ECMWF Workshop on Non-hydrostatic Modelling, 8 – 10 November 2010

Contents

Contents

INTRODUCTION

ECMWF is working on a non-hydrostatic option for its Integrated Forecasting System (IFS) because beyond about 2015, when its operational resolution is planned to be upgraded to 10 km, it is essential to have an efficient, accurate and robust non-hydrostatic model available to prepare for future resolution upgrades. Starting point was the non-hydrostatic formulation developed by ALADIN and Météo-France and made available to ECMWF in the global IFS/ARPEGE model. Further developments are required, in particular to improve efficiency, before this non-hydrostatic model can be used for the Centre's operational applications. It was therefore timely to hold a workshop on non-hydrostatic modelling to bring together experts in this area to discuss and get guidance on the following topics, among others:

- The viability and competitiveness of some of the key elements of the IFS at very high resolution: the spectral transform method, the semi-implicit method and the semi-Lagrangian advection.
- The strengths and weaknesses of alternative approaches taken in the development of non-hydrostatic dynamical cores.
- Alternatives to the fully-compressible Euler equations and their suitability for global NWP in the non-hydrostatic regime.
- Conservation requirements.
- Scaling properties of the various numerical solvers on massively parallel platforms of the future.
- Experience with the unification of general circulation and cloud-resolving models.
- Consistent physics-dynamics coupling for the fully-compressible model.
- Issues arising in the 'grey zone' with partly resolved physical processes.
- Test cases and strategies to assess model performance in the non-hydrostatic regime.

The workshop followed the usual format of one and a half days of invited lectures followed by discussions in three working groups. The working groups made valuable recommendations and suggested potential avenues for ECMWF to explore. These will guide research and development at the Centre in the area of non-hydrostatic modeling in the coming years.

These proceedings contain the reports summarising the discussions and recommendations from the three working groups and the papers presented by the invited lecturers.

The workshop organisers thank all participants for their contribution to a successful and enjoyable event.

Working Group 1: Dynamical solvers (Suitable discretizations and alternatives to global spectral models)

Chaired by Bill Skamarock; Secretary Martin Miller

Others: Terry Davies, Jean Coté, Günther Zängl, Mariano Hortal,
Almut Gamann, Joanna Szmelter

1 Overview

Need for nonhydrostatic global atmospheric solvers for next generation climate and weather models.

Applications focus for group: uniform resolution global nonhydrostatic models - NWP, perhaps up to decadal time scales.

Note: we acknowledge that many applications need/desire local (and perhaps time adaptive) refinement capabilities (i.e unstructured meshes), but this is not our focus here.

Need for conservation (mass, scalar mass, energy, etc) or other solver attributes?

Question - what is sufficient accuracy, i.e. is small-scale phenomenological accuracy at high (NH) resolutions sufficient/needed for global NWP applications?

Weak(?)/Strong(?) scalable codes.

2 Model development goal

Maximize efficiency for applications (flexibility, other needs/constraints) $\longrightarrow accuracy/cost$

2.1 Critical understanding

1. Groups have different applications, accuracy measures, and different tolerance for accuracy versus costs. These differences in large part account for out different model development histories and current efforts.
2. Computer architectures are evolving towards larger number of processors -O(10^4) $-$ (10^6).

This is driving many development efforts in new directions. The computer architecture evolution is creating or exacerbating model problems such as

- memory-cache-processor bandwidth
- inter-node communications
- I/O bottlenecks
- fault tolerance, bit reproducibility, issues on future MPP machines

How do techniques map to these issues: lots of documents/reports available, eg.

References

[1] http://www.mmm.ucar.edu/projects/global_cores/core_assessment.pdf

[2] D. Williamson (2007). The evolution of dynamical cores for global atmospheric models. *J. Meteor. Soc. Japan 85B*, 241-269.

[3] Paul Selwood Nigel Wood et al (2009). Strategies for improving the scalability of the UM in response to changing computer architectures. http://research.metoffice.gov.uk/research/nwp/publications/mosac/doc-2009-10.pdf.

3 Solution techniques

3.1 Horizontal meshes

- Lat-Long - known issues with poles
- Gaussian
- Yin-Yang
- Cubed-sphere - (any SL formulations being developed)?
- Hex - (any SL formulations being developed)?
- Triangles - (any SL formulations being developed)?
- etc...

