

## ***Interactive comment on “Development and performance of a new version of the OASIS coupler, OASIS3-MCT\_3.0” by Anthony Craig et al.***

### **Anonymous Referee #3**

Received and published: 10 May 2017

This paper describes a new version of the OASIS coupler, called OASIS3-MCT, following upon other papers documenting earlier versions.

The key advance here is the inclusion of the MCT (Model Coupling Toolkit) as a layer providing parallel regridding services. The paper convincingly demonstrates the benefits of the MCT layer in computing interpolation weights online and offline. The analysis is quite comprehensive, and covers the technical aspects of coupling and interpolation. It explores the performance aspects of choosing to perform the computationally expensive operations on the source or destination side, or on dedicated core(s). It covers the issues of various coupling timestepping algorithms and parallel layouts, with different arrangements of components in serial and parallel. The tests are well-designed and representative of very high-end models. They have been performed with some toy models, but with resolutions set to the highest-end models of today, and at challenging

[Printer-friendly version](#)

[Discussion paper](#)



core counts.

Links have been provided to ensure public availability of code.

The paper is very well-written with no obvious typos or poor sentence construction.

I recommend publication as-is, as it is a complete description of an existing piece of software. The comments below are what I found to be omissions in the paper, but perhaps should be treated as suggestions for future versions.

1) There is no discussion of OpenMP as an alternative to MPI. Future hardware will require going to more shared memory and less message passing.

2) There is no discussion of GPUs, MICs, etc and plans to port OASIS to novel architectures.

3) I am somewhat taken aback by the extreme cost of providing bfb (bit for bit) reproducing algorithms. In other similar codes this cost ratio is somewhat lower (which could of course mean that the non-reproducing modes in other codes are too slow!) This may require some work.

---

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2017-64, 2017.

Printer-friendly version

Discussion paper

