Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2017-301-RC1, 2018

Geoscientific Model Development Discussions

Interactive comment on "Simulation of the Performance and Scalability of MPI Communications of Atmospheric Models running on Exascale Supercomputers" *by* Yongjun Zheng and Philippe Marguinaud

Anonymous Referee #1

Received and published: 9 February 2018

The authors present work of simulated scaling analysis for different communication algorithms commonly used in atmospheric models using a skeleton codes and a simulation package to examine the scaling performance on possible future supercomputers. This represents significant new information on how these algorithms may perform and is likely to be of interest to the community. The methods used are well described and appear robust. Some of the assumptions made about future architectures, in effect single CPU core nodes, are unlikely to be entirely valid. Whilst these are made the entirely reasonable purpose of make the simulations tractable, they may weaken some of

Printer-friendly version

Discussion paper

the conclusions. For example, almost all CPU base supercomputer are multi-core and multi-socket nodes which then have significant network hierarchy. Moreover, many of the largest machines in the top 500 list have non-CPU architectures such as GPUs and Xeon Phi. These have more complex hierarchies and are unlikely to, or even cannot be, programmed with a single MPI rank bound to single "core". Whilst the authors don't hide this, this is not discussed in the conclusions. Most of the results are presented in the form of graphs. Unfortunately, they are simply too small and it not possible to read the legends, axis labels etc. This makes it difficult to judge the quality of the results and the inferences drawn. These should be reproduced to appear much larger. Moreover, it would appear (although hard to be sure) that some of the plots have number of processors as the x-axis. This is a discrete variable and so line graphs should not be used, a bar chart may be appropriate. Whilst it may be common practice to present scaling data in this way, it is still wrong. This paper has the potential to become an interesting and significant work, but not in its current form. Once some revisions have been made it should be review again. In particular, there are three changes which are necessary. i) The plots must be made bigger so they are legible ii) Plots against discrete variables shouldn't be line graphs iii) The authors should comment on and discuss what conclusions can be drawn from simulations of single core nodes for more complex node architectures and the consequent differences to communication patterns.

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2017-301, 2018.

Printer-friendly version

Discussion paper