Global Atmospheric Dynamical Core Assessment

28 October 2008

This is an assessment of the state-of-the-art in global solvers for atmospheric fluid flow by the participants of the Global Atmospheric Core Workshop held in Boulder on 23-24 September 2008. The workshop was hosted by the National Center for Atmospheric Research and formally was a meeting of the Atmospheric Dynamical Core Working Group that was formed as part of the effort to develop a new combined climate and weather modeling system.

Summary

We are evaluating state-of-the-art approaches for dynamical cores in a new combined weather and climate modeling system. A consensus exists for the necessary attributes of this solver - it must be nonhydrostatic, global, solve the deep atmosphere equations, conserve mass and scalar mass, have local refinement and limited-area capabilities, be capable of relatively uniform resolution over the globe, have positive-definite and monotonic transport options, and must be efficient on planned petascale MPP machines. Exact energy conservation, while potentially desirable, is not viewed as a strict requirement. Existing operational models use latitude-longitude grids. They have some problems with the poles and critically will not scale well to MPP architectures. The development efforts that are moving forward most quickly are converging on two classes of grids for the sphere - cubed-sphere grids and icosahedral grids using finite-volume and finite difference solvers. Efforts have not completely converged - there are problems with all of the approaches but solutions are actively being explored. There are also development efforts centered for hydrostatic solvers based on Galerkin techniques using cubed-sphere grids. Viable nonhydrostatic solvers using these methods are far off. Test cases for nonhydrostatic (and hydrostatic) solvers need further development. There is a relative dearth of discriminating test results for models and, equally important, there are almost no computational costs reported for these potential cores that would allow for a determination of solver efficiency that is necessary for evaluating candidate architectures for applications.

1. Motivation

The primary motivating factors for constructing a new dynamical core for weather and climate applications are:

(1) Existing global climate and weather dynamical models are based on latitude-longitude grids, and these solvers do not scale well to existing and planned MPP architectures using O(10^4) - O(10^6) processors.

(2) As we move to higher resolutions in global simulation we need to solve the nonhydrostatic equations of motion. The few global nonhydrostatic solvers existing today are either on latitude-longitude grids and do not scale well, or possess other problems that need to be addressed.

The major problems faced in developing a global nonhydrostatic solver center around designing a grid and a solver for the equations of motion that scales well on MPP machines, is efficient and produces solutions where the grid is not visible in practical applications.

2. Priority Applications

Priority applications of the atmospheric core need to be identified in order to appropriately design and test the core and in order to make recommendations and/or decisions about any future weather and climate model dynamical core. There are number of possible priority applications, and the most challenging is to produce global, nonhydrostatic cloud-resolving simulations for weather prediction and climate-process studies. At the workshop, it was also noted that any core would have to work well and be efficient for climate applications at resolutions much coarser than nonhydrostatic resolutions, e.g. IPCC simulations (dh ~ O(100 km)).

3. Dynamical Core Requirements

There was a general consensus at the workshop on the following requirements for a dynamical core.

(1) The solver should integrate the fully compressible nonhydrostatic equations of motion. This requirement is driven by the needs of weather, regional climate, and climate-process studies. Nonhydrostatic simulation requires solving the

nonhydrostatic equations. The validity of approximate equation sets (anelastic, pseudo-compressible, etc) is potentially limited, and the discrete solvers for the approximate equations are often not any simpler than compressible Navier-Stokes solvers. During the workshop there was some limited discussion of the development of more-accurate filtered equation sets (filtering out acoustic modes), and that these developments should be watched closely.

(2) The model should integrate the deep-atmosphere equations. In this regard it was noted that the spherical approximation for the shape of the earth was acceptable, but many present and future applications require simulating a deep atmosphere (perhaps hundred of kilometers), hence the shallow atmosphere approximation (dh not a function of r) is not sufficiently accurate.

(3) The model must be applicable to the globe. This is a climate and weather requirement.

(4) The model must be suitable for regional (limited area) modeling using prescribed lateral boundary conditions. This is a regional climate and weather requirement.

