A Fully Implicit Solution Method Capability in CAM-HOMME

Katherine J. Evans*, Damian Rouson‡, Andy Salinger‡, Mark A. Taylor‡, Amik St. Cyr#

*Oak Ridge National Laboratory, ‡Sandia National Laboratory, #National Ctr. Atmospheric Research

Introduction/Motivation

The inclusion of new physics and chemistry and grid refinement of the Community Climate System Model (CCSM) create new algorithmic challenges including coupled nonlinear multiscale processes and enhanced scalability requirements. To maintain scalability, a number of climate models have been returning to fully explicit methods developed several decades ago. However, finer model grids require a superlinear reduction in the time step size to account for the smaller spatial scale and increased multiscale interactions (Keyes et al., 2006). Fully implicit methods are well established as a more accurate and efficent solution method for a range of multiscale applications because they eliminate the stability restriction for time step size selection and solve all the dependent variables consistently. However with the exception of recent linearized oceanic spin-up models (Bernsen et al., 2008; Li and Primeau, 2008), global Earth system models have not implemented a FI solver capability. The current solver implementation utilizes a Fortran interface package within the Trilinos project, which allows fully tested, optimized, and robust code with a suite of parameter options to be included a priori. Presently, a fully implicit (FI) solution method is applied to several shallow water test cases from Williamson et al. (1992) within the High-order Method Modeling Environment (HOMME) component of the CCSM, and early results are presented.

For test case 1, the structure of an anomaly advected by a prescribed wind field is retained even when using a time step much larger than the gravity wave CFL. To evaluate solver performance for steady state analyses, test case 2 is applied using an FI method preconditioned with an adaptation of the existing semi-implicit method, and both solution time and iteration count is reduced appreciably in some cases. A more realistic test modeling the flow of air over a mountain feature highlights the ability to run accurately over the gravity wave CFL but also the limitations of the current preconditioner.

Fully Implicit (FI) Solution Method

First, the equations are written as a nonlinear residual of the state vector so that all the terms are written and solved coherently. For shallow water eq'ns, $\mathbf{x} = \{u, v, h\}^T$

$$F(\mathbf{x}) = 0$$

Taking the first term of a Taylor series expansion at $\mathbf{x}^k$ gives a linearized update for F as the Jacobian, J, times the update. Solving for an $\mathbf{x}^{k+1}$ that satisfies the $F(x)$ within a specified nonlinear tolerance is inexact Newton's method.

$$F(\mathbf{x}^k) = -J(\mathbf{x}^k)\delta\mathbf{x}$$
$$\mathbf{x}^{k+1} = \mathbf{x}^k + \delta\mathbf{x}$$

The update, δx, is found using GMRES, a linear solution method that builds a Krylov vector. Rather than explicitly calculating the Jacobian to get an update for $F(x)$, a finite difference approximation is used. Generally, more iterations are needed, but the end result is more efficient. The efficiency of the linear solver can be enhanced with the use of a preconditioner.

$$J\mathbf{v} \simeq \frac{F(\mathbf{x} + \epsilon\mathbf{v}) - F(\mathbf{x})}{\epsilon}$$

Once a Jacobian-vector multiply is built, it is evaluated to see if it satisfies a linear tolerance requirement, η_l, upon which the latest δx is sent to be evaluated in the nonlinear residual calculation.

$$\frac{|J\delta\mathbf{x} + F(x^k)|_2}{|F(x^k)|_2} \leq \eta_l$$

Semi-Implicit (SI) Solver as a preconditioner

The preconditioner can be written as an operator matrix, $\tilde{M}$. Then, the Jacobian-vector update is multiplied by $\tilde{M}^{-1}\tilde{M}$. Thus $\tilde{M}^{-1}$ is applied once to update the Krylov vector, and then again to return the state vector update δx.

$$J\tilde{M}^{-1}(\tilde{M}\delta\mathbf{x}) = -F(\mathbf{x})$$
$$\text{set} \quad \tilde{M}\delta\mathbf{x} = \delta\mathbf{z}.$$
$$J(\tilde{M}^{-1}\delta\mathbf{z}) = -F(\mathbf{x})$$

In order to minimize code implementation, the existing SI solver is adapted for use as preconditioner. It can be designed to capture the relevant linear physics for the problem of interest and has been used successfully to enhance other shallow water equation models (Mousseau et al., 2002).