  Generally speaking, we do not want to see significant ”grid imprinting”

3.2 Vert Coordinates

- f(pressure) - historical in hydrostatic models, easy to include hydrostatic ”switch” in NH model. Additional eqns to solve.
- f(height) - issues with rigid lid upper b.c.’s? Fixed geometry eases formulation of implicit solvers.
- Lagrangian remap - little experience in NWP models - some (and soon-to-be-more) experience in climate models.
- Discrete Charney-Phillips / Lorenz / finite element

3.3 Terrain

- Terrain following z, p
- z- coordinate: Some research efforts, issues with BL resolution, physics...

3.4 Time integration techniques

- Explicit
- Implicit
- Spectral Element /Discontinuous Galerkin methods
- Dimensional splitting
- Time splitting
- Physics/dynamics coupling - do nonhydrostatic dynamics bring with them new constraints/issues?

4 Main questions

Will implicit techniques scale on future architectures such that they will be efficient? O(10^4) - O(10^4) processors, O(1) km horizontal grid spacing?

Iterative (e.g. CGR) or direct (spectral) implicit solvers - given the need for global sums or transposes, etc... These techniques may have significant problems scaling to new architectures.

Will large timesteps in SLSI schemes (relative to Eulerian schemes) be feasible on O(km) meshes? e.g. how far can timesteps be pushed in SL schemes when deep convection is explicit?

Appropriate test cases and test strategy? Small planet, aquaplanet, other dry/moist tests, etc...

5 Recommendations

1. Level of uncertainty in architectures is high, and continued exploration by the Centre of increased scalability and efficiency of transforms, and of the stability and efficiency of the implicit algorithm, is warranted.
2. Maintain resources to continue investigations of NH IFS dynamics w.r.t. forecast skill.
3. Maintain links and strengthen existing, or establish new, collaborations with NH IFS users/developers outside ECMWF.
4. Keep abreast of outside development efforts using different approaches, and if circumstances dictate, move to develop/adopt a new core, understanding that development of different dynamical cores will require significant new resources.

Working Group 2: Relevant experience with, and issues arising from, the unification of GCMs and CRMs

1 Introduction

A universal "supermodel" which automatically adapts to each resolution is the Holy Grail of Numerical Weather Prediction. Such a model should be able to assure the large scale balances when used with large resolutions but release coherently these balances in order to resolve mesoscale processes at higher resolution while preserving the balances for the averaged states. It should also converge toward a CRM with all the 3D features which may be necessary when the complexities at the scale of convective clouds need to be explicitly computed, with the caveat that it will surely not be with an isotropic mesh.

And of course, it should still "scale" perfectly on the computers of the future and be easy to connect with a data assimilation system.

However, the unification of GCMs and CRMs, if achievable, may take various paths. The replacement of parametrization by "super-parametrizations" can be seen as a synergetic juxtaposition of GCM and CRM. Such a solution needs the co-development of both GCM and CRM and a clever way of interfacing them.

A hierarchy of models especially designed each for a given range of scales but within a common software is an alternative to the Holy Grail and may also be considered at least as a transitory solution.

In any case, the problem of the "grey zones" (convection, turbulence) has to be considered carefully. At the scale of the grey zones, the connection between the dynamical core and the current parametrizations is at the heart of the problem. The problem of the unification of GCM and CRM is then not independent of the questions investigated in the other groups, in particular of the debate about the choice of the adiabatic system of equations for the dynamical core. But in this working group, the discussion was more focused on questions involving physics and its coupling with dynamics.

2 Unification of model physics

The unification of model physics is the main challenge in the development of a "supermodel". The WG agreed that there is still quite a way to go before success.

Complementary approaches are anyway beneficial (e.g. super-parametrization is an integrated auto-evaluation tool for the parametrization path), none of the possible strategies being perfect.

Therefore, ECMWF might seriously consider the trilemma, once having reached a resolution of about 10 km:

- Further increase in resolution if the main goal is a better wind forecast and an improvement in the representation of stable boundary layer processes (but with a tough challenge in order not to deteriorate the previous skills for cloud and precipitation forecasts, while increasing the spatial resolution of scales associated to feed-back loops within the model; see below for the problematic of the "grey-zones")

- Pause around a resolution of 10 km in parallel with R&D on a super-parametrization strategy if the main goal is to improve the hydrological cycle forecast at all resolutions
- Pause around a resolution of 10 km and rely more and more on an EPS system made suitable for this resolution (TL/AD with more physics, stochastic physics…).

Each progress in resolution has its own challenges. For instance the type of balances necessary to be correctly simulated in order to reach a new forecasting target changes and the verification of progress in this new direction requires careful thinking as progress towards new scales should not destroy previously acquired benefits.

