



RESEARCH ARTICLE

10.1029/2022MS003293

Key Points:

- An updated version 3.0.0 of the Kinetic Pre-Processor (KPP) integrator of chemical kinetics for atmospheric models has been developed
- KPP 3.0.0 features an adaptive solver option for reducing chemical mechanisms locally and on the fly where full complexity is not needed
- The adaptive solver implemented in the global GEOS-Chem model shows a 32% speedup with errors less than 1% for key tropospheric species

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

H. Lin,
hplin@seas.harvard.edu

Citation:

Lin, H., Long, M. S., Sander, R., Sandu, A., Yantosca, R. M., Estrada, L. A., et al. (2023). An adaptive auto-reduction solver for speeding up integration of chemical kinetics in atmospheric chemistry models: Implementation and evaluation in the Kinetic Pre-Processor (KPP) version 3.0.0. *Journal of Advances in Modeling Earth Systems*, 15, e2022MS003293. <https://doi.org/10.1029/2022MS003293>

Received 14 JUL 2022

Accepted 21 JAN 2023

© 2023 The Authors. Journal of Advances in Modeling Earth Systems published by Wiley Periodicals LLC on behalf of American Geophysical Union. This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs License](#), which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

An Adaptive Auto-Reduction Solver for Speeding Up Integration of Chemical Kinetics in Atmospheric Chemistry Models: Implementation and Evaluation in the Kinetic Pre-Processor (KPP) Version 3.0.0

Haipeng Lin¹ , Michael S. Long¹ , Rolf Sander² , Adrian Sandu³, Robert M. Yantosca¹ , Lucas A. Estrada¹, Lu Shen⁴, and Daniel J. Jacob¹

¹John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA, ²Air Chemistry Department, Max Planck Institute for Chemistry, Mainz, Germany, ³Department of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA, ⁴Department of Atmospheric and Oceanic Sciences, School of Physics, Peking University, Beijing, China

Abstract Kinetic integration of large and stiff chemical mechanisms is a computational bottleneck in models of atmospheric chemistry. It requires implicit solution of the coupled system of kinetic differential equations with time-consuming construction and inversion of the Jacobian matrix. We present here a new version of the Kinetic Pre-Processor (KPP 3.0.0) for fast integration of chemical kinetics featuring a range of improvements over previous versions in performance, diagnostics, versatility, and community openness. KPP 3.0.0 includes a new adaptive auto-reduction solver to decrease the size of any mechanism locally and on the fly under conditions where full complexity is not needed, by partitioning species as “fast” or “slow” based on their local production and loss rates. Previous implementations of this adaptive solver suffered from excessive overhead in the repeated construction of the local Jacobian matrix or were hard-wired to specific mechanisms. Here we retain the general applicability of the method to any mechanism and avoid overhead by using pre-computed Jacobian matrix terms for the full mechanism and cropping the matrix locally to remove the slow species with no change in memory allocation. We apply this adaptive solver within KPP 3.0.0 to the GEOS-Chem global 3-D model of atmospheric chemistry and demonstrate a 32% reduction in solver time while maintaining a mean error lower than 1% for key species in the troposphere.

Plain Language Summary Calculating chemical evolution in global atmospheric chemistry models is computationally expensive because the chemical mechanisms typically include hundreds of species to account for all conditions from urban to remote. However, the full chemical complexity is not needed under most conditions. Here we have developed an adaptive auto-reduction chemical solver that reduces any mechanism on the fly depending on local conditions and without significant computational overhead. We apply this adaptive solver as an option in a new version 3.0.0 of the Kinetic Pre-Processor (KPP) chemical solver software package that also includes a number of updates relative to previous versions. The adaptive solver achieves a 32% reduction in solver time in a global model simulation while incurring less than 1% average errors for key species.

1. Introduction

Modeling atmospheric chemistry is a grand computational challenge. Current global 3-D models of oxidant-aerosol chemistry use chemical mechanisms that may involve hundreds of coupled chemical species with lifetimes ranging from less than a second to many years. Chemical evolution in such a mechanism is computed by solving a large, stiff system of coupled non-linear ordinary differential equations (ODEs) expressing the chemical kinetics of individual species. Implicit solvers are required to accommodate the coexistence of short and long time constants but are computationally expensive because of the need for repeated construction and inversion of the Jacobian matrix (Brasseur & Jacob, 2017). Chemical integration often dominates the overall computational cost of global 3-D atmospheric chemistry model simulations, even in massively parallel environments or using graphics processing units (GPUs) (Alvanos & Christoudias, 2017; Dawson et al., 2022; Eastham et al., 2018; Zhuang et al., 2020).

Considerable research has gone into devising algorithms to speed up chemistry solvers. A common strategy is to split the mechanism species by time scales in order to decrease the stiffness of the system (Djouad & Sportisse, 2002; Gong & Cho, 1993; Young & Boris, 1977), but this tends to be mechanism-specific and is difficult to apply in global models because of the wide range of conditions that may be experienced. Wholesale reduction of the mechanism, such as in the Super-Fast mechanism used in some climate models (Brown-Steiner et al., 2018), may lead to large errors (Kelp et al., 2022) and incorrect chemical responses to perturbations. Machine learning methods can in principle speed up chemical integration by orders of magnitude but have met with little success because of the large dimensionality of the problem resulting in error growth, in addition to requiring re-training and re-evaluation after even minor updates to the chemical mechanism (Keller & Evans, 2019; Kelp et al., 2020, 2022).

A promising approach for global models is to recognize that the full complexity of the mechanism is not needed everywhere. For example, reactions involving volatile organic compounds (VOCs) and their short-lived products typically account for much of the complexity but may be unimportant outside of continental boundary layers where the VOCs are emitted. Jacobson (1995) thus applied separate mechanisms in a global model for the urban boundary layer, the global troposphere, and the stratosphere. However, fixed geographical separation between domains can result in errors and inefficiencies by not accounting for the interactions at chemical boundaries between domains (Rastigejev et al., 2007) and not allowing for a continuum of chemical regimes from source regions to the remote atmosphere. Santillana et al. (2010) developed an adaptive mechanism reduction method in which the size of the mechanism is adjusted at each grid cell and time step, classifying species as fast (coupled) or slow (uncoupled) on the basis of their total production and loss rates. However, the overhead involved in local definition of the reduced mechanism offset the computational gains. Sander et al. (2019) and Shen et al. (2020, 2022) improved the method by pre-compiling a limited ensemble of chemical sub-mechanisms and selecting the most appropriate sub-mechanism for use based on local conditions at each timestep, thus avoiding the overhead. Shen et al. achieved a 30%–50% reduction in computational cost compared to the full parent mechanism in a global simulation of the troposphere and stratosphere, but the selection of sub-mechanisms had to be customized to the parent mechanism.