Once Kyrlov update sufficient,

$$\tilde{M}^{-1}(\delta\mathbf{z}) = \delta\mathbf{x}$$

When used as a solver, the semi-implicit method solves the advection and background flow terms explicitly and the gravity wave terms implicitly, so a time step above the gravity wave time scale can be taken (Thomas and Loft, 2002). The benefit of splitting this way is that the implicit terms can be rewritten as the Helmholtz equation of the geopotential height, the solutions of which can be found efficiently using the congugate gradient method. The implicit terms are the terms retained in the current application of the preconditioner used here.

$$\frac{\mathbf{u}^{t+1} - \mathbf{u}^{t-1}}{2\Delta t} = -(\omega^t + f)\hat{k} \times \mathbf{u}^t - \nabla\left(\frac{1}{2}(\mathbf{u}^t)^2 + \underline{gh_s} + \underline{\frac{1}{2}(gh^{t+1} + gh^{t-1})}\right)$$
$$\frac{h^{t+1} - h^{t-1}}{2\Delta t} = -\nabla \cdot h^t\mathbf{u}^t - \underline{h_o\nabla \cdot \frac{1}{2}(\mathbf{u}^{t+1} + \mathbf{u}^{t-1})}$$

Implicitly discretized terms (Crank-Nicolson)

Steps for gravity wave preconditioner:

1. Recast u and h in terms of updates δh and δu, and retain only gravity wave terms (underlined above)
2. Solve multiply δu equation by the gradient operator, $\nabla\delta u$
3. Substitute grad u into height equation to get a Helmholtz equation for δh
4. Solve using CG to get δh, then sub into δu equation and return values to FI above

An alternative dycore option within the Community Climate System Model (CCSM): High-Order Methods Modeling Environment (HOMME)

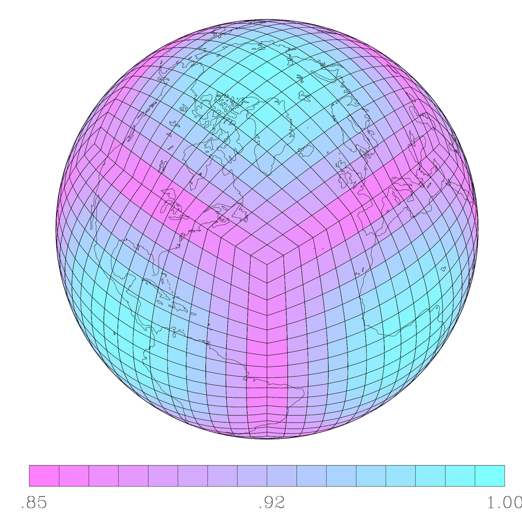

Spectral Element Spatial Discretization

1. Domain: 6 cube faces mapped to the sphere and tiled into lat-lon elements
2. Within each element, variables are approximated by polynomial expansions
3. Communication is only needed at the element edges
4. Mesh refinement: add elements or increase order of spectral degree

Combines favorable aspects of two discretizations onto a cubed sphere grid

Spectral Transform Methods

High order accuracy
High convergence rate

Finite Element Methods

Geometric Flexibility
Minimal Communication

*Developed at NCAR as part of DOE's Climate Chance Prediction Program (CCPP), HOMME has exhibited superior scalability and efficiency. The spectral element formulation as applied in HOMME is the only dycore within CAM that locally conserves mass and energy

