

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 surprising. 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.

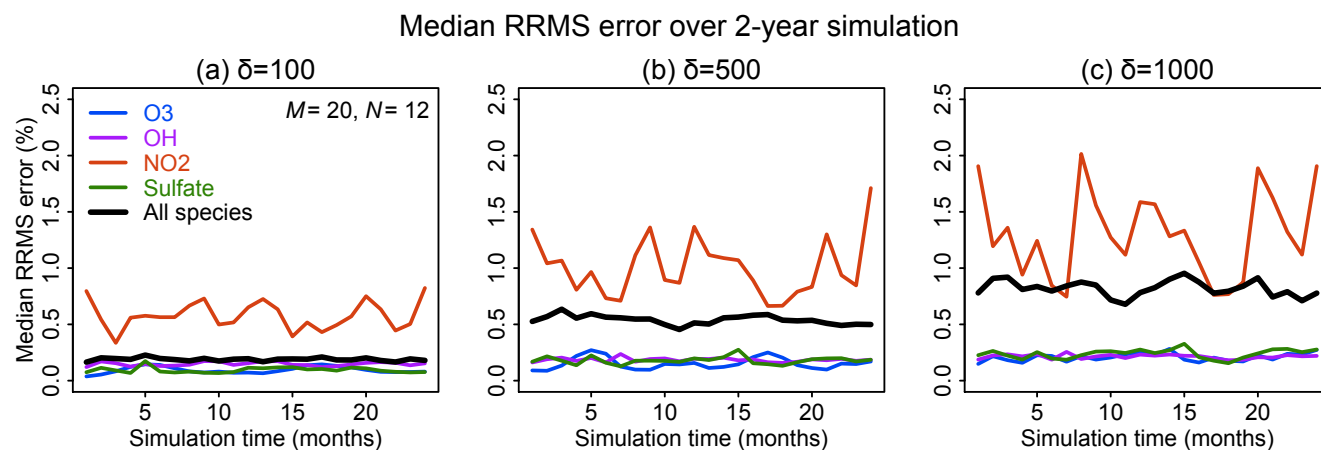


Figure 5. Accuracy of the adaptive reduced chemistry mechanism algorithm over a two-year GEOS-Chem simulation (see text). The accuracy is measured by the 24-hour mean RRMS error on the end day of each month relative to a simulation including the full chemical mechanism. Rate thresholds δ of (a) 100, (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.**

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.

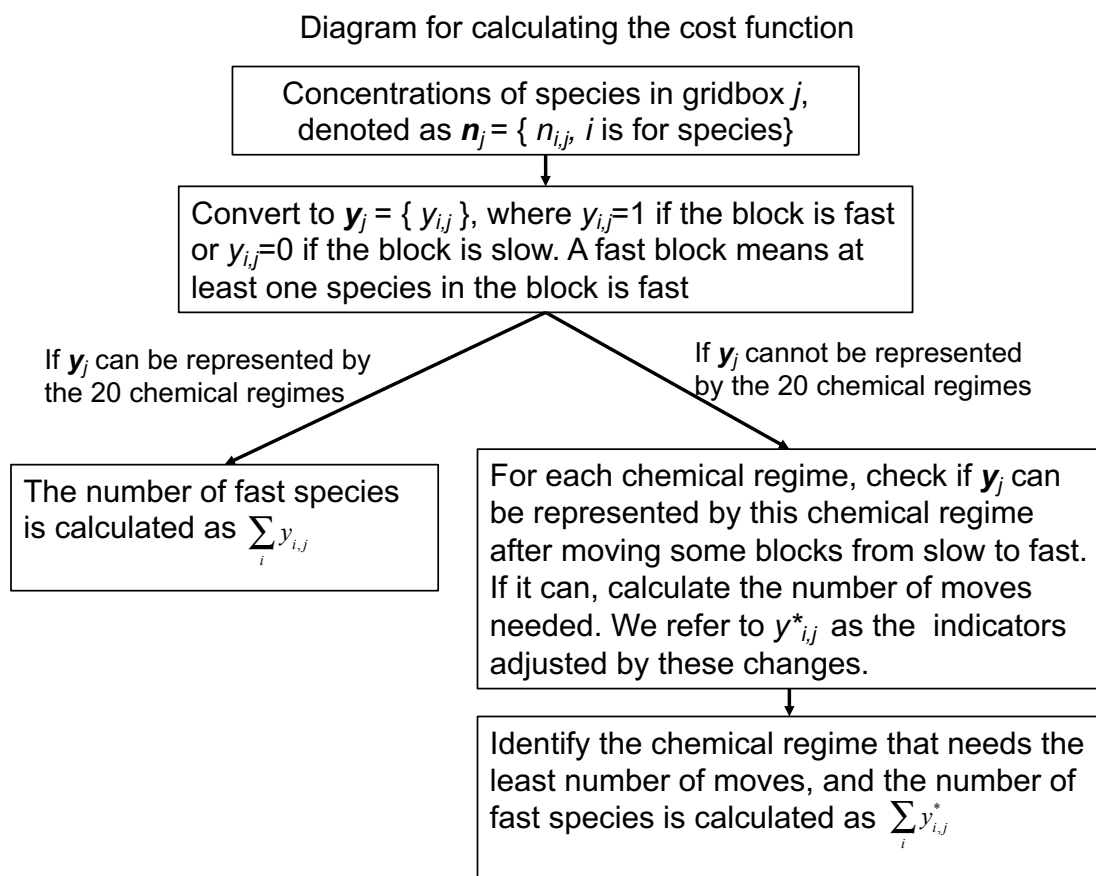


Figure S1. The diagram for calculating the cost function Z_2 . More details can be found in text.

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 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.

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 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.

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.

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

Lu Shen¹, Daniel J. Jacob¹, Mauricio Santillana^{2,3}, Xuan Wang¹, Wei Chen⁴

5

¹John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA

²Computational Health Informatics Program, Boston Children's Hospital, Boston, MA, USA

³Department of Pediatrics, Harvard Medical School, Boston, MA, USA

⁴Department of Physics, Harvard University, Cambridge, MA, USA

10 *Correspondence to:* Lu Shen (lshen@fas.harvard.edu)

Abstract. The major computational bottleneck in atmospheric chemistry models is the numerical integration of the stiff coupled system of kinetic equations describing the chemical evolution of the system as defined by the model chemical mechanism (typically over 100 coupled species). We present an adaptive method to greatly reduce the computational cost of that numerical integration in global 3-D models while maintaining high accuracy. Most of the atmosphere does not in fact
15 require solving for the full chemical complexity of the mechanism, so considerable simplification is possible if one can recognize the dynamic continuum of chemical complexity required across the atmospheric domain. We do this by constructing a limited set of reduced chemical mechanisms (chemical regimes) to cover the range of atmospheric conditions, and then pick locally and on the fly which mechanism to use for a given gridbox and time step on the basis of computed production and loss rates for individual species. Application to the GEOS-Chem global 3-D model for oxidant-aerosol
20 chemistry in the troposphere and stratosphere (full mechanism of 228 species) is presented. We show that 20 chemical regimes can largely encompass the range of conditions encountered in the model. Results from a 2-year GEOS-Chem simulation shows that our method can reduce the computational cost of chemical integration by 30-40% while maintaining accuracy better than 1% and with no error growth. Our method retains the full complexity of the original chemical mechanism where it is needed, provides the same model output diagnostics (species production and loss rates, reaction rates)
25 as the full mechanism, and can accommodate changes in the chemical mechanism or in model resolution without having to reconstruct the chemical regimes.

