

Response to referee comments 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”

We thank the referees for their careful reading of the manuscript and the valuable comments. This document is organized as follows: the Referee’s comments are in *italic*, our responses are in plain text, and all the revisions in the manuscript are shown in blue. **Boldface blue text** denotes text written in direct response to the Referee’s comments. The line numbers in this document refer to the updated manuscript.

Reviewer 2

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.

Response. Thanks for raising these good points. This feedback has significantly improved the manuscript.

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 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.

Response. Thanks. This process is done offline. The 20th chemical regime is the full chemical mechanism, so every gridbox can be matched with a regime after some moves. We have included a diagram in the supplement and mentioned this in text.

Line 194. We check each of the M regimes and select the one that needs least number of moves from slow to fast, and this selection can be pre-defined so it does not add extra computational time. The 20th chemical regime is the full mechanism, so every gridbox can be matched by the M regimes.

Line 199. A diagram for this process can be found in Figure S1.




Figure S1. The diagram for calculating the cost function Z_2 . More details can be found in text.

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.

Response. Thanks for raising this good point. In this study, we have run the optimization multiple times and also tried different temperature parameters. We present the one with lowest cost function. Now we make this clear in text.

Line 186. Throughout this study, we present the results with lowest cost function after running the optimization multiple times and using different temperature parameters.

The set of species blocks are not always exactly the same due to the random processes in simulated annealing. But in general, they can separate the oxidants, anthropogenic VOCs, and biogenic VOCs. Now we say.

Line 216. This algorithm still has shortcomings. There are some unexpected groupings (such as sulfur species and peroxyacetylnitrate) and separations (such as HO_2 and H_2O_2). The blocks are constructed by minimizing the number of fast species in the optimization, so species tend to be in the same block as long as they are fast or slow simultaneously. For example, isoprene products and CFCs are both slow in the stratosphere and clean regions, so they may be assigned into the same group (e.g., block 6). **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). Here we**

chose the set of groupings that minimized the cost function for a number of realizations of the algorithm. These two shortcomings may be addressed by introducing regularization terms in the cost function to enforce known species relationships, which will implement this in a follow-up study.

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.

Response. Now we have a paragraph to discuss this problem.

Line 309. The performance tests presented here were for a single-node implementation of GEOS-Chem using 12 CPUs in a shared-memory Open Message Passing (Open-MP) parallel environment. High-performance GEOS-Chem (GCHP) simulations can also be conducted in massively parallel environments with Message Passing Interface (MPI) communication between nodes and domain decomposition across nodes by groups of columns (Eastham et al., 2018). In principle, the chemical operator scales perfectly across nodes because it does not need to exchange information between columns (Long et al., 2015). However, differences in computational costs between columns (due to differences in chemical regimes) could result in load imbalance between nodes, degrading performance. In the current implementation of GCHP, the MPI domain decomposition is by clustered geographical columns in order to minimize exchange of information across nodes in the advection operator (Eastham et al., 2018). Such a decomposition would penalize our approach since different geographical domains may have different computational loads for chemistry (e.g., oceanic vs. continental regions). This could be corrected by using different MPI domain decompositions for different model operators, and tailoring the domain decomposition for the chemical operator to balance the number of fast species across nodes. Such an approach is used for example in the NCAR Community Earth System Model (CESM) where different domain decompositions are done for advection (clustered geographical regions) and for radiation (number of daytime columns).

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)?

Response. Now we show the results for rate thresholds δ of 100, 500 and 1000 molecules $\text{cm}^{-3} \text{s}^{-1}$. The user can decide which to use based on their needs. Figure 5 is also updated.

Line 262. The best range for δ is between 100 and 1000 molecules $\text{cm}^{-3} \text{s}^{-1}$, where the median RRMS error is below 1% and the improvement in computational performance is in the 30-40% range.

Line 271. Figure 5 shows the time evolution over two years of simulation of the median RRMS error for all species and also for the selected species OH, ozone, sulfate, and NO₂. The median RRMS for all species is 0.2%, 0.5%, and 0.8% for rate thresholds δ of 100, 500, and 1000 molecules $\text{cm}^{-3} \text{s}^{-1}$ respectively. There is no error growth over time. Among the four representative species, the RRMS is highest for NO₂, ranging from 1.0% to 2.0% for δ ranging from 10² to 10³ molecules $\text{cm}^{-3} \text{s}^{-1}$. For OH, ozone and sulfate, the RRMSs are below 0.3% in all cases. Figure 6 displays the spatial distribution of the relative error on the last day of the 2-year simulation, using a rate threshold δ of 500 molecules $\text{cm}^{-3} \text{s}^{-1}$ as an example. The relative errors are below 0.5% everywhere for O₃, OH, and sulfate. The error for NO₂ reaches 1-10% at high latitudes, but this is still well within other systematic sources of errors in estimating NO₂

concentrations (Silvern et al., 2018). Results for rate thresholds δ of 100 and 1000 molecules $\text{cm}^{-3} \text{s}^{-1}$ can be found in Figure S8-9.




Figure 5. Accuracy of the adaptive reduced chemistry mechanism algorithm over a two-year GEOS-Chem simulation (see text). The accuracy is measured by the 24-hour mean RRMS error on the end day of each month relative to a simulation including the full chemical mechanism. Rate thresholds δ of (a) 100, (b) 500 and (c) 1000 molecules $\text{cm}^{-3} \text{s}^{-1}$ are used to partition the fast and slow species in the reduced mechanism. Results are shown for the median RRMS across all 228 species of the full mechanism and more specifically for ozone, OH, NO₂, and sulfate.

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?

Response. The test shows this will bring 4% more computation cost. Now we say

Line 155. This increases the computation cost of chemical integration by only 4% relative to letting the algorithm set them as either fast or slow.

Minor Corrections

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

Response. Fixed, thanks.

Page 7, Equation 7: There is no D1, both Σ are labelled with D2

Response. Fixed, thanks.

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

Response. Fixed, thanks.