



Interactive
Comment

Interactive comment on “Towards convection-resolving, global atmospheric simulations with the Model for Prediction Across Scales (MPAS): an extreme scaling experiment” by D. Heinzeller et al.

Anonymous Referee #1

Received and published: 29 September 2015

1 General Comments

The manuscript addresses a relevant issue in climate and weather forecast modelling world: Will the next generation models, build on advanced numerical methods for unstructured global grids, be able to deliver adequate computational performance on current and future HPC systems? The authors set out to answer this question for the MPAS model (atmosphere only) by testing four different model configurations: One standard resolution case with uniform cell size, two (by today’s standards, at least for

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper



climate models) high-resolution cases with grid refining, and one very high-resolution case, which is clearly beyond the current state-of-the-art in climate research. It is worth mentioning that even the two high-resolution cases represent advanced modelling due to the grid refinement used.

Besides investigating computational performance, the authors try to widen the context by also evaluating the accuracy of the model. This is done by looking at a particular phenomenon, the West African Monsoon. Since I am not at all an expert for Earth system processes, I can not include the corresponding section 3 of the manuscript in my review.

The authors come to the main conclusions, that

- MPAS-A experience "good scaling" across all platforms and all of the standard and high-resolution test cases
- the "number of cells owned by each task" is a "robust" indicator for the breakdown of parallel efficiency
- the computational cost associated with setup and I/O has a strong impact on parallel performance
- the model "scales very well" even for the very high resolution test case
- scaling at very high resolution is governed by the same rules as the high-resolution configurations

Summarising, the authors conclude that "it is possible to conduct global, convection-resolving atmospheric simulations with MPAS on current and future massively parallel systems".

[Full Screen / Esc](#)

[Printer-friendly Version](#)

[Interactive Discussion](#)

[Discussion Paper](#)



2 Specific Comments

The methodology of the study seems well-founded and sufficiently detailed to include most aspects relevant to computational performance. Notably the consideration of I/O (even at very high resolution) is an important aspect that is often overlooked in the analysis of computational performance for weather or climate models. To certify MPAS-A "very good scaling" seems a reasonable conclusion, given the presented results.

The "number of cells owned by each task" as a robust metric for the scalability limit does not appear similarly convincing to me. This could be due to the fact that the "transition zone", which is frequently used in the text, is not clearly defined. It did not become clear to me whether the range of 600–150 cells is, in itself, the definition of the transition zone, or comes as a result of measurements. Of course, it is understood that the parallel efficiency will break down below a certain limit of cells per task, but this is just saying that tasks and cells per task are reciprocal.

The transition zone (or critical cells-per-task metric) is often associated with a limit of 70% parallel efficiency in the text, but that there is a universal (across grids and platforms) limit for that metric remains a bit unclear. I do recognise that the dependency on the platform (e.g. interconnection speed) is mentioned in the text. However, it seems from the results in Figures 4, 6, and 11 that the point where the parallel efficiency drops below 70% is, in fact, both platform and grid dependent. Even Figure 13 does not provide a clear explanation for the transition zone being 600–150 cells.

The level of detail for the performance analysis of MPAS-A is much appreciated. It provides insight to performance relevant characteristics of MPAS-A, which can probably be generalised to other models.

More specific comments:

- pg. 7000, lines 7–9: The text refers to Table D1 but some cases are not present

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper



- in that table (9 nodes ForHLR; 190 grid cells per task). This is a bit confusing.
- pg. 7000, lines 19–21: The statement is not true for 420 tasks
 - pg. 7001, line 1: It is not evident from Figure 8 that the graph becomes non-linear for large numbers of tasks, at least not for platforms other than Juqueen and particularly not for the finest grid.
 - pg. 7002, line 3: The case 35 nodes on Jtest-half is not present in Table E1.
 - pg. 7002, lines 7–12: It can be argued that the case of 420 tasks for the 120km grid is already indication this fact.
 - pg. 7005, line 23: It is concluded (pg. 7004, lines 11–12) that the impact of non-contiguous partitions is negligible. Still it is listed here as a reason for irregular scaling.
 - pg. 7007, lines 18–22: The two test cases to study the break down of parallel efficiency show quite good scaling (about 82% for the first test and still above 70% for the second). The number of cells per tasks (160 and 261) are within the transition zone for both cases. Why where they chosen? Wouldn't one case well above and one case well below the efficiency limit be more illustrative?

3 Technical Corrections

General remark: The use of performance metrics seems to be a bit inconsistent at some places. For example, the speed of the model is given as "Integration in 24h walltime" in Table 1, as "Realtime [s] per 24h" in Table D1, and "Speedup" in Table G2. One suggestion is to use "Simulated years per day (SYPD)" as this seems to become an established metric for climate models.

- pg. 7003, line 3: Not sure, but should it read "... for *less* than 40 owned cells per task, ..."?
- pg. 7005, line 13: "Test runs ..." (not Tests in plural)
- pg. 7018, line 11: The unit of "10Gb" seems to be wrong for transfer speed.
- pg. 7019, line 1: Not sure if "nodes" should actually be "cores".
- pg. 7019, line 6: The formulation "... is to some degree independent..." could be improved.
- pg. 7019, line 7–8: The unit "racks" is not mentioned in the corresponding Table 1.

Interactive comment on Geosci. Model Dev. Discuss., 8, 6987, 2015.

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper