Response to `gmd-2016-114-SC1`:

1. The main paper must give the model name and version number (or other unique identifier) in the title. If the model development relates to a single model then the model name and the version number must be included in the title of the paper. If the main intention of an article is to make a general (i.e. model independent) statement about the usefulness of a new development, but the usefulness is shown with the help of one specific model, the model name and version number must be stated in the title. The title could have a form such as, "Title outlining amazing generic advance: a case study with Model XXX (version Y)".

   In response, we have modified the title as follows: "Coarse-grained component concurrency in Earth System modeling: Parallelizing atmospheric radiative transfer in the GFDL AM3 model using the Flexible Modeling System coupling framework".

Response to `gmd-2016-114-RC1`:

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.

2. I enjoyed the more lyrical style, but in places the style came across as glib, particularly when precision was sacrificed for poetry. Grounding the manuscript better in data and substance would strengthen the exposition greatly.

   We have taken some pains to sharpen the arguments and make them more rigorous at many places through the text, including the specific items highlighted by Reviewer 1 below. We hope we have retained some of the "lyricism" without giving in to "glibness".

3. 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.

   Please see reply to (25) below.

4. 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.

   We agree that that there need to be more studies of the effects of temporal subsampling of radiation. However, there is no case to be made that temporal subsampling is superior for *physical* reasons: at best, we can claim that we do not see marked increase in bias or RMS error as an effect of subsampling, and that it improves time to solution. The Morcrette papers we have cited in response to Reviewer 1's other comments (about spectral and spatial subsampling) also only claim expediency and decreased time to solution as a reason. They indicate error increases are tolerable over short timescales and probably increase model bias on climate timescales. Based on these results, we believe our focus on methods to possibly eliminate temporal subsampling are justified. We have modified our discussion of the Pauluis-Emanuel and Xu-Randall studies to indicate the need for further exploration of the effects of temporal subsampling, including solution convergence as Δt_{rad} and Δt_{atm} converge. See page 4, line 6.

5. 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.

While optimal performance for OpenMP threading is generally achieved when the threadpool is constrained to a single socket, scalability of concurrent components is limited by the number of cores in a shared-memory node. This number is slated to be very much larger than today, $\mathcal{O}(100)$ on MIC architectures. See also reply to (35) below.

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

   We have gone through the text and noted the following acronyms: GPU, MIC, CCC, ENIAC, ESM, IPCC, FAR, AR5, HPC, FMS, SIMD, SPMD, MPMD, SMP, MPI, GFDL, OpenMP/OMP, AOGCM, LBL, PE, AM3/CM3, SST, AMIP, GPCP, CERES EBAF, NOAA, SYPD, CHSY, NPES, CPU. Most of these are quite well-known and commonly used, and cover standard terminology in the computing and climate literature. Some refer to institutions, models, and observational datasets. We hope none of them look "indiscriminate" in the revised text. All have been defined at first use, thanks to the close reading by reviewers: see (8) below.

7. 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.

   We have clarified our use of "ESM" as a generic term for models of any level in the complexity hierarchy. We have also more clearly explained that the current implementation is in an atmosphere-only setting, but readily generalized to more components, which we describe in our plans for future work. This includes components invoked only in ESMs defined in the narrower sense, that of models that include an interactive biosphere. See page 2, line 26; page 16, line 2; and the caption to Figure 2.

8. p1l2: Acronyms

   Fixed, see page 1, line 3

9. p1l24: 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.

   It is mostly believed that the principal cause for the decline of vector computing was economics: commodity clusters were able to match performance of custom hardware at a vastly lower price. We have added some text and citations to bolster this argument, see page 2, line 9.

10. p2l5: 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.

    There was an error: it was Fig 1.4 from the 2007 Report, not the 2013 one we cited. It is now fixed, and a link to the figure itself is provided. See page 2, line 15.

11. p2l7: Would help to explain to the reader the phrase "arithmetic and logic" is a memory fetch logic?

Fixed, see page 2, line 17

12. p2l11: “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.

    Fixed, see page 2, line 24

13. p2l20: “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.

    Fixed, see page 3, line 5

14. p2l30: 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.

    Fixed, see page 3, line 18

15. p3l5: 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.

    This is an important oversight, and we have added a short discussion on spatial and spectral subsampling, see page 3, line 30. The main focus of this paper remains the elimination of temporal subsampling at modest computational expense.

