

DOMAIN DECOMPOSITION FOR THE SHALLOW WATER EQUATIONS

J. G. CHEFTER, C. K. CHU, AND D. E. KEYES

ABSTRACT. Domain decomposition (Krylov-Schwarz) iterative methods are proposed for the parallel implicit solution of the unsteady geopotential equation that arises when semi-implicit, semi-Lagrangian (SISL) methods are employed in the long-time integration of the shallow water equations. SISL methods permit timesteps on the scale of Rossby wave dynamics, in contrast to the small timesteps required to resolve gravity waves, and thus satisfy the stability bound of an explicit method. The price of the semi-implicitness is a global elliptic problem on a multiply-connected semi-periodic domain with variable coefficients that become singular at the poles of latitude-longitude coordinate systems. Elliptic solvers based on domain decomposition offer flexibility in discretization and good algebraic convergence properties. They also provide good data locality with a view towards high-latency coarse- to medium-grained parallelism.

1. THE SHALLOW WATER EQUATIONS

The system of shallow water equations (SWEs) is a hyperbolic problem at the core of many models for the dynamics of the oceans or the atmosphere. In this paper, the SWEs are reformulated for large timestep by putting the computational burden on a scalar elliptic equation addressable through Schwarz-type domain decomposition methods. The three-dimensional Euler equations for a shallow layer of inviscid fluid on a sphere of radius a rotating with angular velocity Ω reduce under the assumptions of constant density, hydrostatic balance in the radial direction, and the Taylor-Proudman hypothesis to:

$$\begin{aligned} (1) \quad \frac{du}{dt} &= (2\Omega \sin \lambda + \frac{u}{a} \tan \lambda)v - \frac{1}{a \cos \lambda} \frac{\partial \Phi}{\partial \varphi}, \\ (2) \quad \frac{dv}{dt} &= -(2\Omega \sin \lambda + \frac{u}{a} \tan \lambda)u - \frac{1}{a} \frac{\partial \Phi}{\partial \lambda}, \\ (3) \quad \frac{\partial \Phi}{\partial t} &= -\frac{1}{a \cos \lambda} \left(\frac{\partial}{\partial \varphi} [(\Phi - \Phi_b)u] + \frac{\partial}{\partial \lambda} [(\Phi - \Phi_b)v \cos \lambda] \right). \end{aligned}$$

1991 *Mathematics Subject Classification.* Primary 65M55, 65Y20, 76C20.

The work was supported in part by Schlumberger-Doll Research (JGC), by National Aeronautics and Space Administration under NASA cooperative grant NCC 5-34 (CKC), by the National Science Foundation under contract number ECS-8957475 (DEK), and by the National Aeronautics and Space Administration under NASA contract NAS1-19480 while the third author was in residence at the Institute for Computer Applications in Science and Engineering.

This paper is in final form and no version of it will be submitted elsewhere.




FIGURE 1. (a) Regular staggered grid employed for (u, v, Φ) . (b) Computational domain.

The unknown fields are u and v , the vertically averaged velocities in the longitudinal (φ) and latitudinal (λ) directions, and the free-surface geopotential, Φ . Potentials $\Phi \equiv gh(\varphi, \lambda, t)$ and $\Phi_b \equiv gh_b(\varphi, \lambda)$, are defined in terms of the height of the top and bottom surfaces, and $\frac{d}{dt} \equiv \frac{\partial}{\partial t} + \frac{u}{a \cos \lambda} \frac{\partial}{\partial \varphi} + \frac{v}{a} \frac{\partial}{\partial \lambda}$ is the horizontal material derivative.

To integrate these equations, we use the semi-Lagrangian scheme first proposed in [7, 8]. The integral of either of the momentum equations (1,2) over one timestep can be approximated to first-order accuracy in time as $\Xi^{k+1} - \Xi_*^k = \xi^{k+1} \tau$, where Ξ^{k+1} is the value of Ξ at the arrival point (φ, λ) of some streamline at time $(k+1)\tau$, Ξ_*^k is the value of Ξ at the departure point (φ_*, λ_*) of the same streamline at time $k\tau$, and $\xi^{k+1} \tau$ is the implicit approximation of the RHS. To find the departure point (φ_*, λ_*) , one needs to integrate backwards along the characteristics, which can be approximated [4] by a sum over N equal intervals $\tau_1 = \frac{\tau}{N}$.

