

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 1

The integration of the chemical differential equations in a chemistry transport model is a significant computational burden. Both climate and air quality forecasting models spend a significant fraction of their computation marching forwards through time, solving these differential equations. Developments in speeding up this code have been not been forthcoming over the last decades and it has only really been the availability of more, faster CPU cores that has allowed us to run increasingly complex chemistry. Algorithmic developments to speed up the integration of this chemistry are to be welcomed.

This paper outlines a method for running different subsets of the differential equations in different geographical regions, so speeding up the solutions. On the one hand, this is a fairly obvious thing to do. We don't need to be integrating all of the chemistry of isoprene over the middle of the Pacific in the same way as we do over the middle of the Amazon. However, attempting this adaptive chemistry has not been readily taken up by the community. This paper attempts a first realistic attempt at such an implementation.

In general, the paper describes a novel new technique with a potentially sound methodology which could be extremely useful to the atmospheric chemistry transport modelling community. At this point though, I am not convinced that the algorithm is working as intended, or if it is, the explanations given in the text are satisfactory for the reader to understand what is going on. If either of these outcomes can be corrected the paper should be published.

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

1) My major concern about the methodology is the split of species into the different blocks. The methodology for doing this is explained but there is little interpretation of the results. In many cases, the blocks seem to have lumped together some fairly random sets of species and this appears to have been hidden away in the SI.

Response. Thanks, we have added more interpretation of the results and one independent paragraph to discuss the shortcoming of our present method. The blocks of species are constructed by minimizing the number of fast species, so it cannot guarantee the groups to be always chemically logical. We have a follow-up project to fix this issue by introducing a regularization term that defines the species' distances as learned from their reactant-product relationships. The preliminary result of the revised method is more chemically logical, but it still has some unresolved issues so we are not able to present it in this work. In this manuscript, we will make it clear that our present method can only generally group species with coherent chemical behaviors but there are unexpected groups.

Block 1: "Aromatics" (Benzene and Toluene but the chemically almost identical XYLE is in Block 7?) but it also contains CH₂I₂. What is the advantage of using the integrator for these together? Does this mean that the aromatic chemistry is also being solved with the integrator in the middle of the oceans where the CH₂I₂ is important? Block 2: "Organic nitrates" This contains some organic nitrates and the N atom. These organic compounds are only really at appreciable concentrations at the surface whereas the N atoms are only really applicable in the upper stratosphere? Block 3: "Isoprene" This seems to contain some isoprene chemistry but also HFCS which seems suprising. Block 6: "Halocarbons" seems to be again mainly isoprene species to me, but it also contains HFC and CFC species which would only be important in the stratosphere.

Is there an explanation for this? Are the blocks in the SI correct? I might be missing an important concept here. But it has not been explained. To me, this feels like an error has occurred somewhere either in the species list given in the SI or in the algorithm.

Response. Our present method cannot perfectly separate these species by their chemical properties. We now discuss this shortcoming in one independent paragraph. Again, we will address this issue in a follow-up paper.

Line 216. This algorithm still has shortcomings. There are some unexpected groupings (such as sulfur species and peroxyacetylnitrate) and separations (such as HO₂ and H₂O₂). 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.

Given the random nature of the annealing algorithm, is this set of species blocks always the same one? What degree of variation is present when running this algorithm multiple times?

Response. Thanks for raising this good point. 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 220. 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.

There needs to be more work done to explain why these blocks are the best ones to use given the variation seen in the species type in each block. I realize this is a result of the optimization algorithm but the situation at the moment appears to be that an algorithm has told us that this is the result and we are going to leave it at that. To my mind, the species within these blocks do not appear to have the properties you would expect given what is trying to be achieved. I'm happy to be convinced otherwise but the text at the moment does not achieve that and it is not possible therefore to be confident that the algorithms are working appropriately.

Response. Thanks for pointing this out. This is a shortcoming of our present method, and we will fix it in the future study.

Line 223. 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.

2) The Supplementary information figures should be contained within the main text. The paper is fairly short and some of the figures are central to understanding the methodology. Figure S2 could be removed by putting the species list into Table 1. The other figures are small enough to be included in the main body without overwhelming the reader

Response. Thanks. We have moved two figures back to the main text. The information of Figure S2 has been included into Table 1. We also have 7 new figures (Figure S1-5, S8-9) in the supplement to better support our discussion.

3) I found the structure of the beginning of the paper a bit confused. We have an Introduction; a section on the chemical operator which includes a very brief description of the GEOS-Chem model and the KPP system. We then have a section on the algorithm being described. The material about the chemical operator should be moved into the introduction as this basically supports the introductory text about the chemical integration. The model description should go into a separate section.