1 Introduction

Accurate representation of atmospheric chemistry is of central importance for air quality and Earth system models (National
30 Research Council, 2016) *but it is computationally expensive.* 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. Computing the kinetic temporal evolution of such systems involves solving a stiff system of N coupled non-linear ordinary differential equations (ODEs) of the form

$$35 \quad \frac{d\mathbf{n}_i}{dt} = P_i(\mathbf{n}) - L_i(\mathbf{n}) \quad (1)$$

where $\mathbf{n} = (n_1, \dots, n_K)^T$ is the vector of species concentrations, expressed typically as number densities (e.g., molecules cm^{-3}), and K is the number of species in the mechanism. $P_i(\mathbf{n})$ and $L_i(\mathbf{n})$ are the production and loss rates of species i that depend on the concentrations of other species in the mechanism. Finite-difference solution of the coupled system of ODEs requires an implicit scheme to avoid limitation of the time step by the shortest lifetime in the system (Brasseur and Jacob, 2017).

40 Implicit schemes involve repeated construction and inversion of the Jacobian matrix ($K \times K$) for the system, and this is computationally expensive for large K . But the full coupled chemical mechanism may not be needed everywhere in the model domain. For example, highly reactive volatile organic compounds (VOCs) have little influence far away from their source regions. Here we show that we can obtain a substantial reduction of computational cost in a global 3-D model by adaptively adjusting the ensemble of species that actually need to be solved as a coupled system in a given model gridbox.

45 We do so with a general algorithm that is readily applicable to any chemical mechanism or numerical solver.

As the simplest example of an implicit scheme, consider the first-order method which approximates Eq. (1) as

$$f_i(\mathbf{n}(t + \Delta t)) = n_i(t + \Delta t) - n_i(t) - s_i(\mathbf{n}(t + \Delta t))\Delta t = 0 \quad (2)$$

where Δt is the time step and $s_i(\mathbf{n}(t + \Delta t)) = P_i(\mathbf{n}(t + \Delta t)) - L_i(\mathbf{n}(t + \Delta t))$ is the net source evaluated at the end of the time step.

This defines a vector function $\mathbf{f} = (f_1, \dots, f_K)^T$ and an algebraic system $\mathbf{f}(\mathbf{n}(t + \Delta t)) = \mathbf{0}$ that is solved iteratively by the Newton-Raphson method. The procedure involves iterative calculation and inversion of the $K \times K$ Jacobian matrix $\mathbf{J} = \partial \mathbf{f} / \partial \mathbf{n}$. Most models use higher-order implicit algorithms designed for accuracy and speed, such as the Gear (Gear 1971; Hindmarsh, 1983) and Rosenbrock (Sandu et al., 1997; Hairer and Wanner, 1991) solvers, but all require iteratively calculating the Jacobian matrix and solving the linear system using a matrix factorization. As a result, the chemical operator that solves for the chemical evolution of species concentrations from Eq. (1) is the most expensive component of atmospheric chemistry models (Eastham et al., 2018), and this computational cost has been a barrier for inclusion of atmospheric chemistry in Earth system models (National Research Council, 2012).

There are various ways to speed up the chemical operator, all involving some loss of accuracy or generality (Brasseur and Jacob, 2017). A general approach is to reduce the dimension of the coupled system of ODEs that needs to be solved implicitly. This can be done by simplifying the chemical mechanism to decrease the number of species (Brown-Steiner et al., 2018; Sportisse and Djouad, 2000), or by isolating long-lived species for which a fast explicit solution scheme is acceptable

(Young and Boris, 1977). Jacobson (1995) used different subsets of their full mechanism to simulate the urban atmosphere, the troposphere, and the stratosphere. 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).

65 Santillana et al. (2010) combined these ideas in an adaptive algorithm for 3-D models that determines locally at each time step (“on the fly”) which species in the chemical mechanism need to be solved in the coupled implicit system. This was done by computing the local production (P_i) and loss rates (L_i) for all species at the beginning of the time step. Species with either P_i or L_i above a given threshold were labeled “fast” and solved with an implicit scheme, while the others were labeled “slow” and solved with an explicit scheme. The complexity of the chemical system to be solved was thus adapted to the local environment. Here ‘fast’ and ‘slow’ refer to the rates in the chemical system, not the species lifetime. For example, short-lived VOCs may be considered slow outside of their source regions because they have negligible influence on other species. The extent of this influence depends on the changing local conditions, hence the need for an adaptive algorithm, The adaptive approach does not pre-judge the local environment, unlike in Jacobson (1995), and instead resolves the dynamic continuum of complexity that may be encountered in the atmosphere. Santillana et al. (2010) applied their algorithm to the
70 GEOS-Chem global 3-D Eulerian chemical transport model (Bey et al., 2001). While the computational savings were promising for the chemical integration within each gridbox, the need to construct a different system in every single grid box and at every time step cancelled out some of the gains and led to only small time-savings when compared to the performance of the standard full-chemistry model.

Here we draw from the approach introduced by Santillana et al. (2010) but use a set of pre-defined chemical regimes to take
80 full advantage of the time-savings from the adaptive reduction mechanism algorithm. We start with the objective identification of a limited number of chemical regimes that encompass the range of atmospheric conditions encountered in the model. These regimes are defined by the subset of fast species from the full mechanism that need to be considered in the coupled system, and we pre-code the Jacobian matrix and its inverse for each. The model then picks the appropriate chemical regime to be solved locally and on the fly. We show that this approach can achieve large computational savings
85 without significantly compromising accuracy when implemented in GEOS-Chem. Our method can be adapted to any mechanism and model, retains the complexity of the full mechanism where it is needed, and preserves full diagnostic information on chemical evolution (such as reaction rates, production and loss of individual species, etc.).

2 Model description

90 ~~The evolution of a chemical system within a 3-D Eulerian atmospheric model is obtained by solving the system of K coupled~~

continuity equations (Brasseur and Jacob, 2017)

95 where U is the wind vector. The first term on the right hand side describes the transport of the species, and the other terms describe chemical production and loss as given by Eq. (1). Emission and deposition may be treated as boundary conditions or added to the chemical production and loss terms. Turbulence not resolved by the wind vector can be represented with additional parameterized transport terms.