Here we develop a mechanism-agnostic, ready-to-use method for adaptive auto-reduction of any chemical mechanism and implement it as an option in a new version 3.0.0 of the Kinetic Pre-Processor (KPP). Originally developed by Damian et al. (2002) and Sandu and Sander (2006), KPP is a software tool that automatically generates code to efficiently integrate chemical mechanisms. KPP takes in a set of human-readable input files describing the mechanism and generates Fortran 90, C, or MATLAB code to solve the corresponding system of ODEs using any of a suite of integration methods. KPP is used in many atmospheric chemistry models including MECCA within MESSy (Jöckel et al., 2010; Sander et al., 2019), WRF-Chem (Fast et al., 2006; Grell et al., 2005), the forward and adjoint GEOS-Chem models (Henze et al., 2007), and the adjoint for the CMAQ model (Hakami et al., 2007; Zhao et al., 2020). Our new version KPP 3.0.0 incorporates several performance and diagnostic updates over the previous version KPP 2.1 (Sandu & Sander, 2006) in addition to the adaptive solver option.

2. Adaptive Solver for Chemical Kinetics

Atmospheric chemistry models alternate chemical integration and transport calculations through operator splitting (Brasseur & Jacob, 2017). The chemistry solver is called for a time interval of length Δt , referred to as the external time step, and returns a vector of updated concentrations $\mathbf{C}$ at the end of that time step to be operated on by transport. The kinetic integration of the mechanism by the chemistry solver is done over internal time steps $h \leq \Delta t$ to reach the desired accuracy.

The chemistry solver integrates a system of N coupled nonlinear first-order ODEs of the form

$$\frac{dC_i}{dt} = P_i(\mathbf{C}) - L_i(\mathbf{C}) \quad (i = 1, 2, \dots, N) \quad (1)$$

where N is the number of coupled species in the mechanism, $\mathbf{C}$ is the vector of species concentrations of dimension N , and $P_i(\mathbf{C})$ and $L_i(\mathbf{C})$ are the production and loss rates of species i that depend on the concentrations of other species in the mechanism through the reaction rate expressions.

The external time step in a global model is typically $\sim 10^3$ s, but many species in the mechanism have lifetimes ~ 1 s or shorter. An explicit solver would require internal time steps shorter than the lifetime of the shortest-lived species in order to achieve stability, but this is not computationally practical. An implicit solver is required. The

simplest such solver is the first-order backward Euler method, which approximates the solution to Equation 1 over the internal time step h with

$$\mathbf{C}(t+h) = \mathbf{C}(t) + \mathbf{s}(\mathbf{C}(t+h))h \quad (2)$$

where we define the net source term $\mathbf{s} = \mathbf{P} - \mathbf{L}$ as a vector of functions ($P_i - L_i$). Solving Equation 2 for the unknown quantity $\mathbf{C}(t+h)$ using the Newton-Raphson method requires the repeated construction and inversion of the $N \times N$ Jacobian matrix $\mathbf{J}$:

$$\mathbf{J} = \frac{\partial \mathbf{s}}{\partial \mathbf{C}} \quad (3)$$

The construction and inversion of this Jacobian matrix is computationally expensive. Higher-order solvers generally used in atmospheric chemistry models, such as Rosenbrock (Hairer & Wanner, 1991; Rosenbrock, 1963; Sandu et al., 1997) or Gear (Jacobson & Turco, 1994), similarly require the repeated construction and inversion of the Jacobian over internal time steps. The Jacobian is typically $\sim 90\%$ sparse allowing for efficient sparse-matrix inversion methods (Sandu et al., 1997), so that the overall cost of construction and inverting the Jacobian scales as $\sim N$ rather than a higher power.

A way to reduce the dimensionality N of the problem is to split the mechanism into “fast” species for which the coupled implicit solution is necessary and “slow” species that may be solved independently over the external time step using a fast explicit method. Young and Boris (1977) and Gong and Cho (1993) separate species into fast and slow based on their lifetimes compared with the integration time step. However, the separation results in non-conservation of mass because the reaction rates are not computed consistently. This may not be of consequence in a regional model (as used in those applications) where the domain is ventilated by the boundary conditions, so that errors do not accumulate, but it is more problematic in a global model. Santillana et al. (2010) separated instead “fast” and “slow” species on the basis of their production and loss rates, with the slow species having sufficiently low rates that their non-conservation of mass would be inconsequential. This is more relevant for global models where concentrations and rates of short-lived VOCs become very small and they have negligible influence on other species outside of their source regions. Shen et al. (2020, 2022) used the same approach to partition species between fast and slow.

Here we also follow the partitioning method of Santillana et al. (2010). At the beginning of each external time step we calculate $\mathbf{P}$ and $\mathbf{L}$ and classify species i as fast if $\max(P_i, L_i) > \delta$ and as slow otherwise, where δ is a user-selected partitioning threshold. Fast species are assigned to the coupled implicit solver as a subset (sub-mechanism) of the full mechanism, while the evolution of slow species over the external time step beginning at t_0 is calculated using an explicit first-order approximation with first-order loss rate constant $k_i(t_0) = L_i(t_0)/C_i(t_0)$:

$$C_i(t_0 + \Delta t) = \frac{P_i(t_0)}{k_i(t_0)} + \left(C_i(t_0) - \frac{P_i(t_0)}{k_i(t_0)} \right) e^{-k_i(t_0)\Delta t} \quad (4)$$

In the Santillana et al. (2010) implementation, the Jacobian had to be reconstructed locally at every external time step for the identified subset of fast species and that incurred large overhead, canceling the benefit of the method. Here we avoid the overhead by taking advantage of the pre-computed Jacobian matrix terms for the full mechanism in the KPP solver to simply remove rows and columns corresponding to the slow species. This is explained in Section 3.2 and is the key new development to make the method computationally practical.

The partitioning threshold δ is set prior to integration and is tuned to balance performance and accuracy. Previous work (Santillana et al., 2010; Shen et al., 2020, 2022) considered that in a typical tropospheric chemistry mechanism, much of the coupling is associated directly or indirectly with cycling of the hydroxyl radical (OH). OH has a daytime concentration of $\sim 10^6$ molecules cm^{-3} and a lifetime ~ 1 s, so its production and loss rates are $\sim 10^6$ molecules $\text{cm}^{-3} \text{ s}^{-1}$. A species with production and loss rates that are several orders of magnitude smaller would not be expected to contribute significantly to the coupling. They found $\delta \sim 10^2 - 10^3$ molecules $\text{cm}^{-3} \text{ s}^{-1}$ to be adequate after testing for performance and accuracy.

