



Implementation of the chemistry module MECCA (v2.5) in the modal aerosol version of the Community Atmosphere Model component (v3.6.33) of the Community Earth System Model

M. S. Long¹, W. C. Keene², R. Easter³, R. Sander⁴, A. Kerkweg⁵, D. Erickson⁶, X. Liu³, and S. Ghan³

¹School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA

²Department of Environmental Sciences, University of Virginia, Charlottesville, VA 22904, USA

³Atmospheric Sciences and Global Change Division, Pacific Northwest National Laboratory, Richland, Washington, USA

⁴Air Chemistry Department, Max-Planck Institute of Chemistry, 55020 Mainz, Germany

⁵Institute for Atmospheric Physics, University of Mainz, 55099 Mainz, Germany

⁶Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA

Correspondence to: M. S. Long (mlong@seas.harvard.edu)

Received: 18 April 2012 – Published in Geosci. Model Dev. Discuss.: 12 June 2012

Revised: 27 November 2012 – Accepted: 8 January 2013 – Published: 22 February 2013

Abstract. A coupled atmospheric chemistry and climate system model was developed using the modal aerosol version of the National Center for Atmospheric Research Community Atmosphere Model (modal-CAM; v3.6.33) and the Max Planck Institute for Chemistry's Module Efficiently Calculating the Chemistry of the Atmosphere (MECCA; v2.5) to provide enhanced resolution of multiphase processes, particularly those involving inorganic halogens, and associated impacts on atmospheric composition and climate. Three Rosenbrock solvers (Ros-2, Ros-3, RODAS-3) were tested in conjunction with the basic load-balancing options available to modal-CAM (1) to establish an optimal configuration of the implicitly-solved multiphase chemistry module that maximizes both computational speed and repeatability of Ros-2 and RODAS-3 results versus Ros-3, and (2) to identify potential implementation strategies for future versions of this and similar coupled systems. RODAS-3 was faster than Ros-2 and Ros-3 with good reproduction of Ros-3 results, while Ros-2 was both slower and substantially less reproducible relative to Ros-3 results. Modal-CAM with MECCA chemistry was a factor of 15 slower than modal-CAM using standard chemistry. MECCA chemistry integration times demonstrated a systematic frequency distribution for all three solvers, and revealed that the change in run-time performance was due to a change in the frequency distribution of chemical integration times; the peak frequency was similar

for all solvers. This suggests that efficient chemistry-focused load-balancing schemes can be developed that rely on the parameters of this frequency distribution.

1 Introduction

The spatial and temporal resolutions of geophysical modeling systems are increasing rapidly. As a result, the need to more explicitly resolve many of the physical and chemical processes that previously operated below the resolution and within the uncertainty ranges of these modeling systems is increasing accordingly. Individually, the computational skill of physical and chemical systems is high, but the computational needs of these systems in combination with dynamical and geophysical models has made coupled investigations prohibitive. The capabilities of current high-performance computing platforms available to geoscientific modeling are beginning to permit the coupling of these systems for scientific research. Of particular interest are the interactions between atmospheric chemistry and climate, particularly with respect to the implications of multiphase processes for tropospheric composition, clouds, precipitation, and radiative transfer.

Multiphase interactions, primarily between gases, aerosols and cloud droplets, represent a highly non-linear set of

processes that significantly impact the processing and lifetimes of many important tropospheric species. Of increasing interest are chemical transformations involving inorganic, halogenated (Cl and Br) compounds and associated influences on the cycling of NO_x, HO_x, S, O₃, CH₄ and non-methane hydrocarbons (NMHCs), Hg, and other species of both natural and anthropogenic origin.

Accurately resolving interactions that control multiphase processes requires they be evaluated explicitly. The computationally difficult solution of the stiff system of ordinary differential equations (ODEs) derives from multiphase processes (e.g., mass transfer). Computational speed must be optimized in order to execute simulations of sufficient duration to provide time for model equilibration (spin up) and generation of a sufficient sample size for analysis.

This manuscript describes a coupled atmospheric chemistry and climate modeling system that leverages an efficient multiphase atmospheric chemistry mechanism, MECCA (Module Efficiently Calculating the Chemistry of the Atmosphere; version 2.5; Sander et al., 2005, 2011) within a 3-mode size-resolving aerosol module (modal aerosol module) version of the National Center for Atmospheric Research's (NCAR) Community Atmosphere Model (version 3.6.33; Gent et al., 2010; Liu et al., 2012; hereafter referred to as modal-CAM). The modal aerosol module in CAM was developed to provide a size-resolving aerosol microphysics package capable of more accurately resolving the direct and indirect impacts of aerosols on climate. Modal-CAM is embedded as the atmosphere component of the NCAR Community Climate System Model (CCSM3.0; Collins et al., 2006). (Note: since completion of the work presented here, CCSM has been renamed the Community Earth System Model, CESM).

