

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

Anonymous Referee #2

Received and published: 23 May 2016

This is a short paper describing work done in the finite element boundary communication in the HOMME model, for both CG and DG discretizations. The authors have refactored the communication routines to allow for overlapping of communication and computation. The paper is well written and clearly explained. The speedups are modest, but still significant and the approach should be adopted by the HOMME model. Larger speedups should be possible.

The authors also describe their use of a MPI micro-benchmark for evaluating the effectiveness of non-blocking communication. Results are presented for a particular machine ("yellowstone"). I assume before undertaking the code refactoring necessary to overlap computation and communication, one should establish how effective this can

C1

be on the target machine. What can researchers expect on modern parallel computers? Are the results on yellowstone typical of an average institutional machine?

page 1, typo, "ocaen" -> "ocean"

page 1, line 20, add: "The SE and DG methods attain this property for arbitrary order, at the expense of a small timestep"

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

C2