100 The system of Eq. (3) must be solved by operator splitting, in which transport and chemistry are solved separately and successively over discrete time steps (Santillana et al., 2016). Operator splitting reduces the dimensionality of the problem because the transport operators do not involve coupling between chemical species, while the chemical operator does not involve spatial coupling. Thus the chemical operator solves the system of ODEs described by Eq. (1) over the operator splitting time step Δt , passes the updated chemical concentrations to the transport operators, which in turn update the concentrations. After the transport operators have been applied, the results are returned to the chemical operator as initial conditions for the next time step.

105 The chemical operator is the most computationally costly component of an “off-line” 3-D atmospheric chemistry model where winds and turbulence parameters are taken as input. It remains a major computational burden even in “on-line” models that solve for atmospheric dynamics as well as chemistry. This holds in massively parallel computing environments despite the near perfect scalability of the chemical operator. For example, Eastham et al. (2018) found in the off-line GEOS-Chem High Performance model at global cubed sphere e180 ($\approx 50 \times 50$ km²) resolution that the chemical operator was responsible for 50% of the overall wall time using 90 cores and 37% using 360 cores.

110 We use the GEOS-Chem 12.0.0 global 3-D model for tropospheric and stratospheric chemistry (<https://doi.org/10.5281/zenodo.1343547>) as demonstration for our algorithm. The model is applied here with a horizontal resolution of $4^\circ \times 5^\circ$ and 72 pressure levels extending from surface to 0.01 hPa. It is driven by MERRA2 assimilated meteorological data from the NASA Global Modeling and Assimilation System (GMAO). The model includes coupled gas-phase and aerosol chemistry as described by Sherwen et al. (2016) and Travis et al. (2016) for the troposphere and Eastham et al. (2014) for the stratosphere. The chemical mechanism has 228 species and 724 reactions. Among these species, 143 are volatile organic compounds (VOCs), 37 are inorganic reactive halogen species, 24 are organic halogen species, and 24 are other inorganic and aerosol species. The chemical reactions are integrated using the Rosenbrock solver (Sandu et al., 1997; Hairer and Wanner, 1991) generated from the Kinetic PreProcessor 2.2.4 (KPP) (Damian et al., 2002) software. The model uses operator splitting between chemistry and transport with a chemistry timestep of 20 minutes (Philip et al 2016). We use
120 12 cores with shared memory in the simulations.

The key processes in the KPP chemical operator are as follows. The operator first updates the reaction rate coefficients on the basis of temperature, actinic flux, etc. It then passes these reaction rate coefficients together with initial species concentrations to the Rosenbrock solver, which solves for the temporal evolution of concentrations over the external timestep Δt . In the process, the Rosenbrock solver approximates the solution at multiple internal timesteps, so it needs to repeatedly recompute the species production and loss rates, construct the corresponding Jacobian matrix, and solve the linear system numerically using a matrix factorization. The bulk of the cost in the overall chemical operator is in the repeated computation of production/loss rates and solving the linear system using a matrix factorization. Reducing the number of species in the system to be solved can significantly reduce the computational cost.

130 3 The adaptive algorithm for the chemical operator

Our adaptive algorithm determines locally what degree of complexity is needed in the chemical mechanism by diagnosing all species in the full chemical mechanism as either “fast” or “slow”, and choosing among pre-constructed chemical mechanism subsets (“chemical regimes”) which is most appropriate for the local conditions. Here we present (1) the definition of fast and slow species and the different treatments for each, and (2) the approach used to pre-construct the chemical regimes.

3.1 Definition of fast and slow species

Following Santillana et al. (2010), we separate atmospheric species as fast or slow based on their production and loss rates in Eq. (1) relative to a threshold δ : fast if either $P_i(\mathbf{n}) \geq \delta$ or $L_i(\mathbf{n}) \geq \delta$, slow if $P_i(\mathbf{n}) < \delta$ and $L_i(\mathbf{n}) < \delta$. Concentrations of the fast species are integrated as a coupled system with the KPP Rosenbrock solver. Concentrations of slow species are integrated by explicit analytical solution of Eq. (1) assuming first-order loss with effective rate coefficient $k_i = L_i/n_i$.

$$\frac{dn_i}{dt} = P_i - k_i n_i \quad (3)$$

$$n_i(t + \Delta t) = \frac{P_i(t)}{k_i(t)} + \left(n_i(t) - \frac{P_i(t)}{k_i(t)}\right) e^{-k_i(t)\Delta t} \quad (4)$$

Solving for $n_i(t + \Delta t)$ by Eq. (4) incurs negligible computational cost, therefore there is considerable advantage in classifying species as slow if this can be done without significant loss in accuracy. We select the threshold δ for species to be classified as fast or slow by numerical testing, as described in Section 4, but some basic chemical reasoning is useful. Consider the OH radical, which is a central species in atmospheric chemistry mechanisms. OH has a daytime concentration of the order of 10^6 molecules cm^{-3} and a lifetime of the order of a second, implying production and loss rates of the order of 10^6 molecules cm^{-3}

150 s^{-1} . Species with production and loss rates that are orders of magnitude lower than 10^6 molecules $cm^{-3} s^{-1}$ are therefore unlikely to influence OH or other species in the coupled mechanism, as these are all to some extent related to OH at least in the daytime. So we may expect an appropriate threshold δ to be of the order of 10^2 - 10^3 molecules $cm^{-3} s^{-1}$. Santillana et al. (2010) recommended $\delta = 100$ molecules $cm^{-3} s^{-1}$ in their algorithm.

One issue with the solution for the slow species by Eq. (4) is that it does not strictly conserve mass, because the loss rate for a given species over the time step does not necessarily match the production rate of the product species. This is usually inconsequential, but we found in early testing that it resulted in the total mass of reactive halogen species growing slowly over time in the stratosphere. To avoid this effect, we treat all 37 reactive inorganic halogen species as fast above 10 km altitude. This increases the computation cost of chemical integration by only 4% relative to letting the algorithm set them as either fast or slow.

3.2 Pre-selecting the chemical regimes

160 Instead of building a local chemical mechanism subset at every time step as in Santillana et al. (2010), we greatly improve the computational efficiency by pre-selecting a limited number (M) of chemical mechanism subsets (chemical regimes) for which we pre-define the Jacobian matrix in KPP. We then determine locally which chemical regime to apply on basis of the ensemble of species classified as fast. This approach reduces the computational overhead of repeatedly allocating and deallocating memory in the method of Santillana et al. (2010).

165 Construction of the chemical regimes can be done objectively by searching for a minimum in the computational cost of the chemical operator over the global domain. But some narrowing of the search is necessary. For the 228-species mechanism in GEOS-Chem, there are in principle $2^{228}-1$ possible combinations of species that would form mechanism subsets. The vast majority of those combinations make no chemical sense, but diagnosing this objectively would be computationally unfeasible. Instead, we start by splitting the mechanism species into N different blocks based on similarity of chemical behavior. Then we classify a block as fast if at least one species in the block is fast, and slow if all species in the block are slow. The chemical regime is defined as the assemblage of fast blocks.