The coupled modeling system used in this study was designed to investigate the role of aqueous processes and inorganic halogen cycling through use of their explicit representation in MECCA combined with the size-resolving modal aerosol physics and atmospheric coupling of modal-CAM. Results will be validated and interpreted in detail in a forthcoming manuscript (Long, M. S., Keene, W. C., Easter, R., Sander, R., Kerkweg, A., Lui, X., Erickson, D. J., and Ghan, S.: Sensitivity of tropospheric chemical composition to halogen-radical chemistry using a fully coupled GCM/size-resolved multiphase chemical system I: Halogen distributions, aerosol composition, and sensitivity of climate-relevant gases).

2 MECCA model description

MECCA version 2.5 is a FORTRAN90 compliant atmospheric chemistry module developed to deploy easily as a submodel within base models using the MESSy interface (Modular Earth Submodel System; see <http://www.messy-interface.org>). Since CAM is not

designed as a MESSy compliant base model, the interface used for this study was designed from scratch to accommodate the complexities of the non-compliant GCM and the needs of the modal aerosol module. MECCA is available at no cost, under the terms of the GNU General Public License (GPL), included within – and not to be confused with – the stand-alone box-model CAABA (Chemistry As A Boxmodel Application).

MECCA contains a comprehensive atmospheric reaction mechanism that includes transformations involving O₃, CH₄, HO_x, NO_x, NMHCs, halogens (Cl, Br, I), and sulfur. In addition to gas-phase reactions, the scheme includes fully integrated multiphase transformations (both aqueous-phase and heterogeneous pathways) involving aerosols and cloud droplets. Mass transfer is calculated dynamically per Schwartz (1986). Photochemical reaction rates vary as a function of solar zenith angle under clear-sky and cloudy conditions based on Landgraf and Crutzen (1998). MECCA is a MESSy-compliant submodel within the ECHAM5/MESSy for Atmospheric Chemistry (EMAC) chemistry climate GCM (CCM). Numerous investigations have been performed using this system: These include evaluation of gas-phase chemistry from the surface to the mesosphere (Jöckel et al., 2006), multiphase cycling of marine-derived halogens (Kerkweg et al., 2008a, b), isotopic composition of the atmosphere (Gromov et al., 2010), and influences of chemical processes on polar stratospheric clouds (Kirner et al., 2011). A full list of EMAC applications can be found on <http://messy-interface.org>. See Table S1 in the Supplement and Sander et al. (2011) for a complete description of the chemical scheme.

MECCA uses the Kinetics PreProcessor (KPP, Sandu and Sander, 2006) to build a solution based on a choice of several predefined numerical methods. KPP was designed to facilitate programming fast and accurate solutions to chemical reaction mechanisms based on user-defined implicit solvers and solver configurations. It relies on sparse linear algebra routines to optimize serial computational performance, and is therefore well suited for atmospheric chemistry problems over a wide range of complexities.

The tropospheric chemical mechanism used in the coupled model was based on a subset of the full MECCA mechanism. Other than the addition of gas-phase reactions for non-methane hydrocarbons (NMHCs; based on von Kuhlmann et al., 2003), the mechanism was identical to that used in Keene et al. (2009), although configured for three rather than eight aerosol size bins. Photochemical rates were calculated using CAABA's JVAL submodel.

3 Modal-CAM atmosphere model

Atmospheric processes were simulated in three dimensions (3-D) using CAM at 1.9° × 2.5° lat-long resolution with 26 vertical levels (Gent et al., 2009). CAM is a FORTRAN90




Fig. 1. Schematic showing the order of relevant chemistry, dynamics and physics routines (light gray boxes) over a single model time iteration (Δt) relative to the MECCA (white horizontal bar) and the modal-CAM (gray dark-gray horizontal bar) tracer arrays. Boxes indicating model operations are oriented vertically across the tracer array bars to indicate whether they interact with one or both tracer arrays. Step indices correspond to those described in the text (see Sect. 4).

compliant general circulation system built upon an extensive set of high-performance computational routines to preserve scalability and performance of the model across changes in resolution and model physics. The high-performance structure relies upon the message passing interface (MPI), or, at the user's discretion, a combination of MPI and shared-memory process routines.

The dynamical core (approximation of the equations of motion on a discrete, spherical grid) is based on a flux-form semi-Lagrangian method (see Lin and Rood, 1996) that is better suited for tracer transport. This approach permits grid-wide stability of the chemistry solution, in contrast to discrete methods that introduce large dispersion and diffusion errors in their approximation of the equations of motion which propagate into and destabilize the chemistry solver.

