Geosci. Model Dev. Discuss.,
doi:10.5194/gmd-2016-126-AC2, 2016

Geoscientific Model Development Discussions

Interactive comment on "A diagnostic interface for the ICOsahedral Non-hydrostatic (ICON) modelling framework based on the Modular Earth Submodel System (MESSy, 2.50)" *by* Bastian Kern and Patrick Jöckel

Bastian Kern and Patrick Jöckel

bastian.kern@dlr.de

Received and published: 2 September 2016

We thank the anonymous **Referee #2** for this concise, however, incisive review. The questions raised are very important. We answer them in the following, and adapt the manuscript to express more precisely our appraisal to these important topics.

General comments

The authors describe an in-situ method of post-processing and claim improved

throughput and reduced disk space requirements when compared to current practices. But they do not appear to actually compare their results to the current practice.

We did not intend to conduct any quantitative measurements w.r.t. throughput and disk space resulting from the application of on-line diagnostics. Numbers would depend on which and how many on-line diagnostics are applied. With the anticipation that data, which are "aggregated" in some way – may it be temporal or spatial average, column integrated values, any combination of variables, ... – have a smaller volume than the raw 3D model output, better throughput and lower disk space requirements are obvious. How this reduction in data volume is achieved, is another question. Currently, there are suggestions to use data compression (on-line or as post-processing) to decrease data volume and reduce the network and storage footprint (Kuhn et al., 2016; Baker et al., 2016).

An in-situ "post-processing" approach can reduce data volume. In practice, we anticipate a combination of the on-line "post-processing" and a limited set of "raw" data, which both could also undergo some on-line or off-line compression (lossy or lossless). As during the simulation checkpointing is anyway required, also a re-calculation of certain periods of the simulation with more output data could be a solution, when interesting features and evolution is detected in the in-situ "post-processing" data.

Does this new method in fact save time or disk space requirements?

If this approach saves disk space depends on which and how many on-line tools are applied, and which raw model fields are written out in addition. But when we can get rid of a large part of raw 3D output this is an opportunity to save huge expenses.

For the GRAGG submodel, as applied in this study, the aggregation has been applied on a 0.5 $\times$ 0.5 target grid, using 10 $\times$ 10 bins for the joint PDF. In this specific case, the data volume of the raw output is about 4 times the output of the jPDF. In case of averaging variables onto the coarser grid, the data volume is reduced by a factor of of 187.

Are scientists willing to reduce the number of model fields written to disk if derived fields are computed in-situ?

We do not know if scientists are willing to do so, but latest with the advent of Exascale computer systems, scientist applying then-state-of-the-art ESMs may have no choice but to reduce the volume of output data in some way. Of course, this will largely change the overall workflow. Today, many simulations are conducted with data output, which is not necessary for the primary scientific question to be answered. However, it offers the opportunity for the scientist or another group to conduct further analyses on these additional data and maybe find some interesting insights.

We do not suggest to do all nowadays post-processing on-line, but we want scientists to think about, which data is needed, which analyses are useful, and which of them could be done already on-line. The tricky part is to find a balance between data reduction by reducing disk output through on-line analyses, and data which need to be written out in full resolution. With checkpointing, which is anyway required, scientists may have the opportunity to restart simulations of a certain period, when they find interesting features. Some additional data may be needed of course for quality checking.

The additional time to be invested for the development and testing of the on-line tools could potentially be saved from reduced efforts in post-processing and data handling (e.g., de-/archiving).

We are convinced that a combination of on-line tools, data compression and re-calculation from checkpoint data will shape the workflow in the next decades.

