



HAL
open science

Development of the adjoint of GEOS-Chem

D. K. Henze, A. Hakami, J. H. Seinfeld

► **To cite this version:**

D. K. Henze, A. Hakami, J. H. Seinfeld. Development of the adjoint of GEOS-Chem. Atmospheric Chemistry and Physics, 2007, 7 (9), pp.2413-2433. hal-00296220

HAL Id: hal-00296220

<https://hal.science/hal-00296220>

Submitted on 18 Jun 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Development of the adjoint of GEOS-Chem

D. K. Henze, A. Hakami, and J. H. Seinfeld

California Institute of Technology, Pasadena, CA, USA

Received: 4 October 2006 – Published in Atmos. Chem. Phys. Discuss.: 19 October 2006

Revised: 6 February 2007 – Accepted: 17 April 2007 – Published: 11 May 2007

Abstract. We present the adjoint of the global chemical transport model GEOS-Chem, focusing on the chemical and thermodynamic relationships between sulfate – ammonium – nitrate aerosols and their gas-phase precursors. The adjoint model is constructed from a combination of manually and automatically derived discrete adjoint algorithms and numerical solutions to continuous adjoint equations. Explicit inclusion of the processes that govern secondary formation of inorganic aerosol is shown to afford efficient calculation of model sensitivities such as the dependence of sulfate and nitrate aerosol concentrations on emissions of SO_x , NO_x , and NH_3 . The accuracy of the adjoint model is extensively verified by comparing adjoint to finite difference sensitivities, which are shown to agree within acceptable tolerances. We explore the robustness of these results, noting how discontinuities in the advection routine hinder, but do not entirely preclude, the use of such comparisons for validation of the adjoint model. The potential for inverse modeling using the adjoint of GEOS-Chem is assessed in a data assimilation framework using simulated observations, demonstrating the feasibility of exploiting gas- and aerosol-phase measurements for optimizing emission inventories of aerosol precursors.

1 Introduction

Chemical transport models (CTMs) enhance our ability to understand the chemical state of the atmosphere and allow detailed analysis of issues ranging from intercontinental pollution transport to the coupling of anthropogenic processes, regional pollution and climate change. Of particular interest in these realms is explicit consideration of the role of aerosols, the importance of which is well documented. Given

Correspondence to: D. K. Henze
(daven@caltech.edu)

the substantial uncertainty that remains in many aspects of detailed aerosol simulations, it is critical to further examine how the numerous parameters in such models steer their predictions, especially estimates of emissions inventories for aerosols and their precursors. The complexity of the thermodynamic and photochemical processes that govern secondary formation of aerosols precludes simple assessment of the dependence of model predictions on such parameters. Working to arrive at CTMs that more reliably reproduce observations, adjoint modeling is often employed as a method for determining the sensitivity of model predictions to input parameters and for optimizing these parameters to enforce agreement between the model predictions and an observational data set.

Several inverse modeling studies have analyzed sources of aerosols and aerosol precursors on regional scales. As of yet, most studies have been fairly coarse, limited to optimization of a few scaling factors for emissions inventories spanning large domains. Park et al. (2003) used multiple linear regression to estimate annual mean sources of seven types of primary carbonaceous aerosol over the United States. A Kalman filter approach was used to estimate improved monthly emissions scaling factors for NH_3 emissions over the United States using observations of ammonium wet deposition in works by Gilliland and Abbitt (2001) and Gilliland et al. (2003, 2006). Mendoza-Dominguez and Russell (2000, 2001) optimized domain-wide emissions scaling factors for eight species over the eastern United States using observations of gas-phase inorganic and organic species and speciated fine particles. Source apportionment models have also been refined using inverse modeling (Knipping et al., 2006; Schichtel et al., 2006).

Data from satellite observations offer tremendous potential for inverse modeling of aerosols (Collins et al., 2001; Kahn et al., 2004). In order to best exploit these, and other, large data sets, it is desired to extend inverse analysis of aerosol models to global scales and to finer decomposition

of the emissions domains. Such goals require consideration of inverse modeling methods designed for large sets of variable parameters. The adjoint method is known to be an efficient means of calculating model sensitivities that afford examination of numerous parameters, where these values can subsequently be used in tandem with an observational data set for data assimilation. First appearing in the field of atmospheric science in the early 1970s (Marchuk, 1974; Lamb et al., 1975), the method later came to be applied extensively in meteorology, e.g., Talagrand and Courtier (1987); Errico and Vukicevic (1992). In the last decade, the adjoint approach has expanded to include ever more detailed CTMs, beginning with the abbreviated Lagrangian stratospheric model of Fisher and Lary (1995) and the Lagrangian tropospheric model of Elbern et al. (1997). Vukićević and Hess (2000) used the adjoint method to perform a sensitivity study of an inert gas-phase tracer over the Pacific, while Elbern and Schmidt (1999) presented the first adjoint of a 3-D Eulerian CTM to include chemistry. These initial works have been followed more recently by similar development and application of adjoint models of several CTMs: CHIMERE (Vautard et al., 2000; Menut et al., 2000; Schmidt and Martin, 2003), IMAGES (Muller and Stavrou, 2005; Stavrou and Muller, 2006), Polair (Mallet and Sportisse, 2004, 2006), TM4 (Meirink et al., 2006), the California Institute of Technology urban-scale model (Martien et al., 2006; Martien and Harley, 2006), and DRAIS (Nester and Panitz, 2006). The adjoint of the regional model STEM also has been developed (Sandu et al., 2005a) and deployed (Hakami et al., 2005, 2006; Chai et al., 2006).

Of all the previous 3-D adjoint modeling studies, none includes detailed treatment of aerosols, likely owing to the difficult prospect of deriving the adjoint of the model routines dealing with aerosol thermodynamics. The study of Hakami et al. (2005) deals only with inert carbonaceous aerosols, and the work of Dubovik et al. (2004), though global in scale, does not include full chemistry or aerosol thermodynamics. Detailed adjoint modeling of aerosols began with the theoretical investigations of Henze et al. (2004) and Sandu et al. (2005b). However, these are preliminary studies performed on idealized box model systems. In the current work we present the first adjoint of a global CTM that includes dynamics, full tropospheric chemistry, heterogeneous chemistry, and aerosol thermodynamics. We demonstrate the potential value of this tool for quantifying and constraining factors that govern global secondary inorganic aerosol formation. In addition, we note the general usefulness of the adjoint model of GEOS-Chem for a wide variety of applications, such as constraining CO emissions using satellite data (Kopacz et al., 2007¹).

¹Kopacz, M., Jacob, D., Henze, D. K., Heald, C. L., Streets, D. G., and Zhang, Q.: A comparison of analytical and adjoint Bayesian inversion methods for constraining Asian sources of CO using satellite (MOPITT) measurements of CO columns, submitted, 2007.

2 Forward and inverse models

The GEOS-Chem model is used to simulate global aerosol distributions (version 6.02.05 with a horizontal resolution of $4^\circ \times 5^\circ$ and 30 layers up to 0.01 hPa, GEOS-3 meteorological fields). This version of the model includes detailed gas-phase chemistry coupled with heterogeneous reactions, inorganic aerosol thermodynamics, and oxidative aging of carbonaceous aerosols (Park et al., 2004). A few of the specific equations for various model processes are given in Sect. 3.3, along with their corresponding adjoints. We note here that gaseous SO₂ and primary sulfate are co-emitted in GEOS-Chem using a single emissions inventory, referred to as SO_x, which is partitioned between the two species on a regional basis, with sulfate comprising 5% of SO_x emissions in Europe, 1.7% in North America, and 3% elsewhere (Chin et al., 2000).

The standard model has been modified to facilitate the specific inverse modeling goals of the present study. We neglect stratospheric chemistry, which over the course of the short simulations considered here should not have a substantial impact. The standard GEOS-Chem tropospheric chemical mechanism comprises 87 species and 307 reactions integrated using the SMVGEARII solver of Jacobson (1995). We retain this standard chemical mechanism; however, we implement a different numerical solver. The details of this are given in Appendix A. To summarize, we implement a 3rd order Rosenbrock solver that not only facilitates construction of the adjoint model, but also improves forward model efficiency. We also consider using offline concentrations of sulfate aerosol for calculation of photolysis rates and heterogeneous reaction probabilities, see Sect. 3.5.

2.1 Inverse modeling

An adjoint model is used to calculate the gradient of a cost function, J , with respect to a set of model parameters, \mathbf{p} , $\nabla_{\mathbf{p}} J$. For data assimilation applications, the cost function is defined to be

$$J = \frac{1}{2} \sum_{\mathbf{c} \in \Omega} (\mathbf{c} - \mathbf{c}_{\text{obs}})^T \mathbf{S}_{\text{obs}}^{-1} (\mathbf{c} - \mathbf{c}_{\text{obs}}) + \frac{1}{2} \gamma_r (\mathbf{p} - \mathbf{p}_a)^T \mathbf{S}_p^{-1} (\mathbf{p} - \mathbf{p}_a) \quad (1)$$

where \mathbf{c} is the vector of species concentrations mapped to the observation space, \mathbf{c}_{obs} is the vector of species observations, \mathbf{S}_{obs} is the observation error covariance matrix, \mathbf{p} is a vector of active model parameters throughout the model domain, \mathbf{p}_a is the initial estimate of these parameters, \mathbf{S}_p is the error covariance estimate of these parameters, γ_r is a regularization parameter, and Ω is the domain (in time and space) over which observations and model predictions are available. We will sometimes use the notation c and p to represent single elements of the vectors \mathbf{c} and \mathbf{p} . Using the variational approach, the gradient $\nabla_{\mathbf{p}} J$ is supplied to an optimization

routine and the minimum of the cost function is sought iteratively. At each iteration, improved estimates of the model parameters are implemented and the forward model solution is recalculated. In this study, the magnitude of each variable parameter is adjusted using a scaling factor, σ , such that $p = \sigma p_a$. We use the L-BFGS-B optimization routine (Byrd et al., 1995; Zhu et al., 1994), which affords bounded minimization, ensuring positive values for the scaling factors.