Modal-CAM incorporates a comprehensive set of processes that control the evolution and coupling of three fixed-width log-normally distributed aerosol modes (Aitken, accumulation and coarse). The modal aerosol treatment is described in detail in Liu et al. (2012). Each mode consists of internally mixed populations of non-sea-salt (nss) SO_4^{2-} , organic matter from primary sources (OM), secondary organic aerosol (SOA) from volatile organic precursors, black carbon (BC), inorganic sea salt, and mineral dust. The nss- SO_4^{2-} is assumed to be in the form of NH_4HSO_4 . OM and BC are treated only in the accumulation mode. SOA is only in the Aitken and accumulation modes, and mineral dust is only in the accumulation and coarse modes. Aerosol number and aerosol water are also calculated for each mode. Aerosol mass and number associated with stratiform cloud droplets are treated explicitly.

The following processes affect aerosols in the model: grid-resolved transport, sub-grid vertical transport by turbulence and convective clouds, emissions (surface and elevated),

sedimentation and dry deposition, cloud droplet activation and subsequent aerosol resuspension, wet removal (in- and below-cloud by stratiform and convective clouds), condensation of $\text{H}_2\text{SO}_4(\text{g})$ and condensation/evaporation of semi-volatile organics and water, cloud chemistry (oxidation of SO_2 to H_2SO_4), transfer (renaming) of particles from Aitken to accumulation mode due to growth via condensation and cloud chemistry, aerosol nucleation, and aerosol coagulation (Aitken and accumulation modes only). Trace gas processes include transport, emission, and dry and wet deposition.

4 MECCA/Modal-CAM coupling

The coupling involves (1) adding MECCA chemical species to CAM, (2) interfacing MECCA gas, aerosol and cloud chemistry routines within CAM (and disabling the corresponding CAM routines), and (3) as needed, modifying CAM routines for processes that affect MECCA and modal-CAM species (e.g., emission of sea-salt species). This initial implementation is not a complete two-way coupling between MECCA and modal-CAM, as indicated in Fig. 1, since some MECCA aerosol species do not interact directly with the modal-CAM physics. This was done to minimize unnecessary modifications to processes that have little impact from one system to another. The impact of this configuration on memory use and model performance was not evaluated.

The MECCA gas and aqueous aerosol species were added to the existing fully-transported trace species in CAM: 96 gas species (4 of which were already treated in modal-CAM), and 31 aqueous aerosol species in each of the 3 size modes. The MECCA aqueous cloud-droplet species (31 species for each mode) were also added to the modal-CAM cloud-borne species, which are not fully transported (Liu et al., 2012). This coupling interfaces the bulk inorganic aerosol

composition considered by microphysical routines in modal-CAM with chemical speciation evaluated for multiphase processes in MECCA. As such, the system uses redundant chemical species to account for nss-SO_4^{2-} and sea salt between two tracer arrays. Gas-phase species are shared between the MECCA and modal-CAM tracer arrays. Since it was necessary to ensure that the impact of model routines on both bulk species in modal-CAM (e.g., NaCl), and corresponding speciation in MECCA (e.g., Na^+ and Cl^-) are proportional, several model routines operate on both tracer arrays simultaneously (see Fig. 1). Modal-CAM stores information about both sets of species throughout a time step, and changes are updated accordingly – either from MECCA to CAM (Fig. 1, step 6) or CAM to MECCA (Fig. 1, step 8). The iteration of one model time step as outlined in Fig. 1 involved 10 discrete steps:

Step 1: This step calculates the advective transport for all chemical species.

Step 2: Vertical transport of gases and interstitial aerosols in both tracer arrays by shallow convective clouds is calculated. Aerosol activation/resuspension in stratiform clouds is then calculated in conjunction with turbulent vertical mixing, acting on both tracer arrays. The aerosol activation utilizes the modal-CAM aerosol composition (e.g., hygroscopicity calculation neglects MECCA chemical species); but tendencies are applied to all aerosol species.

Step 3: Aerosol water uptake is calculated based on modal-CAM's aerosol composition. Resulting aerosol water content is applied to both tracer arrays. Wet deposition of all aerosol species (interstitial and cloud-borne in both arrays) through in-cloud and below-cloud scavenging is then calculated. Next, vertical transport of gases and interstitial aerosols by deep convective clouds is calculated.

Step 4: Below-cloud scavenging by rainwater of all soluble gases occurs here.

