

Interactive comment on “Asynchronous Communication in Spectral Element and Discontinuous Galerkin Methods for Atmospheric Dynamics” by B. F. Jamroz and R. Klöfkorn

Anonymous Referee #1

Received and published: 23 April 2016

In the paper, the authors describe their implementation of asynchronous communication in the High-Order Methods Modeling Environment (HOMME) dynamical core of the Community Atmosphere Model (CAM). Data packing is overlapped with communication for the continuous spectral element method in HOMME at the cost of bit-for-bit reproducibility. Thus the Jablonowski–Williamson baroclinic wave instability test case is used to verify the changes to HOMME were accurate. In the discontinuous spectral element method in HOMME, communication is overlapped with data packing as well as computation and the authors were able to achieve bit-for-bit reproducibility. Scaling results on the Yellowstone supercomputer are shown for both the continuous and discontinuous spectral element methods for the the Jablonowski–Williamson baroclinic

[Printer-friendly version](#)

[Discussion paper](#)



wave instability test case.

Investigations on speeding up HOMME, like what is presented in the paper, do have an impact on the community that uses it in CAM and is thus an interesting topic. The paper describes an implementation of overlapping communication and computation that in practice is well know. It focuses on a detail of this that I don't know if it is known. Specifically, posting the send once the data is packed for a single neighbor and not waiting until all the data is packed into the communication buffers. The approach is demonstrated to yield up to a 16% decrease in run time, depending on the number of elements per process. The methods and assumptions are clearly outlined. The authors fail to report on others attempts to overlap communication and computation to support their claim that they their approach is new. With some work, and access to the Yellowstone supercomputer, I believe the authors work is reproducible. They authors provide a subversion repository and branch where a majority of the code can be found.

Although this change to HOMME is incremental, I believe that improvements in model speed are useful for publication. Please find more comments about issues that the authors should address below.

First, I would suggest that the authors proofread the paper again as there are some typographical errors. Below is a small sample of typographical errors and is not meant as an exhaustive list.

1. *typo on page 5*: “eg.” should be “e.g.”
2. *typo on page 8*: “a auxiliary diagnostic variables” should be “auxiliary diagnostic variables”

Second, I would suggest the authors address the following content concerns.

3. First the authors state :“That is, a process sending a blocking message must wait until the message has been received.” This is technically not true if you are

Printer-friendly version

Discussion paper



- talking about `MPI_Send`. The function `MPI_Send` only blocks until the buffer can be reused. If you are not talking specifically about `MPI_Send` this needs to be clarified.
4. It is unclear how many time-steps were used to compute the numbers in Figure 5 and Tables 1 and 2. Some more detail should be added to the captions.
 5. In Figures 1, 5 and 6 line plots are used for discrete data. Is there a piecewise linear fit between the data? I would suggest to use only symbols where the actual data measured is.
 6. I believe the title needs to include the code name and version.
 7. A literature survey should be given to support the claims that the authors approach is new. I have rarely see this kind of detail presented in the literature, so maybe a review of the most popular finite element methods is in order.

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-23, 2016.

[Printer-friendly version](#)[Discussion paper](#)