175 The partitioning of species into blocks can be optimized by minimizing globally the number of fast species (and hence the computation cost) for a given threshold δ . We use for this purpose a training dataset from a GEOS-Chem simulation for 2013, consisting of the global ensemble of tropospheric and stratospheric gridboxes for the first 10 days of February, May, August, and November sampled every 6 hours. For each gridbox j , we diagnose each individual species i as fast or slow following Section 3.1. We then diagnose the blocks as fast or slow with the indicator $y_{i,j} = 1$ if the block is fast (at least one species in the block is fast) or $y_{i,j} = 0$ if the block is slow (all species in the block are slow). The fraction Z_1 of all species that needs to

be treated as fast over the testing domain is then given by

$$Z_1 = \frac{1}{\Omega} \sum_j \sum_i y_{i,j} \quad (5)$$

where Ω is the total number of gridboxes multiplied by the total number of species (228 in our case). We seek the
180 partitioning of species into blocks that will minimize Z_1 , and we use for that purpose the simulated annealing algorithm
(Kirkpatrick et al., 1983). Starting from an arbitrary partitioning of the 228 species into N blocks, and at each iteration of the
algorithm, we randomly move one species from one block to another. If Z_1 decreases, this transition is accepted; if not, the
transition is accepted with a probability controlled by a parameter named temperature that decreases gradually as the
algorithm proceeds. Among the N blocks, 3 are allocated to the reactive inorganic halogen species, and $N-3$ are allocated to
185 the other species. This forced separation of the reactive inorganic halogen species is because the corresponding blocks are
imposed to be fast above 10 km altitude (see Section 3.1). [Throughout this study, we present the results with lowest cost
function after running the optimization multiple times and using different temperature parameters.](#)

[Once the blocks have been defined in the above manner](#), we define the chemical regimes as different assemblages of blocks.
This yields $2^N - 1$ possible chemical regimes. Individual gridboxes in the model domain may correspond to any of these $2^N -$
190 1 regimes depending on which blocks are classified as fast or slow. We need to limit the number of regimes to a much
smaller number M of most useful regimes in order to keep the compilation of the code manageable, and as we will see the
bulk of conditions in the model domain can effectively be represented by just a few regimes. 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. [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.
Here we refer to \$y^*_{i,j}\$ as the indicators adjusted by these changes.](#) Thus, the fraction Z_2 of species that needs to be treated as
195 fast over the global domain is given by:

$$Z_2 = \frac{1}{\Omega} \left(\sum_{D_1} \sum_i y_{i,j} + \sum_{D_2} \sum_i y^*_{i,j} \right) \quad (6)$$

200 where D_1 are the gridboxes that can be represented by the top M chemical regimes, and D_2 are the gridboxes that are
represented by other regimes and must be matched to the top M regimes. [A diagram for this process can be found in Figure
S1.](#)

We tested a range of values from 3 to 20 for the number N of blocks. In this testing we used a threshold $\delta = 100$ molecules $\text{cm}^{-3} \text{s}^{-1}$ to partition fast and slow species, following Santillana et al. (2010), and a number $M = 20$ of chemical regimes (see next paragraph for choice of M). Figure 1 shows the fraction of fast species in the global domain (Z_2) as a function of N . If N is low such that blocks are large, there is more likelihood that a species in a given block will be fast causing all species in the block to be treated as fast. If N is high, more blocks will need to be moved from slow to fast in order to match the limited number M of chemical regimes. For $M = 20$ we thus find an optimal value $N = 12$ at which only 40% of the species need to be treated as fast.

Table 1 lists selected representative the species of these 12 blocks and a full listing is in Fig. S2. Oxidants such as OH, O₃, and NO₂ are important under all circumstances so block 8 and 9 are fast in most gridboxes. Nonmethane VOCs species often have low concentrations outside of the continental boundary layer, and very low concentrations in the stratosphere, so the dominant VOC blocks 1-7 are fast in fewer than 40% of gridboxes. 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).

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.

We tested different numbers of chemical regimes (M) from 3 to 40 for combining the $N = 12$ blocks, and again selected the regimes to minimize the global fraction Z_2 of species to be included in the implicit solver. Z_2 decreases from 65% to 40% as M increases from 3 to 20 and flattens at higher values of M (Fig. 2a). This is because 88% of the gridboxes can be represented by 20 chemical regimes (Fig. 2b). A larger number of blocks ($N > 12$) would extend the improvement to higher values of M , but the size of M is also limited by considerations of code manageability and compilation speed. We use 20 chemical regimes in what follows.

Table 2 shows the composition of the 20 chemical regimes as defined by the blocks of Table 1. For 72% of the gridboxes, we only need to solve for fewer than 50% of the species as fast. Only 3.6% of gridboxes need to use the full chemistry mechanism, as defined by the 20th regime.

235 Figure 3 further shows the distribution of these 20 chemical regimes globally and for different altitudes, and the corresponding percentage of fast species that needs to be included in the chemical solver. In continental surface air where VOC emissions are concentrated, we find that over 80% of species generally need to be included. This percentage is reduced to 20-60% over the ocean and < 20% over Antarctica. At 5 km altitude, we find a distinct boundary between the daytime and nighttime hemisphere; the daytime chemistry is more active, and the percentage of fast species is higher in the daytime (40-
 240 60%) than at night (10%-30%). At 15 km altitude the extratropics are in the stratosphere, where non-methane VOC chemistry is largely absent, but the model still needs to solve 30-40% species as fast because of the halogens. Deep convection over tropical continents delivers short-lived VOCs and their oxidation products to the upper troposphere, so that a large number of species needs to be treated as fast in the convective outflow where and when it occurs. The importance of deep convective outflow for global atmospheric chemistry has been pointed out in a number of studies (Prather and Jacob,
 245 1997; Bechara et al., 2010; Schroeder et al., 2014), and emphasizes the advantage of reducing the mechanism adaptively on the fly rather than with pre-set geographic boundaries.

4 Error analysis

Here we quantify the errors in our adaptive reduced mechanism method by comparison with a standard GEOS-Chem simulation for the troposphere and stratosphere (version 12.0.0) including full chemistry (228 species). The comparison is
 250 conducted for a 1-month simulation to examine the sensitivity to the rate threshold δ , and for a 2-year simulation to evaluate the stability of the method. In both cases, we use the Relative Root Mean Square (RRMS) metric as given by Sandu et al. (1997) to characterize the error:

$$RRMS_i = \sqrt{\frac{1}{Q_i} \sum_{j=1}^{Q_i} \left(\frac{n_{i,j}^{\text{reduced}} - n_{i,j}^{\text{full}}}{n_{i,j}^{\text{full}}} \right)^2} \quad (8)$$

255 Where $n_{i,j}^{\text{reduced}}$ and $n_{i,j}^{\text{full}}$ are the concentrations for species i and gridbox j in the reduced and full chemical mechanisms, and the sum is over the Q_i gridboxes where $n_{i,j}^{\text{full}}$ is greater than a threshold a . Here we use $a = 1 \times 10^6$ molecules cm^{-3} as in Eller et al. (2009) and Santillana et al. (2010).

