

Interactive comment on “Efficient ensemble data assimilation for coupled models with the Parallel Data Assimilation Framework: Example of AWI-CM” by Lars Nerger et al.

Lars Nerger et al.

lars.nerger@awi.de

Received and published: 26 November 2019

We like to thank for reviewer for the careful review. Please see our response below.

The paper describes the implementation of the software tool PDAF to a coupled ocean-atmosphere model. It discusses essentially the general structure of the PDAF software and how the coupling can be realized on a distributed computing architecture with MPI. While this is interesting, my main issues with this manuscript are the following 4 points:

C1

1. No actual results of the assimilation system are presented. Only the execution time for different settings. It is unclear to me what the role of a reviewer can be in this case. I rather think that the paper should also include the results of such model (see also the following point).

Response: The manuscript was prepared for the particular scope of the Journal Geoscientific Model Development (GMD) as a technical development study. As such the manuscript focuses on the technical aspects and the scalability. Discussing actual data assimilation results would not be in line with the scope. Actually, given that coupled data assimilation is a young approach, and still challenging, we think that GMD would not be the right journal to discuss application results of coupled data assimilation as we would not reach the intended readers. Apart from this, the scalability was only assessed with short experiments over 10 days (i.e. 10 analysis cycles). During this time, the assimilation process is in the initial spin-up phase and the assimilation effect is still small. Significantly longer experiments over a few months or a year would be required to get significant assimilation results. Given limited computing resources, we cannot perform full-length scalability experiments.

To respond to the authors recommendation, we have revised the introduction to better point out the status and challenges of weakly and strongly coupled data assimilation. This should clarify why here we only discuss experiments with weakly coupled assimilation.

2. The manuscript mentions different approaches to implement the assimilation in a coupled system: in a combined state vector spanning the atmospheric and ocean model or separately. The question about which approach is better is still open and it should not be too difficult to the authors to check both approaches. This would help

C2

also to address the previous point and add substantially to the scientific value of this paper.

Response: While technically the strongly-coupled data assimilation is not too difficult, as is actually discussed in the manuscript, the application as such is. Strongly-coupled data assimilation is a very young approach and there are hardly any papers on this topic. By now, we know that the plain application of strongly coupled data assimilation does likely not give optimal results. To this end, we didn't attempt strongly-coupled DA in this manuscript as the assimilation results are most likely not representative. This led to the decision to discuss the model binding AWI-CM-PDAF version 1.0 for weakly coupled data assimilation. This scope is now better clarified in the manuscript.

3. The different time scales of ocean and atmosphere are not discussed and the assimilation is done only in the ocean. To really appreciate the effectiveness of the coupling, data should be assimilated in both the atmosphere and the ocean and the question regarding the assimilation frequency should be addressed. As usual, the models should be validated against independent observations.

Response: As mentioned before, following the scope of GMD, the manuscript discusses the PDAF model binding for a coupled model as a technical development paper. Discussing application results and the question of the assimilation frequency would be a different study, which would certainly not suited for GMD.

4. There is too much overlap between this manuscript and previous manuscripts by the same author concerning the description of PDAF (in particular the memory coupling, general API structure). I think the author should focus this paper on the

C3

coupling aspect and just reference to elements already published before.

Response: We have revised Section 3 (in particular 3.1 and 3.2) to also discuss the particularities of coupling PDAF to the coupled model with multiple executables. For completeness of the manuscript, we prefer to keep aspects like the in-memory coupling or the added subroutines in the manuscript, even though quite a bit of these aspects were already discussed in the previous study (Nerger and Hiller 2012). Even more, aspects like the routine 'Assimilate_PDAF' are new, and its discussion is only possible when also discussing the routines 'Init_parallel_PDAF' and 'Init_PDAF'.

I therefore recommend major revision before this article is published in GMD.

Minor comments:

line 46: tranDAsfers -> transfer

Response: corrected

page 6: MPI Communicators: is this discussion not too technical?

Response: Given that the manuscript is submitted to GMD, we think that the degree of technicality is just right (e.g. see also Kurtz et al., 2016). The configuration of the communicators is actually a core part that makes the online coupled of PDAF with AWI-CM work.

C4

Section 5: How the system scales for a fixed ensemble size?

Response: For a fixed ensemble size, the scalability is determined by the scalability of the models (as discussed in our previous papers on PDAF). As this holds likewise for assimilation with uncoupled and coupled models, we didn't perform systematic scalability tests on this aspect. We now mention the scalability for a fixed ensemble size in Sec. 4.1.

Figure 6: the label mentions relative execution times, but the unit on the axis is [s].

Response: corrected

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2019-167>, 2019.