16. p4l9: 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.

    The reviewer is correct: most of the very high end PE counts are for specialized problems, not full ESMs in climate mode. We have clarified this, and added some references: see page 5, line 9

17. p6l15: Subscript abbreviations are usually written in roman font.

    Fixed, here and everywhere else in the text: see page 7, line 15.

18. 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?

    Noted. We have reworded the caption to Figure 3 to clarify these points.

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

Text clarified: see page 9, line 6.

20. 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.

    Text clarified: see page 9, line 7.

21. p7 [reviewer typo: this is p8] l6: Seconds and hours (line 9) are units, and can be abbreviated.

    Done: see page 9, line 11.

22. p7 [reviewer typo: this is p8] 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.

    Agreed, eliminated last sentence of paragraph, as it's redundant with what come's next: see page 9, line 13.

23. p7 [reviewer typo: this is p8] l19: Spell out PE . . . Processing Elements? A socket?

    Fixed: see page 9, line 26.

24. p7 l14 [reviewer typo: this is p8 l24]: "Architected"? Okay, it can be used as a verb; but I think designed, or constructed would be better.

    Rewritten: see page 9, line 31.

25. 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.

    If the objection is to the use of the Mercator projection: we agree that it introduces distortions, particularly areal distortion as you approach the poles. While this is known, this projection is customary and in very widespread use in the literature. We do not believe the use in this instance to be egregious, as the results are in no way compromised by the projection.

    Similarly, we have used a somewhat standard color scheme: the rainbow in panel 1, and a warm/cool two-color palette for the difference plots. This again is quite customary.

    We are somewhat puzzled by this remark as neither the projection nor the color scheme are in any way misleading or distorting the results.

    That said, we have modified Fig. 5 and 6 to show the model-obs. differences in Panel (b), and the (smaller) model-model differences in Panels (c) and (d). We hope this makes the discussion clearer. See page 11, line 28, and captions to Figs. 5 and 6.

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

Again, we are puzzled by the remark and unsure how (or why) to add land temperatures to a global plot of precipitation. No land or ocean cells have been masked on this plot.

The differences between the simulations are interesting, but cannot be over-interpreted, as in practice we would retune the models to maintain the same top-of-atmosphere net radiative flux. The focus of the paper is on the computational aspects and the coupling algorithm. We expect to visit this issue in greater detail in a future paper describing scientific model results from concurrent-radiation run where $\Delta t_{\mathrm{rad}} = \Delta t_{\mathrm{atm}}$, with appropriate tuning. See also response to (28) below.

27. P10l14: “Remarkably” similar. . . my rule of thumb is that people use the word remarkable when they don’t know what of substance to remark.

    Rephrased, see page 12, line 2.

28. p10l30: I would make the tuning point later, as it seems as though the authors are interpreting the differences as fundamental, rather then simply an illustration of compensating biases in a manner that is to be anticipated.

    The tuning discussion comes at the tail-end of Section 4.1, which is the last discussion of physical results. We agree that the differences are not fundamental, and in the initial draft had indicated that they were “within the margins of the tuning process.” We have expanded on that to make it clearer, as suggested, see page 13, line 26.

29. p11l5: SYPD has a unit, i.e., yr d≏1

    SYPD is itself a unit, in fact it is introduced using the phrase “simulated years per day, or SYPD”, see page 13, line 30. We believe it would be redundant to attach the unit again. Units are mandatory when the same quantity can be expressed in different units, e.g m and cm. That is not so in this case.

    That said, we have added a discussion and definition of SYPD, see page 13, line 30.

30. 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.

    We do indeed believe the CHSY discussion is useful, as both time to solution (SYPD) *and* resource consumption (CHSY) are taken into account in configuring the parallel layout of a production model. See discussion in Balaji et al. (2016). We have also chosen to be consistent with that paper in naming the throughput measure CHSY instead of PHSY, which we agree would have also been a valid choice.

    In regard the second point, the reviewer is correct, and we have made this point now in the text, see page 14, line 7.

31. p13l13: 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.

    It's true that the background material is a bit ungrounded in *this* manuscript, but summarizes the findings of Balaji (2015), which does indeed survey MPMD methods and inexact computing approaches. We have rewritten the paragraph to reflect that these are findings from Balaji (2015). See page 15, line 18.