A critical parameter to select in the algorithm is the rate threshold δ separating fast and slow species on the basis of their production and loss rates. A high threshold decreases the number of fast species and hence speeds up the computation but at

260 the expense of accuracy. We tested different rate thresholds ranging from 10 to 5000 molecules $\text{cm}^{-3} \text{s}^{-1}$ in a 1-month GEOS-Chem simulation starting on August 1 2013. Figure 4 shows the median RRMS error for all species on September 1 and the increased computational performance for different rate thresholds δ . 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.

265 Figure S6 further shows the distribution of RRMS errors over all species for different rate thresholds δ . The 90th percentile RRMS error stays below 5% if $\delta \leq 1000$ molecules $\text{cm}^{-3} \text{s}^{-1}$ but exceeds 10% for $\delta = 5000$ molecules $\text{cm}^{-3} \text{s}^{-1}$. The 99th percentile RRMS error is less than 20% for $\delta \leq 1000$ molecules $\text{cm}^{-3} \text{s}^{-1}$ but rises to 80% for $\delta = 5000$ molecules $\text{cm}^{-3} \text{s}^{-1}$. The largest errors are usually from the tropospheric halogen species (Fig. S7). When near the day-night terminator, the sharp transition of production and loss rates is not properly captured by the first-order explicit equations, resulting in high relative errors.

270 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.

280 5 Conclusions

We have presented an adaptive method to speed up the temporal integration of chemical mechanisms in global atmospheric chemistry models. This integration (“chemical operator”) involves the implicit solution of a stiff coupled system of ordinary differential equations (ODEs) representing the kinetic evolution of individual species in the mechanism. With typical mechanisms including over 100 coupled species, this chemical integration is the principal computational bottleneck in atmospheric chemistry models and hinders the adoption of detailed atmospheric chemistry in Earth system models.

285 Our method takes advantage of the fact that different regions of the atmosphere need different levels of detail in the chemical mechanism, and that greatly reduced mechanisms can be used in most of the atmosphere. We do this reduction locally and on the fly by choosing from a portfolio of pre-selected reduced chemical mechanisms (chemical regimes) on the basis of species production and loss rates, distinguishing between “fast” species that need to be in the coupled mechanism and “slow” species that can be solved explicitly. Our method has six advantages over other methods proposed to speed up the chemical

computation: (1) It does not sacrifice the complexity of the chemical mechanism where it is needed, while greatly simplifying it over much of the world where it is not. (2) It conserves all of the meaningful diagnostic information of the chemical system, such as production and loss rates of species and families, and individual reaction rates. (3) It can be tailored to achieve the level of simplification that one wishes. (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.

We applied the method to the GEOS-Chem global 3-D model for oxidant-aerosol chemistry in the troposphere and stratosphere. The full chemical mechanism in GEOS-Chem has 228 coupled species. We developed an objective numerical method to pre-select the reduced chemical regimes on the basis of time slices of full-mechanism model results. We showed that 20 regimes could cover efficiently the range of atmospheric conditions encountered in the model. We then pick appropriate regimes for the chemical operator on the fly by comparing the local production and loss rates of individual model species to a threshold δ . Values of δ in the range 100-1000 molecules cm^{-3} maintain an accuracy better than 1% relative to a model simulation with the full mechanism and decrease the computational cost of the chemical solver by 32-41%. Comparison testing with a 2-year global GEOS-Chem simulation for the troposphere and stratosphere including the full mechanism shows errors of less than 1% for critical species and no significant error growth over the two years.

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).

Several improvements could be made to our method. (1) The blocks of species used to construct the reduced chemical mechanisms are optimized to minimize the number of fast species but are not always chemically logical, which could be improved by applying prior regularization constraints to the optimization. (2) Optimization in the definition of the reduced mechanisms could take into account not only the number of species but also their lifetimes that affect the stiffness of the system. (3) Separation between fast and slow species could take into account species lifetimes, because species with long lifetimes but high loss rates (such as methane or CO) can be solved explicitly. (4) Mass conservation in the explicit solution could be enforced to enable more species (in particular stratospheric halogens) to be treated explicitly when they play little role in the coupled system. (5) Besides removing the slow species from the implicit chemical operator, we could also remove unimportant reactions, which would reduce the cost in updating the production/loss rates and the Jacobian matrix. These improvements will be the target of future work.

Code availability. The standard GEOS-Chem code is available through <https://doi.org/10.5281/zenodo.1343547>. The updates for the adaptive mechanism can be found at <https://doi.org/10.7910/DVN/IM5TM4>.

Data availability. All datasets used in this study are publically accessible at <https://doi.org/10.7910/DVN/IM5TM4>.

Author contribution. L. Shen and D. Jacob designed the experiments and L. Shen carried them out. L. Shen and D. Jacob prepared the manuscript with contributions from all co-authors.

Competing Interests. The authors declare that they have no conflict of interest.

Acknowledgments. This work was funded by the NASA Modeling and Analysis Program (MAP)

References

- Brown-Steiner, B., Selin, N. E., Prinn, R., Tilmes, S., Emmons, L., Lamarque, J.-F., and Cameron-Smith, P.: Evaluating simplified chemical mechanisms within present-day simulations of the Community Earth System Model version 1.2 with CAM4 (CESM1.2 CAM-chem): MOZART-4 vs. Reduced Hydrocarbon vs. Super-Fast chemistry, *Geosci. Model Dev.*, 11, 4155–4174, <http://sci-hub.tw/10.5194/gmd-11-4155-2018>, 2018.
- Brasseur, G.P. and Jacob, D.J.: *Modeling of atmospheric chemistry*, Cambridge University Press, 2017
- Bechara, J., Borbon, A., Jambert, C., Colomb, A., and Perros, P. E.: Evidence of the impact of deep convection on reactive Volatile Organic Compounds in the upper tropical troposphere during the AMMA experiment in West Africa, *Atmos. Chem. Phys.*, 10, 10321–10334, <https://doi.org/10.5194/acp-10-10321-2010>, 2010.
- Bey, I., Jacob, D. J., Yantosca, R. M., Logan, J. A., Field, B. D., Fiore, A. M., Li, Q., Liu, H. Y., Mickley, L. J., and Schultz, M. G.: Global modeling of tropospheric chemistry with assimilated meteorology: Model description and evaluation, *J. Geophys. Res.*, 106, 23 073–23 095, 2001.