Alternatively, for sensitivity analysis, the cost function can be defined as simply a set of model predictions,

$$J = \sum_{g \in \Omega_s} g(\mathbf{c}) \quad (2)$$

where Ω_s is the set of times at which the cost function is evaluated. The desired gradient values are the sensitivities of this set of model predictions to the model parameters.

2.2 Adjoint modeling

Equations for calculating the desired gradients using the adjoint method can be derived from the equations governing the forward model or from the forward model code. The prior approach leads to the continuous adjoint, while the latter leads to the discrete adjoint (Giles and Pierce, 2000). The continuous adjoint equations for CTMs have been derived previously, using methods based upon the Lagrange duality condition (Vukićević and Hess, 2000; Pudykiewicz, 1998; Schmidt and Martin, 2003) or Lagrange multipliers (Elbern et al., 1997). Continuous adjoint gradients may differ from the actual numerical gradients of J , and continuous adjoint equations (and requisite boundary/initial conditions) for some systems are not always readily derivable; however, solutions to continuous adjoint equations can be more useful for interpreting the significance of the adjoint values. Many previous studies have also described the derivation of discrete adjoints of such systems (Sandu et al., 2005a; Muller and Stavrou, 2005). An advantage of the discrete adjoint model is that the resulting gradients of the numerical cost function are exact, even for nonlinear or iterative algorithms, making them easier to validate. Furthermore, portions of the discrete adjoint code can often be generated directly from the forward code with the aid of automatic differentiation tools. Here we present a brief description of the discrete adjoint method for the sake of defining a self-consistent set of notation for this particular paper; we refer the reader to the cited works for further derivations and discussions of continuous and discrete adjoints.

The GEOS-Chem model can be viewed as a numerical operator, F , acting on a state vector, \mathbf{c}

$$\mathbf{c}^{n+1} = F(\mathbf{c}^n) \quad (3)$$

where \mathbf{c} is the vector of all K tracer concentrations, $\mathbf{c}^n = [c_1^n, \dots, c_k^n, \dots, c_K^n]^T$ at step n . In practice, F comprises many individual operators representing various physical processes. For the moment we will simply let F represent a por-

tion of the discrete forward model which advances the model state vector from step n to step $n+1$.

For simplicity, we consider a cost function evaluated only at the final time step N with no penalty term. We wish to calculate the gradient of the cost function with respect to the model state vector at any step in the model,

$$\nabla_{\mathbf{c}^n} J = \frac{\partial J(\mathbf{c}^N)}{\partial \mathbf{c}^n} \quad (4)$$

We define the local Jacobian around any given step as

$$\frac{\partial \mathbf{c}^{n+1}}{\partial \mathbf{c}^n} = \frac{\partial F(\mathbf{c}^n)}{\partial \mathbf{c}^n} = F_c^n \quad (5)$$

Using the chain rule, we can expand the right hand side of Eq. (4) to explicitly show the calculation of \mathbf{c}^N from \mathbf{c}^n ,

$$\nabla_{\mathbf{c}^n} J = (F_c^n)^T (F_c^{n+1})^T \dots (F_c^{N-1})^T \frac{\partial J(\mathbf{c}^N)}{\partial \mathbf{c}^N} \quad (6)$$

Evaluating the above equation from left to right corresponds to a forward sensitivity calculation, while evaluating from right to left corresponds to an adjoint calculation. When K is larger than the dimension of J , which in this case is a scalar, the adjoint calculation is much more efficient (Giering and Kaminski, 1998).

For the adjoint calculation, we define the adjoint state variable λ_c^n ,

$$\lambda_c^n = \frac{\partial J(\mathbf{c}^N)}{\partial \mathbf{c}^n} \quad (7)$$

This can also be expanded,

$$\lambda_c^n = \left[\frac{\partial \mathbf{c}^{n+1}}{\partial \mathbf{c}^n} \right]^T \frac{\partial J(\mathbf{c}^N)}{\partial \mathbf{c}^{n+1}} \quad (8)$$

$$= (F_c^n)^T \frac{\partial J(\mathbf{c}^N)}{\partial \mathbf{c}^{n+1}} \quad (9)$$

The equation above suggests how to solve for the adjoint variable iteratively. Initializing the adjoint variable at the final time step

$$\lambda_c^N = \frac{\partial J(\mathbf{c}^N)}{\partial \mathbf{c}^N} \quad (10)$$

we solve the following equation iteratively from $n=N, \dots, 1$,

$$\lambda_c^{n-1} = (F_c^n)^T \lambda_c^n \quad (11)$$

The value of λ_c^0 is then the sensitivity of the cost function with respect to the model initial conditions,

$$\lambda_c^0 = \nabla_{\mathbf{c}^0} J \quad (12)$$

The scheme above shows why calculating the adjoint variable is often referred to as “reverse integration” of the forward model, as we step from the final time to the initial time. This should not be confused with simply integrating the forward model equations backwards in time.

In order to calculate the sensitivity of J with respect to other model parameters, such as emissions, similar analysis (see, for example, Sandu et al., 2003) shows that the gradient of the cost function with respect to these parameters,

$$\lambda_p^0 = \nabla_p J \quad (13)$$

can be found by iteratively solving the following equation,

$$\lambda_p^{n-1} = (F_p^n)^T \lambda_c^n + \lambda_p^n \quad (14)$$

where the subscripts c and p indicate sensitivity with respect to c and p , respectively, and

$$F_p^n = \frac{\partial F^n}{\partial p} \quad (15)$$

When a penalty term is included in the cost function, the gradient becomes

$$\nabla_p J = \lambda_p^0 + \gamma_r \mathbf{S}_p^{-1} (p - p_a) \quad (16)$$

3 Constructing and validating the adjoint of GEOS-Chem

Here we present the derivation of the adjoint of GEOS-Chem. While the adjoint of the advection scheme is based upon the continuous approach, the remainder of the adjoint model is based upon the discrete formulation, using automatic differentiation tools for assistance. We use the Tangent and Adjoint Model Compiler (TAMC, Giering and Kaminski, 1998), a freeware multipurpose program, and the Kinetic PreProcessor (KPP, Sandu et al., 2003; Damian et al., 2002; Daescu et al., 2003), a public domain numerical library for constructing the adjoint of chemical mechanisms. Always some, if not significant, manual manipulation of the code is required to use such tools. We often combine automatically generated adjoint code with manually derived discrete adjoint code to improve efficiency and transparency of the adjoint model.

Validation of the adjoint model is an important part of introducing an adjoint model of this size and complexity. Discrete portions of the adjoint code have the advantage of being easily validated via comparison of adjoint gradients to forward model sensitivities calculated using the finite difference approximation. The hybrid approach adopted here (discrete and continuous) requires detailed inspection of the adjoint gradients on a component-wise basis as discrepancies owing to the continuous portion are anticipated to obscure such comparisons for the model as a whole. Additional motivations exist for checking the gradients of subprocesses in the model separately and collectively. For large CTMs, it is not feasible to compare adjoint and finite difference gradients for each control parameter, as the finite difference calculation requires an additional forward model evaluation per parameter. However, component-wise analysis affords simultaneous examination of large numbers of sensitivities throughout the

model domain, a much better approach to revealing potential errors than performing validation checks in only a few locations. Furthermore, as GEOS-Chem has many routines common to other models, it behooves us to consider the adjoint of these routines separately.

Forward model sensitivities, Λ , are calculated using the finite difference (brute force) method. For component-wise tests of nonlinear routines, Λ is calculated using the two-sided formula,

$$\Lambda = \frac{J(\sigma + \delta\sigma) - J(\sigma - \delta\sigma)}{2\delta\sigma} \quad (17)$$

while for testing the full model, the more approximate one-sided finite difference equation,

$$\Lambda = \frac{J(\sigma + \delta\sigma) - J(\sigma)}{\delta\sigma} \quad (18)$$

is used in order to minimize the number of required forward model function evaluations. The latter method is also adequate for testing linear components of the model. We use $\delta\sigma=0.1-0.01$ for most tests, which experience showed to be an optimal balance between truncation and roundoff error. For most of these validation tests, it suffices to use a simplified cost function that does not depend on any observational data set, as in Eq. (2), defining g to be a predicted tracer mass, either gas- or aerosol-phase, in a single grid cell, or the total mass burden over a larger spatial domain.

3.1 Aerosol thermodynamics

The equilibrium thermodynamic model MARS-A (Binkowski and Roselle, 2003) is used to calculate the partitioning of total ammonia and nitric acid between aerosol and gas phases. While it is a relatively simple treatment compared to others such as SCAPE (Kim et al., 1993) or ISORROPIA (Nenes et al., 1998), the MARS-A model is still fairly complex. It uses an iterative algorithm to find equilibrium concentrations, considering two primary regimes defined by the ionic ratio of ammonium to sulfate and several sub-regimes defined by conditions such as relative humidity.

Several factors have historically prevented rigorous treatment of aerosol thermodynamics from inclusion in adjoint modeling studies of CTMs, or even adjoint studies of aerosol dynamics (Henze et al., 2004; Sandu et al., 2005b). Division of the possible thermodynamic states into distinct regimes causes many discontinuities in the derivatives, precluding easy derivation of continuous adjoint equations and raising doubts to the value of such sensitivities. Furthermore, several coding tactics often employed in these types of models render them intractable for direct treatment using automatic differentiation tools.

