original coefficients of the elliptic operator. Although the latter may be satisfactory for smooth problems (which so far comprise the literature) it is absolutely unacceptable for the real orography of the planet.

There is an important lesson in the above-outlined exercise: any preconditioner for applications to natural atmospheric/oceanic flows has to account for small-scale variability of the coefficients. In other words, it is more important to keep \mathcal{P} close to \mathcal{L} than to invert an oversimplified \mathcal{P} exactly. The latter narrows considerably the set of preconditioning methods useful for the applications addressed. Either spectral or multigrid preconditioners appear effective if they can be kept simple. If the majority of the coefficients of \mathcal{L} must be invoked in spectral decomposition or multigrid cycling, then such a preconditioner becomes complex and costly. Then, simply allowing for more iterations in the outer Krylov solver may be more effective. The quest for an ultimate preconditioner will likely never end. At the moment, it seems that following the path of ADI preconditioners, set forth in Skamarock et al. 1997, warrants further investment.

6. REMARKS

In the literature of the Krylov methods, schemes derived based on the minimization of the L_2 norm of the residual error r —such as the minimum or conjugate residual—are sometimes presented as inferior members of the conjugate-gradient family. The accompanying argument is that they are equivalent to schemes that minimize the solution error e but for the squared operator \mathcal{L}^2 , thereby converging much more slowly to the exact solution. Such statements are made out of context of practical applications and fail to recognize that “superior” schemes based on the self-adjointness simply lack robustness.

A potentially confusing statement one can encounter in the literature is that conjugate residual schemes derivable by minimization of the residual error do require self-adjointness. There can be a few special reasons for such statements, which are not true in general. First, some variants may exploit self-adjointness to reformulate coefficients such as β and α in (26); cf. Stoer (1983). Second, the convergence theory may invoke self-adjointness of the operator but then do not require its definiteness (cf. Ashby et al. 1990). In SM94, we have provided a formal proof showing that the definiteness is sufficient for the convergence. Fortunately (for meteorological applications) either self-adjointness or definiteness is sufficient. This flexibility is rarely noted in the mathematical literature where self-adjointness is frequently assumed *a priori*.

Throughout this paper we have advocated those schemes derivable by minimizing the residual error. Here we offer an argument that in atmospheric/oceanic fluid models, it is the residual error whose minimization is of primary concern. Consider a prototype equation for fluids in an arbitrary M -dimensional curvilinear reference frame \mathbf{x}

$$\frac{\partial \rho G \psi}{\partial t} + \nabla \cdot (\rho G \mathbf{V} \psi) = \rho G \mathcal{F}, \quad (46)$$

where $G = G(\mathbf{x})$ is the Jacobian of the coordinate transformation from the Cartesian to the curvilinear framework, ψ is an arbitrary specific variable (e.g., velocity component, potential temperature, water vapor mixing ratio, etc.), $\mathbf{V} = \dot{\mathbf{x}}$, and \mathcal{F} combines all forcings and/or sources (e.g., pressure gradient and Coriolis terms in the momentum equation, heat sources in the potential temperature equation, and $\mathcal{F} = 0$ in the mass continuity equation). The prototype equation (46) can be written equivalently as

$$\frac{D\psi}{Dt} = \mathcal{F} \iff \frac{\partial \psi}{\partial t} + \mathbf{V} \cdot \nabla \psi = \mathcal{F}. \quad (47)$$

The latter form shows that in regions where ψ is locally constant, it can evolve in time only in response to \mathcal{F} . While passing from (46) to (47), we have exploited the mass continuity equation, as the complete form of (47) is

$$\frac{D\psi}{Dt} + \psi \frac{1}{\rho G} \left(\frac{\partial \rho G}{\partial t} + \nabla \cdot (\rho G \mathbf{V}) \right) = \mathcal{F}. \quad (48)$$

The second term on the lhs vanishes in analytic models, since the expression in the parentheses is the mass continuity equation. However, in numerical models, either implicit compressible or anelastic, the entire expression that multiplies ψ in the second term on the lhs,

forms the residual error of the elliptic equation at hand. Thus, Eulerian flux-form schemes for (46) react to spurious sources of the transported field in proportion to the residual error. The semi-Lagrangian schemes are more forgiving in this respect. Since they approximate (47), they are free of the linear source error. However, multiplying either (46) or (47) by ψ and integrating over the entire domain shows that both types of schemes suffer from spurious sink/sources in quadratic integrals. The latter is especially important for the nonlinear stability and accuracy of the long-term integrations; see Smolarkiewicz et al. (1997) for a discussion. In any case, it is essential to assure that the residual error is sufficiently small to keep spurious sources from becoming comparable to those physical sources in \mathcal{F} .

Acknowledgements. Stimulating discussions with John Adams, Jeffrey Scroggs, and Steven Thomas are gratefully acknowledged. National Center for Atmospheric Research is sponsored by the National Science Foundation. Los Alamos National Laboratory is operated by the University of California for the U.S. Department of Energy. This work has been supported in part by the Department of Energy “Climate Change Prediction Program” research initiative. Elements of this work were accomplished during PKS’s collaborative leave at the European Centre for Medium-Range Weather Forecasts, Reading, UK.