(5) The model must have local refinement capability for both regional and global applications. For example, 2-way interactive nesting is available in many modeling systems to meet this requirement (e.g. ARW), but local refinement can be attained in a number of ways. This capability is needed for weather, regional climate, and global climate studies.

(6) The discrete solver must conserve mass. This is important for climate applications, and is the basis for scalar conservation for all other applications.

(7) The discrete solver must conserve scalars and the scalar transport solver should be consistent with the discrete mass conservation equation. This is needed for accurate water, aerosol, and trace gas (chemistry) budgets, i.e. for most applications.

(8) Positive definite transport must be available for scalars. Water, trace gases, etc. are positive-definite (PD) quantities. Most model physics can't handle negative values. PD transport is a physically consistent way to ensure that PD quantities stay PD.

(9) The horizontal grid should be configurable such that it is relatively uniform (grid-cell area should not vary significantly).

(10) The horizontal grid should be configurable such that it is relatively isotropic (dx ~ dy). These last two items express the need for uniform resolution over the globe and isotropy in the numerics.

(11) The model should be reasonably efficient on various existing and proposed supercomputer architectures (time to solution for a given accuracy, understanding that specific efficiency and accuracy measures may be application dependent). This final point just reiterates the need for the solver to be able to run reasonably well on all potential future computing platforms.

It is also desirable that monotonic scalar transport options be available. Some view this as an absolute requirement, but as pointed out by others the level of implicit diffusion associated with these schemes can be substantial.

4. Higher Order Conservation

There has been considerable discussion in the community, and at the workshop, concerning the need for conservation of various other quantities such as energy, various forms of vorticity, etc. Conservation may be desirable, but there are differences of opinion concerning whether or not conservation of a number of higher order moments should be considered a strict requirement. Thuburn (2008) presents an overview of the issues and a framework in which to consider them (see John Thuburn's presentation at http://www.mmm.ucar.edu/projects/global_cores/Thuburn.pdf)

Energy conservation can be achieved a number of ways in a solver. For coupled climate simulations (atmosphere + ocean), energy conservation appears to be crucial for proper climate predictions. In this case it is important to include the atmospheric heating (increase in internal energy) arising from the dissipation of energy and enstrophy (e.g. the explicit and implicit mixing and damping in the momentum equations). Non-conservative solvers handle this problem using energy *fixers* that add in this dissipational heating. Perfect energy conservation may not be obtained, but the combined accuracy and conservation levels appear to be acceptable. Dissipational heating is usually not accounted for in NWP models, but it is thought to be important in tropical cyclone evolution.

Energy conservation in nonhydrostatic solvers can be achieved by integrating a total energy equation and diagnosing the internal energy (i.e., NICAM (Tomita and Satoh, Fluid Dyn. Res., 2004, Vol 35, 357-400), see also George Bryan's workshop presentation (http://www.mmm.ucar.edu/projects/global_cores/Bryan.pdf)). One concern that was raised with this approach is that, if the transfer of energy between potential, kinetic and total energy is not exactly conserved in the model, then an unphysical evolution of the solution enthalpy may occur. There have been examples of this in gas-dynamics simulations, but it is not been observed (or examined) in solvers using this approach to energy conservation. There are similarities between integration techniques that conserve energy and the use of energy fixers. In this regard, it is unclear whether or not solvers that use energy fixers (previous paragraph) might also suffer from this problem.

5. Numerical Methods

Here we consider some aspects of the nature of the numerical integration schemes that are not directly tied to the spatial discretization of the sphere.

Elliptic equation solvers

One aspect of solvers that can impact MPP performance is whether or not global inversions are needed in the solutions process (i.e. solution of a 2D horizontal or 3D elliptic equation, usually as a result of using a temporally implicit time integration technique or to satisfy some constraint, or the need to diagnose some quantity from another, for example, velocity from PV).