We develop the adjoint of MARS-A in pieces, separating the model into several subprograms, the adjoints of which are then created using TAMC. Tracking variables are added

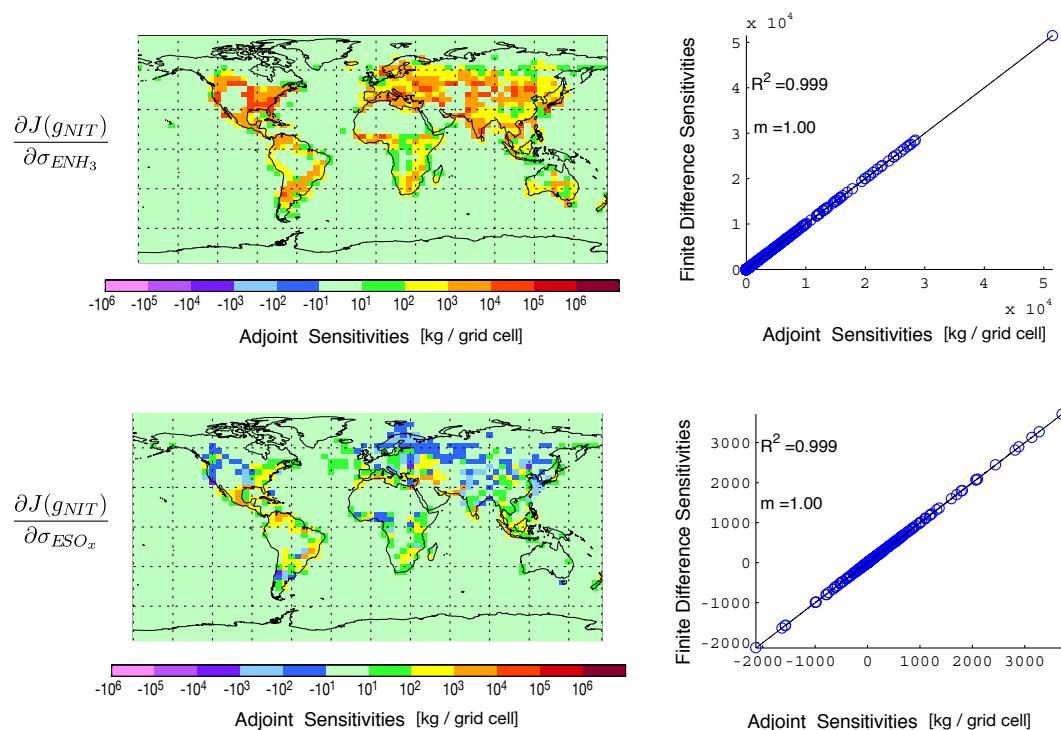


Fig. 1. Thermodynamic adjoint validation. In the left column are the adjoint sensitivities of nitrate aerosol mass at the surface with respect to anthropogenic NH_3 and SO_x emissions scaling factors. In the right column are the adjoint gradients compared to finite difference gradients. The cost function is evaluated once at the end of a week-long simulation that includes only aerosol thermodynamics and emissions of SO_x and NH_3 .

to the forward model routine to indicate which of these sub-routines to call during the adjoint calculation. Initial unequilibrated concentrations at the beginning of each external time step are saved in checkpoint files during the forward calculation. Intermediate values are recalculated from these during the adjoint integration. This type of two-level checkpointing strategy has been shown to optimally balance storage, memory and CPU requirements (Griewank and Walther, 2000; Sandu et al., 2005a).

The accuracy of the resulting adjoint code is tested by comparing adjoint gradients to finite difference gradients calculated using Eq. (17) with $\delta\sigma=0.1$. These comparisons can be made directly throughout the entire model domain by turning off all transport processes. Figure 1 shows comparisons for the sensitivity of surface level nitrate aerosol mass with respect to scaling factors for emissions of surface level anthropogenic SO_x and NH_3 after a week-long simulation. The gradients agree quite well, confirming the accuracy of the thermodynamic adjoint code. Discussion of values of model sensitivities is given in Sect. 4.

3.2 Chemistry

KPP (v2.2) (Sandu et al., 2003; Damian et al., 2002; Daescu et al., 2003) is used to automatically generate code for the

adjoint of the tropospheric chemistry solver, which calculates gradients with respect to the initial species concentrations. We are also interested in the gradient with respect to the emission rates for those species whose emissions are incorporated into the chemical mechanism itself, such as NO_x , (as opposed to those that are simply injected into the model grid cells at intermediate times, such as SO_x). The additional equations for calculating discrete adjoint gradients with respect to reaction rate constants are derived in Appendix B. Though these equations have not been presented previously, KPP does provide the necessary subroutines for solving them.

To assess the accuracy of the adjoints of the chemistry routine, we calculate the sensitivity of the species concentrations at the end of a single chemistry time step (1 h) with respect to the emissions of NO_x (emitted as NO) in a box model test. For this test, the chemical environment is that of a polluted, urban grid cell in the afternoon. Figure 2 shows the ratio $\lambda_{\text{ENO}_x}/\Lambda_{\text{ENO}_x}$ for three separate cases. Using a two-sided finite difference calculation (Eq. 17) with $\delta\sigma_{\text{ENO}_x}=0.1$ leads to agreement within a few percent. The dependence of the internal time step on species concentrations is a feedback not accounted for in the adjoint algorithm; hence, also holding the internal time step fixed at 60 s results in ratios of nearly 1.000 for all species. For comparison, the ratios when

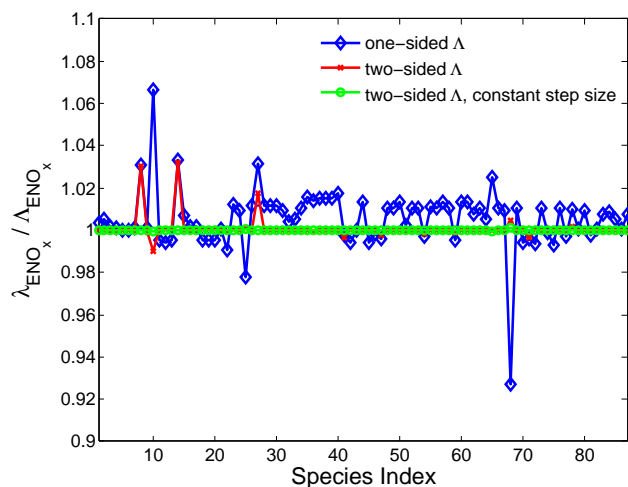


Fig. 2. Chemistry adjoint validation. The ratios of the adjoint to finite difference sensitivities of each species with respect to NO_x emissions are calculated for a 1 h box model simulation. Results are shown for a one-sided finite difference calculation, $\delta\sigma = 0.1$ (blue \diamond 's), a two-sided finite difference calculation (i.e. average of $\delta\sigma = 0.1$ and -0.1 , red \times 's) and a two-sided finite difference calculation with a fixed internal time step of 60 s (green \circ 's).

Eq. (18) is used for Λ_{ENO_x} are also shown, which can differ as much as 8% from unity, demonstrating the nonlinearity of such chemical systems.

The above test was reassuring, yet limited in scope for a global CTM. To test our adjoint model over a wide variety of chemical conditions, we also compare the accuracy of the adjoint derivatives of the chemical mechanism in global simulations over much longer time scales. We turn off all transport related processes in the model and calculate the adjoint and finite difference sensitivities of surface level tracer masses with respect to NO_x emissions in each location after a week-long simulation. As lack of transport leads to unrealistically extreme concentrations, emissions are reduced by an order of magnitude to prevent the chemical systems from becoming too stiff. Many chemical changes associated with aerosols are treated separately from the main tropospheric chemistry mechanism in GEOS-Chem, such as aqueous reactions, dry deposition, chemical aging, and emission of SO_x and NH_3 (Park et al., 2004). The adjoints of these processes are constructed separately (manually and with TAMC) and included in the following tests.

Figure 3 shows the adjoint and finite difference sensitivities of several species with respect to surface level, anthropogenic NO_x emissions scaling factors. We choose to show sensitivities of species such as acetone and methacrolein to NO_x emissions to also highlight the potential value of the adjoint model for analysis of non-aerosol species. We see from these, and similar tests for other active species (not shown), that the sensitivities calculated using the adjoint model consistently agree with those using the finite difference method over a wide range of conditions.

The code generated by KPP allows computation of either the continuous or discrete adjoints of the chemical mechanism. The continuous adjoint equation can be solved faster than the discrete adjoint equation at a given tolerance level, as calculation of the latter requires recalculation of intermediate values from the forward integration and computation of the Hessian during the adjoint integration, see Appendix B. At tight tolerance levels (i.e. very small internal time steps), the results of these methods should converge. However, for tolerance levels appropriate for global modeling, the continuous adjoint is only approximate, as $\lambda + \delta\lambda$, where $\|\delta\lambda\| < C \cdot \text{Tot}$. Given that the computational expense of the Rosenbrock solver increases substantially for tighter tolerance levels (see Appendix A), it is more efficient to use the discrete adjoint, even though this requires an additional forward integration. This is in contrast to the approach of Errera and Fonteyn (2001), who chose to approximate the necessary intermediate values by linearly interpolating from values stored at each external time step, an approach likely more appropriate for their stratospheric chemistry application.

GEOS-Chem accounts for the effect of aerosol concentrations on the radiation available for photolysis reactions and on the available surface area for the heterogeneous reactions included in the main chemical mechanism. The influence of the concentration of sulfate-ammonium-nitrate aerosols on such rates is not currently accounted for in the adjoint model. We assume such an effect is less than 5% (Liao et al., 1999; Martin et al., 2003), especially as the absorbing aerosols (black carbon, mineral dust) are not active variables during these tests. The general agreement between λ and Λ , only the latter of which accounts for this effect, indicates this assumption is adequate, at least for simulations of this length. Further tests indicate that this assumption is valid for most, though not all, cases, see Sect. 3.5.