An example of this situation is the maintenance of the deep convective Quasi-Equilibrium (Arakawa and Shubert, 1974) rule to a sufficient degree of accuracy as resolution is increased and cloud microphysics is more accurately represented and controls some larger scale aspects. It was recognised that all this is mainly a consequence of the fact that the (time and space) scales between "real" updrafts and compensating circulations are very different, and that parametrization cannot thus handle these both fully consistently and within the same time step.

In fact this particular "grey-zone" challenge is so dimensioning that much of the current section is devoted to "unification with respect to deep convection". However opinions about how to best address this particular challenge (when one cannot neglect any more the cloud area fraction) were not consensual within the WG. The following possibilities were discussed:

- "Switching off" the parameterised part when the resolved part is correctly taking over (but how to do it cleanly?)
- Shifting the handling of model subsidence to the dynamics (this requires harmonisation of discretizations, something that might be considered if some rewriting in the dynamics takes place)
- Replacing a diagnostic closure by a prognostic closure in the convection scheme
- Considering "plumes" only as perturbations, on top of the part from the convective drafts which the model already handles within its "resolved" equations

Various types of convection (tropical, post-frontal, breeze-type, etc.) may favour different solutions, since the large scale balance (QE) does not say anything about single occurrences. Hence non-exclusive approaches are recommended.

There was however consensus that a unified modelling of cloud physics (between its various representations) was, even if not a sufficient condition of success, surely a necessary one. This may however raise even more issues of interfacing between numerics and physics. The time-step dependency and grid-point-storm syndromes are for example common problems which are encountered by most of the NWP models when the resolution increases. A good testing strategy is clearly necessary for a better understanding of these questions.

It is recommended that ECMWF proposes organising rather soon a workshop specifically aimed at "grey zone" problems, along the lines suggested in the above summary. This workshop might be co-arranged with WGNE.

3 Representation of observed balances

With the increase in resolution, more processes are supposed to be “resolved”, therefore “built-in” large scale balance or quasi-equilibrium are not necessarily assured anymore in the model. It is then necessary to decide what a model has to be good at (physical thinking) and, by opposition, what still has to be considered as unpredictable adjustment-type degrees of freedom.

The quasi-geostrophic adjustment problem is a simple prototype of this problem. However, in the case of higher resolutions, the problem is more linked to parameterization and interaction between physics and dynamics.

The WG considers that a good solution to handle this problem is to run test-cases, sufficiently simple to help isolating the balance and adjustment-type one wishes to study and realistic enough to stress the parametrization choices in their potential weaknesses (wrong time scales, shift of the stationary solution, too oscillating behaviour, not to speak about the in principle already handled full divergence syndromes).

4 Evolution of phys/dyn coupling at high resolution

4.1 Interplays

As one goes to higher resolution, the interplay between high resolution processes (cloud/precipitation microphysics and unorganised turbulence) and either the dynamics or the macro-physical parametrizations (convection, GWD) is the key hurdle. The interplay in the case the simulation of the atmospheric convection in the “grey zone” range of resolution and for higher resolutions was already discussed (section 2).

But (at least) two other points of view were expressed during the discussion:

- The difficulty in the inherent averaging of highly non-linear processes (Pincus and Klein, 2000). For example, parametrizing mixed phase clouds and, alternatively, preventing the spurious existence in the model of water supersaturation are two basic but challenging tasks, with both numerical and physical aspects.
- The crucial influence of lateral exchanges for grid-boxes which will never be isotropic (Piotrowski et al., 2009) shows at least that the issue of vertical resolution cannot be forgotten, even if it is surely not the only thing to do in order to avoid serious problems.

The WG mentioned that the respect of observed balances (section 3) and consistency of the formulations (section 4.2) are in general separate issues.

4.2 Consistency between dynamical cores and physical packages

Current physical parametrization are usually designed as constant pressure processes. In an hydrostatic model or in an anelastic model, the computation of the pressure is done diagnostically in the dynamics which therefore assure the balance. In a fully compressible model, the pressure is a supplementary degree of freedom. The question of the coupling between the physics and the evolution of the pressure has to be investigated.

Consistency in a fully compressible model is then difficult because physics routines are usually designed only to update the temperature or potential temperature, but either the pressure or the density has to be changed to maintain the equation of state. The Met Office model does not enforce the equation of state for individual processes, but enforces it at the end of the timestep. If this method is not used there is a potential problem with any method for the compressible equations.