- 355 Damian, V., Sandu, A., Damian, M., Potra, F., and Carmichael, G. R.: The kinetic preprocessor KPP – a software environment for solving chemical kinetics, *Comput. Chem. Eng.*, 26, 1567–1579, 2002.
- Eastham, S. D., Long, M. S., Keller, C. A., Lundgren, E., Yantosca, R. M., Zhuang, J., Li, C., Lee, C. J., Yannetti, M., Auer, B. M., Clune, T. L., Kouatchou, J., Putman, W. M., Thompson, M. A., Trayanov, A. L., Molod, A. M., Martin, R. V., and Jacob, D. J.: GEOS-Chem High Performance (GCHP v11-02c): a next-generation implementation of the GEOS-Chem chemical transport model for massively parallel applications, *Geosci. Model Dev.*, 11, 2941–2953, <http://sci-hub.tw/10.5194/gmd-11-2941-2018>, 2018.
- 360 Eller, P., Singh, K., Sandu, A., Bowman, K., Henze, D. K., and Lee, M.: Implementation and evaluation of an array of chemical solvers in the Global Chemical Transport Model GEOS-Chem, *Geosci. Model Dev.*, 2, 89–96, <http://sci-hub.tw/10.5194/gmd-2-89-2009>, 2009.
- 365 Eastham, S. D., Weisenstein, D. K., and Barrett, S. R. H.: Development and evaluation of the unified tropospheric–stratospheric chemistry extension (UCX) for the global chemistry-transport model GEOS-Chem, *Atmos. Environ.*, 89, 52–63, doi:10.1016/j.atmosenv.2014.02.001, 2014.
- Gear C. W.: *Numerical Initial Value Problems in Ordinary Differential Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1971.
- 370 Hairer, E. and Wanner, G.: *Solving Ordinary Differential Equations II. Stiff and Differential-Algebraic Problems*, Springer, Berlin, 1991.
- Hindmarsh, A. C.: ODEPACK: A systematized collection of ODE solvers, *Sci. Comput*, 55-64, 1983.
- Jacobson, M. Z.: Computation of global photochemistry with SMVGEAR II, *Atmos. Environ.*, 29, 2541–2546, 1995.
- Keller, C. A. and Evans, M. J.: Application of random forest regression to the calculation of gas-phase chemistry within the GEOS-Chem chemistry model v10, *Geosci. Model Dev.*, 12, 1209–1225, <https://doi.org/10.5194/gmd-12-1209-2019>, 2019.
- 375 Kirkpatrick, S., Gelatt, C.D. and Vecchi, M.P.: Optimization by simulated annealing, *Science*, 220 (4598), 671-680, 1983.
- Long, M. S., Yantosca, R., Nielsen, J. E., Keller, C. A., da Silva, A., Sulprizio, M. P., Pawson, S., and Jacob, D. J.: Development of a grid-independent GEOS-Chem chemical transport model (v9-02) as an atmospheric chemistry module for Earth system models, *Geosci. Model Dev.*, 8, 595–602, <https://doi.org/10.5194/gmd-8-595-2015>, 2015.
- 380 National Research Council: *A National Strategy for Advancing Climate Modeling*, National Academies Press, Washington DC, 2012.
- National Research Council: *The Future of Atmospheric Chemistry Research: Remembering Yesterday, Understanding Today, Anticipating Tomorrow*, National Academies Press, Washington DC, 2016.
- 385 Prather, M. J. and Jacob, D. J.: A persistent imbalance in HO_x and NO_x photochemistry of the upper troposphere driven by deep tropical convection, *Geophys. Res. Lett.*, 24, 3189–3192, 1997.

- Philip, S., Martin, R. V., and Keller, C. A.: Sensitivity of chemistry-transport model simulations to the duration of chemical and transport operators: a case study with GEOS-Chem v10-01, *Geosci. Model Dev.*, 9, 1683–1695, <http://sci-hub.tw/10.5194/gmd-9-1683-2016>, 2016.
- 390 Sportisse, B., Djouad, R.: Reduction of chemical kinetics in air pollution modeling, *J. Comput. Phys.*, 164, 354–376, 2000.
- Santillana, M., Le Sager, P., Jacob, D.J., and Brenner, M.P.: An adaptive reduction algorithm for efficient chemical calculations in global atmospheric chemistry models, *Atmos. Environ.*, 44(35), 4426–4431, 2010.
- Sherwen, T., Schmidt, J. A., Evans, M. J., Carpenter, L. J., Großmann, K., Eastham, S. D., Jacob, D. J., Dix, B., Koenig, T. K., Sinreich, R., Ortega, I., Volkamer, R., Saiz-Lopez, A., Prados-Roman, C., Mahajan, A. S., and Ordóñez, C.:
395 Global impacts of tropospheric halogens (Cl, Br, I) on oxidants and composition in GEOS-Chem, *Atmos. Chem. Phys.*, 16, 12239–12271, <http://sci-hub.tw/10.5194/acp-16-12239-2016>, 2016.
- Santillana, M., Zhang, L., and Yantosca, R.: Estimating numerical errors due to operator splitting in global atmospheric chemistry models: Transport and chemistry, *J. Comput. Phys.*, 305, 372–386, doi:10.1016/j.jcp.2015.10.052, 2016.
- Silvern, R. F., Jacob, D. J., Travis, K. R., Sherwen, T., Evans, M. J., Cohen, R. C., Laughner, J. L., Hall, S. R., Ullmann, K.,
400 Crouse, J. D., Wennberg, P. O., Peischl, J., and Pollack, I. B.: Observed NO/NO₂ Ratios in the Upper Troposphere Imply Errors in NO-NO₂-O₃ Cycling Kinetics or an Unaccounted NO_x Reservoir, *Geophys. Res. Lett.*, 45, 4466–4474, <https://doi.org/10.1029/2018gl077728>, 2018.
- Schroeder, J.R., Pan, L.L., Ryerson, T., Diskin, G., Hair, J., Meinardi, S., Simpson, I., Barletta, B., Blake, N. and Blake, D.R.: Evidence of mixing between polluted convective outflow and stratospheric air in the upper troposphere during
405 DC3, *J. Geophys. R.*, 119(19), 11–477, 2014.
- Sandu, A., Verwer, J. G., Blom, J. G., Spee, E. J., Carmichael, G. R., and Potra, F. A.: Benchmarking stiff ode solvers for atmospheric chemistry problems II: Rosenbrock solvers, *Atmos. Environ.*, 31, 3459–3472, 1997
- Travis, K. R., Jacob, D. J., Fisher, J. A., Kim, P. S., Marais, E. A., Zhu, L., Yu, K., Miller, C. C., Yantosca, R. M., Sulprizio, M. P., Thompson, A. M., Wennberg, P. O., Crouse, J. D., St. Clair, J. M., Cohen, R. C., Laughner, J. L., Dibb, J. E.,
410 Hall, S. R., Ullmann, K., Wolfe, G. M., Pollack, I. B., Peischl, J., Neuman, J. A., and Zhou, X.: Why do models overestimate surface ozone in the Southeast United States? *Atmos. Chem. Phys.*, 16, 13561–13577, <http://sci-hub.tw/10.5194/acp-16-13561-2016>, 2016.
- Young T. R. and Boris J. P.: A numerical technique for solving stiff ordinary differential equations associated with the chemical kinetics of reactive flow problems, *J. Phys. Chem.*, 81, 2424–2427, 1977.