REFERENCES

- Ashby, S. F., T. A. Manteuffel, and P. E. Saylor, 1990: A taxonomy for conjugate gradient methods. *SIAM J. Numer. Anal.*, **27**, 1542–1568.
- Axelsson, O., 1994: *Iterative Solution Methods*. Cambridge University Press, Cambridge, 654 pp.
- Bernard, R. S., and H. Kapitza, 1992: How to discretize the pressure gradient for curvilinear MAC grids. *J. Comput. Phys.*, **99**, 288–298.
- 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.
- Birkhoff, G., and R. E. Lynch, 1984: *Numerical solutions of elliptic problems*, SIAM, Philadelphia, pp. 319.
- Chorin, A. J., 1968: Numerical solution of the Navier-Stokes equations. *Math. Comp.*, **22**, 742–762.
- Durran, D. R., 1989: Improving the anelastic approximation. *J. Atmos. Sci.*, **46**, 1453–1461.

- Elman, H. C., and D. P. O'Leary, 1998: Efficient iterative solution of the three-dimensional Helmholtz equation. *J. Comput. Phys.*, **142**, 163-181.
- Eisenstat, S. C., H. C. Elman, and M. H. Schultz, 1983: Variational iterative methods for nonsymmetric systems of linear equations. *SIAM J. Numer. Anal.*, **20**, 345-357.
- Frankel, S. P., 1950: Convergence rates of iterative treatments of partial differential equations, *Math. Tables Aids Comp.*, **4**, 65-75.
- Gal-Chen, T., and R. C. J. Somerville, 1975: On the use of a coordinate transformation for the solutions of the Navier-Stokes equations. *J. Comput. Phys.*, **17**, 209-228.
- Glowinski, R., 1992: Ensuring well-posedness by analogy; Stokes problem and Boundary control for the wave equation. *J. Comput. Phys.*, **103**, 189-221.
- Hestenes, M. R., and E. Stiefel, 1952: Methods of conjugate gradients for solving linear systems. *J. Res. NBS*, **49**, 409-436.
- Kapitza, H., and D. Eppel, 1992: The non-hydrostatic mesoscale model GESIMA. Part 1: Dynamical equations and tests. *Beitr. Phys. Atmosph.*, **65**, 129-146.
- Leslie, L. M., and B. J. McAvaney, 1973: Comparative test of direct and iterative methods for solving Helmholtz-Type equations. *Mon. Wea. Rev.*, **101**, 235-239.
- Navon, I. M., 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.
- Richardson, L. F., 1910: The approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stress in a masonry dam. *Philos. Trans. Roy. Soc., London A210*, 307-357.
- Roache, P., J., 1972: *Computational Fluid Dynamics.*, Hermosa Publishers, Albuquerque, 446 pp.
- Saad, Y., and H. Schultz, 1986: GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, **7**, 856-869.
- Shadid, J. N., and R. T. Tuminaro, 1994: A comparison of preconditioned nonsymmetric Krylov methods on a large-scale MIMD machine. *SIAM J. Sci. Comput.*, **15**, 440-459.
- Shewczuk, J. R., 1994: An Introduction to the Conjugate Gradient Method Without the Agonizing Pain. <http://sepwww.stanford.edu/sep/matt/proj/mathematics.html>
- Skamarock, W. C., and J. B. Klemp, 1992: The stability of time-split numerical methods for the hydrostatic and nonhydrostatic elastic equations. *Mon. Wea. Rev.*, **120**, 2109-2127.
- Skamarock, W. C., P. K. Smolarkiewicz, and J. B. Klemp, 1997: Preconditioned conjugate-residual solvers for Helmholtz equations in nonhydrostatic models. *Mon. Wea. Rev.*, **125**, 587-599.
- Smolarkiewicz, P. K., and L. G. Margolin, 1994: Variational solver for elliptic problems in atmospheric flows, *Appl. Math. & Comp. Sci.*, **4**, 527-551.
- Smolarkiewicz, P. K., and L. G. Margolin, 1997: On forward-in-time differencing for fluids: An Eulerian/semi-Lagrangian nonhydrostatic model for stratified flows. *Atmos. Ocean Special*, **35**, 127-152.

SMOLARKIEWICZ AND MARGOLIN: VARIATIONAL ELLIPTIC SOLVERS

- 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.
- Smolarkiewicz, P. K., and L. G. Margolin, 1998: MPDATA: A finite-difference solver for geophysical flows, *J. Comput. Phys.*, **140**, 459-480.
- Smolarkiewicz, P. K., V. Grubišić, L. G. Margolin, and A. A. Wyszogrodzki, 1999: Forward-in-time differencing for fluids: Nonhydrostatic modeling of fluid motions on a sphere. *Proc. 1998 Seminar on Recent Developments in Numerical Methods for Atmospheric Modelling*, Reading, UK, ECMWF, 21-43.
- Stoer, J., 1983: Solution of large linear systems of equations by conjugate gradient type methods. *Mathematical Programming: The State of the Art*, A. Bachem, M. Grötschel, and B. Korte, Eds., Springer-Verlag, pp. 655.