We reformulate and extend the *Introduction* of the manuscript:

(p. 1, l. 23)

"*The performance of present day HPC systems reaches the Petascale, i.e. providing peak performances of more than* 10^{15} *operations per second (see http:// top500.org/). While the increase in computational power still follows Moore's law, this does not hold for the performance of the I/O (sub-)system. I/O operations are limited by the bandwidth the I/O system offers to the compute system. With increasing computational power, the ratio between I/O bandwidth and floating point operations per second decreases. Thus, for many applications the main bottleneck is I/O, potentially impacting the application performance on the next generation of supercomputers (Ali et al., 2009). Furthermore, storage capacity grows at a smaller rate as the computational speed (Kunkel et al., 2014; Kuhn et al., 2016). These problems especially persist for data intensive applications like numerical climate models, relying on fast data throughput and large storage capacity.*

An obvious solution addressing the problems of data handling, the I/O performance gap, and the limited storage capacity, is to reduce the amount of output data. But a simple spatial and temporal reduction (coarse-graining) may prove counterproductive. When we apply post-processing tools to the coarsened datasets, it deteriorates advantages gained through finer resolved numerical simulations.

Recent studies on volume reduction of climate data focus on lossless and lossy compression of data on-line or as post-processing (Kuhn et al., 2016; Baker et al., 2016). We propose an approach of using diagnostic analyses, which are usually applied on the data after they are stored on disk, on-line during the model simulation. Depending on the analyses needed and selected, with the on-line diagnostic tools we might be able to drop parts of the output of variables on the full model resolution. This reduces data transfer over the computer network and disk storage occupation. The application *of on-line diagnostic tools enables us to use the fully resolved data, to perform data reduction, and even to apply advanced diagnostics, which are not applicable off-line. On-line diagnostic tools can also be combined with any kind of data compression.* "

We extend the *Conclusions* Section of the manuscript:

(p. 16, l. 31)

"*The ratio between I/O bandwidth and peak performance of high performance systems decreased in the recent years, and probably will continue to do so for future systems, as the performance gap widens. Additionally the increase in storage capacity has slowed down recently. We have to focus on these problems during the design and provision phase of future HPC systems. Apparently, there will be no quick solutions to the problems, as the widening of the performance gap might persist for decades. For now, we will have to focus on software solutions to bridge this gap by providing efficient I/O operations, making use of parallelisation and appropriate I/O middleware in our codes, and also by reducing the amount of output data. Recent studies focus on data compression as a (transparent) post-processing step, either when the data is stored, or even on-line before data is transferred over the computer network to increase communication performance.*

We suggest the application of on-line diagnostic tools to reduce the volume of data from numerical simulations. This will require a largely modified workflow for scientists, as appropriate analyses have to be selected and on-line tools have to be developed before the simulation starts.

Checkpointing is already applied by model systems to overcome restrictions in compute times of job schedulers, and to restore simulations in case of software or hardware failures. Even when cutting the "full" model output to a minimum, a subsequent re-calculation with increased output volume and frequency is therefore possible. Thus,

scientists should not be afraid of "losing" data, as a combination of all methods should suit their needs.

Currently applied workflows may not retain the same on future HPC systems. A certain degree of development and optimisation of current code and the application of novel methods are required. Here, a generalised interface for on-line diagnostic tools is extremely useful. "

(Additional) References

Baker, A. H., Hammerling, D. M., Mickleson, S. A., Xu, H., Stolpe, M. B., Naveau, P., Sanderson, B., Ebert-Uphoff, I., Samarasinghe, S., De Simone, F., Carbone, F., Gencarelli, C. N., Dennis, J. M., Kay, J. E., and Lindstrom, P.: Evaluating Lossy Data Compression on Climate Simulation Data within a Large Ensemble, Geoscientific Model Development Discussions, 2016, 1–38, doi:10.5194/gmd-2016-146, 2016.

Kuhn, M., Kunkel, J., and Ludwig, T.: Data Compression for Climate Data, Supercomputing frontiers and innovations, 3, doi:10.14529/jsfi160105, 2016.

Kunkel, J., Kuhn, M., and Ludwig, T.: Exascale Storage Systems – An Analytical Study of Expenses, Supercomputing frontiers and innovations, 1, doi:10.14529/jsfi140106, 2014.

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-126, 2016.