32. p13l20: 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.

    We have toned down the speculative statements here, see page 16, line 10.

33. p14l22: 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.

    Fixed, see page 17, line 3

Response to `gmd-2016-114-RC2`:

34. page 2, around line 20: the discussion focuses on performance aspects only. Mention of the implications for energy/power use would also be useful here, as the focus is on extreme-scale systems.

    Fixed, see page 2, line 19

35. page 3, line 12: “how it is achieved without increasing data movement” should, I think, be changed to “and how it is achieved with minimal impact on data movement”. This point is discussed further below [*].

    [*] The following point is my main concern with the paper as it stands. I believe the paper would benefit by being clearer on the use, and limitations, of shared memory threading to implement the concurrent execution of the radiation and the rest of the atmosphere model. It is clear that for a given multi-core processor there will be limits on the number of MPI processes per node and the number of threads each model may use within an MPI process without incurring potentially expensive data movement between caches. The example results given are based on two threads for each model. It would be good to make clear the rationale for this choice. Here are some thoughts on this and suggestions for possible changes which might help achieve this. Currently, the use of threads for executing the radiation in parallel with the rest of the atmosphere is described as not incurring any communication costs. While it is true there will be no MPI communication incurred between cores running the atmosphere and cores running the radiation, I believe there may well be some extra remote data accesses (i.e. cache misses) incurred between cores running the atmosphere and cores running radiation. The magnitude of this effect is, of course, architecture dependent and also depends on the number of threads used for each model and the mapping of the threads to cores. The results presented are for AMD Interlagos processors with two threads used for the atmosphere model and two threads used for the radiation model (within each MPI process). The Interlagos processor chip consists of eight 2-core modules. Two threads executing on the same module share an L2 cache. All 2 core modules share a large L3 cache. So, if one atmosphere and one radiation thread share a module, they can share data in the L2 cache (as well as the L3 cache). If two Atmosphere threads share a module and two radiation threads share a module, they will communicate through the L3 cache, which is more expensive in terms of cycles to access. If threads are on separate processor chips, there will be (even more expensive) data movement within the shared memory node.

    The cache behaviour in either of the above cases is likely to be different to that of a single thread running first the atmosphere and then the radiation. If more than two threads were used for each model, some sharing would have to take place via the L3 cache. Total thread numbers are clearly limited by the core count of a shared memory node. I would suggest to the authors that some clarification of these issues be made. For example: - in Figures 3 and 4, the images depicting MPI and OpenMP could be re-drawn to illustrate the relationship of threads within MPI in each case. In Figure 3, this would simply show multiple threads in an MPI task and multiple MPI tasks. In Figure 4, MPI tasks could be shown with both atmosphere and radiation threads or with ocean threads. For Figure 4, this might be something like: [AARR] [AARR]... [O] [O]... (where [] here represent MPI processes and letters represent threads and their models) - In Section 3 (perhaps?), a brief description of a multi-core processor (like the Interlagos) could be given along with the implications of thread-to-core mapping. This description would help to explain the benefit of using OpenMP to exploit parallelism between the atmosphere and radiation models (i.e. no MPI communication) and pave the way for a discussion of the potential for sharing data between caches in the specific configurations presented in the results section.

The concurrency software design was formulated to ensure the radiation component was truly independent from the remaining atmosphere components with a single data synchronization point (copy) at the end of each time step. The data copying, itself implemented with OpenMP threading, can be completely removed via index flipping, but this might result in non-local accesses or extra cache invalidations - so a simple data copy is preferable. Additionally, each thread immediately calls a subroutine (either atmosphere or radiation) and is working with either explicitly blocked data structures or pseudo-blocked data created via copy-in variables. So not only are the atmospheric and radiation components completely independent, but within each component the data for each thread is isolated as well. If a radical paradigm shift in computing occurs, this forethought in the design allows us to quickly re-cast the radiation threads as separate MPI processes with a needed MPI communication synchronization point replacing the current data copy. See also (5) above.

36. page 5, Figure 1: The role of A_t should be depicted in the top figure (to be consistent with equations 1 and 2), I feel. Also, in many models, the atmosphere is executed on more processors than the ocean (because it scales better). Is this diagram consistent in this respect with the FMS model being described? Also, the bottom figure in Figure 1 implies that the ocean is executed on fewer processors in the concurrent set up. Is the intention to simply show a deployment utilising the same number of processors in total? If so, that should be made clear in the caption and text.

