

Interactive comment on “Coarse-grained component concurrency in Earth System modeling” by V. Balaji et al.

Anonymous Referee #1

Received and published: 28 June 2016

General Comments:

This is a well composed manuscript written in a lively style to describe experiments in finer grained component concurrency for at atmospheric model. The novel aspect is the use of component concurrency to improve scalability in an atmospheric code. In general the manuscript is appropriate for GMD, but revisions would strengthen the presentation.

The suggested revisions are substantial, but of a nature that the editor can readily adjudicate.

Specific Comments:

1. I enjoyed the more lyrical style, but in places the style came across as glib, particu-

[Printer-friendly version](#)

[Discussion paper](#)



larly when precision was sacrificed for poetry. Grounding the manuscript better in data and substance would strengthen the exposition greatly.

2. The figures are substandard and not sufficiently quantitative. In particular differences between different configurations of the same simulation are preferred over differences to the observations, as the former, not the latter, is the point of the manuscript. Also attention to map projections and color scales is required.

3. The frequent reference to Pauluis and Emanuel was insufficiently discriminating and in places misleading. The reference gives the impression that infrequent coupling to radiation is a substantial source of bias and also model instability. This is not what the manuscript is about. Moreover the Pauluis Emanuel study, while noteworthy, has not been shown to generalize. It might, but the literature is not there. In the grand scheme of things, the trade-off in accuracy and stability of calling radiation more frequently, versus simulating at higher resolution, is not well understood.

4. Some of the hard-core computational issues are insufficiently addressed. In particular, component concurrency probably will affect the hard scaling floor, and the trade-off between communication vs processing in codes like ESMs with low arithmetic intensity. I guess there are also trade-offs that arise because of the need for the concurrency to be in a shared memory implementation. The manuscript would be strengthened if these issues were discussed in a more thorough manner.

5. There are too many acronyms, and some seem indiscriminately chosen.

6. The manuscript does not distinguish between AGCMs and ESMs, frequently discussing the results in terms of ESMs but then presenting results for the atmospheric GCM alone. Do the results generalize to problems that already have much more concurrency? I guess so but this is a separate point and the manuscript should discriminate between what was done and what is inferred based on what was done.

Details:

p112: Acronyms

p1124: Is that true, or did the technology also drive things. it makes it sound like everything was possible and we just chose something, rather than necessity driving development.

p215: I am not sure what figure the authors are trying to say here. Actually resolution has not kept pace with computing as far as I can tell, and the reference to the figure does not make sense.

p217: Would help to explain to the reader the phrase “arithmetic and logic” is a memory fetch logic?

p2111: “The state of play of climate computing in the face of these challenges” this phrase comes across as a bit of a throw away. Did the Balaji (2015) paper make a point that is important for the present discussion? If so what was it.

p2120: “and there is constant churn of operations “ missing an article. . . lyrical, but it gives the idea that the computer is kept busy computing rather than moving information around. Most codes are memory bandwidth limited . . . which you get to shortly. But this intro sentence did not prepare me well.

p2130: Here: “Of the many factors of 10 increase in performance needed to get to the promised land of “exascale computing”, we believe at least one can come from component organization. “ I would prefer precision over poetry.

p315: Weather centers also run spatially coarse-grained radiation. Also some have proposed a form of coarse graining in the spectral domain, i.e., Monte-Carlo Spectral Integration.

p419: A reference is needed here. My intuition suggests that such high processor counts have only been applied to more traditional Atmosphere, or Ocean or Atmosphere Ocean Problems, but not simulations of the carbon cycle, i.e., the use of ESM here is misleading as earlier it implied biology.

p6l15: Subscript abbreviations are usually written in roman font.

Figure 3: I spent some time on this and I am not sure I understood it. The vertical dimension denotes sequence, first to last from top to bottom. The boxes indicated either the legend or a component process? The thickness of the box denotes? A figure should be illustrative, not a riddle. Also a bit more structure might help me understand what atmosphere up is. The meaning of some colors seems to be specific, others decorative?

p7 l2: 'interact strongly with atmospheric chemical species and clouds' this could exclude water vapor, why not say, couple strongly to composition.

p7 l4: By this definition the ocean does not have tracers. Aren't tracers really just scalar quantities that are transported with the flow subject to source and sink processes.

p7 l6: Seconds and hours (line 9) are units, and can be abbreviated.

p7 l9: This makes it sound worse than it may be, some codes rescale the radiative heating rates at each timestep by the insolation, or the surface temperature, in a sense linearizing about the state defined every 3 hrs. Something you mention on the next page, but it comes late.

p7 l19: Spell out PE ... Processing Elements? A socket?

p7 l14: 'Architected'? Okay, it can be used as a verb; but I think designed, or constructed would be better.

Figure 5: This needs redrafting, first for the colour scale (no rainbows); second to show the common color scales in those panels where it is appropriate. The highly distorted projection should either be motivated or replaced.

Figure 5: I would like also to see land temperatures. And differences among the simulations are far more interesting than differences with GPCP

P10l14: 'Remarkably' similar. ... my rule of thumb is that people use the word remark-

Printer-friendly version

Discussion paper



able when they don't know what of substance to remark.

p10130: I would make the tuning point later, as it seems as though the authors are interpreting the differences as fundamental, rather than simply an illustration of compensating biases in a manner that is to be anticipated.

p1115: SYPD has a unit, i.e., yr d^{-1}

Table 1: Maybe spell out the acronyms; for example does the introduction of CHSY really help anything? And if so why not PHSY, processor hours per simulated year.

The computational cost of radiation appears small. If increasing the frequency of radiation nine-fold increases the computational cost by 50% this implies that the cost of radiation is about 5% of the total computational cost in the default configuration. Is this correct? If so this is rather small compared to some other models, suggesting that the proposed approach might be even more beneficial for other centers, or offer the possibility of more exact representations of radiative transfer. Here some clear numbers would be useful.

p13113: This paragraph is a bit ungrounded in the manuscript, which does not evaluate MPMD approaches. Certainly the GPU rewrite of COSMO has a factor of 3.6 speed up on a first implementation ... so there is some room for efficiency gains through reprogramming, also the inexact hardware approaches (Dübben and Palmer) merit mention if this paragraph were to be retained and better grounded in the manuscript.

p13120: Where does the order ten components come from. I think "probably not more than ten" would be more accurate, but in either case if this comes at the end it should also be better ground in the manuscript.

p14122: By data do the authors mean model output? Or the performance data, gleaned from the benchmarks? The use of "data" suggests the latter, but the former should also be addressed.

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

[Printer-friendly version](#)

[Discussion paper](#)