Here we include the option to dynamically define δ instead of specifying a uniform value over the domain, to account for rates varying with local conditions. This is done by identifying a target species which is considered central to the mechanism and scaling its production and loss rates to define a local partitioning threshold:

$$\delta = \alpha_{\text{target}} \max(P_{\text{target}}, L_{\text{target}}) \quad (5)$$

```
#INTEGRATOR    rosenbrock_autoreduce
#LANGUAGE      Fortran90
#MINVERSION    3.0.0
#AUTOREDUCE    on

#FAMILIES
POx : O3 + NO2 + HNO3;
LOx : O3 + NO2 + HNO3;

#DEFVAR
O3    = IGNORE;
NO    = IGNORE;
...

#EQUATIONS
O3  + hv = 2OH  + O2 : PHOTOL(1);
NO  + O3 = NO2  + O2 : ARR(1.8d-12, -1370.0);
CO  + OH = CO2  + H2O : 2.4d-13;
HO2 + NO = OH   + NO2 : ARR(3.7d-12, 240.0);
NO2 + hv = NO   + O3 : PHOTOL(2);
HO2 + HO2 = H2O2 + O2 : ARR(1.2d-13, 749.0);
H2O2 + hv = 2OH   : PHOTOL(3);
CH4 + OH = CO   + H2O : ARR(3.1d-12, -187.0);
...
```

KPP Options
Integrator, target language, and other code generation options

Production and Loss Families (Optional)

Definition of species
"IGNORE" if mass balance checks are not required
Otherwise specify stoichiometric composition

List of reactions
Reactants = Products : Rate constant
Fortran 90 expressions are supported in rate constants

Figure 1. Example KPP input file. The KPP input file includes options for code generation, adaptive auto-reduction, diagnostics for production and loss rates of chemical families, and a list of species and reactions. Reaction rate constants can be specified as expressions in the target language (in this case Fortran 90), here showing calls to functions calculating photolysis and Arrhenius rate expressions.

where P_{target} and L_{target} are the local production and loss rates of the target species, and $\alpha_{\text{target}} \ll 1$ is a user-selected coefficient that depends on the target species but is otherwise fixed for the model domain. For example, a model may use OH as target species for daytime, and NO₂ or the nitrate radical (NO₃) for nighttime. When using OH as a target species and with $\max(P_{\text{OH}}, L_{\text{OH}}) \sim 10^6 \text{ molecules cm}^{-3} \text{ s}^{-1}$, a value $\alpha_{\text{OH}} \sim 10^{-4} - 10^{-3}$ would correspond to the criteria for δ used by Santillana et al. (2010). But $\max(P_{\text{OH}}, L_{\text{OH}})$ can in fact vary over orders of magnitude depending on pressure, UV flux, and other factors, and our local specification of δ accounts for this variability. We find that this specification of dynamic threshold improves accuracy with no significant overhead since P_{target} and L_{target} are computed at the beginning of each external time step in any case.

We also include two new options in the algorithm. First is to force individual species to remain in the coupled implicit solver even if $\max(P_i, L_i) \ll \delta$. As we will see, this may be helpful for inorganic halogen species that cycle between radical and non-radical forms across sunrise/sunset. Second is to include an "append" functionality in the algorithm so that a species initially diagnosed as slow at the beginning of the external time step can be transferred into the coupled implicit solver if it becomes fast over the course of the integration. This increases accuracy with minimum overhead.

3. Kinetic Pre-Processor (KPP) Version 3.0.0

3.1. KPP 3.0.0 Overview

KPP (Damian et al., 2002; Sandu & Sander, 2006) generates code for solving the chemical kinetics for a given chemical mechanism defined by a list of species, reactions, and rate constants. It is designed for speed by exploiting sparse matrix algebra and pre-computation of terms. We have made several improvements to KPP relative to the previous version 2.1 (Sandu & Sander, 2006), including the adaptive solver option. We present these improvements as KPP version 3.0.0, available for download from <https://github.com/KineticPreProcessor/KPP/> (<https://doi.org/10.5281/zenodo.7308373>) with detailed documentation at <https://kpp.readthedocs.io>. Existing features in the previous versions of KPP are maintained in KPP version 3.0.0. Our new version has zero numerical differences compared to KPP version 2.1 when using the same configuration as verified using the "saprc2006.kpp" example.

KPP takes in as input one or more text files, with an example shown in Figure 1. The input is not necessarily in one single file; it may be split into several files for readability. The files describe the chemical mechanism, the choice of numerical solver (e.g., Rosenbrock), target language (i.e., Fortran 90, C, or MATLAB), floating point

type (single- or double-precision), production and loss diagnostics for selected chemical families (optional), and whether to use the adaptive solver option (optional). Some inputs are not part of the KPP code generation and are instead left for users to adjust at runtime, including convergence criteria (absolute and relative tolerance), numerical order of the solver, and adaptive solver options, in order to enable the user to experiment with different thresholds for performance and accuracy.

Based on the specifications in the KPP input file(s), KPP creates files in the target language containing a description of the system of ODEs (number of species, reactions, and a numbered list of species), code to calculate the time derivatives, Jacobian matrix, and solution by back-substitution, along with a copy of the numerical solver (such as Rosenbrock) and supporting routines for sparse linear algebra. This set of files can be either run standalone as a box model or can be included in a 3-D model to update concentrations locally over external time steps.

Auto-reduction of the mechanism with the adaptive solver described in Section 2 is an option in KPP 3.0.0 using the Fortran 90 language, enabled in the configuration file by `#AUTOREDUCE` on and specifying the corresponding integrator (e.g., `#INTEGRATOR rosenbrock_autoreduce`). This sets up the capability for the user to reduce the mechanism locally through specification of the partitioning threshold δ between fast and slow species, listing any species for which this partitioning should not be applied. These specifications are done at runtime for flexibility. Even when the auto-reduction solver is used, mechanism auto-reduction can be disabled at runtime, defaulting to the behavior of the original integrator, to check for accuracy. A test case box model for auto-reduction is included as part of KPP 3.0.0 documentation.

In addition to the option for adaptive mechanism auto-reduction, several additional features and improvements were added to KPP 3.0.0 relative to the previous version 2.1 (Sandu & Sander, 2006). These include:

1. *Redeployment of KPP source code, continuous integration, and documentation for community development.* KPP 3.0.0 source code has been redeployed on GitHub for community development. The GitHub repository incorporates continuous integration (CI) tests which automatically compile the KPP source code to build and run combinations of sample chemistry mechanisms and integrators into box models at every code revision. This helps to ensure that new features and updates added to KPP do not break existing functionality and that numerical results for existing configurations are not affected. The documentation has also been relocated to <https://kpp.readthedocs.io> where it is automatically built with each code revision on GitHub.
2. *New diagnostics.* The following diagnostic features have been added to KPP 3.0.0:
 - *Production and loss rates for chemical families.* KPP 3.0.0 allows for the definition of families of chemical species for computing production and loss for that family, ignoring interconversion reactions within the family. This is useful in models for example, to keep track of odd oxygen (Bates & Jacob, 2020) or nitrogen oxides. Families are defined in the `#FAMILIES` section of the KPP input file.
 - *Stoichiometric numbers.* The stoichiometric numbers of all reactions in the mechanism are now available in the `CalcStoichNum` subroutine in the KPP-generated code. This feature is used to calculate the importance of chemical species in a mechanism for the skeletal mechanism reduction in Sander et al. (2019).
 - *Individual reaction rates and time derivatives.* The reaction rates and their time derivatives are now available as optional outputs from the KPP-generated code.
3. *Addition of new solvers.* Several solvers have been added as options for integrating chemical kinetics including VODE (Brown et al., 1989), SDIRK (Hairer & Wanner, 1991), 3-stage Runge-Kutta, and forward and backward Euler methods.
4. *Addition of new rate law functions.* Rate law functions for three-body reactions using the formulas proposed by JPL (<https://jpldataeval.jpl.nasa.gov>) and IUPAC (<https://iupac.aeris-data.fr/>) have been added to the built-in rate laws in KPP 3.0.0. Rate law functions are not limited to those built in KPP and can be added by including extra source code files in the KPP input.
5. *Miscellaneous performance improvements in Fortran 90.* KPP 3.0.0 optimizes for Fortran 90 performance by applying several guidelines in coding, particularly in reaction rate calculations that are computed repeatedly in loops:
 - *Unifying number precision.* Previous inputs to KPP used both single- and double-precision numbers. KPP input files and code now do not mix number precision to avoid conversion, which loses precision and costs computational time.