The relative balance of PE count between atmosphere and ocean depends on many factors, and within the FMS system, we have examples of both. However, in Figure 1 this is mostly schematic, indicating that some PEs may idle in serial processing.

We have adjusted the numbering of the component timesteps to be consistent with Eqs. 1-4.

37. page 9, line 2: “chosen to offer optimal load balance” could be extended to “chosen to offer optimal load balance and data sharing”, for example. [I generally have a concern over the use of the work “optimal”, which has a formal sense of “provably best”. The word “good” might be better unless the load balance is provably optimal?]

Fixed, see page 10, line 1

38. page 11, line 9: The above arguments are tied up with the statement that “All runs use the optimal processor/thread layout for a given PE count”. Some explanation about what this layout is and how it was chosen could be added.

Done, see page 13, line 33

39. page 1, line 4: I suggest changing “based on marginal increases in clock speed” to, for example, “based on, at best, marginal increases in clock speed” since it is likely clock speeds may decrease in future in some systems.

Done, see page 1, line 4

40. page 1, line 14: Define the acronym CCC here.

    Done, see page 1, line 16

41. page 1, line 15: is a little ambiguous about what is running in parallel (“and all other atmospheric physics components”. I would suggest making it clear that there are only two concurrent components (i.e. not all “other atmospheric components” are executed in parallel with each other!

    Done, see page 1, line 18

42. page 2, line 3: perhaps provide a reference to the IPPC assessments.

    See page 2, line 15.

43. Section 2: page 4, line 4: needs a closing bracket after “example”.

    Done, see page 4, line 33

44. page 8, line 24: “Individual” should be “individual”.

    Done, see page 9, line 32

45. page 9, Figure 4: In this figure, the Land and Ice models are shown as executing concurrently but this is not mentioned in the text. This should be explained (or made consistent with Figure 3).

    Fixed. It should indeed have been consistent with Fig. 3. See new Fig. 4.

46. page 10, line 9: “that that” should be “that”.

    Done, see page 11, line 21

47. page 10, line 11: This sentence would benefit from having a reference added.

    Done, see page 11, line 23

48. page 10, line 13: Expand the acronym GPCP here as a definition.

    Done, see page 11, line 26

49. page 10, lines 15-17: The point here is, I think, that this result is counter intuitive. If that is correct, it would be worth stating.

    As noted above in (25) and (28), the differences between the runs are small and within the bounds of the tuning process. Differences should not be over-interpreted as they are likely to vanish upon tuning. See page 13, line 26 also.

50. page 11, line 3: “less expensive as...” should be “less expensive as the...”.

    Done, see page 13, line 25

51. page 12, line 4: the figures for processor count and SYPD given in this line are rounded versions of those in Table 1. Those in the previous sentence are not rounded. Please use the precise figures for consistency.

    Done, see page 15, line 4

52. page 13, line 18: It would be worth giving the definition of CCC again here to remind the reader.

    Done, see page 16, line 2

Response to Topical Editor's comments:

53. My main concern relates to the issue #35 (referee 2) which is also linked to issue #5 (referee 1). In fact, I do not really understand your replies to these two important issues and I understand that you did not make any modifications in the document to address them. Thank you in advance for a revised version of the document taking into account these issues.

    Section 3 has been largely rewritten in a manner which I hope will address your concerns, as well as those raised in #5 and #35.

54. #6: IPCC, AR4, GFDL, CM3, NOAA, PE, GFDL, CPU are still not defined

    Fixed, see page 1, line 16; page 2, line 13; page 2, line 15; page 11, line 5; page 1, line 1; page 9, line 23; page 16, line 22.

55. #25 - #26. :

    (a) I think the reviewer was not asking to add land temperatures to Fig 5 and 6 but just add an extra plot similar to Fig.5 and 6 for land temperatures. Is that possible?

        With regards to the request to show the response of surface temperature to the radiative time step and the serial vs. concurrent coupling, we prefer not to do this since these integrations prescribe the ocean surface temperature, which prevents land temperatures from changing much as well. The changes in top of atmosphere radiative fluxes in an AMIP model that we are showing are more relevant to the temperature changes in a coupled model than are the temperature changes in an AMIP (prescribes sea surface temperature) model themselves.

