

Interactive comment on “HOMMEXX 1.0: A Performance Portable Atmospheric Dynamical Core for the Energy Exascale Earth System Model” by Luca Bertagna et al.

Anonymous Referee #2

Received and published: 29 November 2018

"HOMMEXX 1.0: A Performance Portable Atmospheric Dynamical Core for the Energy Exascale Earth System Model"

Authors: Luca Bertagna, Michael Deakin, Oksana Guba, Daniel Sunderland, Andrew M. Bradley, Irina K. Tezaur, Mark A. Taylor, and Andrew G. Salinger

The authors describe porting the HOMME dynamical core (used by the atmospheric component in both E3SM and CESM) to HOMMEXX. HOMMEXX is a single source code that supports execution on both CPUs and GPUs. The work required to transform this sizable code base to be (mostly) platform independent was likely quite substantial and was a rather bold undertaking. Because climate simulations are performed

Printer-friendly version

Discussion paper



by the community on a variety of architectures, the single code base solution is appealing - especially in preparation for exascale computing. While HOMMEXX may not get the performance on GPU that one would aspire to, it's frankly impressive that the code transformation did not result in worse performance on the CPU. Also I appreciate that many architectures were evaluated. That said, I do think that the manuscript itself could use some improvements, which I detail below.

————— Comments/suggestions/questions —————

(1) The authors do say that there was substantial work involved in this refactoring, but I am interested in them quantifying it in some way (if possible) - e.g., people hours, or lines of code touched, or percentage of HOMME code that is unchanged in HOMMEXX ?

(2) On the same line of thinking, can the authors speculate how this effort would compare to the effort that would be required to use OpenACC to port HOMME to GPUs. (Section 5 indicates that the authors are considering this undertaking as well.)

(3) I would have liked to see some performance metrics in the paper that are in more common use by the climate modeling community, e.g. simulated years per compute day or CPU-core hours per simulation year.

(4) Because the stated motivation of this work is running at exascale, it seems that the Dennis et al 2017 paper on this subject ("Preparing the Community Earth System Model for Exascale Computing") should probably be cited and mentioned in Section 1.

(5) Page 4: text refers to Figure 4 before Figure 3 (swap order of these figures?)

(6) When referring to a resolution for the first time in the text via n_e (e.g., $n_e = 240$), please also mention the corresponding grid degree (as done for $n_e = 30$ on page 5 and in Fig. 4 caption).

(7) Section 2: Why only use the outer threading for HOMME? Please explain/justify.

Printer-friendly version

Discussion paper



(8) In general, I'd like to see more detail given in Section 3, given that this is a journal paper (as opposed to length-constrained C.S. conference submission) that should be of interest to climate modelers. It may be necessary to break up section 3 if significant more detail is added.

–Section 3.1: I really would like more justification for the choice of using Kokkos. And I'd like to understand more of what was involved to use it - how about a simple example?

–Section 3.2: What is the correctness-testing build? (p.8. line 15). Is this what lines 16-17 are describing? I assume that this is not what you use for performance testing? Please clarify and consider expanding the correctness discussion.

– Section 3.3: this section could benefit from some code snippets/pseudocode to clarify (9) page 11, line 9: What is a subview?

(10) I believe that Section 4 (results) could be improved quite a bit. It was a bit frustrating at times:

–section 4.1: I'd like it to be easier to get the info on the various test machines and figure out which was which in the figures - so maybe put them in a bulleted list or large table. I'd suggest that the name of the platform - as referred to in the figures should be in bold and appear first. For those readers not familiar with these DOE machines, it was frustrating that Fig. 5 listed "Edison", which was alternatively referred to by its processor (IB) in Figure 6. I kept having to search through information in paragraphs on pages 12 and 13 while looking at the plots. Similarly in section 4.3 (page 16, line 17), when a reference was made to the Xeon machines as a group, I had to search through these paragraphs. I think that not all readers are familiar with these machines, so please make it more accessible.

–page 14: this discussion of power consumption feels stuck in here. The rough guideline numbers on power use (lines 10-14) should probably be included in the bulleted list or table of machines. Or maybe this power discussion should be in a new section

[Printer-friendly version](#)[Discussion paper](#)

4.5 that happens after the reader has seen the other results - making it easier to follow. As it is, figures 5-8 are referred to in line 2 before they have really been presented in the text (which happens in later subsections).

–page 15, lines 10-11: this info feels like it should be in the intro of section 4 (page 12) - not stuck at the end of 4.1

–section 4.1: I also find it unhelpful to refer to the figures before "presenting" them. For example, page 15, line 1 and page 15, line 3: These comments about the plot attributes should go in the sections where the plots are described (4.2 and 4.3)

–Consider combining 4.2 and 4.3 into a single "strong scaling" section. Also make the machine labels more consistent between plots 5 and 6, for example. (Even though 5 has fewer platforms.)

–Section 4.4: Overall, interpreting the results could be easier (more readable) by referring in the text to specific examples in the figures. For example, in page 16 line 29, say which platform in Figure 8 is the one that "does not always perform as well" with a large number of threads per rank.

(11) Table 1: consider naming the kernels for those familiar with HOMME (rather than kernel 1, kernel 2, ...)

(12) section 4.5, lines 14-15: It seems that Figure 6 does not indicate that V100 is strictly faster than HSW, though this text suggests that (the 1.2x to 3.8x)

(13) page 13, line 8: I don't know whether to be concerned about what else may be hardware-specific in this single source code. What percentage of code is different? What are the types of code constructs that are problematic for Kokkos? What are the broader implications for other codes?

————— Minor items: —————

(1) page 7, line 6: "around" => "on"

Printer-friendly version

Discussion paper



(2) page 9, line16: "the the" => "the"

(3) page 9, line 19: "nested loops," => "nested loops"

(4) page 17, line 12: I'd assume that the reader is not necessarily aware of what nvprof is...should at least cite this.

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2018-218>, 2018.

Printer-friendly version

Discussion paper

