

Interactive comment on “An adaptive method for speeding up the numerical integration of chemical mechanisms in atmospheric chemistry models: application to GEOS-Chem version 12.0.0” by Lu Shen et al.

Anonymous Referee #2

Received and published: 2 December 2019

Shen et al describes the implementation of a method of reducing the computational complexity of solving a chemical mechanism within GEOS-chem. The paper is interesting, although further revisions are required before it can be considered for publication.

Major Comments

I found the discussion in Section 3.2 very hard to follow, specifically how blocks are grouped into regimes and then the subsequent changing of blocks from slow to fast if

[Printer-friendly version](#)

[Discussion paper](#)



a gridbox does not correspond to any of the regimes. The sentence in question is

"Gridboxes that do not correspond to any of the M regimes need to be matched to one of the M regimes by moving some blocks from slow to fast, which will change the values of the corresponding indicators $y_{i,j}$ from 0 to 1."

Could the authors explain just how the mapping of species to blocks to regimes to these re-matched regimes is done? A diagram or pseudocode would be useful here. This crucial step is not explained well, and I'm not sure if this step is done online or not. How is the regime determined during a model run, and how is it ensured that the regimes are correctly matched (and what happens when they do not match)? This information is required to adequately understand the method presented.

I would be interested to know how robust the particular organisation determined from the Simulated Annealing algorithm is. Were multiple simulated annealing simulations performed? Was the rate of reduction of "temperature" changed to see if this affected the results? As with any global optimisation technique it is possible to get stuck in local minima, and a single run-through this algorithm will not be sufficient to determine whether the true minima has been found.

While not essential for this manuscript, I would be interested to know if this classification has any load-balancing implications. I can imagine that for codes with MPI parallelisation across many nodes of a HPC, this method will increase the imbalance between different MPI tasks (while still decreasing the overall run-time). This could then lead to further speed improvements if the load-balancing is improved.

Most of the discussion and plots presented use a δ of 100 molecules $\text{cm}^{-3} \text{s}^{-1}$ (or a range is presented), except when $\delta = 500$ is used for Figures 4 (the equivalent plot for $\delta = 100$ is Figure S5) and 5 and the discussion surrounding the 2-year runs in Section 4. Given that the $\delta = 100$ results seem noticeably better, why were the $\delta = 500$ presented in the main text? Are there equivalent plots (especially the Figure 5 equivalents) for the other values of δ used (100,1000)?

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

Given the errors associated with halogen species presented in Figure S4, would there be a large drop in performance if these species were always treated as fast?

Minor Corrections

Page 6, Line 165: ",fast blocks,"

Page 7, Equation 7: There is no D_1 , both Σ are labelled with D_2

Page 16, Figure 1: The X-axes for the panels are slightly off-set. This can be clearly seen in the downward dotted lines.

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2019-279>, 2019.

Printer-friendly version

Discussion paper