A minimum consensus was reached under the following lines:

- Solutions with either hydrostatic primitive equations in mass-type coordinates or Eulerian equations in height based coordinates are easier to make consistent, from first principles.
- The difficulty for the current NH-IFS set-up comes from its mixed character (mass based coordinate but elastic adjustments in the dynamics). This is a possible disadvantage of the said solution, when one wants to project diabatism in a more sophisticated way than the "quasi-anelastic" coupling advocated by Thurre and Laprise (1992). The latter solution seems however to deliver very correct total impacts in all studied case up to now, to the price of course of a filtering hypothesis in the translation of the full equations.

If not using this solution, a reviewing of the way to design the parametrization schemes may be needed, without guarantee of success in this quite special case.

This problem is connected with the issues concerning the organisation of the time-step. In the current IFS organisation (physics after the explicit dynamics but before the SI correction in spectral space), adjustment-type parametrizations may suffer from SI important corrections. It may then be better handled with a physics called before the dynamics (ARPEGE-ALARO-AROME solution) in which case the "true" end of the time step is equivalent to the very beginning of the next one. But there may be other drawbacks of this choice.

Predictor-corrector methods permit much better physics-dynamics coupling by allowing the physics to interface with correctly time-averaged dynamics. Though more expensive, such methods may well improve the accuracy enough to make them cost-effective.

The working group agreed that more research is needed to better document this problem.

4.3 Conservations

The discussion was even less consensual than on the previous aspects, but something nevertheless came out of it.

One may distinguish three successive levels of consistency if one wishes good conservation properties in the physics:

- In the physically correct computation of thermodynamic parameters when the model (dynamics included) needs to resolve the multiphasic composition of the air parcels (c_p, c_v, R, L_v/L_s).
- In the computation of the physical tendencies of the model prognostic variables from the output of the parametrizations (but with which budget-type conservation law?).
- In ensuing consistency with respect to the previous constraints in each parametrization algorithm.

Given the associated workload, it is recommended to first assess the magnitude of the impacts of inconsistencies with sensitivity studies (e.g. by intentional degradation of available modelling solutions having reached a higher degree of "conservation" than IFS), for each step subsequently.

5 Evolution of physics at high resolution

5.1 Evolution toward a 3D physics for high resolution

An evolution towards a 3D physics in the IFS is not only a scientific challenge but also a technical one. In the current code, the physic computations are done along independent columns. The parallelization of the physics is based on this simplification. In the present IFS, 3D computations are possible only in the s.-Lag. stencil space, i.e. at the origin points of the trajectories.

5.1.1 3D Radiation

The working group came quickly to a consensus here, with direct recommendations: ECMWF should only consider the impact of slopes and shadows on the surface radiative fluxes and leave the "cloud to cloud" problem, relative to vertical flux divergences, for a later perspective, if ever. The impact is likely to appear around 5km meshes (or a bit earlier) but the implications are big enough for ECMWF to start thinking about the topic soon. One prerequisite step would be to complexify the partial radiative computations which are done at each time step and each grid-point which are currently too crude with respect to the complexity of the full radiative computations done with a lower time and space resolution.

5.1.2 3D Turbulence

The WG was less unanimous on this issue, but a majority consensus seemed to emerge for recognising that a full 3D treatment would likely be too sophisticated for what the anisotropy of the meshes and the horizontal resolutions would allow to be meaningful. The use of a predictor-corrector scheme means that the physics has a 3d element which may well be sufficient at resolutions lower than a kilometer.

A reasonable compromise is currently explored by the ALADIN community. In this approach, the vertical diffusion is complemented by a consistent "3D-minus-1D" scheme built to simulate a true lateral diffusion but with some 3D flavour. The additional part, completed for efficiency reasons in the 3D space of the SL stencil, is mostly reduced to its isotropic terms but with time- and space-varying coefficients. This necessary simplification is dictated by the highly non-isotropic computational grid used in NWP models. One must be careful that there is no gap left by the parametrizations through which the model dynamics would develop its own way to mitigate too strong gradients. In principle, in the above-mentioned framework, the fact that the diffusion physical operator is ultimately combined with its numerical counterpart (for avoiding energy accumulation at the end of the spectrum) should prevent this risk.

More exploratory work is needed but the LAM applications in the IFS galaxy can be seen as useful precursors towards a better treatments of the horizontal subgrid mixing and energy cascade at the end of the spectrum.