3.3 Convection, turbulent mixing, and wet removal

Wet removal of tracers in GEOS-Chem is generally treated as a first-order process, leading to discrete forward model equations of the form,

$$c_k^{n+1} = c_k^n e^{-r_{w,k}\Delta t} \quad (19)$$

Since the loss rate $r_{w,k}$ for most species does not depend on any active variables (Jacob et al., 2000), the corresponding adjoint is simply

$$\lambda_k^n = \lambda_k^{n+1} e^{-r_{w,k}\Delta t} \quad (20)$$

The adjoints of these routines are generated using hand-created code, retaining efficiency and legibility. However, the in-cloud formation and cycling of sulfate aerosol from SO_2 is decidedly nonlinear, as the soluble fraction of SO_2 is limited by availability of H_2O_2 , and a fraction of the SO_2 is reintroduced into the gas phase as sulfate when droplets evaporate (Park et al., 2004). Such nonlinearities that span

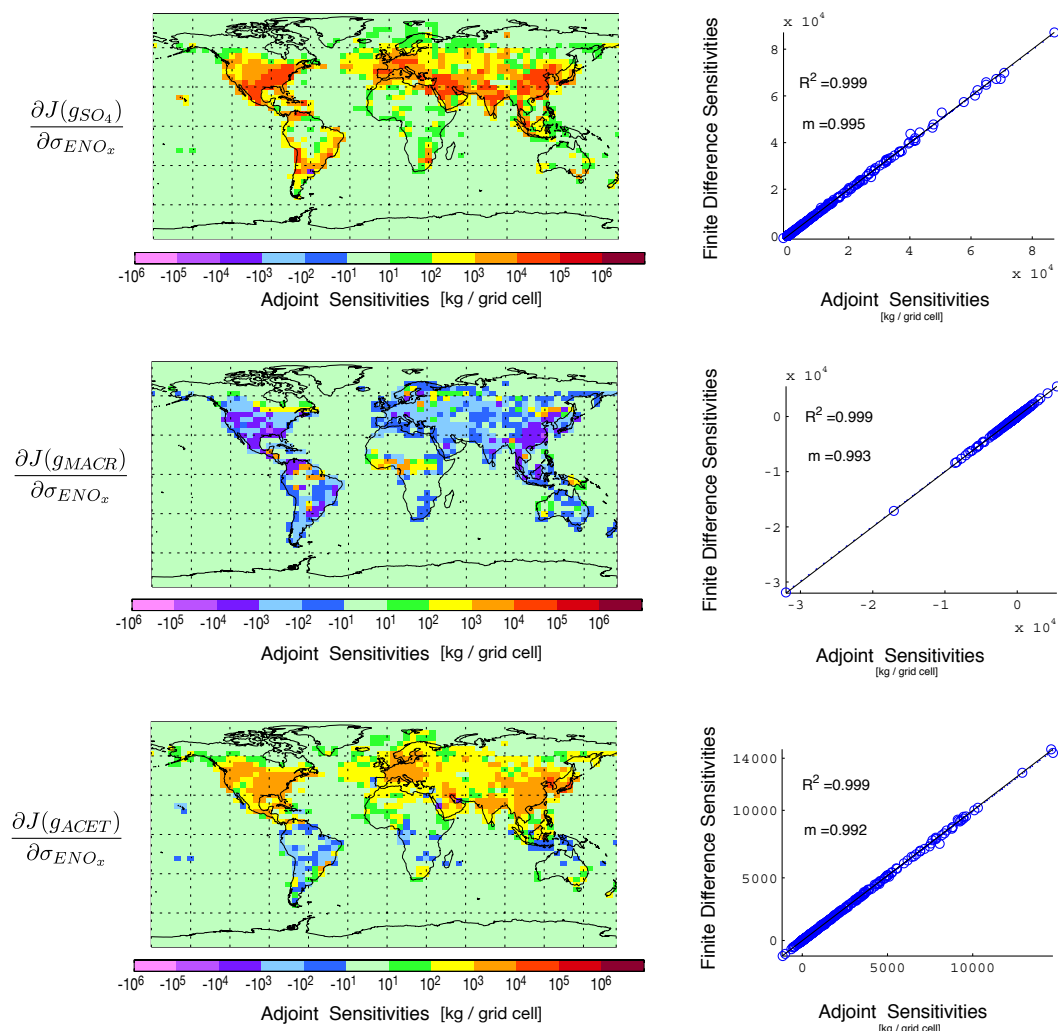


Fig. 3. Chemistry adjoint validation. In the left column are the adjoint sensitivities of sulfate (SO_4), methacrolein (MACR), and acetone (ACET) at the surface with respect to surface level anthropogenic NO_x emissions scaling factors. In the right column are the adjoint gradients compared to finite difference gradients. The cost function is evaluated once at the end of a week-long simulation with only chemistry and emissions $\times 0.1$.

multiple program modules are treated both manually and with the help of TAMC, requiring additional recalculation and checkpointing of intermediate values.

Turbulent mixing in the boundary layer in the forward model is calculated according to a mass-weighted mixing algorithm applied every dynamic time step (30 min for our case),

$$\mu_{k,j}^{n+1} = \frac{\sum_{l=1}^L m_l \mu_{k,l}^n}{m_T} \quad (21)$$

where $\mu_{k,j}$ is the mixing ratio (c/ρ , ρ is the density of air) of tracer k in layer j , m_l is the air mass in a single layer l , m_T is the total air mass in the boundary layer column, and L is the number of layers in the boundary layer. Rewritten in

matrix form, this equation reads,

$$\begin{bmatrix} \mu_{k,1} \\ \vdots \\ \mu_{k,L} \end{bmatrix}^{n+1} = \begin{bmatrix} \frac{m_1}{m_T} & \cdots & \frac{m_L}{m_T} \\ \vdots & \ddots & \vdots \\ \frac{m_1}{m_T} & \cdots & \frac{m_L}{m_T} \end{bmatrix} \cdot \begin{bmatrix} \mu_{k,1} \\ \vdots \\ \mu_{k,L} \end{bmatrix}^n \quad (22)$$

Direct application of Eq. (11) yields the corresponding adjoint equation,

$$\begin{bmatrix} \lambda_{\mu_{k,1}} \\ \vdots \\ \lambda_{\mu_{k,L}} \end{bmatrix}^n = \begin{bmatrix} \frac{m_1}{m_T} & \cdots & \frac{m_1}{m_T} \\ \vdots & \ddots & \vdots \\ \frac{m_L}{m_T} & \cdots & \frac{m_L}{m_T} \end{bmatrix} \cdot \begin{bmatrix} \lambda_{\mu_{k,1}} \\ \vdots \\ \lambda_{\mu_{k,L}} \end{bmatrix}^{n+1} \quad (23)$$

which can be simply written as,

$$\lambda_{\mu_{k,j}}^n = \frac{m_j \sum_{l=1}^L \lambda_{\mu_{k,l}}^{n+1}}{m_T} \quad (24)$$

Deep convection is calculated in the forward model using cumulus cloud fluxes and an RAS type algorithm, see Appendix A of Allen et al. (1996). We calculate the discrete adjoint of this scheme using TAMC, noting that TAMC initially generates code that is accurate, yet several orders of magnitude slower than necessary due to several superfluous loops that have to be removed manually. The numerical scheme for the forward calculation iteratively solves a set of essentially linear equations, with an internal time step of five minutes. If we neglect a single conditional statement that checks only for rare floating point exceptions, then storage or recalculation of the intermediate values is not required for the adjoint calculation.

The adjoint model performance for a simulation including convection, turbulent mixing, and wet deposition is tested by comparison of finite difference sensitivities to the adjoint sensitivities of concentrations of a soluble tracer with respect to its initial concentrations in a location exhibiting strong convection, deposition, and mixing. Horizontal transport, chemistry, and aerosol thermodynamics are turned off for these tests. We use a perturbation of one percent for the finite difference calculation. The ratio λ_c/Λ_c for simulations that are 6 h, 1 d and 3 d in length are 0.9998, 1.0002 and 1.0003, from which we see consistent satisfactory agreement between the two methods. Performance is similar in other tested locations.

3.4 Advection

We implement the adjoint of the continuous advection equations. GEOS-Chem nominally employs a monotonic piecewise parabolic (PPM) advection routine (Colella and Woodward, 1984; Lin and Rood, 1996). Below we briefly show how this scheme can be used to solve the continuous adjoint advection equations and afterwards address some of the issues wedded to this approach. We consider the 1-D example of the advection equation for a tracer in mass concentration units,

$$\frac{\partial c}{\partial t} = -u \frac{\partial c}{\partial x} \quad (25)$$

where u is the wind velocity in the x -direction. The forward numerical model actually solves the flux form of Eq. (25) in terms of the mixing ratio (Lin and Rood, 1996),

$$\frac{\partial(\rho\mu)}{\partial t} = -\frac{\partial(\rho\mu u)}{\partial x} \quad (26)$$

Assuming that the continuity equation for ρ is satisfied, this can be rewritten in the advection form,

$$\frac{\partial\mu}{\partial t} = -u \frac{\partial\mu}{\partial x} \quad (27)$$

Applying the adjoint variable as a Lagrange multiplier and integrating by parts (see, for example, Appendix A of Sandu et al., 2005a), the continuous adjoint of Eq. (27) is

$$-\frac{\partial\lambda_\mu}{\partial t} = \frac{\partial(\lambda_\mu u)}{\partial x} \quad (28)$$

where λ_μ is the adjoint of the mixing ratio. Note that we have assumed that the winds (or any other met fields) are not active variables; taking the adjoint with respect to the meteorology is another task in itself (see, for example, Giering et al., 2005). Applying the simple transform $\hat{\lambda}_\mu = \lambda_\mu/\rho$, and substituting this into Eq. (28), we arrive at the following adjoint equation,

$$-\frac{\partial(\rho\hat{\lambda}_\mu)}{\partial t} = \frac{\partial(\rho\hat{\lambda}_\mu u)}{\partial x} \quad (29)$$

which is similar in form to Eq. (26). If we assume that ρ is relatively constant over a single dynamic time step and that the advection is linear, then we can simply solve Eq. (29) using the same numerical code that was used to solve Eq. (26) in the forward model, scaling the adjoint by $1/\rho$ before and re-scaling by ρ afterwards, which is equivalent to solving Eq. (28).