- *Switch to control update of reaction rate constants.* Previously, KPP called subroutines to update reaction rate constants at every internal time step, but this is computationally expensive and may not be needed in most model applications where the variables affecting rate constants are not updated between internal time steps. An optional switch has been added so that rate constants are updated only at the beginning of the external time step.
 - *Optimized rate law functions.* Rate law functions have been split to avoid computation of unnecessary terms. For example, previously a single function with Arrhenius temperature dependence was used for all reactions: $\text{ARR}_{\text{abc}}(A, B, C) = A * \exp(B/T) * (300/T)^C$. Many reactions may have $B = 0$ or $C = 0$ but the expressions $300/T$ or $\exp(0)$ are still needlessly computed. Separate rate functions such as $\text{ARR}_{\text{ab}}(A, B) = A * \exp(B/T)$ have been added, leading to a 44% performance improvement of reaction rate computations in a full-chemistry GEOS-Chem model run.
 - *Avoiding conditionals and optional arguments.* Conditional clauses such as IF, ELSE, and SELECT CASE, and testing for optional arguments add significant computational cost if called thousands of times. If a conditional clause or optional argument is present in a frequently called subroutine, the subroutine is split into different functions for each case.
 - *Thread safety for generated code.* The code generated is now thread safe, so that calls for updating rate constants and running the integrator can be placed in an OpenMP parallel loop for parallelization.
 - *Improved expressions for vector and array functions.* Several functions for basic vector and array operations originally used reference BLAS (Basic Linear Algebra Subprograms) implementations. These have been replaced with Fortran 90 expressions to allow for compilers to better optimize the code.
6. *Miscellaneous code improvements in KPP.* The C source code of KPP has been improved so that no compiler warnings are generated. A more consistent memory allocation helps to avoid buffer overflow problems. The KPP language is now parsed by `bison` instead of `yacc`.

3.2. Adaptive Solver Implementation in KPP 3.0.0

We implemented the adaptive solver as an option within KPP's Fortran 90 version of the Rosenbrock solver using sparse matrix algebra. KPP is computationally efficient because the functions to compute the time derivatives for each species and the Jacobian matrix (expressed in terms of reaction rates and species concentrations), along with all sparse matrix algebra routines, are pre-generated and use fixed indices to access species vectors and matrices. This means that the problem size and memory space is fixed at compile time, avoiding expensive memory allocation operations. However, this also yields a fixed problem structure that is difficult to manipulate. One major source of overhead in locally defined sub-mechanisms as in Santillana et al. (2010) is associated with repeatedly re-allocating and de-allocating memory to accommodate changing problem sizes for each sub-mechanism (Shen et al., 2020). Our adaptive solver implementation in KPP 3.0.0 uses a mapping operation to project the full mechanism into sub-mechanisms, thus reusing the same memory space to avoid expensive resizing operations.

Figure 2a shows the sparse Jacobian data as stored within KPP. In this example mechanism, there are three species (NVAR) and eight non-zero entries in the Jacobian (LU_NONZERO) of the full mechanism. The row and column indices of these eight non-zero entries in Jacobian matrix are correspondingly specified in LU_IROW, LU_ICOL in row-compressed form. At the beginning of every external time step, the production and loss rates of each species are calculated and compared to the partitioning threshold to separate species into fast and slow. Figure 2b shows an example where species B is slow. The entries in the Jacobian corresponding to B no longer need to be computed, and a mapping operation is performed: JVS_MAP corresponds to the non-zero Jacobian matrix entries still present in the sub-mechanism consisting of fast species. The smaller sub-mechanism can now be described by two species (rNVAR) and four non-zero entries in the reduced Jacobian (cNZERO) with indices described by cLU_IROW and cLU_ICOL. The data structure of the smaller sub-mechanism is identical to the full mechanism, and the same routines are used to solve it, without the need to generate extra code, or resizing memory.

The mechanism auto-reduction is performed once at the beginning of every external time step. The set of fast and slow species are established according to the runtime options for the partitioning threshold δ and the list of species to be excluded from partitioning and kept in the fast subset under all conditions (`keepSp-`




Figure 2. Mechanism auto-reduction in KPP. Panel (a) shows the data structure of the sparse Jacobian data within KPP for a sample 3-species mechanism. If species B is diagnosed as slow, then the corresponding rows and columns of the Jacobian are no longer calculated (panel (b)), and the indices pointing to the sparse Jacobian data (LU_IROW, LU_ICOL) are adjusted to remove the slow species (cLU_IROW, cLU_ICOL) through a mapping array (JVS_MAP) while preserving the sparse matrix in row-compressed form. The result is a reduced mechanism with the same data structure as the original one, but with smaller dimensions.

cActive). Based on the list of species in the fast set, the mapping (JVS_MAP) from the full mechanism to the fast sub-mechanism is created. Because KPP generates hard-coded source code to compute each term of the Jacobian matrix and back-substitution for computational efficiency, two logical control vectors, DO_JVS and DO_SLV, are created to skip computation of terms corresponding to slow species as these are no longer necessary.

The separation between fast and slow species is controlled by the initial conditions at the beginning of the external time step, but an optional “append” function is added to account for an initially slow species becoming fast over the course of the internal time steps. This function appends new species to the fast sub-mechanism and adjusts the logical control vectors if these species are initially partitioned as slow but their production or loss rate exceed the partitioning threshold over the course of the internal time stepping. Diagnosing this has little overhead, because the production and loss rates of all species are already computed at every internal time step.

We used a box model integration of the GEOS-Chem chemical mechanism (described below) to analyze the overhead of the adaptive solver implemented within KPP. By forcing the adaptive solver routines to run with a threshold of 0, we determined that the KPP overhead added by the auto-reduction is 10%–16%. The main source of overhead is the copying of data between the full and reduced sub-mechanism (Figure 2), where the worst-case scenario is when all species are partitioned as fast and all species' data need to be copied between the full to the “reduced” data structures. Profiling tests show that other steps such as the partitioning of species between fast and slow, or the first-order approximation for slow species, contribute negligible overhead.

We also used the box model to verify that our adaptive solver introduces minimal mass non-conservation compared to the full solver. Integrating the box model over four sets of initial conditions from the GEOS-Chem model over 10,000 time steps shows that the total masses of odd oxygen, reactive nitrogen, chlorine, and bromine (defined in Table S1 in Supporting Information S1) are conserved within 0.30% in the standard configuration of the adaptive solver as compared with the full mechanism solver. Figure S2 in Supporting Information S1 shows

the geographical distribution of mass non-conservation for these four families at a snapshot in time in a global GEOS-Chem model simulation, again comparing the adaptive solver in its standard configuration to the full solver. Mass conservation is within 0.01% for odd oxygen and reactive nitrogen and at worst within 0.2% for bromine.

