Numerical Methods of the Atmosphere and Ocean: Present Status and Roadmap for Future Research and Development Activities

Jean Côté (University of Quebec at Montreal) and Christiane Jablonowski (University of Michigan)

Atmosphere and ocean models, and in particular their dynamical cores, are at an important crossroads. One motivation for numerical design changes is the emergence of a new generation of computer architectures that promises to be as disruptive as the transition from vector to parallel computing in the 1990s. The new supercomputer architectures with millions of (most often) quite diverse processing elements demand that computations stay local in memory and minimize the number of global communication steps. In addition, the unleashed power of new computer generations enables modelers to run multi-scale simulations at unprecedented, fine grid resolutions that approach the km-scale. This in turn impacts the choice of the underlying equation sets and the numerical discretization methods that allow a more faithful and detailed numerical representation of atmospheric and oceanic phenomena at all scales. The anchors of new model developments are the operational meteorological offices, national climate centers and laboratories as well as the university research community. We group the current and projected future research activities in various categories and center the discussion around the design of atmospheric General Circulation Models (GCMs). Most categories will similarly apply to ocean models.

1. **Equation Sets:**
   Traditionally, most operational weather and climate GCMs were built upon the hydrostatic approximation with spherical coordinates and constant gravity. More recently, the global modeling community has moved towards non-hydrostatic equation sets (either with the shallow-atmosphere approximation or the deep-atmosphere variant) that allow for grid spacing's below 10 km. Some groups consider filtered approximations for the non-hydrostatic equations to eliminate fast vertically propagating sound waves that are meteorologically insignificant, others consider filtering or damping the sound waves at the level of the numerical approximation. Today, active research is underway to replace the traditional spherical-Earth approximation with a non-spherical elliptically shaped Earth. A long term goal would be to move towards a truly 3D representation of the Earth and its gravitational anomalies.
2. **Time discretization:**
   Advection: Time discretization methods are closely linked to the equation sets, which can be written in Eulerian or Lagrangian form. The latter allows for point-wise or flux-form semi-Lagrangian (SL) methods. The mass-conservative flux-form SL discretization has recently regained popularity for the tracer advection component of GCMs due to its projected computational efficiency gains in the presence of many tracers. SL discretizations allow long, efficient time steps with Courant-Friedrichs-Lewy (CFL) numbers greater than one. However, it is likely that future computer architectures will favor Eulerian or SL methods with CFL not much larger than 1. This will keep computations local and reduce the parallel communication costs, thereby off-setting the costs of a shorter time step.
   Waves: Explicit versus (semi-)implicit methods: The minimization of global communication will favor a splitting between slow and fast waves which can be treated differently. A popular split is HEVI that stands for horizontally-explicit and vertically-implicit. Active research is

also underway to explore other IMplicit/EXplicit (IMEX) combinations and non-traditional time integrators such as exponential time integration methods for stiff problems.

Long term research: Parallel-in-time time integrators would exploit untapped parallelism in GCMs, utilize the newest computing architectures more efficiently and thereby reduce the wall-clock execution time of GCMs. New efficient time integration algorithms are needed for success with the GCM equations.

3. **Horizontal discretization:**

   Methods and grids: Traditionally, the global spectral-transform method on latitude-longitude (or Gaussian/reduced-Gaussian) computational grids has been the dominant choice for weather and climate GCMs. Today, there is a general move towards local spatial discretizations to avoid global communication. Examples are finite-difference, finite-element, spectral element and discontinuous Galerkin methods that utilize almost uniform polyhedral grids on the sphere. Examples of these emerging computational meshes are cubed-sphere, hexagonal, triangular or spherical Voronoi grids that avoid the convergence of the meridians near the pole points.

   High-order: The expected computer architecture with fast computations and relatively slow memory movement favors the development of high-order numerical approximations by making extra local computations almost free. One also has to take into account the fact that the amount of memory per processing element will be smaller than today. In a finite element discretization one has the choice of reducing the element size (h-convergence) or increasing the discretization order in each element (p-convergence). The global error can then be minimized with a gain in efficiency but without forgetting that numerical spatial errors have to balance the errors that come from all other sources.

   Variable resolution: Variable resolution for weather forecasting is making a comeback with the development of the cubed-sphere and hexagonal grids that allows a more gradual change in resolution without the extreme aspect ratios of past attempts. The use of variable resolution for climate applications is more heavily debated.

   Dynamic grid adaptivity: Dynamic adaptivity of the mesh (Adaptive Mesh Refinement/AMR) has also been a topic of interest for several years and progress has been made towards the development of global models that can track objects and refine resolution as the objects evolve smaller scales. A potential bottleneck for AMR and variable-resolution in general is the need of scale-aware physical parameterizations.

   Model adaptivity: Model adaptivity, where the model changes locally in a region of the domain, will also play an important role in future computing. It will be most effective when combined with adaptive mesh and algorithm refinement techniques. These models can describe the same physics at different levels of fidelity at the same location or can describe different physics. Model adaptivity can potentially exploit different levels of parallelism, asynchrony, and mixed precision and can minimize communication across layers. Model adaptivity (as well as AMR) could be tied to error control and uncertainty management to apply the finer-grained models only in those regions where the extra expense improves the solution accuracy. There are clearly opportunities in developing scalable adaptive algorithms, but more research is needed.