Shallow Water: Test Case 1

Advection of a cosine bell curve using prescribed velocities

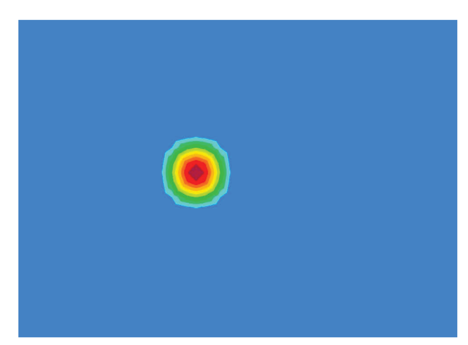

Initial, t=0

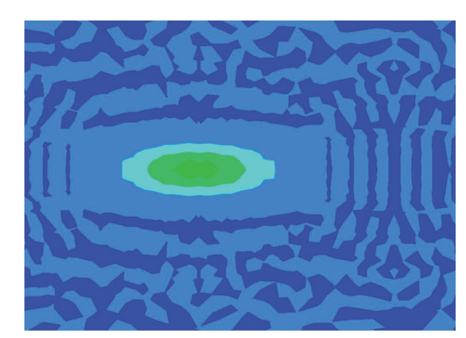

1st order, Fully Implicit Solve

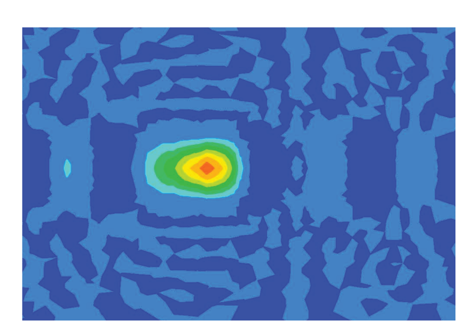

2nd order Filtered Leapfrog

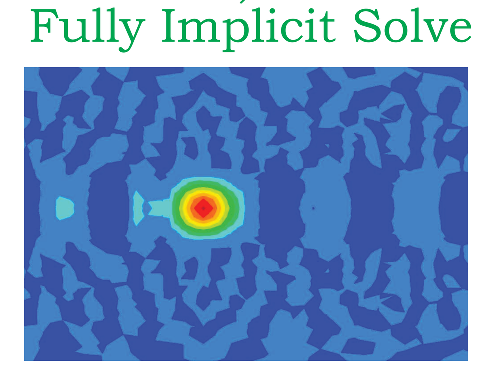

2nd order, CN Fully Implicit Solve

Plots of the cosine bell anomaly on a course grid (~500km spacing) with a 1 hour time step, which is close to the limit at which explicit leapfrog is stable. Notice that the 1st order method (BE) is overly diffusive. Both leapfrog (LF) and Crank Nicolson (CN) are 2nd order accurate, but leapfrog must be filtered to maintain stability, which affects accuracy. All the methods experience some spectral 'ringing' given such a coarse grid. However this configuration highlights the accuracy issues with various time discretization strategies.

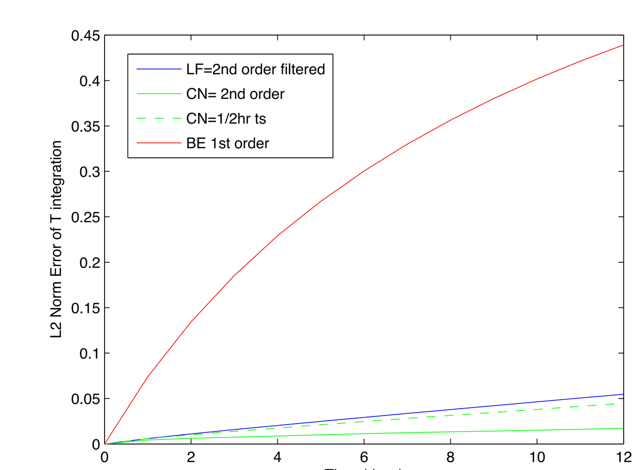

The L2 norm of error for runs using the three methods mentioned above for a finer grid (~140km) after the anomaly has been advected around the sphere. The leapfrog (LF), backward euler (BE), and Crank Nicolson (CN) are plotted with a time step size of 16 minutes (solid lines), close to the limit at which explicit leapfrog is stable. Not surprisingly, the first order method (BE) has the highest error. A run with a time step size of 1/2 hour is also performed using CN (dashed green line), and still the L2 error is less than LF.