Scalability on MPP machines is a major issue for these solvers, and multigrid methods appear to show some promise on MPP machines in this regard (see Ross Heike's workshop presentation at http://www.mmm.ucar.edu/projects/global_cores/Heikes.pdf). The cost of multigrid methods scale with the number of points discounting parallelization, and parallelization using common domain decomposition techniques may be adequate to maintain efficiency and allow for good scaling. Problems with this approach include poor performance when one no longer has a simple coefficient structure in the elliptic problem (for example, in the Helmholtz equation for pressure found in solvers using terrain following coordinates in semi-implicit nonhydrostatic formulations).

When the simple coefficient structure is lost (full Jacobian), algebraic multigrid solvers could be used. There was some discussion concerning whether or not this is viable for a full atmospheric model, and the primary question centers around the applicability of the existing algebraic multigrid software packages and if they possess sufficient flexibility and robustness for these applications. Another alternative that was mentioned is to use multigrid as a preconditioner for conjugate gradient schemes.

Semi-implicit solution techniques that use finite volume and finite difference methods for integrating the nonhydrostatic equations are fairly well developed - they are used in a number of operational NWP and climate models. There has been some development of semi-implicit solution techniques for use with spectral element global models, but these efforts are still in their early stages. Nonhydrostatic solvers using Garlerkin methods are still being developed - first using fully explicit methods; semi-implicit formulations have yet to be explored.

Combined explicit and semi-implicit solvers

Solvers for the nonhydrostatic compressible equations have been constructed using finite volume or finite difference methods using some combination of implicit and explicit techniques to overcome the timestep restriction arising from the fast acoustic modes. All the solvers use implicit integration techniques for the vertically propagating acoustic modes because the vertical grid spacing is always much smaller than the horizontal grid spacing.

Some models use a fully explicit step to integrate the horizontally propagating modes (acoustic and gravity waves, and their advection). Examples of these formulations presented at the workshop include the FIM model (Jin Lee's presentation), the OLAM model (Bob Walko's presentation), and the nonhydrostatic version of the S.J. Lin's FV model (Bill Putnam presented results from a hydrostatic version in his presentation).

Some nonhydrostatic models use what is known as a split-explicit integration technique whereby the horizontal acoustic and gravity-wave modes are integrated with a small explicit timestep while the advection terms and other terms associated with slow processes are integrated with a larger explicit timestep. An example of this approach is given in the description of the global WRF model integration scheme in George Bryan's presentation, and in the NICAM model.

Other nonhydrostatic models use a semi-implicit time integration scheme whereby both horizontally and vertically propagating acoustic modes, and horizontally propagating gravity waves, are integrated with an implicit scheme, usually time centered. Transport and other processes are integrated explicitly, hence their description as semi-implicit. Models using this approach need to solve a 3D Helmholtz equation (see the discussion in the previous section). An example of this approach is the Unified Model model developed by the UK Meteorological Office.

Semi-Lagrangian (SL) methods

A number of the operational global NWP and climate models use semi-Lagrangian (SL) solvers. There are active development efforts in this area, but there were no presentations about them at workshop. For SL solver, areas of concern are conservation properties and accuracy. SL methods are naturally cast in advective (non-conservative) form, hence conservation of first-order quantities does not directly fall out of the formulations. Cell-integrated SL methods represent one approach to gaining mass and scalar conservation. The accuracy of simulated gravity waves in SL simulations has also been of some concern. SL methods are usually used with timesteps that are much larger that those used in explicit Eulerian models (this timestep difference is often as large as a factor of 5 to 10). The time integration accuracy with respect to the trajectory calculations and the integration of forcing terms along the trajectories has been the focus of some development efforts.