While the continuous approach was in part adopted for reasons of practicality (the discrete advection algorithm in the forward model not being directly amenable for use with automatic differentiation tools), subsequent investigation indicates that the continuous approach is suitable, if not preferable. This is not surprising, as it is well documented that discrete adjoints of sign preserving and monotonic (i.e. non-linear and discontinuous) advection schemes are not well behaved and can contain undesirable numerical artifacts, see for example Thuburn and Haine (2001), Vukićević et al. (2001), and Liu and Sandu (2006)².

To illustrate the benefits of the continuous adjoint approach for our system, the following numerical test is performed. The sensitivity of aerosol concentrations with respect to concentrations in a neighboring cell six hours earlier are calculated for a meridional cross section of the northern hemisphere. To afford simultaneous calculation of finite difference and adjoint sensitivities throughout this domain, only horizontal advection in the E/W direction is included in these tests. Figure 4 shows finite difference sensitivities calculated using Eq. (18) for several values of $\delta\sigma$ as well as the adjoint gradients. The undesirable nature of the finite difference sensitivities is indicated by negative sensitivities that have no physical meaning. That negative values become more prevalent as $\delta\sigma \rightarrow 0$ indicates such values are caused by discontinuities in the discrete algorithm (Thuburn and Haine, 2001). We can expect that adjoint sensitivities of the discrete advection algorithm would contain similar features, which, despite being numerically precise gradients of the cost function, can result in convergence to undesirable local minimums for data assimilation (Vukićević et al., 2001). Given the importance of transport for analysis of aerosols, use of the continuous approach is deemed preferable to implementing a linear transport scheme with well-behaved discrete adjoints at the cost of forward model performance.

²Liu, Z. and Sandu, A.: Analysis of Discrete Adjoints of Numerical Methods for the Advection Equation, Int. J. Numer. Meth. Fl., submitted, 2006.

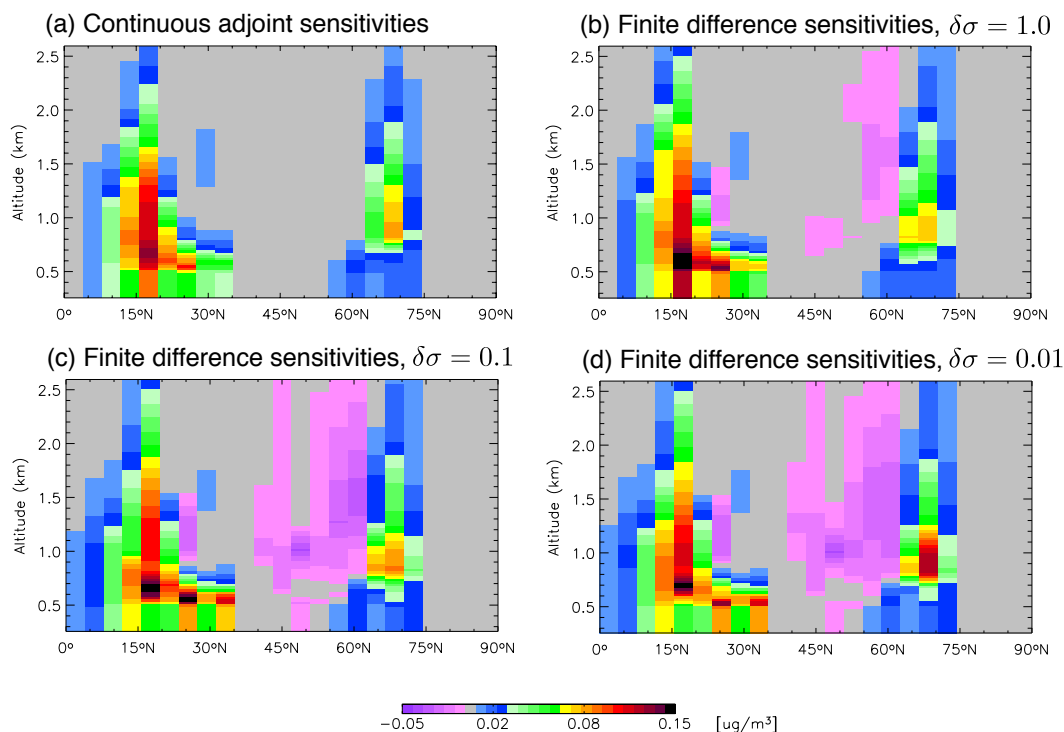


Fig. 4. Sensitivities of aerosol concentrations with respect to concentrations in adjacent cells 6 h earlier considering only E/W advection. Sensitivities are calculated using: (a) continuous adjoint equation and (b)–(d) one-sided finite difference method with perturbations of $\delta\sigma$. The finite difference sensitivities contain more extreme values, including physically meaningless negative sensitivities that become more prevalent as $\delta\sigma \rightarrow 0$.

3.5 Combined performance

Again we compare the gradients calculated using the adjoint model to those calculated using the finite difference method, this time including all model processes. We calculate the sensitivity of global aerosol distributions of sulfate, ammonium, and nitrate to surface emissions of anthropogenic SO_x , NO_x and NH_3 in select locations. As noted previously, such comparisons are quite time consuming to perform on a global scale owing to the expense of the finite difference calculations. Attempting to cover a wide range of conditions, while keeping the number of required calculations within reason, we choose to analyze ten locations for each set of emissions considered, see Fig. 5. The simulations are one day in length, and the cost function (Eq. 2) is evaluated only once at the end of the day. We use a perturbation of $\delta\sigma=0.1$ and Eq. (18) for the finite-difference calculations.

Figure 6 shows the adjoint gradients compared to the finite difference gradients for each of nine relationships. From visual inspection of the scatter plots, it is clear that the agreement is generally within reason given the fact that using a continuous adjoint for advection is expected to cause some amount of discrepancy. Regression lines, slopes, and R^2 values are given for each set of comparisons. The absolute difference between the two methods is often more substantial

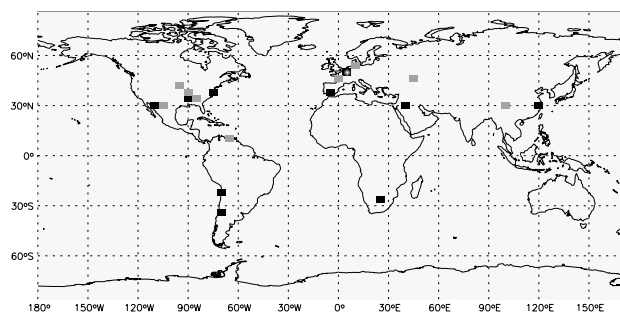


Fig. 5. Select points for accuracy tests. Black locations used for anthropogenic emissions of SO_x and NO_x , grey points for NH_3 , with one overlapping pair in Europe.

for the larger values. As the gradients in a given set usually span several orders of magnitude, many of the slopes are biased by a few such larger values and are not representative of the overall fit. However, accounting for such heteroscedasticity by re-scaling the gradients by $1/p$ or performing weighted regressions that place less emphasis on the larger values still leads to the same general results. Picking twice as many test cells, different test cells, or a different value of $\delta\sigma$ also was not found to substantially alter the overall comparisons.

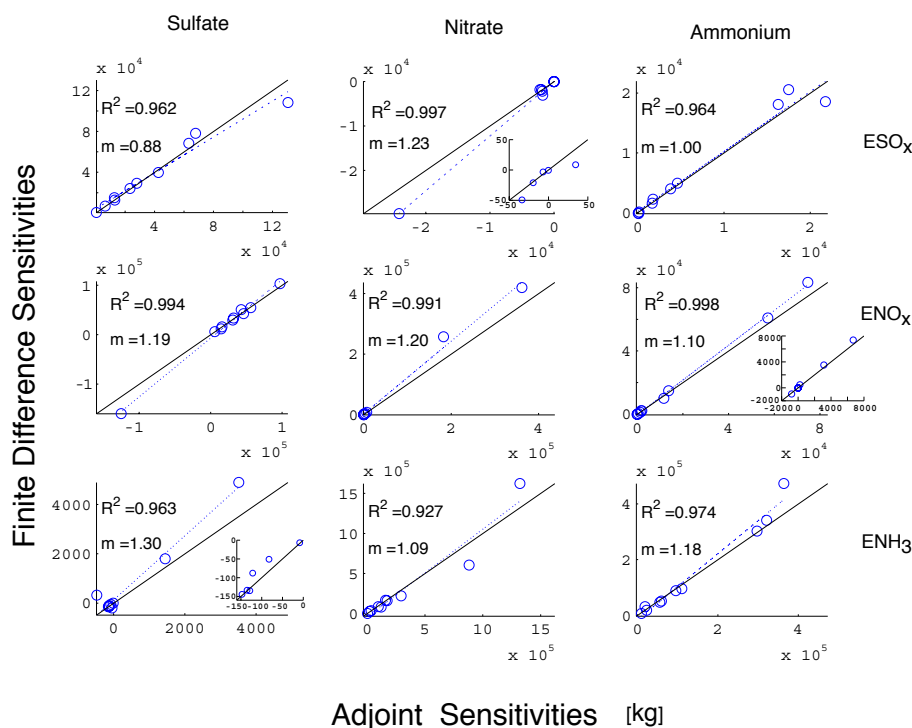


Fig. 6. Full model performance. Comparison of sensitivities of global aerosol burdens (kg) to anthropogenic precursor emissions scaling factors calculated using the adjoint method vs. the finite difference method. A few of the plots contain insets with magnified views of a cluster of points.