5.2 Slow versus fast processes

The WG considered this to be a minor issue, except perhaps for 1D vertical turbulent fluxes that may have to be treated as a slow process as the resolution increases and the time step decreases. However, the parametrization of deep convection which mixes processes with differing time scales may show some surprises when one goes away from quasi-equilibrium. ECMWF should be ready to monitor the situation if necessary, this closing the loop of this WG discussions by a return to the multiscale aspects of the representation of convection (section 2).

6 Conclusion

The discussions of the WG2 showed that the path to be chosen for the transition from GCM toward GCRM is not clearly marked. Most of the problems are complex and need more research and sensitivity studies.

But even if the definition of a “universal” model is subject to various interpretation and may be inaccessible at the end, the idea of a unified multiscale system has to be behind each new development. This is a condition of progress for all possible options. All the recommendations which have been made by the WG2 are following this line.

Working Group 3: Numerical solution procedures

Chaired by Dale Durran; Secretary Nils Wedi

Others: Piotr Smolarkiewicz, Zavisa Janjic, Rupert Klein, Celal Konor, Jin Lee, Michael Baldauf, Hirofumi Tomita, Pierre Bénard, Jozef Vivoda

The working group was to consider numerical solution procedures for the fully-compressible equations and the suitability and potential advantages/disadvantages of alternative model equations for global NWP in the non-hydrostatic regime. The group started with a discussion of the viability of some of the key elements of the existing IFS model, namely the spectral transform method, the semi-implicit solution procedure, and the semi-Lagrangian advection scheme.

The key question to address for ECMWF is whether the **spectral representation** is still affordable at ultra-high resolution. If ECMWF is to continue with the spectral transform method, disadvantages could be the parallel efficiency of the Legendre transforms and the efficiency of the transpositions (i.e. the necessary communications involved in a massively parallel processing (MPP) environment when changing from spectral to gridpoint representations and vice-versa). A further disadvantage could be the problem of spectral ringing when representing "less smooth" fields. Advantages are the possibility of a trivial (direct) 2D elliptic solver and spectral accuracy in particular of the associated derivatives. Given the success, progress and investment made by ECMWF in spectral methods, the group did not suggest abandoning the spectral approach without careful consideration and comparison tests.

The second question addressed was if it is okay to slow acoustic waves in the context of a **semi-implicit solution procedure of the compressible Euler equations** or if it was better to apply alternative algorithms that either filter acoustic waves a-priori or handle acoustic waves in a physically more appropriate way. The background to this question arises from the fact that semi-implicit numerics with relatively large time-steps (chosen to satisfy the advective CFL) tend to slow the phase and group velocity of acoustic waves to zero and thus artificially modify adjustment processes (i.e. physical parametrizations) driven by these waves. Similar arguments, however, apply to the adjustment processes in semi-implicit hydrostatic models (Janjic and Wiin-Nielsen, 1977). In contrast, a-priori filtering of acoustic waves enforces infinitely fast adjustment as phase and group velocities of sound waves become infinite. Notably, while hybrid models such as the Arakawa-Konor system appear to preserve the Lamb mode, this is not the case in some classical anelastic systems and there was no consensus regarding the influence of infinitely fast adjustment processes on the overall solution in the context of NWP.

The third question addressed the relative efficacy of the **semi-Lagrangian advection** at high advective CFL numbers and if advantage was preserved in the limit of ultra-high resolution. Generally, the group's consensus was that in the regime of resolved/permitted convection a local CFL of less than one is necessary for acceptable accuracy. We note that current global semi-implicit semi-Lagrangian models typically run with local CFL numbers of 3-5 near strong jet stream cores, but local CFL numbers smaller that unity throughout the vast majority of the domain. Let $u_{\max}$ be a characteristic maximum jet-stream wind and $w_{\max}$ a characteristic maximum convective vertical velocity. The utility of semi-Lagrangian approximations to the advective operators in a global convection-permitting model will then hinge on the relative magnitudes of $u_{\max}\Delta t/\Delta x$ and $w_{\max}\Delta t/\Delta z$, where Δx and Δz are the local horizontal and vertical grid spacings. Note that if $w_{\max} = 5\ m/s$ and $\Delta z = 200$ m, the time step required to keep the local CFL number less than unity in a region of resolved convection would be sufficient to keep the local CFL less than unity at a spot where a 150 m/s jet was represented with horizontal grid spacing of 6 km. In such a case, if the local CFL for the vertical convection was set less than one there would appear to be no need for a semi-Lagrangian treatment of the horizontal advection. Furthermore, concerns were expressed regarding the efficiency of the semi-Lagrangian halo in a MPP environment, the issue of ensuring conservation (of mass) and positivity of moisture at minimal cost, and the accurate treatment of vertical boundary conditions in the presence of strong vertical motions.

