Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2019-102-AC1, 2019

Geoscientific Model Development Discussions

Interactive comment on "WAVETRISK-1.0: an adaptive wavelet hydrostatic dynamical core" *by* Nicholas K.-R. Kevlahan and Thomas Dubos

Nicholas K.-R. Kevlahan and Thomas Dubos

kevlahan@mcmaster.ca

Received and published: 12 June 2019

1. (i) The 3D code has exactly the same structure for horizontal adaptivity as the 2D code (there is no vertical adaptivity). It has the same parallelization and same data structure. It therefore inherits the scaling inherits the numerical properties of the 2D code, including the fact that cpu time per grid point is independent of the compression ratio for given coarsest and finest grids. The relevant graph is shown in figure 7(a) of Aechtner et al (2015) cited in the manuscript. Note that this property does not depend on the particular flow considered, or the number of vertical levels (although the actual compression does). For a given compression ratio the cpu time increases proportionally to the number of vertical levels.

(ii) Figure 4 illustrates the parallel scaling, not absolute performance, and so 'speed up' is the right measure (especially since the absolute speed is highly machine dependent). We address the absolute speed of the code compared with the non-adaptive version of the code, DYNAMICO, on pp 27 and 28 where we point out that the wavetrisk is about 3-4 times slower per adaptive grid point than DYNAMICO on the same machine with the same number of cores.

(iii) WAVETRISK-1.0 uses the same hybrid patch-quad tree adaptive memory management structure as the 2D shallow water code described in Aechtner et al (2015), as described on pp 7-8. We have checked that for large numbers of cores the memory used per core is proportional to the number of cores. Memory is therefore not a limiting factor since large problem are run on larger numbers of cores.

As an indication, the total memory used per active grid point is about (30*NZ+2)*8 bytes/1e3 (bytes/kb) = 7.2 kb per grid point for 30 fields on NZ=30 vertical levels. Note that because memory is not a limiting factor, we have not made an effort to minimize memory use per active grid point.

2. (i) Thank you for the reference, we will discuss it in the paper. We think the claims are not contradictory: filtering the wavelet coefficients does provide an objective criterion for grid refinement based on the local polynomial interpolation error of variables and/or trend. (Note that the non-dimensional tolerance is scaled by the relevant norms for each variable.) Apart from refining on the wavelets of the variables or trend, our comments on p 29 refer to other factors determining the accuracy and stability of the method, such as the vertical re-gridding algorithm and use of hyperdiffusion.

(ii) Thank you for letting us know about the unpublished 2012 talk by Popinet et al showing some preliminary results for Held and Suarez, that we were previously unaware of. We will note that his group, as well as the CHOMBO group, have made some (unpublished) significant steps towards developing and evaluating AMR for complex 3D atmospheric flows.

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2019-102, 2019.