

1 General remarks

First of all we wanted to thank you for taking your time to go through the manuscript in detail. Your contribution is very much appreciated. Answers to the specific comments are given at the second part of this response letter.

- *This manuscript describes progress to couple an explicit solver for three-dimensional radiative transfer with a large-eddy simulation LES hydrodynamic code. Two numerical choices are explored (the iterative solver and the preconditioner) with emphasis on both strong and weak scaling efficiency. Solving the three-dimensional radiative transfer problem (rather than one-dimensional problem in which every model column is treated independently) is incompatible with one of the algorithmic choices underlying the current treatment of radiation in the LES code, namely the Monte Carlo Spectral Integration algorithm in which the spectral interval for each column is chosen randomly. The authors perform simulations to assess whether a weaker version of the MCSI (where spectral points are held constant across the domain but are chosen randomly in time) is still a viable approach to coupling radiation to LES. This is technical work that will enable some potentially very interesting research on whether three-dimensional radiative transfer effects systematically affect large-eddy simulations.*

Great, right to the point.

- *It's not clear how useful it is to report this work in isolation. The particular performance results for preconditioners and matrix solvers are specific to the problems (including domain size and the amount of cloudiness) and to the computer systems used for testing, while the results on the weak form of MCSI will no doubt need to be revisited for new problems. Personally, I'd advise students of my own to include this material in the subsequent papers describing results, and if GMD has explicit editorial standards for novelty and relevance the editors may want to look closely at whether this manuscript is adequate.*

We agree that this work focuses on more technical aspects of the model but we understood that this is well within the scope of GMD, “development and technical papers, describing developments such as new parameterizations or technical aspects of running models”. We also

think that it is important to show that a parameterization is actually running not only in theory, but also in reality, and in particular that the code scales well in a real-world application on a multi-processor machine. Also, the weak form of the MCSI is a new application and we think that it is valuable to show that it works, even if it hasn't been tested for all possible resolution and cloud scenarios. We showed the parallel scaling behavior on three distinctly different architectures and periphery setups. As for the applicability to a range of different model states, the point of the strong scaling experiments is to examine and understand how a particular solver/preconditioner depends on the scene. The fact that the performance of an iterative solver does depend on the complexity of the simulation makes rigorous testing and documenting ever more interesting and important. We feel confident that the presented edge cases put sensible boundaries on the increase in runtime that is to be expected for atmospheric simulations.

- *With that caveat, the manuscript is generally successful at what it attempts to do. It would be improved most by a little pruning and reorganization aimed at more cleanly separating the two classes of issues (weak MCSI vs. algorithmic choices) and providing a more consistent level of detail to support the argument, keeping in mind that readers will come from both the LES and radiative transfer communities. Some general guidance is provided below. The authors might consider a modified hierarchy for the manuscript to reflect the different concepts being explored. To this reader the top-level ideas/headings might be: Introduction LES and Ten-stream models Weak MCSI Numerical scaling Conclusions The introduction might be similarly reorganized to reflect the separate concepts.*

We agree concerning the structure of the manuscript and made a few modifications accordingly. In particular, the Monte-Carlo-Spectral-Integration part was given an own section, as suggested.

- *The discussion of the broad motivation for the work that radiation influences cloud development, and that three-dimensional radiative transfer is normally neglected could be expanded by three or five sentences so readers understand why the problem is relevant.*

We added two sentences to elaborate on first order 3D effects.

- *It will likely to be easier to discuss three-dimensional issues, including the need for efficient algorithms (ten-stream) and implementations (numerical issues to be explored here) before MCSI issues because the motivation for examining weak MCSI comes from wanting to use three-dimensional RT. The general motivation for MCSI (that is, most of*

the discussion on 9023) should be deferred to the section describing the tests of weak MCSI.

We agree that the introduction as to why we care about MCSI is rather broad but we feel that the wide introduction helps readers not accustomed to radiative transfer.

- *Section 2.3, and again on 9033: it should be more clear if the experiments with weak MCSI use one- or three-dimensional radiative transfer calculations. On a related note it would be useful to explain why one set of experiments is used to test weak MCSI and another set used to assess performance of the ten-stream solver. Pincus and Stevens 2009 included an experiment in which the mean radiative driving was suppressed and only the noise remained. It would be useful to repeat these experiments with weak MCSI.*

Yes, we added a sentence stating that the MCSI simulations were done using 1D solvers. The DYCOMS simulation were used in analogy to the work of Pincus and Stevens (2009), whereas the scaling experiment setups are deliberately kept as simple as possible to investigate how the complexity of cloud dynamics influences the radiative transfer solver computationally (see also next point). To put the impact of the increased noise of the uniform MCSI into perspective, we added a more elaborate discussion on the MCSI (see review response letter #1 and MCSI section).

- *What is the point of the clear-sky experiment? One would think that three-dimensional radiative transfer would be irrelevant in the absence of significant scattering, so its not clear what is being tested or learned with these experiments.*

The strongly forced warm bubble and the clear-sky experiment are the limits how the cloud field may change between calls to the radiation routines. In a real application, both cases occur and with these two setups we test the assumption that reusing an earlier solution may help convergence. We added a second sentence at the first paragraph of strong scaling section to highlight that.

2 Specific comments:

- 9022, line-24: *Surely the idea of radiation coupling to cloud dynamics predates Muller and Bony 2015.*

Of course, the fact that heating drives convective motion is basic physics. We wanted to highlight relevant work on cloud radiative interaction. We added a sentence putting the work in context.

- 9024, line 9–10: *formatting of references is incorrect*

Corrected.

- 9024, line 13: *It would be kind to add one sentence explaining how the ten-stream solver works for those not familiar.*

Added a short description to clarify the key concepts and steps of the TenStream solver.

- 9026, line 21: *This is an abrupt transition. It also sounds a bit like advertising.*

Rewrote the iterative solver paragraph with a more general introduction.

- 9029, lines 16–26: *The explanation of weak and strong scaling is valuable but could be 50 percent shorter.*

Thanks, we also feel that a thorough introduction is necessary for the parallelization novice. We trimmed two sentences at the beginning and the end.

- 9031, line 3: *“Retrieving the transport coefficients from the look-up table” ... what transport coefficients? what lookup table? Readers who don’t know the ten-stream model well are here left behind.*

Added a brief explanation what the transport coefficients are and why we need them.

- 9032, line 6: *What is the Mistral computer?*

Added a reference to the table that lists the respective computing machines.

- 9032, line 15: *Pure speculation about the causes for reduced efficiency is not particularly helpful.*

In order to really know what causes sub-optimal performance on a machine it is imperative to have a rigorous FLOP and memory model. This however is a huge undertaking for complex algorithms. While we can not separate the individual reasons for inefficiencies we feel that it is nevertheless helpful for the reader to know what the possible mechanisms are.

Many thanks,

Fabian Jakub

References

Pincus, R. and Stevens, B.: Monte Carlo spectral integration: A consistent approximation for radiative transfer in large eddy simulations, *Journal of Advances in Modeling Earth Systems*, 1, doi:10.3894/JAMES.2009.1.1, 2009.