Table 1. Partitioning of GEOS-Chem chemical species into $N = 12$ blocks^a.

Block	Type of species ^b	Number of species	Species	% gridboxes where fast ^c
1	Aromatics	21	CH2I2 LBRO2H LBRO2N LTRO2H LTRO2N SO4H2 IMAE BENZ TOLU TRO2 BRO2 CH2CI2 IMAO3 RA3P RP PP IPMN GLYX A3O2 PO2 R4N1	33.4
2	Organic nitrates	7	INDIOL SO4H1 PPN IONITA N RCO3 R4N2	39.3
3	Isoprene, terpenes	30	CH2ICI LISOPH LISOPNO3 MONITA OCS CHBr3 CHCI3 HCFC22 PRPN HPALD HONIT RIPB RIPA LIMO MONITS ISOPNB CH3CHOO MVKN PRN1 MONITU CH2OO PROPNN ISOP OLND OLNN HC5OO ISN1 HC5 RIO2 INO2	13.9
4	Alkanes, alkenes, acetone	12	MSA MAP ETP SO4 ATOOH C2H6 ATO2 ACTA ACET ETO2 PRPE ALD2	41.4
5	Higher alkanes, methyl ethyl ketone	14	CH3I RB3P CH3CI ALK4 R4P C3H8 EOH B3O2 KO2 MGLY R4O2 HAC RCHO MEK	36.5
6	Halocarbons, isoprene products	55	CH2IBr ISN1OA ISN1OG LVOCOA LVOC PYAC SOAMG DHDN CH3CCI3 H1301 H2402 PMNN CC14 CFC11 CFC12 CFC113 CFC114 CFC115 H1211 IEPOXD CH2Br2 HCFC123 HCFC141b HCFC142b CH3Br DHPCARP IAP HPC52O2 MOBA ISNP MAOP MRP RIPD ETHLN ISNOHOO NPMN MOBAOO DIBOO ISNOOB INPN MACRNO2 MVKOO GAOO MGLYOO MGLOO MAN2 ISNOOA ISOPNDO2 MACROO MACRN MAOPO2 LIMO2 ISOPNBO2 ISOPND NMAO3	10.2
7	Secondary organic aerosol	25	LXRO2H LXRO2N SOAGX SOAIE SOAME DHDC IEPOXA IEPOXB XRO2 XYLE PIP HC187 VRP DHMOB MTPA MTPO ROH IEPOXOO HCOOH PIO2 GLYC VRO2 MRO2 MACR MVK	15.6
8	Sulfur, peroxyacetylnitrate	15	CO2 N2O DMS HNO4 HNO2 PAN MP H CH4 H2O2 MCO3 SO2 CO O1D O	95.9
9	Oxidants	12	MPN N2O5 HNO3 CH2O MO2 O3 NO HO2 NO3 NO2 H2O OH	100.0
10	Iodine reservoirs	13	AERI ISALA ISALC I2O4 I2O3 IBr INO HI ICI CINO2 BrSALC BrSALA I2	69.5
11	Bromine and chlorine inorganic species	11	CIOO BrCl Br2 BrNO3 HOBr HOCl CINO3 Cl HBr ClO HCl	99.9
12	Bromine and iodine radicals	13	I2O2 BrNO2 Cl2O2 IONO OIO OCIO HOI IONO2 Cl2 I IO BrO Br	85.0

^aThe full GEOS-Chem mechanism has 228 species. The full names of these acronyms can be found at http://wiki.seas.harvard.edu/geos-chem/index.php/Species_in_GEOS-Chem. Results in Column 2-4 are obtained using data from the first 10 days of February, May, August, and November sampled every 6 hours.

420 ^bQualitative descriptor of the most important species in the block.

^cGlobal percentage of GEOS-Chem model gridboxes in the troposphere and stratosphere where the block is treated as fast. Values are for August 1 2013 sampled every 6 hours.


425 **Table 2.** Composition and frequency of the 20 chemical regimes in the adaptive algorithm^a.

Regime #	Block												% fast species ^b	% gridboxes ^c
	1	2	3	4	5	6	7	8	9	10	11	12		
1	0	0	0	0	0	0	0	0	1	0	0	0	5.6	0.1
2	0	0	0	0	0	0	0	0	1	0	1	0	10.3	3.9
3	0	0	0	0	0	0	0	0	1	0	1	1	15.8	0.1
4	0	0	0	0	0	0	0	1	1	0	1	0	18.4	5.4
5	0	1	0	0	0	0	0	1	1	0	1	0	21.4	2.2
6	0	0	0	0	0	0	0	1	1	0	1	1	23.9	0.5
7	0	1	0	0	0	0	0	1	1	0	1	1	26.9	0.2
8	0	1	0	1	0	0	0	1	1	0	1	0	26.9	1.1
9	0	0	0	0	0	0	0	1	1	1	1	1	29.5	46.3
10	0	0	0	1	0	0	0	1	1	0	1	1	29.5	0.5
11	0	0	0	1	0	0	0	1	1	1	1	1	35.0	3.3
12	0	0	0	1	1	0	0	1	1	0	1	1	35.5	0.7
13	0	1	0	1	1	0	0	1	1	0	1	1	38.5	2.4
14	1	1	0	1	1	0	0	1	1	0	1	1	47.4	5.2
15	1	1	0	1	1	0	0	1	1	1	1	1	53.0	12.7
16	1	1	0	1	1	0	1	1	1	0	1	1	58.1	1.7
17	1	1	1	1	1	0	1	1	1	1	1	1	76.5	3.7
18	1	1	1	1	1	1	1	1	1	0	1	0	88.9	2.3
19	1	1	1	1	1	1	1	1	1	0	1	1	94.4	4.4
20	1	1	1	1	1	1	1	1	1	1	1	1	100.0	3.6

^aThe chemical regimes are defined by the ensemble of fast species that need to be treated as a coupled system with implicit solution in the chemical operator. The species are assembled into blocks as listed in Table 1, and here we identify the blocks treated as fast in the chemical regime (1 ≡ fast, 0 ≡ slow).

430 ^bPercentage of the 228 species in the GEOS-Chem chemical mechanism treated as fast in the chemical regime.

^cGlobal percentage of GEOS-Chem tropospheric and stratospheric gridboxes for which the chemical regime is selected. Values are for August 1 2013 sampled every 6 hour.