4. Adaptive Mechanism Auto-Reduction in the GEOS-Chem Model Using KPP Version 3.0.0

We demonstrate the adaptive solver capability in KPP version 3.0.0 with the GEOS-Chem global 3-D atmospheric chemistry model version 13.4.0 (Bey et al., 2001; <https://doi.org/10.5281/zenodo.6511970>). The chemical mechanism includes comprehensive oxidant-aerosol chemistry in the troposphere and the stratosphere with 291 chemical species and 913 reactions. Recent updates to the mechanism include Cl-Br-I tropospheric halogen chemistry (Wang et al., 2021), isoprene chemistry (Bates & Jacob, 2019), aromatic chemistry (Bates et al., 2021), hydroxymethanesulfonate cloud chemistry (Moch et al., 2020), and NO_y cloud and aerosol chemistry (Holmes et al., 2019). Heterogeneous sulfur chemistry that was previously simulated with a separate module was brought into KPP in GEOS-Chem version 13.4.0 with the addition of new rate functions.

GEOS-Chem has been structured to interact with the KPP-generated solver code through the FlexChem interface which prepares data for the solver code, executes the code, and retrieves concentrations at the end of the time step. FlexChem allows GEOS-Chem to use modules outside of KPP for computing reaction rates, for example, interfacing with Fast-JX for photolysis rates (Bian & Prather, 2002; Mao et al., 2010). FlexChem also includes a derived type object, `State_Het`, which passes state variables from GEOS-Chem model for calculating heterogeneous chemistry reaction rates including cloud liquid water content, aerosol size distribution, pH, and alkalinity. This derived type object holds common intermediate quantities necessary for heterogeneous reaction rate computations, such as aerosol area, avoiding repeated computation and memory use. FlexChem also allows GEOS-Chem users to modify the chemical mechanism input into KPP without modifying the GEOS-Chem source code.

We evaluate the accuracy and computational performance of the adaptive solver in KPP version 3.0.0 using global GEOS-Chem simulations at 2° × 2.5° resolution with 72 vertical levels extending up to 0.1 hPa. The model is driven by the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2) meteorological fields from the NASA Global Modeling and Assimilation Office (GMAO). The external time step for the chemistry solver is 20 min. All simulations are conducted on the same single-node hardware with 24 Intel Cascade Lake physical cores (Intel(R) Xeon (TM) Platinum 8268 CPUs with a base clock speed of 2.90 GHz, no hyper-threading logical cores), 100 GB of RAM, and a high-performance Luster parallel file system. The model was compiled using Intel(R) Fortran Compiler (ifort) version 2021.2.0.

We select as the standard configuration of the adaptive solver within the GEOS-Chem model a dynamically defined threshold with OH as target species during daytime, and NO₂ during nighttime, with coefficients $\alpha_{\text{OH}} = 5 \times 10^{-5}$ and $\alpha_{\text{NO}_2} = 1 \times 10^{-4}$. We find that using NO₃ as an alternative nighttime target species results in higher accuracy but lower speed-up (Table S3 in Supporting Information S1). We do not force any species to remain in the implicit KPP solver as fast and we do not use the append functionality. We find that we achieve a net 32% reduction in integration time in this standard configuration.

Figure 3 shows the percentage of species partitioned as fast (and hence retained in the KPP integration) using the adaptive solver's standard configuration. Starting from the ensemble of 291 species in the GEOS-Chem full mechanism, we find that over 60% are partitioned as fast in surface air over land, reflecting VOC sources, whereas only 10%–50% are partitioned as fast over the ocean. The fraction of retained species decreases with altitude, and fewer than 40% are partitioned as fast in the stratosphere where VOC chemistry is mainly limited to methane. Fewer species are partitioned as fast at night than in daytime. These results are consistent with Shen et al. (2020).

We evaluate the accuracy of the adaptive solver (AS) relative to the full mechanism solver over the global GEOS-Chem domain using the relative root mean squared (RRMS) error metric. For each species i , the RRMS error is:

$$\text{RRMS}_i = \sqrt{\frac{1}{N_i} \sum_{j=1}^{N_i} \left(\frac{C_{i,j,\text{full}} - C_{i,j,\text{AS}}}{C_{i,j,\text{full}}} \right)^2} \quad (6)$$




Figure 3. Percentage of GEOS-Chem species retained in the fast sub-mechanism when using the adaptive solver. The full GEOS-Chem mechanism has 291 species to describe tropospheric and stratospheric oxidant-aerosol chemistry. Only a fraction of species is retained as fast in the KPP solver, while the other species are solved individually using Equation 4. Results are shown for a snapshot in time on 1 July 2014, 12:00 UTC at different altitudes. The adaptive solver uses a dynamically defined threshold (Equation 5) with target species OH in daytime ($\alpha_{\text{OH}} = 5 \times 10^{-5}$) and NO_2 at night ($\alpha_{\text{NO}_2} = 1 \times 10^{-4}$).

where $C_{i,j,\text{full}}$, $C_{i,j,\text{AS}}$ are the concentrations of species i in grid box j for simulations without and with the adaptive solver. The RRMS is computed over the ensemble N_i of ordered grid boxes that account for 99% of the total mass of species i in the boundary layer (surface to PBL height from MERRA2), free troposphere (boundary layer height to tropopause), and stratosphere, respectively, and where $C_{i,j,\text{full}}$ is greater than 10 molecules cm^{-3} .

Figure 4a shows the mean errors at the end of a 1-year simulation for species in different chemical categories (Table S4 in Supporting Information S1) and Figures 4b–4d shows the time evolution of errors over the 1-year simulation for all species in the standard configuration of the adaptive solver within the GEOS-Chem model, respectively. The mean error for most categories is below 1% within the boundary layer and the free troposphere. There is no error growth in the troposphere but there is some in the stratosphere. Figure 5 shows the geographical distribution of relative errors (absolute errors shown in Figure S6 in Supporting Information S1) for ozone, OH, and NO_2 at the surface and 500 hPa for the end of the 1-year simulation. Errors are generally lower than 1% except for OH at high latitudes and NO_2 over the Southern Ocean, where errors are up to 3% for OH and 6 ~ 10% for NO_2 . Errors over land for these species are minimal.

The largest errors in Figure 4 are found for inorganic halogen radicals and their reservoirs. Halogen radicals cycle rapidly during the day but become locked in reservoirs at sunset, to be released again at sunrise (Wang et al., 2021). Partitioning between fast and slow species at sunrise/sunset leads to large errors and mass balance issues associated with the use of the first-order approximation (Equation 4). These problems were previously identified by Shen et al. (2020, 2022). Long residence times in the stratosphere compound the problem and drive slow error growth for other species. The impact is much less in the troposphere where the inorganic halogens are removed by wet deposition.