Grid staggering in spatial discretizations

In schemes using finite volume and finite difference approaches, choices must be made concerning where variables are defined on the grid (staggering). Three different horizontal staggered schemes that appear in nonhydrostatic formulations - the unstaggered (A) grid, where all variables are defined at the cell center, the C-grid staggering, where the prognostic velocities are defined at cell faces and are normal to the faces (and all other variables are defined at the cell center), and the D-grid staggering, where the prognostic velocities are defined at cell faces and are tangential to the cell faces. C-grid formulations are most popular for nonhydrostatic cloud and mesoscale models - where horizontally divergent modes are important (The UKMO/Hadley Centre nonhydrostatic model is a example). A-grid (unstaggered) formulations (NICAM being an example of a nonhydrostatic formulation, the FIM model also uses an unstaggered grid) perform well for rotational modes because the velocities are collocated and the Coriolis terms require no averaging, but perform less well for divergent modes. The D grid (e.g. S.J. Lin's FV formulation) is not optimal for either divergent or rotational modes, but it the vertical vorticity (PV) is naturally defined at the mass (scalar) points, hence it is straightforward to maintain correlations between scalars and PV. A general discussion of grid staggering and the linear response of these grids using second-order centered discretizations is given in Randall (1994). The choice of grid staggering may have a significant impact on the filtering necessary for robust model performance. The impacts of the grid staggering is at the truncation level, so for higher-order methods the errors should be reduced.

Filtering

All models needs filters because of the downscale cascade of enstrophy and energy. Discussions of various aspects of filtering arose throughout the workshop.

For the spectral element (Galerkin) based solvers, it was pointed out by Mark Taylor that is is important to formulate any spatial filters in physical space because of the uneven spacing of nodal points on the elements - failure to do so results in strong grid-imprinting in the solutions.

Many models use dissipation that is implicit in their formulation (e.g. damping characteristics in advection schemes), while other models parameterize the dissipation explicitly using some eddy viscosity or hyperviscosity formulation. The main problems are that there is no theoretical justification for either approach in applications using grid spacings coarser than a hundred meters, and one has little control of the damping when it is implicit in the schemes.

Different filters are being used in various nonhydrostatic finite-volume models. Monotonicity constraints in transport schemes can produce a significant amount of damping. Phil Collela (see his workshop talk at http://www.mmm.ucar.edu/projects/global_cores/Collela.pdf) pointed out that care must be taken in the limiting process so as not to damp isolated maxima and minima (one of the motivations behind WENO type schemes). The FV core (see Bill Putnam's presentation at http://www.mmm.ucar.edu/projects/global_cores/Putnam.pdf) uses monotonicity in the transport operators (that are also used as interpolators) and horizontal divergence damping as filters in the D-grid formulation. Moving to higher-order divergence-damping filters was shown to improve the accuracy of the model. Kinetic energy spectra was shown in presentations by Bill Putnam and Mark Taylor as a way of examining filter characteristics in addition to demonstrating the physical correctness of the solutions. It was also pointed out that the A-grid NICAM model is using a 4th order 3D divergence damping filter (to filter acoustic modes as opposed to gravity waves that are filtered with horizontal divergence damping). D-grid and A-grid primitive-variable models need more filtering to control divergent modes (2D and 3D) than the C-grid models, which explains the preponderance of C-grid formulations in cloud models.

Vertical coordinates

There are a large number of options for the vertical coordinates and discretization. In his workshop presentation, Dave Randall outlined the rational for a hybrid isentropic/terrain-following vertical coordinate. The isentropic coordinate results in the horizontal coordinate surfaces being material surfaces, hence the solutions are not degraded by errors introduced by vertical advection. In the FV core, the vertical integration is Lagrangian - horizontal surfaces are material surfaces. The coordinate surfaces are occasionally remapped to prevent severe distortions or singularities from forming.

An area of significant research effort concerns the question of how to handle the lower boundary. Terrain-following coordinates are used in many models, and their formulation is well understood. The existence of large errors in calculation of horizontal pressure gradients above steep terrain has led many groups to develop techniques where the physical terrain cuts through grid cells. These approaches go by a variety of names, including shaved-cell, cut cell, immersed boundary methods, etc.

The literature has numerous papers expounding on the strengths and weaknesses of these approaches.

6. Dynamical Solvers

There are a number of ongoing efforts to develop dynamical cores based on various spatial grid structures. In this section we consider some existing solvers, and solvers under development, based on their horizontal discretization of the sphere.

Latitude-Longitude grids

No presentations were given at the workshop concerning development efforts using latitude-longitude grids.

Advantages: The grids are orthogonal and regular. Fast/accurate solvers are available in a variety of formulations. Local refinement is straightforward. A considerable experience base exists using these approaches.