435

Figure 1. Minimum of cost function Z_2 (global fraction of chemical species treated as fast) as a function of the number N of blocks used to group the species for mechanism reduction. Values were computed using the GEOS-Chem troposphere + stratosphere simulation on the first days of February, April, August and November 2013, over 24 hours and sampled every 6 hours. Shaded area shows the standard deviation of the cost function minimum computed for each sample.

440




Figure 2. Speed-up of the chemical computation as a function of the number M of chemical mechanism subsets (chemical regimes) used in the coupled implicit solver of the GEOS-Chem model for adaptive simulation of the troposphere and stratosphere. Top: Minimum of cost function Z_2 (global fraction of chemical species treated as fast) as a function of the number of chemical regimes. Bottom: Percentage of model gridboxes that can be represented by the M chemical regimes without adjustment (see Equation 5 and related text). Dashed lines show the values for $M=20$. 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).

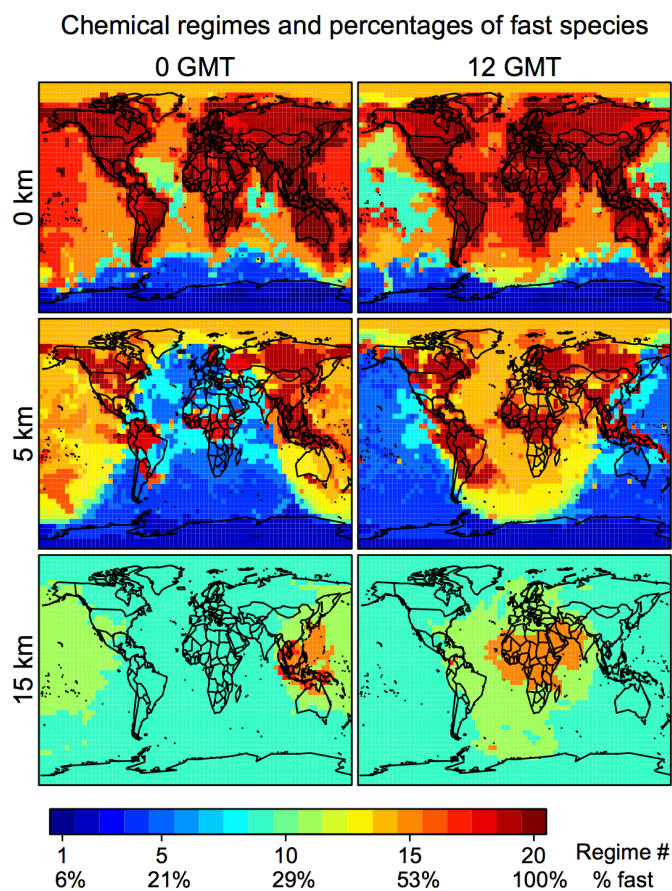



Figure 3. Chemical mechanism complexity needed in different regions of the atmosphere. The Figure identifies the chemical regime from Table 2 needed to simulate a given GEOS-Chem gridbox on August 1 2013 at 0 and 12 GMT. The percentage of the 228 species treated as fast (requiring coupled implicit solution) in that chemical regime is shown on the colorbar and more details are in Tables 1 and 2. Results are shown for different altitudes and using a threshold δ of 100 molecules $\text{cm}^{-3} \text{s}^{-1}$.



455 **Figure 4.** Performance and accuracy of the adaptive chemical mechanism reduction method for different rate thresholds δ
 (molecules $\text{cm}^{-3} \text{s}^{-1}$) to separate fast and slow species. The performance is measured by the reduction in computing processor
 unit (CPU) time for the chemical operator, and the accuracy is measured by the median relative root mean square (RRMS)
 error for species concentrations relative to a global GEOS-Chem simulation for the troposphere and stratosphere using the
 full chemical mechanism (228 species treated as fast). The second x axis gives the global fraction of species that need to be
 460 treated as fast depending on the value of δ . The number of blocks (N) is 12 and the number of chemical regimes (M) is 20.

Median RRMS error over 2-year simulation


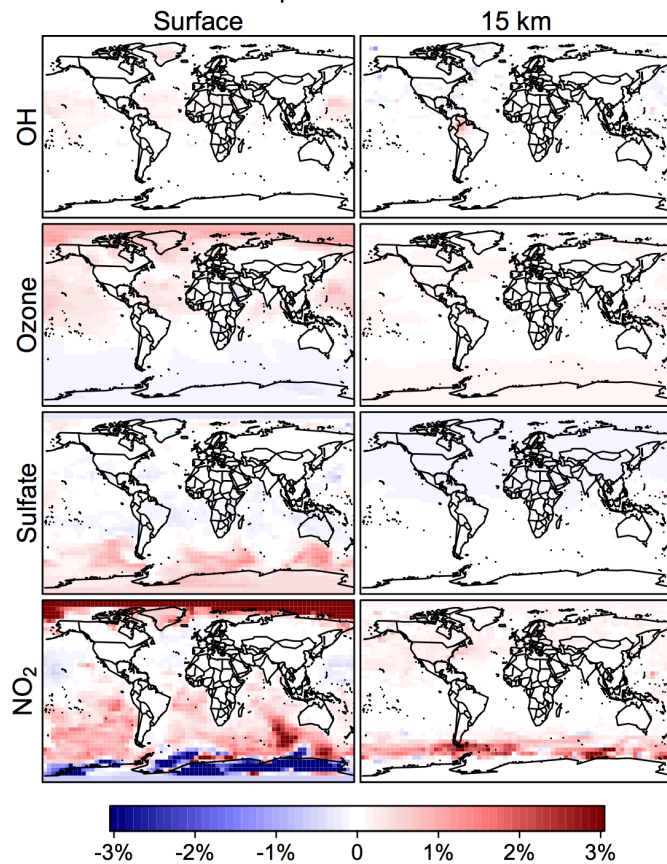


Figure 5. Accuracy of the adaptive reduced chemistry mechanism algorithm over a two-year GEOS-Chem simulation (see text). The accuracy is measured by the 24-hour mean RRMS error on the end day of each month relative to a simulation including the full chemical mechanism. Rate thresholds δ of (a) 100, (b) 500 and (c) 1000 molecules cm⁻³ 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.

Relative error in the adaptive mechanism reduction method



470 **Figure 6.** Relative error from the adaptive mechanism reduction method after two years of simulation in the GEOS-Chem
 global 3-D model for tropospheric-stratospheric chemistry. The figure shows relative differences of 24-h average OH, ozone,
 sulfate and NO₂ concentrations relative to the full-chemistry simulation on the last day of the two-year simulation (2013-
 2014). The relative error for surface NO₂ can be up to $\pm 10\%$ in polar regions. The calculation uses a rate threshold $\delta = 500$
 molecules $\text{cm}^{-3} \text{s}^{-1}$ to partition the species between fast and slow. The number of blocks (N) is 12 and the number of chemical
 475 regimes (M) is 20.