Step 5: Gas, aerosol, and photo chemistry act only on the MECCA tracer array (see Sect. 2). Total overhead stratospheric O_3 necessary for photochemical rate calculations in MECCA's *JVAL* routine was prescribed. Ion balance is maintained in MECCA by adjusting an inert dummy cation tracer representing the combined charges of Na^+ , Ca^+ , and Mg^+ , which was not coupled to modal-CAM NaCl mass. With the exception of that involving SOA, all gas–aerosol exchange was calculated by MECCA.

Step 6: Cloud chemistry includes MECCA-only gas/cloudwater exchange of soluble gas species, equilibrium, and aqueous chemistry in cloud droplets. Cloud chemistry was only activated above a grid-box cloud-fraction threshold of 1×10^{-5} .

Step 7: nss-SO_4^{2-} is passed to modal-CAM after completion of MECCA chemistry. To differentiate total SO_4^{2-} in MECCA, which includes sea-salt SO_4^{2-} from nss-SO_4^{2-} in modal-CAM, only the net change in nss-SO_4^{2-} due to MECCA aerosol chemistry ($\Delta\text{SO}_{4,\text{Chemistry}}$ from aqueous reaction and H_2SO_4 vapor uptake) was considered where

$$\text{nss-SO}_4^{2-}(t + \Delta t) = \text{nss-SO}_4^{2-}(t) + \Delta\text{SO}_4^{2-},\text{Chemistry}. \quad (1)$$

nss-SO_4^{2-} , as passed to modal-CAM, is the sum of MECCA's $\text{H}_2\text{SO}_4(\text{aq})$, HSO_4^- , and SO_4^{2-} species calculated from Eq. (1). nss-SO_4^{2-} for each mode, H_2SO_4 vapor, and corresponding net changes per time-step were calculated here for use by the modal-CAM microphysical routines (step 8).

Step 8: The aerosol microphysical processes of condensation (SOA only), intermodal transfer (renaming) after particle growth, nucleation, and coagulation are calculated in modal-CAM routines. Intermodal transfer and coagulation are now applied to both modal-CAM and MECCA aerosol species. Since mass-transfer from the gas to aqueous phase is included in the MECCA chemical ODE, modal-CAM gas–aerosol exchange and condensation routines are switched off for all species except for SOA.

Step 9: The only net source of nss-SO_4^{2-} in step 8 was through nucleation of $\text{H}_2\text{SO}_4(\text{g})$. This increase in nss-SO_4^{2-} due to modal aerosol processing was passed to MECCA as addition of $\text{H}_2\text{SO}_4(\text{aq})$ to the Aitken mode. In this configuration, both total nss-SO_4^{2-} and H^+ are conserved.

Step 10: Emissions of gases, black carbon, primary organic matter, and NH_4HSO_4 aerosol are driven by offline datasets, while sea salt and dust emissions are calculated online. NH_4HSO_4 aerosol are emitted directly as NH_4^+ and HSO_4^{2-} into the MECCA tracer array. The transfer of nss-SO_4^{2-} into the modal-CAM array occurs later (step 6). Sea salt aerosol is emitted as both NaCl in the modal-CAM array and speciated in MECCA as Na^+ , Cl^- , SO_4^{2-} , CO_3^{2-} , and Br^- . Sea-salt derived SO_4^{2-} is excluded from modal-CAM (see step 6). Vertical turbulent mixing is applied to all gases. (This is done in step 3 for aerosols.) Dry deposition includes all aerosol and gas-phase species in both tracer arrays.

Since impacts on aerosol physical properties are limited due to small changes in abundance of inorganic aerosol species other than Na^+ , NO_3^- , NH_4^+ and SO_4^{2-} , and to simplify the modal-CAM aerosol size and inter-modal exchange routines, mass and density of any species specific to MECCA

only were not considered in calculations of particle mass and size (i.e., density of aged and fresh sea salt are the same). As a result they only interacted with particle dry diameter through changes in nss-SO_4^{2-} , NH_4^+ , and NO_3^- . As well, the volume-weighted hygroscopicities and refractive indices of aerosol modes were calculated using modal-CAM species: bulk NaCl, nss-SO_4^{2-} , dust, BC, POM, and SOA in the 3-mode version, plus NH_4^+ and NO_3^- in the 7-mode version.