Summary

The FI solution framework acts as an accelerator for existing methods in the HOMME model that require smaller time steps on relatively fine grids, and allows new time stepping methods to be developed. The solution is verified to be as accuate as other methods even with larger time steps, and can produce solutions more quickly than the explicit (or a semi-implicit formulation, which published results show 2-3 times increased efficiency compared to explicit (Thomas and Loft, 2002)) for early results with a preconditioner. Initial work with test case 5 has produced solutions with the FI method as well, but the preconditioner does not effectively reduce iterations and increase efficency to make it a competitive option. A high fidelity preconditioner that addresses the advection terms is needed and is currently being developed.

Citations:

Bernsen, E., Dijkstra, H.A., and F.W. Wubs, (2008). "A method to reduce the spin-up time of ocean models." J. Comp. Phys. **130**: 92-108.
Keyes, D.E., Reynolds, D.R., and C. Woodward, (2006). "Implicit solvers for large scale nonlinear problems." J. Phys. Conf. Series **46**: 433-442.
Li, X. and F.W. Primeau, (2008). "A fast Newton-Krylov solver for seasonally varying global ocean biogeochemistry models." Ocean Modelling **23**: 13-20.
Mousseau, V.A., Knoll, D.A., and J. Reisner, (2002). "An Implicit Nonlinearly Consistent Method for the Two-Dimensional Shallow-Water Equations with Coriolis Force." Mon. Wea. Rev. **130**: 2611-2625.
Thomas, S.J., and R. Loft, (2002). "The NCAR Spectral Element Climate Dynamical Core: Semi-Implicit Formulation." J. Sci. Comp. **25**: 307-321.
Williamson, D.L., and J. B. Drake, J. J. Hack, R. Jakob, and P.L. Swarztrauber, (1992). "A Standard Test Set for Numerical Approximations to the Shallow Water Equations in Spherical Geometry." J. Comp. Phys. **102**: 211-224.

Acknowledgements:

Research sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the US Department of Energy under Contract No. DE-AC05-00OR22725 and the DOE/BER SciDAC Project: A scalable and extensible Earth system model. HOMME development is funded as part of DOE's Climate Change Prediction Program (CCPP) program at the National Center for Atmospheric Research (NCAR).

LDRD

Shallow Water: Test Case 2

Steady state nonlinear zonal flow

(1) 4x12x12* (avg grid=227km, min grid=69km)

Time Integration Method	Time Step Size	Simulation Time (Relative to Explicit)	L2 Error	Solver Stats (see text) NL/Lin per NL
Explicit Leap Frog (2nd order)	120s	≡1	2.1e-12	N.A.
Fully Implicit CN (2nd order)	86400s (1 day)	0.18	5.7e-14	1/62
Fully Implicit CN (w/precon)	86400s (1 day)	0.19	2.8e-16	2/4

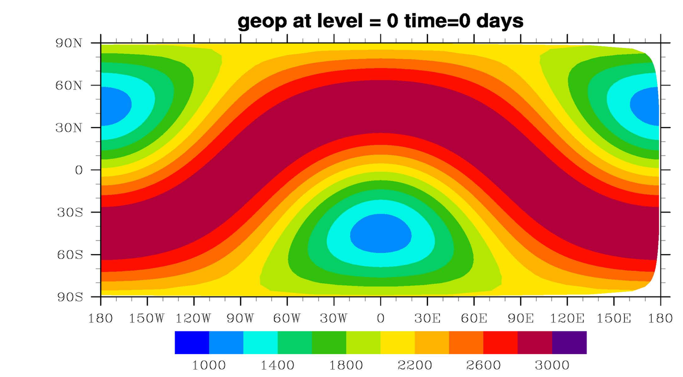

(2) 15x4x4* (avg grid=222km, min grid=184.5km)

