

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.

D. Heinzeller et al.

heinzeller@kit.edu

Received and published: 13 October 2015

First we would like to thank Anonymous Referee #2 for the evaluation of our manuscript and for the detailed comments, suggestions and corrections. While we will be waiting for the comments of referee #3 before providing an updated version of our manuscript, we would like to answer to the questions and concerns raised by referee #2:

(1) General Comments

We appreciate the concerns raised by the referee regarding the significance of the

C2473

conclusion made, based on a single 11-month comparison of MPAS with and without grid refinement. We suggest to put this section in the context of an exemplary study, to rephrase the conclusions drawn from it in the abstract and in the main text accordingly and to highlight the limitations of our comparison. We will also look into the possibility of shortening Section 3 rather than taking it out of the manuscript, since this is the first application of MPAS that focuses on the (West) African Monsoon. We would also like to bring to the referee's attention that we have already applied for computational resources to conduct an ensemble modelling experiment with various MPAS meshes for West Africa with the focus on seasonal predictions.

With regards to the spin up time for soil, we have conducted several experiments with WRF and with WRF-Hydro (coupled atmospheric-hydrological model) to investigate the required spinup times with focus on West Africa. For example, experiments were conducted in which the model was initialised in January 1 of five subsequent years (e.g., 2003, 2004, 2005, 2006, 2007) and for which the soil moisture/temperature was compared in 2007, using the 2003 run as reference. While the error was large for the run initialised in 2007 (i.e., basically comparing the forcing data after running it through WPS to the 4-year spinup run), it was reduced to only slightly above the interannual variability for the 2006-run and of the same order as the interannual variability for the 2005- and 2004-run. These and our other tests suggest that two years of spinup time are sufficient in this region and that even one year might be acceptable. We believe that these comparably short spinup times are due to the extreme conditions in the area: With an intense rainy season and a prolonged dry season, the soil conditions are basically reset after completing one cycle. Attached to this response is a small document with plots from the WRF 12km run (used as reference run in the manuscript), which supports our findings regarding the spinup times. We will extend the corresponding discussion in the manuscript to reflect your concerns and our results regarding the spinup time.

(2) Specific Comments

C2474

Section 2.1: We agree with the referee's suggestion to condense the important information into one table and to provide additional information in the appendix, if necessary.

Section 2.2: A model top of 30km with 41 vertical levels on a hybrid height-based terrain-following coordinate is used in this study. Being developed out of a regional climate modelling context, the version of MPAS-A used here has not been adapted to extend higher into the atmosphere. In more recent versions, the model top has been extended to 1mb (around 42km). We agree that a higher model top (and the implementation of a coupler between the three MPAS components) would be required if MPAS was to be used as a full earth system model in the future. It is probably interesting to note that regional climate models follow very different approaches regarding the model top, even though they are all used for long-term climate simulations: While WRF works with a static and simple model top at a specified pressure (usually 10-50mb) and does not incorporate any information from the forcing data set, the Cosmo-CLM model uses a model top derived from forcing GCM data at a specified level. We will add the missing information on the vertical levels and a short note on its limitations to this section. We will also add information on the number of 2D and 3D variables written to disk to Table G1, as suggested by the referee. We will also note in table G1 that the sets of fields written to output files are completely configurable at run-time.

Section 2.3: Following the referee's suggestion, we fitted a modified Amdahl law to the three test cases and separately for the Intel architecture (Jtest-full/half, ForHLR1, Curie) and the Bluegene architecture (Juqueen). We will add this information to the manuscript in form of a table, and we will also plot the so-obtained fitting curves in Figure 12 and the new equivalent plots for the 120km and the 100-25km test cases.

The performance data was derived from two to three test runs, depending on how close the first two measurements were: For sufficiently close results from the first two runs, the average was taken from these runs. Otherwise, a third run was conducted and the average was taken from all three runs. We will add this information to the manuscript.

C2475

Page 7000, line 27: Indeed we mistakenly stated that the relationship is linear. We fitted a power law $\text{commvol} = A * \text{tasks}^B$ to the data for each test case and obtained an exponent $B=0.52$, independent of the grid. This exponent of 0.52 agrees with what we calculated from the uniform mesh plot. We will add these fits to the three panels in Fig. 8 and correct the text accordingly.

Use of figures: We will follow the referee's suggestion and merge figures 1-3. Figure 4, 6 and 11 will be combined with Figure 12 and equivalent figures for the other two test cases for an easier comparison. Figure 5 and 10 will be merged as suggested.

General comment on colour maps and line plots: The first author is particularly grateful for this hint, since he was not aware of these limitations. We will update the contour plots with appropriate colour maps. Regarding the line plots in our paper, we assume that this concerns mainly Figs. 16 and 17, which we will improve accordingly.

Section 2.4: We will shorten the discussion and drop Fig. 9, but we would like to keep Fig. 7. We feel that the aspect of contiguous and non-contiguous partitions and the variability in the communication properties for individual patches (not existent for regular grids) are important to mention, even though they have no significant influence in this particular case.

Section 2.6: We will drop panels (d)-(f) of Figure 13, as suggested, and modify the text accordingly. Regarding Fig. 14, we agree that a table containing absolute timings and percentages can provide the same information in a clearer way. As suggested by Anonymous Referee #1, rather than presenting the details for two model runs within the transition zone, we will keep the 2024-task run for the 60-12km mesh and add a second run on the same mesh with a larger number of tasks to highlight which parts of the dynamical solver are responsible for the breakdown of the parallel efficiency.

Section 3 and model spin up time: see above

Orography in Fig. 15: This was indeed an issue with the display and will be rectified in

C2476

the updated version of the manuscript,

Absolute values in figures: We believe that the referee's suggestion, applied to all 2D plots in Section 3, will make it difficult for the reader to understand and judge the results of the models. For instance, to understand the impact of the model resolution on the orography (Fig. 15), absolute values are much more suitable. Figure 17, which displays the July 1982 rains and thus visualises the monsoon dynamics in form of the shape and northern extent of the rainband, is much less informative if only differences are plotted. To some extent, this also applies to Figure 16, 18, 20, in which the reader can identify the Saharan Heat Low through the position of the maximum temperature and the depression in mean sea level pressure. We believe that by choosing an appropriate colour bar rather than a rainbow colour bar (see also discussion above), the 2D plots will become easier to understand. We will label axes and colour bars with the appropriate units as suggested by the referee.

Section 4.3, NaNs: The default compiler flags used in MPAS for Bluegene systems do not enable floating-point error trapping via the "-qflttrap" option. Without this flag, floating-point exceptions like a division by zero or the use of a NaN as an operand do not generate signals that would cause the program to halt. We will update the discussion in Section 4.3 to emphasise that for this reason, the model does not simply halt upon encountering a NaN.

Following the referee's suggestion, we will use nodes in Section 4 as the unit of choice and we will mention this difference to Section 2 at the beginning of Section 4. Also, we have adapted Figure 21 so that each column is labelled by category and time spent in seconds.

(3) Technical Corrections

Both typos will be corrected in the revised version of the manuscript.

Please also note the supplement to this comment:

C2477

<http://www.geosci-model-dev-discuss.net/8/C2473/2015/gmdd-8-C2473-2015-supplement.pdf>

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