5 Computational configuration and performance

In a global simulation grid that includes the breadth of atmospheric chemical scenarios at any one time point in the simulation, the use of implicit methods for the multiphase chemistry solution disrupts the scalability of the MPI-based parallel system. In particular, the stiffness of the chemical mechanism, and thus the time and resources needed to reach a solution for a given grid box, varies geographically in the 3-D domain (e.g., see Kerkweg et al., 2007). Proximity to large sources and sinks of highly-reactive species or large gradients in physical or chemical conditions complicate the implicit solution. In CAM, systematic, non-random decomposition and allocation of column subsets of the 3-D grid to the available computational processes segregates a disproportionately large chemical-solution burden into a small subset of processes. Since CAM's time-stepping routines rely on an MPI AllGather routine, model performance is limited by the speed of the slowest column subset. Model load-balancing optimizations (available since CAM version 3.6) permit the building of MPI column subsets and allocating them to processes in ways which enhances the distribution of "difficult" columns across available computational process units (see Mirin and Worley, 2011). The option used for this study (`phys_loadbalance = 2`) builds column subsets from north/day–south/night grid-point pairs. For example, a grid point at 45° N, 0° E is paired with the point at 45° S and 180° E. Consequently, since most land area is in the Northern Hemisphere, this procedure load-balances based across day/night, season, and, to a large extent land/ocean. The set of paired points are then combined into column subsets and assigned to processes. The maximum number of column subsets that can be obtained (thus, the number of independent computational processes that can be used simultaneously) is controlled by dynamic limitations and the horizontal grid size. CAM has been designed to permit allocating additional processes to solve model physics separately from the dynamics routines, which allows a much faster computation of the coupled system.

To evaluate the computational performance of the coupled system, decomposed as described, three positive-definite, adjustable-timestep Rosenbrock methods were tested for accuracy and performance metrics. Sander et al. (2005) found that, for the MECCA chemical mechanism, 2nd and 3rd order solvers performed best in terms of both stability and

Table 1. Comparison of one-month benchmark simulations of the coupled modal-CAM/MECCA system for Ros-2 and RODAS-3 solvers versus Ros-3.

	Species Line	Regression	R^2 (%)	RMSE
Ros-2	O ₃	0.98x + 0.56	0.99	6.5 %
	OH	$1.0x + 1.5 \times 10^{-4}$	0.98	17 %
	Br ₂	0.91x + 0.87	0.94	42 %
	H ⁺ (Coarse Mode)	2.35x + 0.048	0.93	410 %
RODAS-3	O ₃	0.98x – 0.36	0.99	5.1 %
	OH	$0.99 + 7.3 \times 10^{-5}$	0.99	11 %
	Br ₂	0.95x + 0.98	0.97	29 %
	H ⁺ (Coarse Mode)	0.92 + 0.011	0.97	120 %

computational speed. Other studies have investigated the stability and efficiency of the Rosenbrock solvers in KPP across a range of chemical scenarios (Henze et al., 2007; Verwer et al., 1999; Sandu et al., 1997). To our knowledge, this study is the first in which KPP's Rosenbrock solvers were tested against such a complex chemical mechanism including gas, multiphase, and photochemistry through the entire troposphere.

For stability reasons, the Ros-3 (3-stage, order 3(2), L-stable) solver was employed preferentially in past MECCA simulations (e.g., Keene et al., 2009). Consequently, for these tests, coupled simulation results using Ros-3 are considered the benchmark against which results using Ros-2 (2-stage, order 2(1), L-stable) and RODAS-3 (4-stage, order 3(2), stiffly accurate) solvers are compared (see Hairer and Wanner, 1991). The coupled system was run for 5 yr with the Ros-3 solver to stabilize chemistry in the troposphere – defined as a net change in year-to-year total global O₃ mass of less than 1 % (actual net O₃ change between years 4 and 5 was 0.16 % versus 4.7 % between years 3 and 4 of the equilibration period). One-month (January) simulations were then executed using the three solvers. Ros-2 and RODAS-3 were compared to Ros-3 for computational speed and reproducibility of several species. Absolute and relative tolerances were set to 10 cm^{-3} and 0.01, respectively.

In the implicit solution to the multiphase mechanism, the main sources of instability and stiffness involved complex, fast, multiphase chemistry in the near-surface layers. In addition to high liquid water contents in these layers relative to others, there were large wind and geography driven 3-D gradients in reactive species and trace intermediates due to reactions in neighboring grid regions, emissions, deposition, microphysical processing, and scavenging. Thus, it is in close proximity to the surface that the limitations of each numerical method – whether in computational stability or accuracy of the solution – were best evaluated.




Fig. 2. Comparison of O_3 (a and b), OH (c and d), Br_2 (e and f), and coarse mode H^+ (g and h) at grid boxes between the surface and 900 mb from one-month benchmark simulations using RODAS-3 (left column y-axis) and Ros-2 (right column y-axis) solvers versus Ros-3 (x-axis) over the same time period. Black lines depict least-squares standard linear regressions (see Table 1).