Disadvantages: The grids are highly anisotropic near the poles, the pole is a singularity, the grids have major resolution variance, and most solvers need polar filters of some sort. The polar filters (and the non-local basis functions used in spherical-harmonics-based models) present major scaling problems for MPP architectures.

Numerics: We could fill a page with references - spherical harmonics, Semi-Lagrangian Semi-Implicit (SLSI), finite difference and finite volume. These are mostly hydrostatic formulations, but there are a number of nonhydrostatic models that have appeared over the last decade.

Examples: Almost every production climate and NWP model is based on lat-long grid.

Future: Most development groups are looking at other approaches (grids) given the difficult-to-overcome disadvantages listed above. There are some groups using reduced grids (where cells are removed as one approaches the pole in such a way as to keep the longitudinal grid spacing from becoming too small).

Icosahedral grids (triangles)

Bob Walko presented recent results from his experience developing the OLAM model. OLAM is based on an icosahedral grid using triangles (see his workshop presentation at http://www.mmm.ucar.edu/projects/global_cores/Walko.pdf). In the discussion the main advantages and disadvantages were considered (see below). Bob Walko pointed out that selective refinement, easily implemented with triangles (dividing triangles is straightforward), is also straightforward to implement for an icosahedral-grid model based on hexagons by just moving to the dual grid (the hexagons) from the refined triangle-based grid. The advantage of using triangles is their simplicity, whereas using hexagons allows for greater solution accuracy (e.g. solution isotropy).

Advantages: The grids can be very isotropic and possess very uniform resolution. They are easy to refine globally and locally.

Disadvantages: High-order discretization is difficult but can be accomplished by using multidimensional polynomial fitting. Solutions are usually less isotropic than either those from rectangular grids or hexagonal grids, and the timestep more restricted than for rectangular or icosahedral grids of the same resolution. With regards to coding and software engineering, more complex coding and data structures are needed to take advantage of arbitrary local refinement capabilities.

Numerics: Finite-volume and finite-difference and are being investigated.

Examples: OLAM (Robert Walko, Duke Univ.) - full 3D solver for the sphere, with local refinement (papers in press); OMEGA model (David Bacon, SAIC) - nonhydrostatic, in the private sector.

Future: The OLAM model is being actively developed, and is at a stage where it can be tested at high resolution on the sphere. The OMEGA model has existed for quite some time, but it is not in the public domain and little documentation exists. There are a few other development efforts going on over in Europe that are still in an early prototype (SW) stage.

Icosahedral grids (hexagons)

At the workshop, Jin Lee (NOAA/GSD) gave a presentation concerning the FIM (Flow-following finite volume Icosahedral Model) and NIM (Nonhydrostatic Icosahedral Model) models (see his workshop presentation at http://www.mmm.ucar.edu/projects/global_cores/Lee.pdf). Joe Klemp presented a few slides showing SW tests and 3D nonhydrostatic cloud model results, both on perfect hexagonal C-grids (see http://www.mmm.ucar.edu/projects/global_cores/Klemp.pdf).

Advantages: The grids are very isotropic and have very uniform resolution. Global refinement is relatively straightforward.

Disadvantages: High-order discretization is difficult (using higher-order multidimensional polynomial fits, etc). 12 special points (pentagons) need special handling. As with the triangle-based grids, more complex coding and data structures are needed to take advantage of arbitrary local refinement capabilities.

Numerics: Finite difference and finite volume models exist, and they are low-order schemes (2nd order at best). Groups are working on higher-order transport and better 2nd order schemes.

Examples: Primary example is nonhydrostatic NICAM model (Tomita and Satoh, Fluid Dyn. Res., 2004). The DWD has used a hydrostatic model on this grid for global forecasting for some time. Dave Randall's group at Colorado State University and Jin Lee at NOAA Global Systems Division are developing models using this grid (hydrostatic at present, nonhydrostatic solvers are in design stage). Todd Ringler, John Thuburn, Joe Klemp and Bill Skamarock are developing C-grid discretizations (currently testing limited-area cloud model and global shallow-water model formulations). They have developed a solution for the zero-frquency geostrophic mode that alleviates the major impediment to using the hexagonal C grid.