4. **Vertical discretization:**

   Staggered grids: The choice of the vertical staggering seems to be less important if higher-order numerical discretizations in the vertical are used. However, care must still be taken to control computational modes.

Orography: Orography is most often represented via terrain following vertical coordinates that can lead to large discretization errors in the presence of steep terrain. The trend towards high horizontal resolutions will further steepen the mountain slopes and worsen this problem. More research is needed to assess the pros and cons of alternative representations, such as the step-mountain, cut-cell or embedded-boundaries approach.

5. **Issues that involve all dimensions**

   Stability: Numerical stability is always a requirement and a constraining factor.

   Accuracy: In Numerical Weather Prediction (NWP) it is essential to operate at the correct level of accuracy for each component in the simulation. As soon as the required accuracy is reached for a component there is no need to further improve it because the improvement will be invisible in the result.

   Dispersion relations: It is important to have the proper sign of the group velocity to avoid a conflict between the physical and the numerical direction of propagation of physical properties.

6. **Built-in physical constraints:**

   Mass, total energy and potential vorticity (PV) conservation: Conserving these properties is important for climate models, but less so for forecast models that are integrated for shorter time periods. It is desirable to choose numerical discretizations that guarantee conservation without a-posteriori mass or total energy fixers. Conservation of PV is not standard.

   Mass-conserving and positive-definite tracer advection: Mixing ratios need to stay positive during a simulation to obey their physical principles. Mass-conserving and positive-definite numerical discretizations are paramount.

7. **Physical parameterizations:**

   Column physics: Currently, all subgrid-scale parameterizations only operate in vertical columns. In the future, more horizontal coupling becomes necessary, especially for increasing resolution. More properties, like rainfall, will need to become prognostic instead of diagnostic quantities. It will also be possible that multi-dimensional radiation schemes are needed. Another area that needs attention is the traditional modular implementation of the physics forcings. It artificially separates processes like boundary-layer turbulence and shallow convection despite the fact that such processes are tightly linked. More holistic approaches are needed to model subgrid-scale phenomena simultaneously and thereby avoid double-counting.

   Physics-dynamics coupling: The physics package and dynamical core are most often coupled in a first-order time-split way. Within the physics package a time-split approach is also used which makes the results dependent on the order of the operations. Active research is underway to compute the dynamics and physics forcings on different grids, or to embed a 2D cloud-resolving model in a GCM grid cell. The latter is called super-parametrization. The implications of all these choices need to be better understood.

   Numerical discretizations: Numerical discretizations in physics routines are often of low-order and not consistent with the numerical discretization in the dynamical core. Physical parameterization packages also do not converge at their expected order when time steps are reduced.

7. **Computational aspects:**

   Code development and maintenance: This becomes a problem as we go to extreme parallelism and heterogeneous machines. A solution would be machine-generated code from a high-level language. It is already exploited with the finite-element formalism. The code generated eventually will call highly optimized kernels.

Parallelization strategy and scaling: Parallel computing is essential. Applications are needed that scale well in the strong scaling sense as the problem size is generally fixed. GCMs will need to exploit all forms of parallelism with a hybrid MPI/OpenMP and Graphical Processing Units (GPU) accelerator strategy. Domain decompositions need to minimize the number of ghost cells to minimize communication.

Mixed precision: To minimize the size of data movement and memory footprint it is preferable to store in single precision and compute in double precision only where necessary. Since the GPUs that will be used in these machines have a much lower performance in double precision one would need to compute a first solution in single precision and refine if necessary.

Resilience: For a large problem distributed on a large machine the usual strategy of checkpointing/restarting from a file will have to be revised. Local recovery mechanisms are required that leverage the mathematical properties of the algorithms in the application.

Reproducibility: The nondeterministic nature of failure and recovery in large machines, if occurrences are sufficiently frequent, will lead to nonreproducibility and make code correctness difficult to assess.

Related references and links:

Lauritzen, P. H., C. Jablonowski, M. A. Taylor and R. D. Nair (Eds.) (2011): Numerical Techniques for Global Atmospheric Models, Lecture Notes in Computational Science and Engineering, Springer, Vol. 80, 564 p.
Link:
http://www.springer.com/mathematics/computational+science+%26+engineering/book/978-3-642-11639-1

J. Dongarra, J. Hittinger, J. Bell, L. Chacón, R. Falgout, M. Heroux, P. Hovland, E. Ng, C. Webster, S. Wild, (2014): Applied Mathematics Research for Exascale Computing, U.S. Dept. of Energy (DOE) Advanced Scientific Computing Research (ASCR), Exascale Mathematics Working Group Report, March 2014.
Link: http://science.energy.gov/~/media/ascr/pdf/research/am/docs/EMWGreport.pdf

D. Ham, M. Piggott, T. Ringler, H. Weller and N. Wood: Isaac Newton Institute for Mathematical Sciences “Multiscale Numerics for the Atmosphere and Ocean”, 22 August - 21 December 2012
Link: http://www.newton.ac.uk/programmes/AMM/

P. Lauritzen, D. Williamson, P. Ullrich, J. Behrens, H. Weller, R. Nair, B. Wingate: Partial Differential Equations on the Sphere 2014 Workshop, 7-11 April.
Link: http://www2.cgd.ucar.edu/events/workshops/pdes2014