Initial comparison (not shown) of gradients for five of the 90 tests showed underestimation of adjoint sensitivities by more than an order of magnitude. Four of these tests were for the sensitivity of sulfate with respect to NH_3 emissions while one was for the sensitivity of nitrate with respect to SO_x emissions. Using offline concentrations for calculation of the contribution of sulfate aerosol to photolysis rates and heterogeneous reaction probabilities in the main tropospheric chemical mechanism for these tests alleviated the discrepancy, demonstrating that while this feedback is generally negligible, it is occasionally quite strong. Future work will extend the adjoint model to account for this feedback.

Napelenok et al. (2006) performed a complementary analysis on a regional scale, calculating the sensitivities of local aerosol distributions with respect to domain-wide precursor emissions over the United States with a forward sensitivity method (DDM-3D), using finite-difference calculations to check their results. While they found similarly good agreement for the more direct relationships (such as sensitivity of sulfate with respect to SO_2 emissions, or ammonium with respect to NH_3 emissions), they had difficulty verifying the variability in the sensitivities of some of the more indirect relationships (such as the sensitivity of sulfate to NH_3 emissions or nitrate to SO_2 emissions). Granted, they used the more complex and rigorous thermodynamic model ISOR-ROPIA; they suggested that such discrepancies were due to

numerical diffusion, with spatial oscillations of the sensitivities indicative of errors due to transport.

In our tests, transport does not drastically degrade the consistency of the correlation between the two approaches; all of the R^2 are near unity. There is, however, some amount of bias in the comparisons, as indicated by slopes ranging from 0.8 to 1.3, and this does appear to be a result of transport. Figure 7 contains scatter plots of the sensitivities of sulfate with respect to NO_x emissions for several additional tests. Panel (a) shows the results when advection is turned off. This leads to improved agreement, $m=1.03$, compared to the center left panel of Fig. 6; hence, the source of this bias is presumably advection. As shown in Fig. 4, the adjoint gradients are likely smoother and more physically meaningful than the finite difference sensitivities.

To assess the extent to which using the continuous adjoint of advection hinders this approach to validating the adjoint model as a whole, we perform additional tests, the results of which are shown in Fig. 7. Including advection, but evaluating the cost function only in a single location, rather than globally, leads to a very unsmooth adjoint field and triggers many nonlinear and discontinuous aspects of the numerical scheme in a manner inconsistent with advection of the relatively smooth concentration field in the forward model; hence, agreement between adjoint and finite difference gradients under these conditions is worse, see panel (b). All of the

tests so far have been based on a single evaluation of the cost function at the end of a day-long simulation. The effects of changing the assimilation window (the time between consecutive evaluations of the cost function) and the total simulation length are shown in panels (c) and (d). Doubling both the simulation length and the assimilation window to two days leads to an increased discrepancy, panel (c). Again, such behavior is likely owing to discrepancies between the finite difference and adjoint sensitivities of the advection scheme that can accumulate when integrating such sensitivities over several other nonlinear processes. Doubling only the simulation length but maintaining a one-day assimilation window improves the agreement, panel (d), as forcing from additional observations outweighs spurious discrepancies from advection.

Finally, we consider a more realistic example. Model predictions are compared to measurements of aerosol nitrate from the IMPROVE network of monitoring stations (<http://vista.cira.colostate.edu/improve/>). The sensitivities of the error weighted squared difference between predicted and observed nitrate aerosol with respect to natural NH_3 emissions scaling factors are shown in Fig. 8. The cost function is evaluated regionally only on the U.S. East Coast (72.5°W – 82.5°W), and the model is run for ten days starting 1 January 2002. Daily average measurements are assimilated during three of the ten days. Also shown is a comparison between the adjoint sensitivities and finite difference sensitivities evaluated for the same domain. That the overall discrepancy is not much different from the simple 24 h tests (Fig. 6, or Fig. 7, panel b) increases our confidence in the ability of short tests to diagnose the model's performance in practical applications.

Overall, we find the accuracy of the adjoint gradients to be satisfactory. The adjoint model clearly captures the dependence of inorganic aerosol burdens on the chemical and thermodynamic interactions that lead to their formation. While using the continuous adjoint of advection makes this verification process more laborious, we have characterized the discrepancies for future reference.

3.6 Computational efficiency

Here we report computational resource requirements for running the adjoint model of GEOS-Chem on a Linux workstation with dual Intel Itanium 1.5 GHz processors and 4 GB of RAM. The adjoint model utilizes multiple processors on shared memory architectures as efficiently as the forward model. It requires 16 KB of checkpoint storage space per simulated day per grid cell; this amounts to 11 GB of storage space per week with the current model configuration. This is comparable to the storage requirements of other adjoint models of CTMs such as STEM, 40 KB per day per cell (Sandu et al., 2005a), or the CIT model, 100 KB per day per cell (Martien et al., 2006), taking into account that the time step is 30 min in GEOS-Chem (for this study), 15 min for STEM,

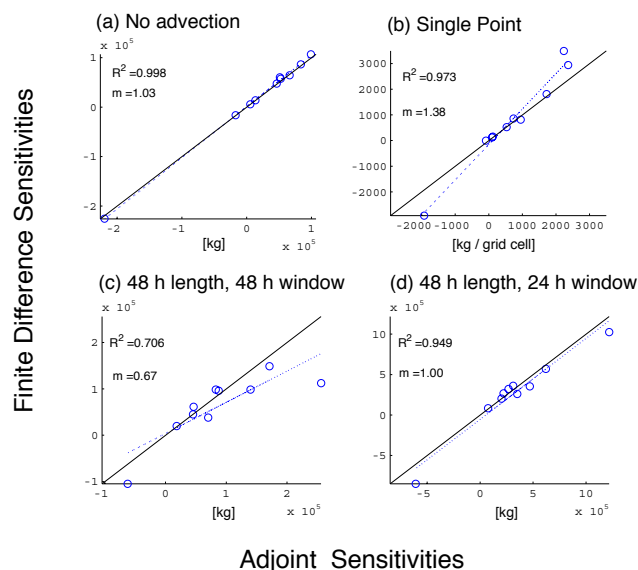


Fig. 7. Effects of advection. Comparison of sensitivities of sulfate burdens to NO_x emissions scaling factors calculated using the adjoint method vs. the finite difference method. The base case (center left panel of Fig. 6) employs the standard PPM advection scheme, and the cost function is evaluated globally once at the end of a 24 h simulation. These cases differ from the base case in the following manner: (a) advection is turned off; (b) the cost function is evaluated in only a single region; (c) both the assimilation window and total simulation length are increased to 48 h; (d) the simulation length is increased to 48 h while the cost function is evaluated every 24 h.

and 3 min for the CIT model. The computational cost of the adjoint model (backward only) of GEOS-Chem is 1.5 times that of the forward model, requiring 2.5 h for a week long iteration (forward and backward). Adjoint models of other CTMs report this ratio as: STEM: 1.5, CHIMERE: 3–4, IMAGES: 4, Polair: 4.5–7, CIT: 11.75. We see that the adjoint of GEOS-Chem is quite efficient; in general, adjoint codes that are derived by hand or use specialized tools such as KPP are most efficient. Such efficiency is the trade-off for the labor involved in manually constructing an adjoint model of this size and complexity.

4 Sensitivity analysis

In this section we demonstrate how the adjoint model can be used as an efficient method of investigating the sensitivity of modeled aerosol concentrations to their precursor emissions. Sensitivity calculations for the full model are performed for a week-long simulation. Figure 9 shows the sensitivity of global burdens of sulfate, nitrate and ammonium aerosol to surface level emissions of anthropogenic SO_x , NO_x and NH_3 . The cost function is evaluated once daily. Other results retrieved from the same calculations (not

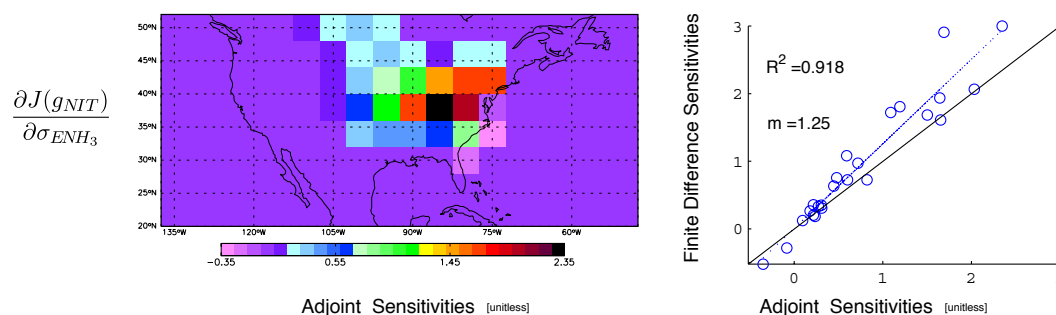


Fig. 8. Sensitivities with respect to the error weighted squared difference between predicted and observed nitrate aerosol from the IMPROVE network for the first ten days of January, 2002. The cost function is evaluated only on the U.S. East Coast (72.5° W–82.5° W). Shown are the sensitivities of the cost function with respect to natural NH₃ emissions scaling factors. On the right are the same quantities compared to finite difference sensitivities.