    (b) Fig. 6 captions do not read OK to me now. Please add "(a)." after "CERES EBAF v2.8 climatological top-of-atmosphere radiation budget" and change ", and (b) Model ..." for "Panel (b) shows model ..."

        Fixed, see new caption to Figure 6

56. #23, p.9, l.23: I am not sure I understand your "PE/PElist" concept, as you are mixing hardware (core) and software (communicator) concepts. If a PElist is equivalent to the communicator concept in MPI, you refer here to the software concept that I would call a list of MPI tasks. But then in Table 1 legend, you write "NPES is the total PE count (MPI*OMP)" so it looks like there is no equivalence between a PE and an MPI task. If two processes share one core (multitasking), would you have two PEs (two processes) or only one PE (one core) in your PElist? Can you clarify?

    You are correct: a PE is a hardware unit, and a PElist is an MPI communicator. We have clarified the text, see page 9, line 26.

57. #36 I think the new numbering of the component timesteps in the top panel of Fig. 1 is wrong and the previous was right. Eq. (1) states that $A^{t+1} = f(A^t, O^t)$, and $O^{t+1} = f(O^t, A^{t+1})$. The top panel of Fig. 1 would imply $A^{t+1} = f(A^t, O^{t+1})$ and $O^{t+1} = f(A^t, O^t)$.

    I think I have got it right this time, see new Figure 1!

58. #49.: I understand your argument about differences probably vanishing with tuning, but as referee 2, I still think that the results that are counter intuitive and this should be stated. Also the sentence "Although this difference is small we feel it is robust and likely related to a sensitivity to the radiative time step discussed below." seems contradictory to me. If they are robust, they would not vanish with tuning?

    We have rewritten the discussion on tuning, which we hope is clearer, see page 11, line 25; page 12, line 2; page 13, line 2; page 13, line 21. The main point is that the magnitude is within the bounds of tuning, so we perhaps not pay too much attention to features within those bounds.

59. Fig.3 and associated text: even if the captions help understand the FMS architecture, I still agree with the referees that some aspects are misleading.

    (a) First, I would remove the green openMP box and the dark blue MPI box on the figure (as the meaning of the colours is now clearly explained in the text. Second, for coherency, I would replace "Atmos Down" by "atmosphere-down" and "Atmos Up" by "atmosphere-up" in the figure, or put "(Atmos Down)" after "atmosphere-down" and "(Atmos Up)" after "atmosphere-up" when they appear in the text.

        Fixed: see page 8, line 5 and new Figures 3 and 4.

    (b) In the text, you write "The implicit coupling requires . . . (land and ocean) ..." but I understand that FMS implements the implicit coupling only over the land and not over the ocean. Can you clarify this?

        Actually, it is implicit also over the ocean surface, which includes sea ice. The explicit coupling is with the 3D ocean, see Balaji et al. (2006) for details.

60. p.2, l.13 please change "... assessments, and their complexity (the number of feedbacks and phenomena simulated), exhibits ..." for "... assessments and their complexity (the number of feedbacks and phenomena simulated) exhibit ..."

    Fixed, see page 2, line 13

61. p.2, l.14 please change "Figure 1.2[1] Figure 1.4[2]" for "Figure 1.2[1] and Figure 1.4[2]"

    Fixed, see page 2, line 15

62. p.3, l.2: please change "operatorss" for "operators"

    Fixed, see page 3, line 4

63. p.3, l.2: why do you qualify "operators and operands" as "new"?

    They are "new" because new instructions keep arriving in the processor during each context switch. I have changed it to "fresh", which I hope is clearer: see page 3, line 4

64. p.3, l.3: please change "... and locality and reuse hard to achieve." for "... and locality and reuse are hard to achieve."

    Fixed, see page 3, line 5.

65. p.3, l.23: please move “(temporal subsampling)” right after “coarser timestep than the rest of the atmosphere”

Fixed, see page 3, line 25

66. p.3, l.25: I do not understand the “as well as the temporal”. Do you mean “as well as in the temporal domain”?

Fixed, see page 3, line 27

67. p.4, l.7: please change “... AMIP (The Atmospheric ...” for “... AMIP (the Atmospheric ...”

Fixed, see page 4, line 10

68. In Table 1 legend, please replace “MPI/OMP” by “MPI*OMP”

Fixed, see legend to Table 1.