Figure 2 compares mass mixing ratios of the one-month benchmark for O_3 , OH, Br_2 , and coarse-mode aqueous H^+ for all model layers between the surface and 900 mb. These species were selected to reflect climate relevance, source of stiffness, halogen cycle reproducibility, and relevant aqueous processes. Regression statistics are given in Table 1. Ros-2 is able to reproduce O_3 and OH with reasonable confidence,

whereas Br_2 and to a much greater extent, H^+ were less precisely reproduced. The reason for the systematic over-prediction of H^+ by Ros-2 is not clear, but may reflect stiffness associated with the aqueous (acid–base) reactions and mass transfer. Conversely, results based on RODAS-3 were more similar to those based on Ros-3 in terms of both absolute (regression slope near 1) and relative differences (higher correlation coefficient; Table 1). The H^+ root mean square error (RMSE; normalized against mean Ros-3 mixing ratios) was still high for the RODAS-3 results. The scatter at higher H^+ mixing ratios generally corresponded to continental regions where sources of atmospheric acids are relatively greater and sea-salt Cl^- and associated regulation of aerosol acidity via HCl phase partitioning is relatively less important. H^+ is highly sensitive to changes in chemistry and circulation in these regions. Circulation changes may also be reflected in the other species due to radiative forcing by O_3 over the benchmark time period. The comparisons demonstrate that RODAS-3 performs markedly better than Ros-2 for all four species.

Relative to Ros-3, completion of the one-month benchmark simulation with RODAS-3 was 9% faster and Ros-2 was 18% slower. This is in agreement with a study of KPP solvers in the GEOS-Chem chemistry transport model (Henze et al., 2007; Eller et al., 2008), although GEOS-Chem uses KPP only for gas-phase calculations and is driven by offline circulation. The frequency distributions of average integration times (or waiting time for completion of one chemistry timestep) for all grid cells varied among the solvers tested (Fig. 3). Relative to RODAS-3, Ros-2 and, to a lesser extent, Ros-3 were skewed towards relatively longer integration times, though there was no systematic change in the peak integration time frequency. These results indicate that the performance gain is due primarily to reduction in frequency of large waiting times and suggest that chemistry-centric grid decomposition and column subsetting that leverages this frequency distribution may yield better model performance. The distribution of waiting times across the global grid demonstrates a physical dependence. While not shown here, chemistry waiting times are inversely dependent upon altitude – the maxima occur in the model surface layer. Further, data show a weak but positive correlation to a combination of total aerosol liquid water and solar zenith angle (as a measure of photochemistry). Based on the benchmark simulation inter-comparison, MECCA chemistry for the fully-coupled simulations was solved using the RODAS-3 solver.

MECCA, as the chemical operator in CAM, had a substantial impact on model runtime prohibiting the use of this configuration for long-term (century-scale) simulations without a large cost in computational resources. Incorporation of the MECCA species and chemistry routines increased CAM's runtime by a factor of 15 relative to modal-CAM configured with the standard chemical module. Replacing modal-CAM's chemical module with MECCA chemistry slowed overall computational speed by a factor of 8. The transport




Fig. 3. Histogram of the percent frequency distribution of per grid box chemistry integration times (in milliseconds) using MECCA in the modal-CAM global atmosphere for the one-month benchmarks using three different Rosenbrock solvers.

routines were a factor of 7 slower due to an increase from 25 to 205 active tracers. Further, the data storage needs of a system this extensive were large enough that considerations of input/output (I/O) frequency and number of diagnostic quantities were necessary. Monthly-mean output from a 10-yr simulation of the coupled system required nearly 850 GB storage, which added an additional computational burden due to the system I/O. I/O is often a limiting factor in high-performance system scalability, though it was not a large factor in this system.

6 Summary

A coupled atmospheric chemistry and climate system model was developed to investigate the details of multiphase processes and associated impacts on chemistry and climate. The computational needs of the chemical system required that performance of individual modules be enhanced. Comparison of three implicit Rosenbrock solvers revealed substantial differences in computational performance for coupled simulations that were distinct from similar investigations based on box models alone. This is likely due to the effect of fixed versus variable physical conditions in 0-D versus 3-D global models, combined with the impact of load-balancing methods on the net system runtime (solver performance in individual grid boxes was not evaluated). Overall the RODAS-3 solver provided the best performance for the current computational configuration.