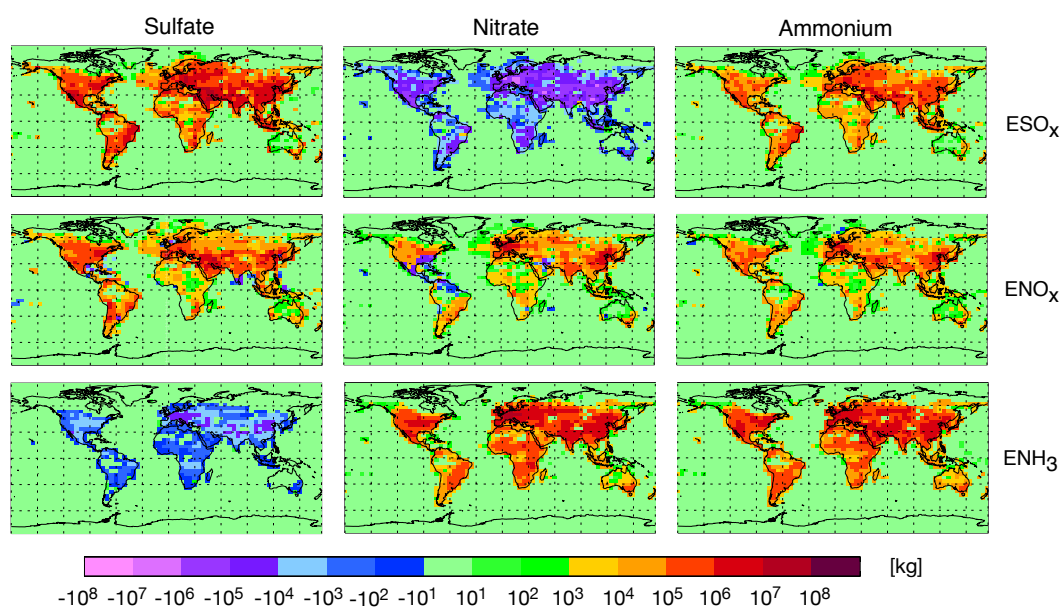


Fig. 9. Sensitivities of global burdens of sulfate, nitrate and ammonium aerosol to anthropogenic SO_x, NO_x and NH₃ emissions scaling factors calculated using the adjoint model for a week-long simulation.

shown) are sensitivities of these species with respect to the following emissions: stack SO_x, stack NO_x, biofuel SO₂, biomass burning SO₂, ship SO₂, biofuel NH₃, biomass burning NH₃, and natural NH₃.

The sensitivities in Fig. 9 encompass a wide range of relationships between aerosols and their primary precursors. Some of these relationships are practically intuitive, such as the sensitivities of sulfate to SO_x emissions or of nitrate to NO_x emissions, both of which are generally large and positive. The sensitivity of ammonium to emission of NH₃ is also positive, and the sensitivities of ammonium to SO_x and NO_x emissions are always positive, owing to uptake of NH₃ on inorganic aerosol by sulfate and nitrate.

Some of the relationships in Fig. 9 are less obvious, such as the negative sensitivity of sulfate to emissions of NH₃. This effect is smaller in magnitude than some of the others,

because the relationship between NH₃ emissions and sulfate aerosol concentrations is less direct. As total sulfate is conserved in the MARS-A aerosol equilibrium model, this effect is not due to thermodynamic interactions between ammonium and sulfate. The only species directly affected by NH₃ or ammonium concentrations are nitrate and nitric acid, via thermodynamic interactions. Therefore, the relationship between NH₃ and sulfate is dictated by the interactions between sulfate and nitrate, and, hence, NO_x. The sensitivity of nitrate to SO_x is largely negative, owing to thermodynamic competition between nitrate and sulfate for ammonium. The sensitivity of nitrate to NH₃ is entirely positive, due to the necessary presence of excess NH₃ for HNO₃ to condense. The combination of these two effects explains the overall negative relationship between sulfate and emissions of NH₃.

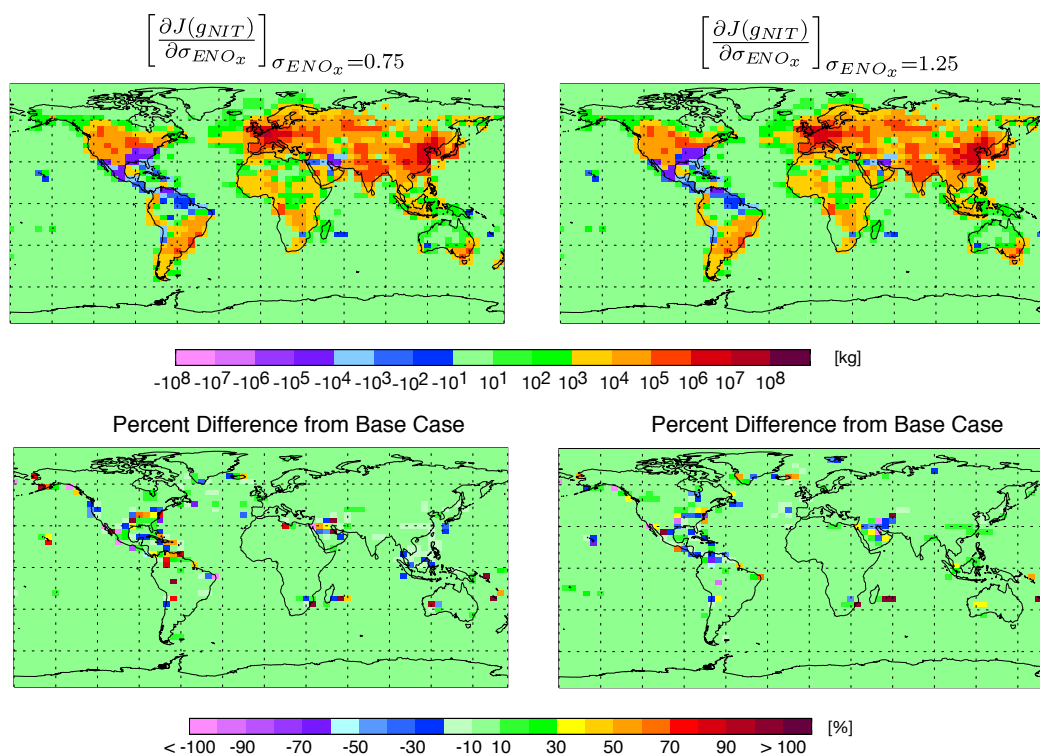


Fig. 10. Sensitivities of nitrate aerosol to emissions of anthropogenic NO_x when the emission inventories are scaled by factors of 0.75 and 1.25, and the percent difference between these sensitivities and those calculated with the base case ($\sigma_{\text{ENO}_x}=1.0$), shown in Fig. 9.

Within the global trends noted above, there is also much discernible local variability. For example, there are a few locations where the sensitivity of sulfate to NH_3 emissions changes abruptly from predominantly negative to locally positive. Some of these actually correspond to similarly abrupt shifts between areas that are sulfate-poor to areas that are sulfate-rich, such as the tip of South America and immediately west of the Iberian Peninsula. In other conditions or times of the day, emission of NO_x can actually lead to a decrease in nitric acid, and, hence, nitrate.

While the adjoint model accounts for nonlinearities in the relationships between emissions and aerosols, the results of the adjoint calculation are still merely tangent linear derivatives (gradients) which are likely to be valid over only a limited range of values for the parameters (emissions). We explore the robustness of the aerosol sensitivity calculations with respect to the magnitude of the emissions. Figure 10 shows the sensitivity of nitrate with respect to NO_x emissions calculated when the emissions are multiplied by uniform scaling factors of 0.75 and 1.25; the relative differences between these values and base case sensitivities shown in Fig. 9. The sensitivities can differ substantially on a point to point basis ($>50\%$), particularly near boundaries between the positive and negative sensitivities or in areas where the sensitivities are very small. The differences are generally much less ($<20\%$) in areas with the largest sensitivities such

as Europe, Eastern Asia and the Eastern United States. Despite these relative differences, the sensitivity field, viewed on the global (log) scale, remains nearly identical to the base case values. While individual sensitivities may be valid only over a limited range, the sensitivity field as a whole appears fairly robust.

Overall, the adjoint model is a promising tool for examining the dependence of aerosol concentrations on emissions. We note that the time required to calculate all of these sensitivities was less than 10 times the cost of a single forward model evaluation, while obtaining these results using the finite difference method would have required >5000 times the cost of a forward run.

5 Inverse modeling tests

Several inverse modeling tests are performed to assess the capabilities of the adjoint model in a data assimilation application. Using the twin experiment framework, pseudo observations, c_{obs} , are generated with the forward model using a base set of emissions parameters, $p=p_a$. An active subset of the parameters used to generate these observations is then perturbed using scaling factors, $\sigma=p/p_a$, each of which is allowed to vary independently in every grid cell for each emitted species. The inverse model uses the pseudo-observations

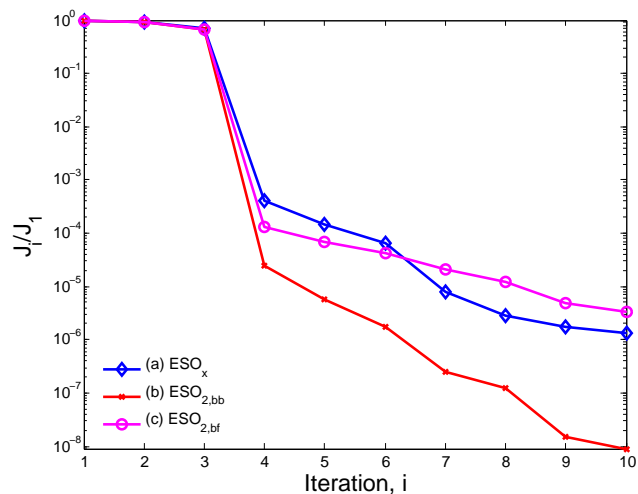


Fig. 11. Cost function reduction for tests DA1. A uniform perturbation is applied to emission inventories of (a) SO_x (b) biomass burning SO_2 (c) biofuel SO_2 . Complete daily measurements of sulfate aerosol are utilized for the data assimilation during a week-long simulation.

to recover the original unperturbed values of these active parameters.