Our results in Figure 4 indicate that halogen species should be kept in the implicit solver as fast species if the primary interest of the simulation is the stratosphere. Figure S5 in Supporting Information S1 shows the mean errors and time evolution of errors over the 1-year simulation when halogen species are kept as fast in the adaptive solver. In this configuration, the performance improvement of the adaptive mechanism is 23%, as compared to 32% in the standard configuration. Keeping the halogen radicals and their reservoir species as fast prevents large errors from developing in the stratosphere, in addition to avoiding error growth and spikes in errors as shown in Figure 4d. A similar problem though not of the same magnitude is found for hydrogen oxide radicals controlling the concentration of OH in the stratosphere, and this causes high errors in stratospheric sulfate exceeding 10% above 85 hPa. The list of species to be kept in the implicit solver in the GEOS-Chem model can be configured in the *FlexChem* interface and passed to KPP 3.0.0 as a runtime option. An alternative to customizing the list of species to be kept in the implicit solver is to disable the adaptive solver in the stratosphere while retaining it in the troposphere. In this case, the performance improvement from using the adaptive mechanism decreases from 32% to 20%. Such ad hoc adjustments to the fast/slow species partitioning would be mechanism-specific and can be




Figure 4. Accuracy of the adaptive solver in a 1-year GEOS-Chem simulation. Panel (a) shows the mean RRMS errors for species in different categories at the end of the 1-year simulation starting on 1 July 2014. The categories are as defined in the standard GEOS-Chem benchmarking output diagnostics (Table S4 in Supporting Information S1). Panels (b–d) show the time evolution of RRMS errors for all species in the boundary layer, free troposphere, and stratosphere, respectively, with colored lines for species of particular interest.




Figure 5. Accuracy of the adaptive solver for ozone, OH, and NO_2 global distributions. The figure shows the percent errors relative to the full solver for the daily mean concentrations on 1 July 2015, the last day of the 1-year simulation, in surface air and at 500 hPa. The absolute error values corresponding to this figure are shown in Figure S6 in Supporting Information S1.

experimented with by users. We find in GEOS-Chem that they are not needed for standard applications focusing on tropospheric chemistry.

We find in this application that the append functionality (allowing species to switch from slow to fast over internal time steps) does not provide significant error reduction, in particular for the halogen species, and degrades the performance improvement of the adaptive mechanism to 24% instead of 32%. We retain it as an option in the adaptive solver code as it may be helpful for longer external time steps (here 20 min) or other chemical mechanisms.

The adaptive solver uses a dynamically defined threshold (Equation 5) with target species OH in daytime ($\alpha_{\text{OH}} = 5 \times 10^{-5}$) and NO_2 at night ($\alpha_{\text{NO}_2} = 1 \times 10^{-4}$), and does not force any species below that threshold to remain as fast.

Our analysis of the performance and accuracy of the adaptive solver is based on 20-min external time steps for the chemistry solver in the GEOS-Chem standard $2^\circ \times 2.5^\circ$ configuration, but regional models with finer grid resolution would use finer external time steps for operator splitting with transport (Philip et al., 2016). We tested the sensitivity of our results to the size of the external time step within the GEOS-Chem $2^\circ \times 2.5^\circ$ environment as shown in Table S3 in Supporting Information S1. We find that finer external time steps result in higher accuracy but lower speed-up, which makes sense in terms of the overhead required at the beginning of each external time step to define the sub-mechanism and cull the Jacobian, but also the opportunity to update the separation between fast and slow species more frequently for better accuracy. Ultimately, this sensitivity to external time-stepping will need to be tested by users in their own model environment (and their own mechanism).

5. Conclusions

We presented an updated version of the (KPP 3.0.0) to integrate stiff chemical mechanism kinetics typical of atmospheric chemistry models. KPP was originally designed for flexibility and speed. KPP 3.0.0 features several improvements for performance, diagnostics, choice of solvers, and code openness. It includes an adaptive solver capability for mechanism auto-reduction where and when the full mechanism is not needed.

The adaptive solver performs auto-reduction of the chemical mechanism locally and on-the-fly at runtime, by comparing the local production and loss rates of each species with a partitioning threshold (δ). Species with

production and loss rates higher than the threshold are considered fast and are solved as a coupled sub-mechanism within KPP, while other species are considered slow and solved individually by an explicit method. Previous application of this adaptive solver method suffered from large overhead due to the need for local reconstruction of the reduced Jacobian matrix and the associated memory allocation and deallocation. We solved this problem here by using pre-computed Jacobian terms for the full mechanism with a mapping operation to crop rows and columns corresponding to the slow species without changing the memory allocation.

KPP 3.0.0 features additional improvements for performance, diagnostics, versatility, and openness. Improved performance includes more efficient calculation of reaction rates from the KPP rate functions and thread safety for parallelization. New diagnostics include individual reaction rates, production and loss rates for chemical families, and stoichiometric numbers. Improved versatility includes expanded choice of chemical solvers. KPP 3.0 is now hosted on GitHub (<https://github.com/KineticPreProcessor/KPP>) to enable community access, development, and testing.

We evaluated the adaptive solver implemented in KPP 3.0.0 by conducting a 1-year simulation with the global 3-D GEOS-Chem atmospheric chemistry model including 291 species in the full mechanism for oxidant-aerosol chemistry in the troposphere and stratosphere. Results show that a 32% performance improvement in the solver can be achieved with a target error of 1% for key species in the troposphere. Errors in the stratosphere can be larger, driven by halogen chemistry. Lower errors especially in halogen species can be achieved by keeping these species within the fast sub-mechanism but this reduces the performance improvement to 23% and is mainly beneficial in the stratosphere.

The release of KPP 3.0.0 introduces improvements in development infrastructure, diagnostics, and performance, particularly in Fortran 90 applications. However, one of the strengths of the KPP software is the capability to generate code for different programming languages. Development directions for future versions include (a) adding support for modern languages such as Python and Julia; (b) refactoring of the generated code to avoid global data structures for easier parallelization; (c) streamlining inputs and outputs of all integrators for consistency; (d) supporting the adaptive solver option in other integrators and programming languages; and (e) improving interaction and compatibility with the Master Chemical Mechanism (<http://mcm.york.ac.uk>). These improvements will allow KPP to better serve the community as a versatile tool for solving chemical kinetics within atmospheric chemistry models and other applications.

Data Availability Statement

Model Code Availability: Source code for KPP 3.0.0 is available at <https://github.com/KineticPreProcessor/KPP> (Sandu et al., 2022; <https://doi.org/10.5281/zenodo.7308373>). The adaptive solver box model is available at <https://github.com/KineticPreProcessor/KPP-AR-boxmodel/> (Long & Lin, 2022; <https://doi.org/10.5281/zenodo.6791657>). The adaptive solver implementation within the GEOS-Chem atmospheric chemistry model used in this work is available at <https://github.com/jimmiein/geos-chem/tree/staging/autoreducekpp> (The International GEOS-Chem Community, 2022; <https://doi.org/10.5281/zenodo.6791655>).