Future: There are very active ongoing development efforts (see Examples), most using finite-volume techniques. The efforts are focused on increasing the accuracy of the discretizations and handling the pentagons, with some effort towards incorporating local refinement.

Cubed sphere grids

There were four presentations concerning cubed-sphere grids: Mark Taylor presented results for a spectral element (Galerkin) hydrostatic solver, Bill Putnam presented some results from the hydrostatic FV (finite-volume) core, Ram Nair presented approaches using discontinuous Galerkin methods, and Phil Collela discussed high-order finite-volume methods. Phil Collela brought up the point that cubed-sphere methods can be considered as one type of mapped multiblock algorithms. He argued that it was important to use high-order methods (at least of order 4) so that the inevitable reduction in accuracy at the cube boundaries is tolerable.

Advantages: The grid can be relatively isotropic with relatively uniform resolution. Local refinement is possible.

Disadvantages: Possibly non-conformal and non-orthogonal. The 8 special points (the corners of the cube) and the boundaries between faces need special handling except for the case of Galerkin approaches that are designed for unstructured grids (these handle all elements identically).

Numerics: Finite-volume, finite-difference and various Galerkin-based approaches are being investigated.

Examples: There are two nonhydrostatic implementations - the MIT GCM (John Marshall's group), and S.J. Lin's core. Both use finite-volume based solvers and both use a vector-invariant form of the momentum equations. There are no Galerkin-based nonhydrostatic solvers. There are a number of hydrostatic implementations.

Future: There are significant ongoing development efforts in the US using a variety of approaches. Only those using finite-volume techniques integrate the nonhydrostatic equations. Quantifying the accuracy of (overcoming) the low-order nature of the schemes and treating the corner points are the focal points of these efforts.

Composite grids (e.g. Yin-Yang grids)

Bill Skamarock briefly reviewed the main issues confronting development efforts using this approach (see http://www.mmm.ucar.edu/projects/global_cores/Skamarock.pdf). One must construct a method to stitch together the component grids, usually involving some interpolation of fluxes and values between the grids in or next to the overlap regions. The interpolation of fluxes can be designed in such a way as to globally and locally conserve first-order quantities in the solver, but no one has yet developed methods that are more accurate than first order. The overlap region is obvious in the solutions, and higher-order approaches are needed. It is not clear if a general (viable) solution to this problem exists.

Advantages: relatively uniform resolution, conformal and orthogonal. Local refinement is possible. Existing solvers could be used.

Disadvantages: boundaries (overlapping regions) need special treatment, refinements through these boundaries would also need special treatment.

Numerics: Finite-volume and finite-difference discretizations have been used with these grids.

Examples: A nonhydrostatic Yin-Yang model is being developed at the Earth Simulator Center in Japan. Another Yin-Yang model variant is the global version of MM5 that uses 2 polar-stereographic projections stitched together at the equator; this formulation is not conservative.

Future: Development at the Earth Simulator Center continues. The low-order nature of the overlap conditions leads to noise at the interface. It is not obvious how to formulate transparent conservative overlap conditions. Bill Skamarock stated that he has been able to develop conservative high-order flux interpolation schemes but that instabilities exist when using solvers that are neutral (non-damping).

7. Test Cases

Christiane Jablonowski presented a number of test cases for solvers on the sphere (testing only hydrostatic scales) and some test cases for nonhydrostatic solver using two dimensions (x,z). Her presentation can be found at http://www.mmm.ucar.edu/projects/global_cores/Jablonowski.pdf A reasonable set of test cases exists for 2D shallow water models on the sphere. However, good test cases are lacking for global 3D models in general, and they are also lacking for nonhydrostatic global models. Test cases for 3D global hydrostatic solvers are under development, and presently center around simulations of unstable baroclinic waves. Outside of using a very-small radius planet, nonhydrostatic tests for the sphere are prohibitively expensive. For this reason we strongly recommend that limit-area versions of all global nonhydrostatic solvers be constructed so that rigorous testing of nonhydrostatic performance can be carried out. The 3D limited area test codes can also be used to perform 2D (x,z) nonhydrostatic testing. For the limited-area configurations, simple boundary conditions are often preferred (periodic, walls) given the difficulties of posing clean (non-arbittrary) open boundary conditions for any given solver and equation set.

