

Interactive comment on “Compact Modeling Framework v3.0 for high-resolution global ocean-ice-atmosphere models” by Vladimir V. Kalmykov et al.

Anonymous Referee #2

Received and published: 16 April 2018

This paper presents versions 2.0 (based on MPI) and 3.0 (based on PGAS) of the Compact Modeling Framework, a software infrastructure offering different services (coupling, I/O, data assimilation) to implement numerical geophysical models. Tests on different architectures are performed and show the scalability of the framework for relatively high-resolution problems. The paper contains interesting information on CMF that would deserve publication in GMD but I consider important parts are missing and the manuscript certainly needs major revisions before it can be considered for publication.

A) Major comments:

C1

A.1) The description of CMF2.0, results of Test I and II and related conclusions need to be clarified:

- First, I agree with Referee #1 that figure 1 is misleading as it suggests a 1:1 connection between the models processes and the coupler processes. I don't understand either the "which means data locality" on P.4, L.50, as there is certainly some exchange of data needed between the component processes and the coupler processes.

- P.5, L.63-66: these two paragraphs are not clear at all. What does "a subset of component's cores works only with individual master core in the coupler" mean? What does "for the two cases, of the source and destination type" mean and why do you put a reference to Craig et al. 2005 here?

- P.5., L. 69-70: You write "All necessary links are initialized at the beginning of run and are used at the calculation stage as persistent (Jacob et al., 2005). I suppose that the links you mention here are not the SCRIP links? You should be using another word as this is confusing. Also, why do you mention Jacob et al. 2005 here?"

- P.5, L.72-73: I don't understand this sentence: "It is worth noting, that links are not sent directly, but as sorted unique cells vectors which allow one to avoid sending duplicated data." Again what "links" are you talking about here?

- In Figure 2, 3 & 4 captions and on P.6, L.88, you should recall what is included in the timing. I suppose this is the time in seconds for the whole 10 model days for the whole ocean-atmosphere and atmosphere-ocean exchanges through the coupler. (Same remark for Figure 6).

- On p.6, L.92, you write: "It is clear that 20-40 coupler cores provide satisfactory speed for such problems, because ~10 seconds costs ..." It is not clear how you get these numbers. I see at best, i.e. with the MVS-10p_16, something between ~15 and ~20 seconds. Same remark for the number presented on P.7, L.4-5. This re-joins referee #1's comment about the fact that the log scale does not allow one to get the detailed

C2

information mentioned in the text.

- P.7, L.6: again here I don't understand why you write "perform only local communication"

- Section 3.4 on I/O should be completely revised. The long theoretical section on P.8 with detailed formula is not useful here, especially as you finally simply state "Asynchronous scheme was incorporated in the latest version our framework" without giving finally any results! The theoretical section should be cut and numbers obtained for the grid sizes you list for INMIO World ocean model should be provided.

- The different utility modules available should be briefly described or a reference to a documentation or User Guide should be provided.

A.2) The description of CMF3.0 needs to be extended

This paper is supposed to mainly describe CMF3.0, or at least this is what the title implies but very little is written about it. It looks like the author was in a hurry to finish the paper. More details should certainly be given on how the I/O service work (1st paragraph on P.11) and on the Data Assimilation part (currently only 2 lines, P.11, L8-9). These improved descriptions should be backed up with performance results (as is done for the interpolation).

Also the discussion of the interpolation results should be extended and detailed. How do the author get to 2-3 seconds per modelling day on 20-50 CPL cores? (This is mentioned also in the conclusions P.15, L.28.) Why is it expected that results would be worse than for CMF2.0 (and "worse" should not be used here because it implies that results for CMF2.0 are bad and that results for CMF3.0 are even worse)? Is it because of the shift from MPI to PGAS? Or because the tests were performed on a different platform?

Also, in the conclusions, you write: "The key part of it, coupler, has a sufficiently small code size for such programs (about 5000 lines of code with unit tests) and is able to

C3

manage the main parallel problems of the coupled modeling - synchronization, regridding and I/O." I don't understand why you write that the coupler manages the I/O as this is not the case in CMF3.0.

A.3) The whole text needs reviewing by a native English speaker. The style and wording needs revision as some sentences are simply not understandable (at least by me), e.g.:

- P.3, L.16: "Unquestionable advantage of non-coupler design is the absence of interference in the user code": why "non-coupler design"? Also, with OASIS, there is some interference in the user code but the objective is to minimize it.

- P.5, L.76: "Several ping-pong tests were carried out for interpolation system using coupled ocean-atmosphere model."

- P.6, L.95-96: "work of the sequential algorithm is only possible with restriction that memory is allocated only for interpolation block, which is impossible in practice": I am not sure what this sentence means exactly and why you write that it is impossible in practice while you do get some results on 1 core; do you mean it would be impossible with real models as the interpolation per se would require all the available memory?

- P.13, L.63: "but we are more interested in scalability of the program on perspective sizes of computational resources"

- P. 13, L.71: "Time evolution of the sea-ice surface temperature is described in the same way as in prescribed ocean experiments.":

- P. 13, last paragraph: please rephrase, the current sentence with the "min." is too difficult to follow.

- P.2, L25: "Coupling through shared file or sequential component is acceptable ...": could be "Coupling through shared file with components executing sequentially is acceptable ..."

