Re-review of Shen et al: An adaptive method for speeding up the numerical integration of chemical mechanisms in atmospheric chemistry models: application to GEOS-Chem version 12.0.0

Many thanks to the authors for addressing mine and the other Referee's points. The changes have helped clarify matters greatly. I have a few further questions/suggestions before I can recommend publication.

Major comments

On the Simulated Annealing Algorithm used, I am concerned by the statement:

"In addition, there are still noticeable changes of species groups if we run the simulated annealing algorithm with different initializations and choices of the temperature parameter, even though the optimized blocks can generally separate the oxidants, anthropogenic VOCs, and biogenic VOCs (Table S1)."

To me, this means that the energy landscape is full of local minima and/or is possibly degenerate. Re-running the algorithm should give the same global minimum multiple times, if the algorithm is robust. Table S1 only lists two other potential groupings – why show only these two in the supplementary information? How many times was this algorithm run for each value of N? Given that when calculating N, the value of δ is fixed at 100 molecules cm^{-3} s^{-1}, but later simulations change δ to higher values, how robust is the categorisation to the value of δ?

It may be that the groupings don't actually matter that much – if there are lots of local minima then each could give similar performance. However, this should ideally be tested if this is the case to see if the errors remain the same. Simulated Annealing is not a great global optimisation technique really, and others such as Basin Hopping or Genetic Algorithms have shown better performance for problems with a large number of potential solutions.

Diagram S1 should be placed in the main text. If I understand this method correctly, you use a training dataset from 4 GEOS-Chem simulations that have been run for 10-days and use output from every 6 hours. Using this you have categorised the species into 12 different distinct blocks (using simulated annealing), which are then combined together into 20 different regimes, and you have assigned each gridbox a regime. Please clarify further if this is not correct.

Does the regime that a gridbox has change in time during the simulation, and if so how? In your response you state that this is calculated offline, so how does time of day or season affect things? If the emissions were changed, would everything need to be re-calculated again? Similarly, if you are wanting to run a pre-industrial or future scenario, what would need to be changed or re-calculated? Figure 3 and S2-S5 show that things are changing in time, but I am not clear how this is determined given your response that this is calculated offline.

Minor Comments

Line 184: Here you state that “Among the N blocks, 3 are allocated to the reactive inorganic halogen species, and N-3 are allocated to the other species.”, and on line 203 you state that “We tested a range of values from 3 to 20 for the number N of blocks”. Does this mean that you just shuffled the halogen species between these three along with everything else, or were you testing 6 to 23 blocks?

Line 190: “~~selected representative~~” and “~~and a full listing is in Fig. S2~~”

Line 254: you use “10^2” and “10^3” here, but 100 and 1000 elsewhere.

The colourbars on figures S2-S5 are completely redundant.