Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2018-280-RC1, 2018

Geoscientific Model Development Discussions

Interactive comment on "Improving climate model coupling through a complete mesh representation: a case study with E3SM (v1) and MOAB (v5.x)" *by* Vijay S. Mahadevan et al.

Anonymous Referee #1

Received and published: 9 December 2018

General Comments

The paper describes with accurate details the work accomplished to implement an efficient workflow for remapping and communications tasks in E3SM. The new workflow being based on the MOAB data structures and mesh libraries and on the TempestRemap algorithms, it is identified by the MBTR acronym.

Thourough details on the algorithm and on the implementation are provided and a complete performance analysis is carried on for two test cases.

The work is interesting and deserves publication, yet, in the current formulation, it risks to be a mixture between a technical report and a research paper.

The potential of the proposed workflow should be better situated w.r.t. to other coupled infrastructures than E3SM.

Some suggestions are provided in the "Specific Comments" and "Needed clarifications" sections.

Specific Comments

The scheme at p.4 ll.7-17 and the following comparison in section 2 should clearly distinguish what features are available in distributed softwares or have just been presented as conceptual algorithms (e.g. the advanced clipping: is it in Portage?) and what has been practically tested by the authors or just inferred from documentations (suggestion: avoid sentences like "It is also unclear whether" unless you add the source of your information. User guide, publications, application cases, ...)

From the user point of view, it is important to know beforehand the amount of information needed to describe the meshes, the decompositions, the fields and the treatments. In the comparison of the coupling approaches this point should be stressed. The MCT paradigm requires a very agile data description (in it's OASIS3 implementation, it is a commitment to be able to work without the connectivity description - at the price of being "oblivious" of some structures). Please assess somehow the user friendliness of the MOAB API's (in particular in their fortran version). A good anchor could be p.8 ll.18-19 where you mention the need of introducing extra calls to describe the details of the mesh to MOAB.

Please include considerations on the memory requirements for storing the MOAB data structures and the supermesh informations. Is there any extra-memory to be accounted for on the source and target processes if adaptive or moving meshes have

to be enrolled runtime in MBTR? This assessment could make the last paragraph of section 3.1 more useful, since its aim is not very clear in the current paper. Refer also to step 5: of Algorithm 1.

Description in section 3 is fluctuating between the hub-and-spoke and the MOAB workflow (e.g. p.10 l.8) please state clearly what's the starting point, the reference for comparison and the new proposal.

The potential of hybrid parallel implementations (MPI processes + threaded tasks) is not always consistently addressed neither in MBTR (e.g. for intersection computation) nor for comparison. Check that the use of process and task is coherent through the whole paper (in particular section 3.4), please.

In order not to restrain the scope of this paper to the replacement of ESMF in E3SM, how would you assess and compare the overall efficiency w.r.t. to a "non hub-and-spoke" coupler interleaving computation and remapping on the same sets of processors (e.g. YAC, OASIS3-MCT) and with couplers already addressing the issue of online weights updates (YAC, C-COUPLER2)?

Needed clarifications

———————

p.4 l.29 explain "consistently respecting the underlying discretization" or remove the sentence.

p.6 l.6 define (or cite) "component architecture". Versus what?

p.6 l.21 does "Fig. 1 (right)" apply to OASIS3-MCT also? Or does the mere-library approach (no separate N_x for the coupler) defines another workflow?

p.7 l.25 what does "field [...] aware" mean?

p.7 l.29 "during the setup phase" is in contradiction with the aim of allowing for adaptive and moving meshes. Indicate whether it is just a practical choice in the current

implementation.

p.7 ll.30-32 in what exactly is the MBTR stack an improvement w.r.t. the MCT view, since MCT is able to handle decomposed meshes?

p.8 l.3 indicate under which scheduling assumptions the N_x processes can share with the N_{c,l} processes part of the processor ressources as implied later by Fig. 4.

p.8 l.14 please define a "DoF": since it is not a word used for cell-centered couplers, it is not a common term for all the readers.

p.8 whole section 3.1 (and following) please include references or links for HOMME, MPAS, VisIt and in general do so for all mentioned models, libraries and other software tools (Zoltan, ParMetis, Eigen3, etc).

p.8 l.26 why "replicated" meshes. Isn't it rather "partitioned"?

p.8 l.29 "in terms of a 'Tag'." is a useless statement unless you make it clear to the reader.

p.8 l.29 n_p has not been defined and is not trivial.

p.9 Algorithm 1. The formulation is too compact and missing some previous definition. Insert references to following sections for details.

p.11 ll.2-3 state here (or anticipate) the rational for replacing MCT as a broker.

p.12 l.8 Kd-tree is a relatively common technique (already mentioned at p.4) BVH-tree deserves a reference here (only provided at p.20).

p.12 l.11 why "unique" ?

p.12 l.17 the same consideration as for p.7 l.29 applies.

p.12 ll.26-29 Fig 6. is not immediate to read without some further "step to step" details in the text.

pp.12-13 subsection 3.3.1 does the seed determination can be fully automated or its efficiency depend on user tuning?

p.13 l.11 what does the sentence "without approximations" refer to (especially w.r.t what alternative)?

p.14 l.6 computing a meaningful bounding box is not trivial in polar or periodicity regions for lon/lat grids.

p.14 l.7 does "to all tasks" refer to tasks (or rather processes) on the source side?

p.14 l.8 "Cells [...] are sent": how are they represented? What's the size of the communications? Is any packing strategy used to avoid latency in separate small communications?

p.14 l.10 please clarify the term "superset": a superset usually refers to inclusion of similat objects. Does it imply that the after representation in MOAB - through the definition of the supermesh - the source and the target side share the same spatial discretisation?

p.14 l.14 is the "crystal" router explained in Tautges et al. (2012) [N.B. reference not freely available] or does it need an extra reference?

p.15 l.5 how expensive can be the communication of ghost intersection elements on highly distributed components?

p.15 l.13 does "has the potential to" mean that is just an idea or is there a prototype?

p.16 l.28 "it is non-trivial to": did you find a way?

p.21 l.8 reference to NE11 configuration not known to the reader.

p.26 Fig.14(b) provide an explanation for the difference of behaviour when going beyond 64 processors

Technical Corrections

---

p.2 l.5 vs l.31 (and elsewhere) make the use of "donor" or "source" consistent

p.2 l.9 should probably be "conservation for critical quantities"

p.3 l.20 the subject of "that nonlinearly couple" should not be "solution fields" (they are just exchanged in the nonlinear coupling process)

p.6 l.6 unclear (if not useless) reference "Section (1)"

p.6 l.32 the "GLL acronym" is used before definition which is given a few lines later

p.7 l.29 "an in-memory" instead of "a in-memory"

p.8 l.26 remove "a" before "replicated SE and MPAS" meshes

p.9 Algorithm 1.

- Step 1: if l can only be s or t - as in step 4: - indicate l \in [s,t] also in step 1: otherwise if the formulation is generic for more than one mesh for component, the naming should be consistent.

- Step 2: if you indicate W_{ij} instead of W_{st} you should not define i,j as a mesh pair. Later at step 20: i takes a specific meaning.

p.12 l.1 "partitioner" instead of "repartioner"

p.12 l.23 remove "is" before "results"

p.12 l.26 "each" instead of "Each"

p.14 Fig.7 caption: "fully covers" instead of "fully cover"

p.15 l.3 "the intersection vertices [...] need" instead of "needs"

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2018-280, 2018.