References

- Alvanos, M., & Christoudias, T. (2017). GPU-accelerated atmospheric chemical kinetics in the ECHAM/MESSy (EMAC) Earth system model (version 2.52). *Geoscientific Model Development*, 10, 3679–3693. <https://doi.org/10.5194/gmd-10-3679-2017>
- Bates, K. H., & Jacob, D. J. (2019). A new model mechanism for atmospheric oxidation of isoprene: Global effects on oxidants, nitrogen oxides, organic products, and secondary organic aerosol. *Atmospheric Chemistry and Physics*, 19(14), 9613–9640. <https://doi.org/10.5194/acp-19-9613-2019>
- Bates, K. H., & Jacob, D. J. (2020). An expanded definition of the odd oxygen family for tropospheric ozone budgets: Implications for ozone lifetime and stratospheric influence. *Geophysical Research Letters*, 47(4), e2019GL084486. <https://doi.org/10.1029/2019GL084486>
- Bates, K. H., Jacob, D. J., Li, K., Ivatt, P. D., Evans, M. J., Yan, Y., & Lin, J. (2021). Development and evaluation of a new compact mechanism for aromatic oxidation in atmospheric models. *Atmospheric Chemistry and Physics*, 21(24), 18351–18374. <https://doi.org/10.5194/acp-21-18351-2021>
- Bey, I., Jacob, D. J., Yantosca, R. M., Logan, J. A., Field, B. D., Fiore, A. M., et al. (2001). Global modeling of tropospheric chemistry with assimilated meteorology: Model description and evaluation. *Journal of Geophysical Research*, 106(D19), 23073–23095. <https://doi.org/10.1029/2001JD000807>
- Bian, H., & Prather, M. J. (2002). Fast-J2: Accurate simulation of stratospheric photolysis in global chemical models. *Journal of Atmospheric Chemistry*, 41(3), 281–296. <https://doi.org/10.1023/A:1014980619462>
- Brasseur, G. P., & Jacob, D. J. (2017). *Modeling of atmospheric chemistry*. Cambridge University Press. <https://doi.org/10.1017/9781316544754>

Acknowledgments

This work was supported by the US EPA Science to Achieve Results (STAR) Program, by the NASA Modeling, Analysis, and Prediction (MAP) Program, and by the NASA Atmospheric Composition Modeling and Analysis Program (ACMAP).