The diagram in figure 1 suggests **potential avenues for ECMWF to explore**: The group discussed the advantages and disadvantages of explicit versus semi-implicit methods for the compressible Euler equations for the particular application of numerical weather prediction. **Explicit schemes** are naturally implemented in a MPP environment and they avoid a three-dimensional elliptic solve. However, disadvantages are that time-split procedures are linearly unstable without a filter and unsplit methods require a relatively small time-step, which may not be a problem if non-iterative time-differencing is used for horizontal advection where appropriate. There was no consensus in the group if either semi-implicit or explicit methods are competitively efficient. In contrast, **semi-implicit methods** do require a 3D elliptic solver. Issues were suspected with scalability (competitive elliptic solvers should scale linearly with the number of gridpoints), potentially a lack of efficient pre-conditioners and some problems were mentioned with the handling of orography in the context of a multigrid solver. However, given an efficient 3D elliptic solver, some members in the group strongly advocated this option, as the potential for exploring alternative equations that require the solution of a 3D elliptic problem is greater and it provides an easy switch between sound-proof and semi-implicit compressible solvers.

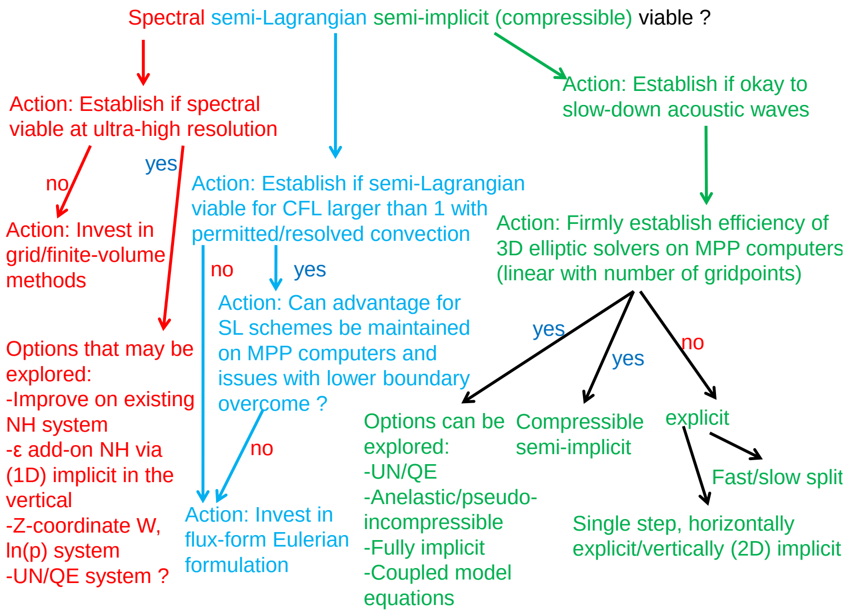

Figure 1: Sketch of proposed options for ECMWF to explore. Arrows indicate potential directions pending the outcome of research into the previous level.

The working group then discussed aspects of **conservation** with a general recommendation for ECMWF to explore ways to upgrade a-posteriori forced global (domain-averaged) conservation to conservation due to local flux cancellation. While there was no consensus regarding all the quantities that should be conserved, it was suggested to ensure mass conservation, and examine the feasibility of ensuring conservation of enstrophy, total energy and angular momentum. The choice of the quantities to conserve may also depend on whether the finite-difference/finite-volume or the semi-Lagrangian approach is taken.

Finally, given the research directions presented in some of the lectures towards alternative grid structures such as hexagonal C-grids, the chairman reminded the group of some recent work in Reinecke and Durran (2009). In this work mountain waves excited by an $8\Delta x$-wide isolated ridge appear qualitatively correct, but are in serious error due to differences in numerical phase speeds arising from the different centered, second-order spatial derivative stencils used on the staggered mesh to evaluate advection, and to evaluate pressure gradients and divergence. Switching to fourth-order differencing of the advection terms provided one way to eliminate the problem.

References

Janjic, Z., and A. Wiin-Nielsen (1977). On geostrophic adjustment and numerical procedures in a rotating fluid. *J. Atmos. Sci. 34*, 297–310.

Reinecke, P.A., and D. R. Durran (2009). The over-amplification of gravity waves in numerical solutions to flow over topography. *Mon. Wea. Rev. 137*, 1533–1549.