In addition to optimizations discussed above, such as chemistry-centric load balancing, several strategies can be pursued to further increase the coupled system's performance. First, chemical species with atmospheric lifetimes

shorter than residence times in a given grid box (so called short-lived species such as $O(^1D)$) can be ignored by the dynamics routines. Prior to including MECCA into CAM for this study, the cost of including additional tracers was the largest factor impacting the system's computational burden. Second, reduction of the size of the chemical mechanism in combination with load balancing will likely have the greatest impact on runtime. A systematic approach to determining the smallest mechanism necessary to constrain the behavior of a specific subset of chemical species (e.g., O_3 and sulfur) is currently being developed using this system. Lastly, the adoption of optimized or parallel-capable linear algebra routines has the potential to significantly speed up the implicit chemistry, but we are not aware of any successful studies showing this. Doing so would require substantial changes to the existing parallelization strategy in CAM. The development of hybrid systems using stream and conventional processors provides a good opportunity to examine this approach.

MECCA, CAM, and the CESM are available for download. The code used here can be made available to users upon request.

Supplementary material related to this article is available online at: <http://www.geosci-model-dev.net/6/255/2013/gmd-6-255-2013-supplement.pdf>.

Acknowledgements. Financial support was provided by the US Department of Energy's (DOE's) Office of Science through the Office of Biological and Environmental Research (BER, grant numbers DE-FG02-07ER64442 and DE-SC0007120 to the University of Virginia), a Global Change Education Program Graduate Research Environmental Fellowship, and the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by DOE's Office of Science (BER) under contract DE-AC05-00OR22725. The CESM project is supported by the National Science Foundation and the DOE's Office of Science (BER). PNNL authors were funded by the US Department of Energy, Office of Science, Scientific Discovery through Advanced Computing (SciDAC) Program. The Pacific Northwest National Laboratory is operated for DOE by Battelle Memorial Institute under contract DE-AC06-76RLO 1830.

Edited by: V. Grewe

References

- Collins, W. D., Bitz, C. M., Blackmon, M. L., Bonan, G. B., Bretherton, C. S., Carton, J. A., Chang, P., Doney, S. C., Hack, J. J., Henderson, T. B., Kiehl, J. T., Large, W. G., McKenna, D. S., Santer, B. D., and Smith, R. D.: The Community Climate System Model Version 3 (CCSM3), *J. Climate*, 19, 2122–2143, doi:10.1175/JCLI3761.1, 2006.
- Eller, P., Singh, K., Sandu, A., Bowman, K., Henze, D. K., and Lee, M.: Implementation and evaluation of an array of chem-