C4

- P.2, L33: “. . . and their representation in the interfaced style understandable . . .” could be “. . . and their adaptation to the interface understandable . . .”
- P.2, L41: “GFDL FMS (Balaji, 2012) system additionally suggests fully parallel data storage with file post processing at the end of the run” could be “In the GFDL FMS (Balaji, 2012) system, fully parallel data storage with file post processing at the end of the run is offered”
- P.3, L.21: “According to proposals of Earth System Modeling conference, (Valcke et al., 2012), . . .” could be “According to the analysis of coupling technologies for Earth System Modeling by Valcke et al. 2012, . . .”
- P.6, L86: “Performance is based on a standard Intel Fortran compiler” should be moved to the next paragraph and could be “On all supercomputers, the coupled system was compiled standard Intel Fortran compiler.”
- P.6, L.90: “Increasing number of coupler size”; should be “Increasing number of coupler processes” or “Increasing the size of the coupler communicator”
- P. 11, L. 12: “Optimizations regarded to ignore repeated cell requests are preserved” could be “Optimization regarding repeated cells are preserved”.
- P. 13, L.56: “Latest version of INMIOWOM model was fully integrated to CMF” could be “CMF2.0 was fully integrated in the latest version of INMIOWOM”
- P.14, 3rd line: “Ice model was built into the ocean model , land model – into the atmosphere model” could be “The ice model is integrated in the ocean model and the land model in the atmosphere model”
- P.14, L.6: “structurize” could be “structure”

B) Other comments:

B.1) References:

C5

- P.2, L.6: The reference to OASIS3 should be Valcke 2013 (i.e. p18, L45)
 - P.2, L.36: The reference to OASIS3-MCT should be Craig et al 2017: A. Craig, S. Valcke, L. Coquart, 2017: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, Geosci. Model Dev., 10, 3297-3308, <https://doi.org/10.5194/gmd-10-3297-2017>, 2017.
 - P.3, L.7 (as annotated in the manuscript): Reference to ESMF should be the more recent one: Theurich, G., Deluca, C., Campbell, T., Liu, F., Saint, K., Vertenstein, M., Chen, J., Oehmke, R., Doyle, J., Whitcomb, T., Wallcraft, A., Iredell, M., Black, T., Da Silva, A. M., Clune, T., Ferraro, R., Li, P., Kelley, M., Aleinov, I., Balaji, V., Zadeh, N., Jacob, R., Kirtman, B., Giraldo, F., McCarren, D., Sandgathe, S., Peckham, S., and Dunlap IV, R.: The Earth System Prediction Suite: Toward a Coordinated U.S. Modeling Capability, B. Am. Meteor. Soc., 97, 1229–1247, <https://doi.org/10.1175/BAMS-D-14-00164.1>, 2016.
- B.2) P.2, L.35-36: OASIS3-MCT proposes coupled system not only as single executable, so this sentence is misleading; the important feature is that the coupling functions are not provided by a standalone coupler but by a coupling library linked to the component models. Please correct.
- B.3) P.2, L.44-50: This list mixes functionalities (1. and 2.) and characteristics (3. and 4.); please reorganise.
- B.4) P.3, L.11: It is not really fair to write that the computational costs of CESM coupler are quite significant 20%, as the CESM coupler does not only perform coupling and remapping but also performs the surface flux computation.
- B.5) P.3, L. 21-24: these 3 lines do not give an appropriate summary of the analysis provided by Valcke et al. 2012. Please correct.
- B.6) P.4, L.42-43 & P.9, L.47: Either provide details on what interceptor or Template methods are or don't mention them; one should not have to read the reference (Gamma

C6

et al., 1995) to understand the sentence.

B.7) P.5, L.77: References here are misleading as they seem to imply that “Test I condition” are explicitly defined in Valcke et al., 2012 or in Craig et al. 2012, while they are not.

B.8) P.6, L.93: You should not use the word “failure” here as the test does not fail, it is just very slow.

B.9) Using very technical coding terms along the text does not help understanding it (e.g. P.4, L.45: “Component class”; P.9, L.63:“resulted in class Communicator”; p.10, L.77-78:“since all services in CMF3.0 inherit base class Service it also allows one to easily add new Åž ; P.10, L.80 : ÅŕIt receives data using Communicator Åž ; P.12, L.25 : Åŕ NormalEvent Åž or Åŕ SyncVarEvent Åž : P.12, L.27 : ÅŕGenerators realize abstract class EventGenerator, so new specific generator subclasses could be easily added); I think it would be better to explain the concept that using these abstract terms.

B.10) P.12, L.26: Reference to Griffies et al. is not useful here.

B.11) P.12, section 5.1: What is the resolution of the INMIO World Ocean model in these tests?

B.12) P.15, L.25-27: please specify if these numbers apply to CMF2.0 or CMF3.0. On line 27, change “CPL3.0” for “CMF3.0”!

B.13) In general, I think the section 6 on Conclusions and future work could be fleshed out.

B.14) I think the “Code availability” section is not satisfactory regarding GMD standard but I will let the Topical Editor decide on this point.

C) Minor comments:

C.1) In the abstract, you write “As addition a parallel realisation of the EnOI (Ensemble Optimal Interpolation) data assimilation method as program service of CMF3.0 is

C7

presented.” but this is not the only example presented in section 5.

C.2) P.3, L.13: It is not right to write that OASIS3 is the most popular version of OASIS as most groups are using OASIS3-MCT today.

C.3) P.3, L.16: In OASIS3-MCT, MCT procedures are executed on all component model cores

C.4) P.4, L.33: Define SOA the first time it appears in the text

C.5) P.2, L.53: Define WOM the first time it appears in the text.

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2017-294>, 2018.