With regard to the nonhydrostatic tests, there are still very few tests with simple moist physics. Cloud-model development experience indicates that model formulations are most severely stressed in simulations of moist unstable convection. Standard tests need to the developed in this area.

Tests for the sphere

These tests are used to evaluate large-scale flow solutions. The simplest tests are for the 2D shallow water equations. Much recent work has been done to develop 3D tests, usually centered around unstable jets and baroclinic wave development. One looks for solution accuracy, convergence rates, any pathological behavior, and whether or not one can "see the grid" in the solution. A test suite can be found at http://www-personal.umich.edu/~cjablono/

Test for a nonhydrostatic solver

These tests verify that the solver is indeed solving the nonhydrostatic equation set, and provide some insight into the level of dissipation in the schemes and phase error characteristics. One looks for solution accuracy (where analytic or converged solutions exist) and solver robustness, the latter being especially important in the simulation of moist convection. A test suite has been proposed, see http:// www.mmm.ucar.edu/projects/srnwp_tests/ As noted above, development of a moist unstable convection test cases needs to be done.

8. Discussion

All of the methods discussed in sections 5 and 6 represent some approach to optimization based on the cost and accuracy of solver components and assumptions concerning what is most important to resolve accurately and conserve in a model formulation. In the workshop presentations, as in other workshop presentations and also in much of the literature, very little has been presented concerning the cost of these schemes. Thus, very little can be said about the relative efficiency of one scheme compared to another. One of the problems with providing runtime or CPU time statistics for various tests is that the costs will be affected by the machine it is run on and the level of effort that has gone into optimizing a specific code for the architecture, thus scientists are often reluctant to provide these statistics. Other measures of cost, such as operations per grid volume per time-step, could serve to give at least some indication of expected cost. Everyone agreed that some measures are needed if any meaningful evaluation is to occur.

Even without having any timing statistics or other measures of computational efforts, it is obvious that we have been examining a full range of scheme formulations - from low-order schemes (e.g. OLAM, FIM), to higher-order schemes (e.g. HOMME). Tests results shown by Christiane Jablonowski indicate that the higher-order schemes do give better solutions at a given resolution (see her presentation at http://www.mmm.ucar.edu/projects/global_cores/Jablonowski.pdf), but which schemes are more efficient is not possible to determine. There are two other factors to consider in this regard, the parallel scaling of the schemes, and the cost of scalar transport. In future applications the transport of scalars will dominate the cost of the fluid flow solver. However, various dynamics options will affect the efficiency of scalar transport, especially if consistency with the dynamics is required.

9. Working Group Members and Workshop Participants

Working group members

Bill Skamarock, *Convener* (NCAR)

Amik St-Cyr (NCAR)
Bill Putnam (NASA)
Christiane Jablonowski (Univ. Michigan)
Dave Williamson (NCAR)
George Bryan (NCAR)
Henry Tufo (NCAR, CU)
Jin Lee (NOAA/GSD)
Joe Klemp (NCAR)
Joe Tribbia (NCAR)
Joern Behrens (Alfred Wegener Institute)
John Thuburn (Exeter Univ.)
Mark Taylor (Sandia)
Peter Lauritzen (NCAR)
Phil Colella (LBL)
Piotr Smolarkiewicz (NCAR)
Ram Nair (NCAR)
Robert Walko (Duke Univ.)
Todd Ringler (LANL)
Richard Rood (Mich., NCAR)
John Drake (ORNL)
Dale Durran* (Univ. Washington)
Phil Rasch* (PNNL)
Bill Collins* (LBL)
Jean-Michel Campin* (MIT)

(* were not able to attend the workshop)

Also attending the workshop

Dave Randall (CSU)
Ross Heikes (CSU)
Eigil Kaas
Pat Worley (ORNL)
Matthew Norman (NCAR)