Implicit approximation of the integrals leads to better stability properties of the scheme in comparison with explicit time integration. In fact, explicit schemes for shallow water equations must respect the CFL condition [10]: $\tau < \delta / (|u| + |v| + \sqrt{gh})$, where δ is a cell diameter. The speed of the gravity wave \sqrt{gh} usually poses the most restrictive condition on allowed timestep. For high latitudes this may lead to timesteps of the order of minutes, which are unreasonably small for desired simulation periods. Semi-implicit schemes remove the gravity wave speed from the CFL condition. Used with semi-Lagrangian integration of the material derivative, they often result in unconditionally stable schemes.

Applying semi-Lagrangian integration for momentum equations (1,2) we get a system of two linear equations. Solving it for u^{k+1} and v^{k+1} , substituting them into the implicitly integrated continuity equation (3), and cancelling cross-derivative terms, a self-adjoint elliptic equation for Φ^{k+1} results:

$$\begin{aligned}
 (4) \quad & \Phi^{k+1} \cos \lambda - \eta^2 \cos \lambda \frac{\partial}{\partial \varphi} \left(W^k \frac{\partial \Phi^{k+1}}{\partial \varphi} \right) - \zeta^2 \frac{\partial}{\partial \lambda} \left(W^k \cos \lambda \frac{\partial \Phi^{k+1}}{\partial \lambda} \right) \\
 & = \Phi^k \cos \lambda - \eta \cos \lambda \frac{\partial}{\partial \varphi} (W^k (u_*^k + F^k v_*^k)) - \zeta \frac{\partial}{\partial \lambda} (W^k \cos \lambda (v_*^k - F^k u_*^k)),
 \end{aligned}$$




FIGURE 2. FLOPs versus number of unknowns for $\|r\|/\|r_0\| < 10^{-8}$. 1-ILU(4), 2-ASM, 3-MSM. Both Schwarz methods are based on four longitudinal strips.

and simple formulae for u^{k+1} and v^{k+1} in terms of Φ^{k+1} :

$$(5) \quad u^{k+1} = \frac{1}{1 + F^{k2}} \left(u_*^k - \eta \frac{\partial \Phi^{k+1}}{\partial \varphi} + F^k \left(v_*^k - \zeta \frac{\partial \Phi^{k+1}}{\partial \lambda} \right) \right)$$

$$(6) \quad v^{k+1} = \frac{1}{1 + F^{k2}} \left(v_*^k - \zeta \frac{\partial \Phi^{k+1}}{\partial \lambda} - F^k \left(u_*^k - \eta \frac{\partial \Phi^{k+1}}{\partial \varphi} \right) \right),$$

where $F^k = (2\Omega \sin \lambda + \frac{u^k}{a} \tan \lambda) \tau$, $\eta = \frac{\tau}{a \cos \lambda}$, $\zeta = \frac{\tau}{a}$, and $W^k = \frac{\Phi^k - \Phi_b}{1 + F^{k2}}$.

Our staggered mesh is shown in Fig. 1(a). Using centered five-point space differencing, we get a first-order in time, second-order in space numerical scheme. In matrix form, for naturally ordered $\Phi_{i,j}^{k+1}$,

$$(7) \quad A\Phi = g,$$

where A is a symmetric positive definite block-tridiagonal matrix. Linear stability analysis shows that the numerical scheme is unconditionally stable.

The computational domain with one island is shown in Fig. 1(b). For the simplest case of boundaries parallel to coordinate lines, the momentum equations at the boundaries are reduced to the following mixed-derivative implicit conditions for (4), for boundaries of constant λ and constant φ , respectively:

$$(8) \quad \zeta \frac{\partial \Phi^{k+1}}{\partial \lambda} = -F^k \left(u_*^k - \eta \frac{\partial \Phi^{k+1}}{\partial \varphi} \right), \quad \eta \frac{\partial \Phi^{k+1}}{\partial \varphi} = F^k \left(v_*^k - \zeta \frac{\partial \Phi^{k+1}}{\partial \lambda} \right).$$

By treating the surface height gradients in the momentum equations explicitly we could get simple Neumann boundary conditions, but in our tests these conditions led to instabilities on the boundaries. Boundary conditions (8) are stable, but they make the system (7) slightly nonsymmetric.

Equation (4) with Dirichlet or Neumann boundary conditions is of a form for which overlapping domain decomposed preconditioners exist such that the convergence rate is independent of the resolution and the number of subdomains [3].