Time Integration Method	Time Step Size	Simulation Time (Relative to Explicit)	L2 Error	Solver Stats (see text) NL/Lin per NL
Explicit Leap Frog (2nd order)	240s	≡1	6.6e-7	N.A.
Fully Implicit CN (2nd order)	21600s (6 hr)	6.7	5.3e-7	6/84.7

*For spectral element grids, MxNxN notation is used, where M is the number of elements in each direction on each cube face and N refers to the spectral degree order within each element. Grids (1) and (2) to the left have similar average grid spacing, but (2) has a larger spectral discretization and thus higher accuracy but smaller time step contstraints for the explicit method. NL is the # outer Newton iterations, and Lin per NL refers to the # of linear Krylov iterations per NL.

Test case 2 solves the steady state nonlinear flow and is designed to evaluate performance of the dycore and solution method. It also tests methods that would assist with model spin up and parameter continuation studies, where a number of steady state simulations are desired. Because the solution is a steady state, FI can run at an arbitrary time step size without sacrificing error, and thus complete the simulation relatively more efficiently upon grid refinement. For example, in configuration (1), the implicit method can take 1 time step over the length of the simulation and thus save 80% time over the extra cost of iteration and forming a preconditioner. The number of iterations in this case due to the preconditioner is reduced from 62 to 8 (per time step). Because this is a simple test problem, the time of set up is most of the work and thus the time is not reduced appreciably.

Future Plans for Implicit Methods in HOMME

Shallow Water: Test Case 5

More realistic case: Zonal flow over an isolated mountain

This case uses an equatorally aligned zonal flow in the presence of an isolated mountain feature. As with the other test cases, the FI method does not require a smaller time step with grid refinement to maintain stability. For a 4x12x12 grid (see above for description), a 20 minunte time step (verses 2 min for explicit) can be taken without a loss resolution of the physical processes of the system. However, early runs of this test case show that unlike test case 2, the Helmholtz preconditioner does not improve the implicit method, possibly because of the advection dominated behavior, which is not included in the linear inversion of the Helmholtz. Several strategies for an improved preconditioner are being pursued. For current numerical methods in HOMME, there is an imposed numerical diffusion to keep the KE spectrum reasonable. This issue will be also addressed with the ongoing accuracy and efficiency assessment for the FI method.

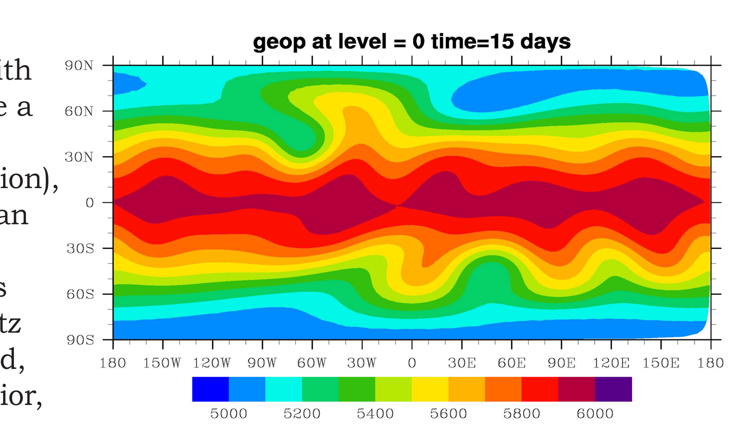

Improved Preconditioning strategies

1. Include more linear physics in the SI preconditioner as model solves more complicated problems
2. Use Restrictive Additive Schwarz preconditioner; used in spectral element models with success
3. Use Multi-level preconditioner; complementary to implementation of time-parallel capability

Further Acceleration: Parallel-in-time Integration

1. Take large implicit time steps (t_N) to span entire simulation quickly
2. Assign each processor one time subset covering a large time step
3. In parallel, integrate at a finer time scale (ticks) to resolve relevant scale
4. Correct discrepancy between fine and coarse simulations (Δ_N)
5. Translate corrections along time trajectory (starred points)
6. Repeat until convergence (typically 2-3x)