Response. Thanks for pointing this out. Since the introduction is already very long, so we decide to remove the first two paragraphs about the chemical operator. And we have also changed the section name to 'Model description'.

4) Page 1 Line 30. The number of reactions thought in play in atmospheric chemistry is significantly more than the "hundreds" described in the text. The MCM has 10s of thousands and mechanisms produced by GECKO-A produces millions. Hundreds are used in the simplified mechanisms for atmospheric chemistry transport models. The text should be clarified here.

Response. Now we say this.

Line 30. The complete Master Chemistry Mechanism (MCM, version 3.3, <http://mcm.leeds.ac.uk/MCMv3.3.1/>) consists of 5,832 species and 16,701 reactions. Atmospheric chemistry models use greatly simplified mechanisms, which still include hundreds of species coupled through production and loss pathways and with lifetimes ranging from less than a second to many years.

5) Page 2 lines 55. There are now some other approaches to speed up the chemical integration using "Machine Learning" approaches they could be cited here.

Response. Now we say.

Line 62. Machine learning algorithms have been developed to replace the role of the conventional chemical solver; but these methods have only been applied to simple scenarios and are subject to error growth as simulation time progresses (Keller and Evans, 2019).

6) Page 6 line 155 the reference to Santillana is 210 rather than 2010.

Response. Fixed, thanks.

7) Page 6 line 166. Do the authors mean 'fast blocks' rather than ,fast blocks,?

Response. It is a typo. Now fixed, thanks.

8) Figure 1. When asking the reader to "see text" can this be more specific? What does the shaded area represent? The SD between what, the monthly values?

Response. Now we say

Line 446. See Equation 5 and related text.

Line 447. For both panels, results are for the first 10 days of February, May, August, and November sampled every 6 hours (shaded area denotes one standard deviation of results sampled every 6 hours).

9) Figure 2. Although using this approach does provide some information it would be useful to split the dataset in another way. Could there be a figure which shows a map of the world indicating whether each block is switched on at that location. This need only be done at the surface for 0 GMT and 12 GMT but it would give some confidence that the approach is working. isoprene block should only be on over continental regions etc. It is very hard to get this level of information from the figures as presented. Why is the value of delta of 100 used in this figure and 500 used in other figures?

Response. We have four supplementary figures to display the regions where anthropogenic and biogenic VOCs are treated as fast.

Line 213. Anthropogenic VOC species (blocks 4 and 5) are found to be fast in boundary layer and daytime mid-troposphere (Figure S2-3). Biogenic VOC species have shorter lifetimes, so they are found to be fast only in lower and middle troposphere over the land (Figure S4-5).

Now we show the results for rate thresholds δ of 100, 500 and 1000 molecules $\text{cm}^{-3} \text{s}^{-1}$. 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 NO_2 . 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 NO_2 , ranging from 1.0% to 2.0% for δ ranging from 10^2 to 10^3 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 O_3 , OH, and sulfate. The error for NO_2 reaches 1-10% at high latitudes, but this is still well within other systematic sources of errors in estimating NO_2 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, NO_2 , and sulfate.

10) Figure 4. Can the figure caption give more information here? What actually is being compared? Is this the RRMS in the monthly mean fields, or in the hourly values averaged to a monthly mean? Is this all of the species in the Jacobian?

Response. Now we say.

Line 465. The accuracy is measured by the 24-hour mean RRMS error relative to a simulation including the full chemical mechanism on the end day of each month.

Line 467. Results are shown for the median RRMS across all 228 species of the full mechanism and more specifically for ozone, OH, NO₂, and sulfate.

11) *It's not obvious that the code for the annealing algorithm is included in the repository. I've had a look but can't find it.*

Response. We have uploaded the code. Please check.

12) *Conclusions. a. It would be useful to discuss whether this algorithm could be used within the adjoint framework for data assimilations, inversion studies? b. The authors discuss the suitability of this approach to minor mechanistic changes. However, if the algorithm is to be useful it needs to be sustainable within the software lifecycle of the chemistry transport model. Could this be spelt out in more detail? Presumably, if a new species was added the training algorithm (which species into which block and how many blocks etc) would need to be re-run with new data, but a small change in species lifetime would not lead to a re-running. It would be useful to have the conditions which are required for the training to be updated to be described.*

Response. Now we say this.

Line 294. (4) It is robust against small mechanistic changes, as these may not alter the choice of chemical regimes or may be accommodated by minor tweaking of the **regimes (new species may be assigned to their most appropriate groups on the basis of chemical logic)**. (5) It is robust against increases in model resolution, where source gridboxes (e.g., urban areas) will simply default to the full mechanism. **(6) If an adjoint is available for the full chemical solver, then it can also be used in our method since the software code of the full chemical solver (e.g. KPP) is retained.**