- Brown, P. N., Byrne, G. D., & Hindmarsh, A. C. (1989). VODE: A variable-coefficient ODE solver. *SIAM Journal on Scientific and Statistical Computing*, 10(5), 1038–1051. <https://doi.org/10.1137/0910062>
- Brown-Steiner, B., Selin, N. E., Prinn, R., Tilmes, S., Emmons, L., Lamarque, J.-F., & Cameron-Smith, P. (2018). 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. *Geoscientific Model Development*, 11(10), 4155–4174. <https://doi.org/10.5194/gmd-11-4155-2018>
- Damian, V., Sandu, A., Damian, M., Potra, F., & Carmichael, G. R. (2002). The kinetic preprocessor KPP—a software environment for solving chemical kinetics. *Computers & Chemical Engineering*, 26(11), 1567–1579. [https://doi.org/10.1016/S0098-1354\(02\)00128-X](https://doi.org/10.1016/S0098-1354(02)00128-X)
- Dawson, M. L., Guzman, C., Curtis, J. H., Acosta, M., Zhu, S., Dabdub, D., et al. (2022). Chemistry across multiple phases (CAMP) version 1.0: An integrated multiphase chemistry model. *Geoscientific Model Development*, 15(9), 3663–3689. <https://doi.org/10.5194/gmd-15-3663-2022>
- Djouad, R., & Sportisse, B. (2002). Partitioning techniques and lumping computation for reducing chemical kinetics. APLA: An automatic partitioning and lumping algorithm. *Applied Numerical Mathematics*, 43(4), 383–398. [https://doi.org/10.1016/S0168-9274\(02\)00111-3](https://doi.org/10.1016/S0168-9274(02)00111-3)
- Eastham, S. D., Long, M. S., Keller, C. A., Lundgren, E., Yantosca, R. M., Zhuang, J., et al. (2018). GEOS-chem high performance (GCHP v11-02c): A next-generation implementation of the GEOS-chem chemical transport model for massively parallel applications. *Geoscientific Model Development*, 11(7), 2941–2953. <https://doi.org/10.5194/gmd-11-2941-2018>
- Fast, J. D., Gustafson, W. I., Jr., Easter, R. C., Zaveri, R. A., Barnard, J. C., Chapman, E. G., et al. (2006). Evolution of ozone, particulates, and aerosol direct radiative forcing in the vicinity of Houston using a fully coupled meteorology-chemistry-aerosol model. *Journal of Geophysical Research*, 111(D21), D21305. <https://doi.org/10.1029/2005JD006721>
- Gong, W., & Cho, H.-R. (1993). A numerical scheme for the integration of the gas-phase chemical rate equations in three-dimensional atmospheric models. *Atmospheric Environment, Part A: General Topics*, 27(14), 2147–2160. [https://doi.org/10.1016/0960-1686\(93\)90044-Y](https://doi.org/10.1016/0960-1686(93)90044-Y)
- Grell, G. A., Peckham, S. E., Schmitz, R., McKeen, S. A., Frost, G., Skamarock, W. C., & Eder, B. (2005). Fully coupled “online” chemistry within the WRF model. *Atmospheric Environment*, 39(37), 6957–6975. <https://doi.org/10.1016/j.atmosenv.2005.04.027>
- Hairer, E., & Wanner, G. (1991). *Solving ordinary differential equations II. Stiff and differential-algebraic problems*. Springer.
- Hakami, A., Henze, D. K., Seinfeld, J. H., Singh, K., Sandu, A., Kim, S., et al. (2007). The adjoint of CMAQ. *Environmental Science & Technology*, 41(22), 7807–7817. <https://doi.org/10.1021/es070944p>
- Henze, D. K., Hakami, A., & Seinfeld, J. H. (2007). Development of the adjoint of GEOS-Chem. *Atmospheric Chemistry and Physics*, 7(9), 2413–2433. <https://doi.org/10.5194/acp-7-2413-2007>
- Holmes, C. D., Bertram, T. H., Confer, K. L., Graham, K. A., Ronan, A. C., Wirks, C. K., & Shah, V. (2019). The role of clouds in the tropospheric NOx cycle: A new modeling approach for cloud chemistry and its global implications. *Geophysical Research Letters*, 46(9), 4980–4990. <https://doi.org/10.1029/2019GL081990>
- Jacobson, M. Z. (1995). Computation of global photochemistry with SMVGear II. *Atmospheric Environment*, 29(18), 2541–2546. [https://doi.org/10.1016/1352-2310\(95\)00194-4](https://doi.org/10.1016/1352-2310(95)00194-4)
- Jacobson, M. Z., & Turco, R. P. (1994). SMVGear: A sparse-matrix, vectorized gear code for atmospheric models. *Atmospheric Environment*, 28(2), 273–284. [https://doi.org/10.1016/1352-2310\(94\)90102-3](https://doi.org/10.1016/1352-2310(94)90102-3)
- Jöckel, P., Kerkweg, A., Pozzer, A., Sander, R., Tost, H., Riede, H., et al. (2010). Development cycle 2 of the modular Earth submodel system (MESSy2). *Geoscientific Model Development*, 3, 717–752. <https://doi.org/10.5194/gmd-3-717-2010>
- Keller, C. A., & Evans, M. J. (2019). Application of random forest regression to the calculation of gas-phase chemistry within the GEOS-Chem chemistry model v10. *Geoscientific Model Development*, 12(3), 1209–1225. <https://doi.org/10.5194/gmd-12-1209-2019>
- Kelp, M. M., Jacob, D. J., Kutz, J. N., Marshall, J. D., & Tessum, C. W. (2020). Toward stable, general machine-learned models of the atmospheric chemical system. *Journal of Geophysical Research: Atmospheres*, 125(23), e2020JD032759. <https://doi.org/10.1029/2020JD032759>
- Kelp, M. M., Jacob, D. J., Lin, H., & Sulprizio, M. P. (2022). An online-learned neural network chemical solver for stable long-term global simulations of atmospheric chemistry. *Journal of Advances in Modeling Earth Systems*, 14(6), e2021MS002926. <https://doi.org/10.1029/2021MS002926>
- Long, M. S., & Lin, H. (2022). KineticPreProcessor/KPP-AR-boxmodel: GC_v13.4-AR-BoxModel [Software]. Zenodo. <https://doi.org/10.5281/zenodo.6791657>
- Mao, J., Jacob, D. J., Evans, M. J., Olson, J. R., Ren, X., Brune, W. H., et al. (2010). Chemistry of hydrogen oxide radicals (HO_x) in the Arctic troposphere in spring. *Atmospheric Chemistry and Physics*, 10(13), 5823–5838. <https://doi.org/10.5194/acp-10-5823-2010>
- Moch, J. M., Dovrou, E., Mickley, L. J., Keutsch, F. N., Liu, Z., Wang, Y., et al. (2020). Global importance of hydroxymethanesulfonate in ambient particulate matter: Implications for air quality. *Journal of Geophysical Research: Atmospheres*, 125(18), e2020JD032706. <https://doi.org/10.1029/2020JD032706>
- Philip, S., Martin, R. V., & Keller, C. A. (2016). Sensitivity of chemistry-transport model simulations to the duration of chemical and transport operators: A case study with GEOS-chem v10-01. *Geoscientific Model Development*, 9(5), 1683–1695. <https://doi.org/10.5194/gmd-9-1683-2016>
- Rastigejev, Y., Brenner, M. P., & Jacob, D. J. (2007). Spatial reduction algorithm for atmospheric chemical transport models. *Proceedings of the National Academy of Sciences of the United States of America*, 104(35), 13875–13880. <https://doi.org/10.1073/pnas.0705649104>
- Rosenbrock, H. H. (1963). Some general implicit processes for the numerical solution of differential equations. *The Computer Journal*, 5(4), 329–330. <https://doi.org/10.1093/comjnl/5.4.329>
- Sander, R., Baumgaertner, A., Cabrera-Perez, D., Frank, F., Gromov, S., Grooß, J.-U., et al. (2019). The community atmospheric chemistry box model CAABA/MECCA-4.0. *Geoscientific Model Development*, 12(4), 1365–1385. <https://doi.org/10.5194/gmd-12-1365-2019>
- Sandu, A., & Sander, R. (2006). Technical note: Simulating chemical systems in Fortran90 and matlab with the kinetic PreProcessor KPP-2.1. *Atmospheric Chemistry and Physics*, 6(1), 187–195. <https://doi.org/10.5194/acp-6-187-2006>
- Sandu, A., Sander, R., Long, M. S., Yantosca, R. M., Lin, H., Shen, L., & Jacob, D. J. (2022). Kinetic Pre-Processor/KPP: The Kinetic Pre-Processor (KPP) 3.0.0 [Software]. Zenodo. <https://doi.org/10.5281/zenodo.7308373>
- Sandu, A., Verwer, J., Blom, J., Spee, E., Carmichael, G., & Potra, F. (1997). Benchmarking stiff ode solvers for atmospheric chemistry problems II: Rosenbrock solvers. *Atmospheric Environment*, 31(20), 3459–3472. [https://doi.org/10.1016/S1352-2310\(97\)83212-8](https://doi.org/10.1016/S1352-2310(97)83212-8)
- Santillana, M., Le Sager, P., Jacob, D. J., & Brenner, M. P. (2010). An adaptive reduction algorithm for efficient chemical calculations in global atmospheric chemistry models. *Atmospheric Environment*, 44(35), 4426–4431. <https://doi.org/10.1016/j.atmosenv.2010.07.044>
- Shen, L., Jacob, D. J., Santillana, M., Bates, K., Zhuang, J., & Chen, W. (2022). A machine-learning-guided adaptive algorithm to reduce the computational cost of integrating kinetics in global atmospheric chemistry models: Application to GEOS-chem versions 12.0.0 and 12.9.1. *Geoscientific Model Development*, 15(4), 1677–1687. <https://doi.org/10.5194/gmd-15-1677-2022>
- Shen, L., Jacob, D. J., Santillana, M., Wang, X., & Chen, W. (2020). An adaptive method for speeding up the numerical integration of chemical mechanisms in atmospheric chemistry models: Application to GEOS-chem version 12.0.0. *Geoscientific Model Development*, 13(5), 2475–2486. <https://doi.org/10.5194/gmd-13-2475-2020>

- The International GEOS-Chem Community. (2022). jimmielin/geos-chem: AR-GC-archival_13.4 [Software]. Zenodo. <https://doi.org/10.5281/zenodo.6791655>
- Wang, X., Jacob, D. J., Downs, W., Zhai, S., Zhu, L., Shah, V., et al. (2021). Global tropospheric halogen (Cl, Br, I) chemistry and its impact on oxidants. *Atmospheric Chemistry and Physics*, 21(18), 13973–13996. <https://doi.org/10.5194/acp-21-13973-2021>
- Young, T. R., & Boris, J. P. (1977). A numerical technique for solving stiff ordinary differential equations associated with the chemical kinetics of reactive-flow problems. *Journal of Physical Chemistry*, 81(25), 2424–2427. <https://doi.org/10.1021/j100540a018>
- Zhao, S., Russell, M. G., Hakami, A., Capps, S. L., Turner, M. D., Henze, D. K., et al. (2020). A multiphase CMAQ version 5.0 adjoint. *Geoscientific Model Development*, 13(7), 2925–2944. <https://doi.org/10.5194/gmd-13-2925-2020>
- Zhuang, J., Jacob, D. J., Lin, H., Lundgren, E. W., Yantosca, R. M., Gaya, J. F., et al. (2020). Enabling high-performance cloud computing for Earth science modeling on over a thousand cores: Application to the GEOS-chem atmospheric chemistry model. *Journal of Advances in Modeling Earth Systems*, 12(5), e2020MS002064. <https://doi.org/10.1029/2020MS002064>