2. KRYLOV-SCHWARZ ALGORITHMS

Schwarz-preconditioned Krylov solvers for linear systems, $A\Phi = g$, find the best approximation of the solution Φ in a small-dimensional subspace that is built up from successive powers of the preconditioned matrix on the initial residual. A variety of parallel preconditioners, whose inverse action we denote by B^{-1} , can be induced by decomposing the domain of the underlying PDE, finding an approximate representation of A on each subdomain, inverting locally, and combining the results. Generically, we seek to approximate the inverse of A by a sum of local inverses:

$$(9) \quad B^{-1} = \sum_k R_k^T A_k^{-1} R_k,$$

where R_k is a restriction operator that takes vectors spanning the entire space into the smaller dimensional subspace in which A_k is defined.

The simplest domain decomposition preconditioner is block Jacobi, which can be regarded as a zero-overlap form of additive Schwarz [5]. The convergence rate of block Jacobi can be improved, at the price of a higher cost per iteration, with subdomain overlap and (for many problems) by solving an additional judiciously chosen coarse grid system. Our tests show that even a relatively small overlap of two mesh widths make preconditioner (9) comparable to the popular incomplete LU on a serial computer. In serial, the most natural form of Schwarz iteration is multiplicative, which improves the convergence rate of the algorithm, at the price of coarser parallel granularity, by enforcing sequentiality between the subdomain solves, just as Gauss-Seidel improves on Jacobi. Results are given for both additive (ASM) and multiplicative (MSM) versions of the preconditioner.

3. NUMERICAL RESULTS

The SISL code is built on top of the Argonne PETSc library [6]. Results below are from Sparc10 and Intel Paragon implementations. Figure 2 compares the performance of ILU, ASM, and MSM as left-preconditioners for GMRES. ILU(4) is used in the comparison, since four levels of fill led to the best ILU results. Strip subdomain problems (with overlap of $2h$) are solved exactly with a nested dissection ordering within each subdomain for the Schwarz methods. Even without a coarse grid, exact subdomain Schwarz methods performed better than ILU for this modest granularity, and improve upon ILU not only in the number of FLOPs, but also in the rate of growth with problem dimension.

P	Pois.-Dirichlet		Pois.-Neumann		SISL-with c.g.		SISL-w/o c.g.	
	add.	mult.	add.	mult.	add.	mult.	add.	mult.
4	13	6	51	16	16	7	13	7
16	11	5	30	12	17	8	15	8
64	9	4	18	9	16	8	17	8
256	8	3	15	7	15	8	18	10

TABLE 1. Number of iterations with number of subdomains P , for four problems, additive or multiplicative preconditioning.

We experimented with a two-level Schwarz preconditioner by adding a coarse grid with one degree of freedom per subdomain, with standard bilinear grid transfer

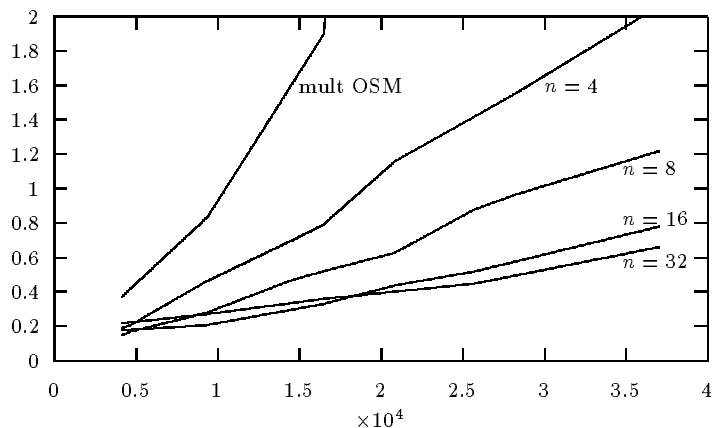


FIGURE 3. Wall-clock time (sec) vs. number of unknowns for $\|r\|/\|r_0\| < 10^{-8}$ on the Paragon, for different numbers of subdomains (processors), n .

operators. To verify its correct operation (convergence rate bounded independent of mesh resolution and subdomain diameter, asymptotically), the Poisson Dirichlet problem in a unit square was solved as in [2] for a range of coarse grid granularities, and the results appear in Table 1, together with convergence results for a Poisson operator with periodic BCs in x and Neumann BCs in y and potential equation (4) with explicitly integrated boundaries, with and without a coarse grid. (The Poisson Neumann problem has a null space of the constant vector, which was eliminated by employing a Dirichlet condition at one point common to both fine and coarse grids.) A box decomposition is used, for a 128×128 grid, iterated until $\|r\|/\|r_0\| < 10^{-5}$.