- ical solvers in the Global Chemical Transport Model GEOS-Chem, *Geosci. Model Dev.*, 2, 89–96, doi:10.5194/gmd-2-89-2009, 2009.
- Gent, P. R., Yeager, S. G., Neale, R. B., Levis, S., and Bailey, D. A.: Improvements in a half degree atmosphere/land version of the CCSM, *Clim. Dynam.*, 34, 819–833, doi:10.1007/s00382-009-0614-8, 2010.
- Gromov, S., Jöckel, P., Sander, R., and Brenninkmeijer, C. A. M.: A kinetic chemistry tagging technique and its application to modelling the stable isotopic composition of atmospheric trace gases, *Geosci. Model Dev.*, 3, 337–364, doi:10.5194/gmd-3-337-2010, 2010.
- Hairer, E. and Wanner, G.: *Solving Ordinary Differential Equations II, Stiff and Differential-Algebraic Problems*, Springer, 1991.
- Henze, D. K., Hakami, A., and Seinfeld, J. H.: Development of the adjoint of GEOS-Chem, *Atmos. Chem. Phys.*, 7, 2413–2433, doi:10.5194/acp-7-2413-2007, 2007.
- Jöckel, P., Tost, H., Pozzer, A., Brühl, C., Buchholz, J., Ganzeveld, L., Hoor, P., Kerkweg, A., Lawrence, M. G., Sander, R., Steil, B., Stiller, G., Tanarhte, M., Taraborrelli, D., van Aardenne, J., and Lelieveld, J.: The atmospheric chemistry general circulation model ECHAM5/MESy1: consistent simulation of ozone from the surface to the mesosphere, *Atmos. Chem. Phys.*, 6, 5067–5104, doi:10.5194/acp-6-5067-2006, 2006.
- Keene, W. C., Long, M. S., Pszenny, A. A. P., Sander, R., Maben, J. R., Wall, A. J., O'Halloran, T. L., Kerkweg, A., Fischer, E. V., and Schrems, O.: Latitudinal variation in the multiphase chemical processing of inorganic halogens and related species over the eastern North and South Atlantic Oceans, *Atmos. Chem. Phys.*, 9, 7361–7385, doi:10.5194/acp-9-7361-2009, 2009.
- Kerkweg, A., Sander, R., Tost, H., Jöckel, P., and Lelieveld, J.: Technical Note: Simulation of detailed aerosol chemistry on the global scale using MECCA-AERO, *Atmos. Chem. Phys.*, 7, 2973–2985, doi:10.5194/acp-7-2973-2007, 2007.
- Kerkweg, A., Jöckel, P., Pozzer, A., Tost, H., Sander, R., Schulz, M., Stier, P., Vignati, E., Wilson, J., and Lelieveld, J.: Consistent simulation of bromine chemistry from the marine boundary layer to the stratosphere – Part I: Model description, sea salt aerosols and pH, *Atmos. Chem. Phys.*, 8, 5899–5917, doi:10.5194/acp-8-5899-2008, 2008a.
- Kerkweg, A., Jöckel, P., Warwick, N., Gebhardt, S., Brenninkmeijer, C. A. M., and Lelieveld, J.: Consistent simulation of bromine chemistry from the marine boundary layer to the stratosphere – Part 2: Bromocarbons, *Atmos. Chem. Phys.*, 8, 5919–5939, doi:10.5194/acp-8-5919-2008, 2008b.
- Kirner, O., Ruhnke, R., Buchholz-Dietsch, J., Jöckel, P., Brühl, C., and Steil, B.: Simulation of polar stratospheric clouds in the chemistry-climate-model EMAC via the submodel PSC, *Geosci. Model Dev.*, 4, 169–182, doi:10.5194/gmd-4-169-2011, 2011.
- Landgraf, J. and Crutzen, P. J.: An Efficient Method for online calculations of Photolysis and Heating Rates, *J. Atmos. Sci.*, 55, 863–878, 1998.
- Lin, S.-J. and Rood, R. B.: Multidimensional flux form semi-Lagrangian transport schemes, *Mon. Weather Rev.*, 124, 2046–2070, 1996.
- Liu, X., Easter, R. C., Ghan, S. J., Zaveri, R., Rasch, P., Shi, X., Lamarque, J.-F., Gettelman, A., Morrison, H., Vitt, F., Conley, A., Park, S., Neale, R., Hannay, C., Ekman, A. M. L., Hess, P., Mahowald, N., Collins, W., Iacono, M. J., Bretherton, C. S., Flanner, M. G., and Mitchell, D.: Toward a minimal representation of aerosols in climate models: description and evaluation in the Community Atmosphere Model CAM5, *Geosci. Model Dev.*, 5, 709–739, doi:10.5194/gmd-5-709-2012, 2012.
- Mirin, A. and Worley, P. H.: Improving the performance scalability of the community atmosphere model, *Int. J. High Performance Comput. Appl.*, 26, 17–30, doi:10.1177/1094342011412630, 2011.
- Sander, R., Baumgaertner, A., Gromov, S., Harder, H., Jöckel, P., Kerkweg, A., Kubistin, D., Regelin, E., Riede, H., Sandu, A., Taraborrelli, D., Tost, H., and Xie, Z.-Q.: The atmospheric chemistry box model CAABA/MECCA-3.0, *Geosci. Model Dev.*, 4, 373–380, doi:10.5194/gmd-4-373-2011, 2011.
- Sander, R., Kerkweg, A., Jöckel, P., and Lelieveld, J.: Technical note: The new comprehensive atmospheric chemistry module MECCA, *Atmos. Chem. Phys.*, 5, 445–450, doi:10.5194/acp-5-445-2005, 2005.
- Sandu, A. and Sander, R.: Technical note: Simulating chemical systems in Fortran90 and Matlab with the Kinetic PreProcessor KPP-2.1, *Atmos. Chem. Phys.*, 6, 187–195, doi:10.5194/acp-6-187-2006, 2006.
- 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, doi:10.1016/S1352-2310(97)83212-8, 1997.
- Schwartz, S. E.: Mass-transport considerations pertinent to aqueous phase reactions of gases in liquid-water clouds, in: *Chemistry of Multiphase Atmospheric Systems*, NATO ASI Series, edited by: Jaeschke, W., Springer Verlag, Berlin, Vol. G6, 415–471, 1986.
- Verwer, J. G., Spee, E. J., Blom, J. G., and Hundsdorfer, W.: A Second-Order Rosenbrock Method Applied to Photochemical Dispersion Problems, *SIAM J. Sci. Comp.*, 20, 1456–1480, doi:10.1137/S1064827597326651, 1999.
- von Kuhlmann, R., Lawrence, M. G., Crutzen, P., and Rasch, P.: A model for studies of tropospheric ozone and nonmethane hydrocarbons: Model description and ozone results, *J. Geophys. Res.*, 108, 4294, doi:10.1029/2002JD002893, 2003.