The effectiveness of the coarse grid for the constant-coefficient operator in the first two pairs of columns can be seen in the decrease in the number of iterations as the granularity of the decomposition increases for both boundary types and both additive and multiplicative preconditioning. Comparison of the last two pairs of columns (for the same SISL problem) shows that the coarse grid does not make a large difference, though the difference widens for large problems. A time step of one hour is used in (4), which makes matrix A only slightly diagonally dominant, so the effect of the coarse grid is not buried in any parabolicity of the problem [1]. The ratio of the diagonal term to the biggest off-diagonal term is about 3 or 4 near equator and goes down to unity near the poles. This aspect of the problem remains under investigation.

4. PARALLEL RESULTS

Parallel implementation of the one-level Schwarz-based preconditioners was carried out on the Paragon. Vectors and matrices are distributed row-wise across processors and all subdomain solves are processed concurrently up to the number of available processors. Unfortunately, Semi-Lagrangian integration turns out to be poorly parallelizable, because the relative location of (ϕ_*, λ_*) in the processor array

is solution-dependent and generally irregular. At present, it is done sequentially, which limits the scalability of the computation.

Figure 3 shows the dependence of elapsed time on the number of the unknowns N for various numbers of subdomains (processors). For zonally dominant (i.e., longitudinally oriented) flows, zonal strip decomposition gave better performance than box decomposition for the problem sizes considered. Sequential multiplicative OSM, which in our experiments has proved to be the most efficient algorithm of all tested, is used for comparison. Its high rate of growth with N for large N reflects thrashing of the memory hierarchy.

5. CONCLUSIONS

By concentrating data dependencies locally, domain decomposition preconditioners exploit the two-level memory hierarchy of high-latency distributed memory architectures. Low-communication zero or small overlaps between the preconditioner blocks are feasible with small convergence rate penalty, at least for intermediate granularities. Demonstrating the applicability of elliptic-based domain decomposition preconditioners to the shallow water equations opens the door to a variety of parallel implicit models in long-time integration geophysics applications.

REFERENCES

1. X.-C. Cai, *Some Domain Decomposition Algorithms for Nonselfadjoint Elliptic and Parabolic Partial Differential Equations*, Ph.D. thesis, Courant Institute, NYU, 1989.
2. X.-C. Cai, W. D. Gropp and D. E. Keyes, *A Comparison of Some Domain Decomposition and ILU Preconditioned Iterative Methods for Nonsymmetric Elliptic Problems*, Numer. Lin. Algebra Applics. (1994), to appear.
3. X.-C. Cai and O. B. Widlund, *Domain Decomposition Algorithms for Indefinite Elliptic Problems*, SIAM J. Sci. Stat. Comput. **13**(1992), 243–258.
4. V. Casulli, *Semi-implicit Finite Difference Method for the Two-dimensional Shallow Water Equation*, J. Comp. Phys. **86**(1990), 56–74.
5. M. Dryja and O. B. Widlund, *An Additive Variant of the Alternating Method for the Case of Many Subregions*, Courant Institute, NYU, TR 339, 1987.
6. W. D. Gropp and B. F. Smith, *Simplified Linear Equations Solvers Users Manual*, ANL-93/8, Argonne National Laboratory, 1993.
7. O. Pironneau, *On the Transport Diffusion Algorithm and Its Applications to the Navier-Stokes Equations*, Num. Math. **38**(1982), 309–322.
8. A. Robert, *A Stable Numerical Integration Scheme for the Primitive Meteorological Equations*, Atmosphere-Ocean **19**(1981), 35–46.
9. Y. Saad and M. H. Schultz, *GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems*, SIAM J. Sci. Stat. Comp. **7**(1986), 865–869.
10. E. Turkel and G. Zwas, *Explicit Large Time-step Schemes for the Shallow Water Equations*, in *Advances in Comp. Meths. for Partial Diff. Eq. III*, R. Vichnevetsky and R. S. Stepleman, eds., IMACS, Brunswick, 1979, 65–69.

DEPARTMENT OF APPLIED PHYSICS, COLUMBIA UNIVERSITY, NEW YORK, NY, 10027-0029.
E-mail address: jgc@appmath.columbia.edu

DEPARTMENT OF APPLIED PHYSICS, COLUMBIA UNIVERSITY, NEW YORK, NY, 10027-0029.
E-mail address: chu@apne.columbia.edu

COMPUTER SCIENCE DEPARTMENT, OLD DOMINION UNIVERSITY, NORFOLK, VA 23529-0162
 AND ICASE, MS 132C, NASA-LARC, HAMPTON, VA 23681-0001
E-mail address: keyes@icase.edu