We begin by generating a week-long set of observational data using the forward model with all scaling factors set equal to unity. For these initial tests, we perturb one set of emissions by re-scaling the emissions in every cell by a factor of two, and we use observations in every grid cell once every 24 h to force the data assimilation. As there is no error in these observations, equal weight is ascribed to each (S_{obs}^{-1} is the identity matrix), and the error covariance of our initial (perturbed) estimate of the emissions scaling factors is infinite (S_p^{-1} is zero). Such conditions are unrealistic and serve only to test the adjoint model under the most ideal conditions possible.

In the first set of tests (DA1), we perturb the emission inventories of (a) surface level anthropogenic SO_x , (b) biomass burning SO_2 and (c) biofuel SO_2 . We assimilate observations of sulfate for the week of 1–7 July 2001. Figure 11 shows the progression of the normalized (divided by the initial value) cost function at iteration i during the optimization procedure, J_i/J_1 . The cost function quickly reduces by at least five orders of magnitude in each case. The correct emissions inventories are essentially entirely recovered.

In the next test (DA2), we perturb the emission inventory of NH_3 from anthropogenic sources, and assimilate observations of aerosol ammonium. This is a slightly more difficult inversion as ammonium measurements alone do not fully constrain NH_3 emissions (Gilliland et al., 2006). As demonstrated in Sect. 3.5, ammonium is indirectly, yet appreciably, coupled to gas-phase oxidants. Utilizing observations of O_x (O_3 , NO_2 and NO_3) in conjunction with ammonium observations noticeably increases the convergence rate over

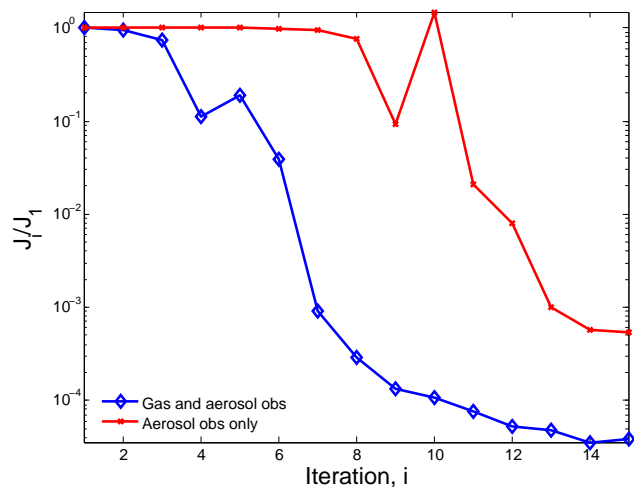


Fig. 12. Cost function reduction for tests DA2. A uniform perturbation is applied to emission inventories of anthropogenic NH_3 . Complete daily measurements of (red-crosses) ammonium aerosol and (blue-diamonds) ammonium aerosol and gas-phase O_x are utilized for the data assimilation during a week-long simulation.

using either type of observations alone, see Fig. 12. This demonstrates, albeit in a highly idealized fashion, the potential for exploiting multi-phase measurements as constraints for aerosol modeling.

The final test (DA3) attempts to mimic a slightly more realistic scenario than the previous tests: improving estimates of global anthropogenic SO_x and NO_x emission inventories using surface measurements of sulfate, nitrate, and ammonium aerosol. In this case, the emissions inventories are perturbed regionally by 5–30% with an additional random factor of order 5%. For example, the anthropogenic SO_x and NO_x emissions in North America are perturbed by factors of $0.8+r$ and $0.85+r$, respectively, while emissions in Asia are perturbed by factors of $1.2+r$ and $1.3+r$, where r is a random number uniformly distributed between 0 and 0.05. The error covariance matrix S_p is calculated using an ascribed error of 100% and is assumed to be diagonal. Observations are used once per day in only half of the land-based surface grid cells. The reduction of the cost function after 15 iterations is shown in Fig. 13. The difference between the true emission inventories for SO_x and NO_x and the estimated inventory at the first and final iterations are shown in Fig. 14. While there are substantial improvements in the SO_x emissions and the NO_x emissions in Europe and Asia, the NO_x emissions in North America have yet to converge. Although the cost function has reduced by nearly two orders of magnitude, the optimization procedure has clearly yet to reach a minimum. In applications of this type, the procedure is often halted according to an appropriate convergence criteria. Further iterations might be justified; however, care must be taken to avoid overly minimizing the predictive error component of the cost function at the sake of generating noisy solutions.

6 Summary and conclusions

The derivation of the adjoint model of GEOS-Chem has been presented in a piecewise fashion. We have implemented the first adjoint of an aerosol equilibrium thermodynamic model (MARS-A, Binkowski and Roselle, 2003), derived using the automatic differentiation tool TAMC (Giering and Kaminski, 1998), which required significant manual pre- and post-processing owing to the structure and complexity of the code. To facilitate construction of the adjoint of the GEOS-Chem gas-phase chemical mechanism, we implemented a Rosenbrock solver using the KPP numerical library (Sandu et al., 2003). This has allowed for automatic generation of the adjoint of the chemical mechanism and also improved forward model performance (see Appendix A). The adjoints of wet removal, deep convection, and turbulent mixing were derived manually and with the aid of TAMC. We have used the continuous adjoint method to treat advection, wherein the same numerical algorithm is used to solve the continuous adjoint advection equation as was used for tracer advection in the forward model.

All aspects of the adjoint model have been tested both separately and together by comparing the adjoint gradients to finite difference gradients. Each individual discrete adjoint routine showed satisfactory performance over a wide range of conditions. The adjoint gradients of the cost function evaluated using the full model are well correlated with the numerical gradients, as measured using finite difference calculations, with most $R^2 > 0.95$. The hybrid approach adopted here avoids physically unrealistic noise associated with discrete adjoints of nonlinear and discontinuous advection schemes and does not entirely preclude validation of the adjoint model as a whole via comparison to finite difference gradients. Such comparisons are understandably unrevealing when considering sparse or infrequent data; however, in both ideal test calculations with smooth adjoint forcings and realistic tests of week-long sensitivities of predictions of actual aerosol observations, the comparisons are consistent enough to ensure proper derivation of the adjoint. Nevertheless, this treatment necessitated additional inspection of model performance on a component-wise basis. While these benchmarks set the standard for further use and development of this adjoint model, future applications may require additional testing.

The adjoint model clearly demonstrates the importance and relative strengths of many complex nonlinear relationships connecting concentrations of aerosol species and their precursor emissions. Though indirect, relationships such as the dependence of sulfate aerosol concentrations on emission of NH_3 or NO_x are captured by the adjoint model and can be determined globally in an efficient manner. The sign and magnitude of many of these sensitivities exhibit a rich array of features owing to the influence of environmental factors, such as the sulfate to ammonium ratio, cloud processing of SO_2 , and variability in the NO_x and O_x levels.

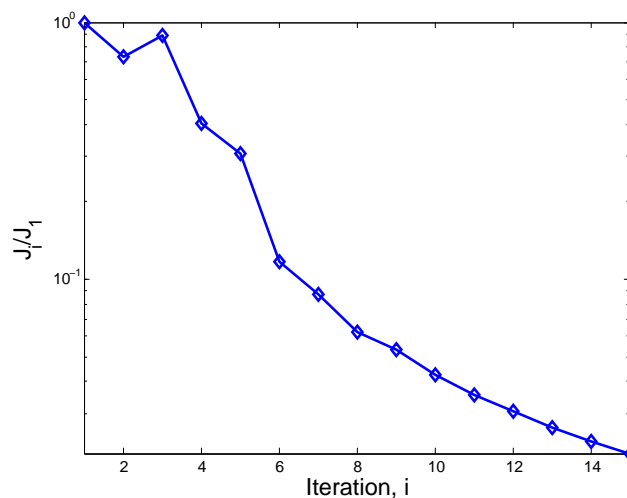


Fig. 13. Cost function reduction for tests DA3. Emissions inventories of anthropogenic SO_x and NO_x emissions are perturbed regionally and optimized simultaneously utilizing sparse daily measurements of aerosol sulfate, ammonium, and nitrate during a week-long simulation.

We have also demonstrated the capabilities of the adjoint model in mock data assimilation applications. An adjoint model of this type allows for the possibility of exploiting multi-phase observations to constrain emissions of aerosol precursors. Here we have focused on regional variability of the emissions inventories, though the emissions can also be adjusted on a temporal basis. For real data assimilation projects, many application specific issues inherent in this type of inverse modeling have yet to be resolved, such as specification of the error covariance matrices \mathbf{S}_{obs} and \mathbf{S}_p . The dependence of adjoint model performance is known to depend strongly on such factors (Chai et al., 2006), proper formulation of which is necessary to ensure scaling of the inventories that are physically realistic (Stavrakou and Muller, 2006). Real world application will also likely require conditioning of the cost function to improve convergence rate (Meirink et al., 2006) and tuning of the regularization parameter (Hakami et al., 2005).

Subsequent studies will focus on expanding the adjoint model to capture feedbacks such as the effect of sulfate aerosol concentrations on photolysis rates and heterogeneous reaction probabilities, seen here to occasionally be quite important. Work on the adjoint of the aerosol equilibrium model ISORROPIA (Nenes et al., 1998) is also in progress. Further application of the GEOS-Chem model will focus also on the exploitation of multi-phase measurements from sources such as surface stations, aircraft, and satellites as model constraints. The adjoint of GEOS-Chem has already been used to constrain emissions of carbon monoxide from Asia using satellite (MOPITT) measurements (Kopacz et al., 2007¹), demonstrating the potential for addressing a wide range of scientific questions with this type of inverse model.

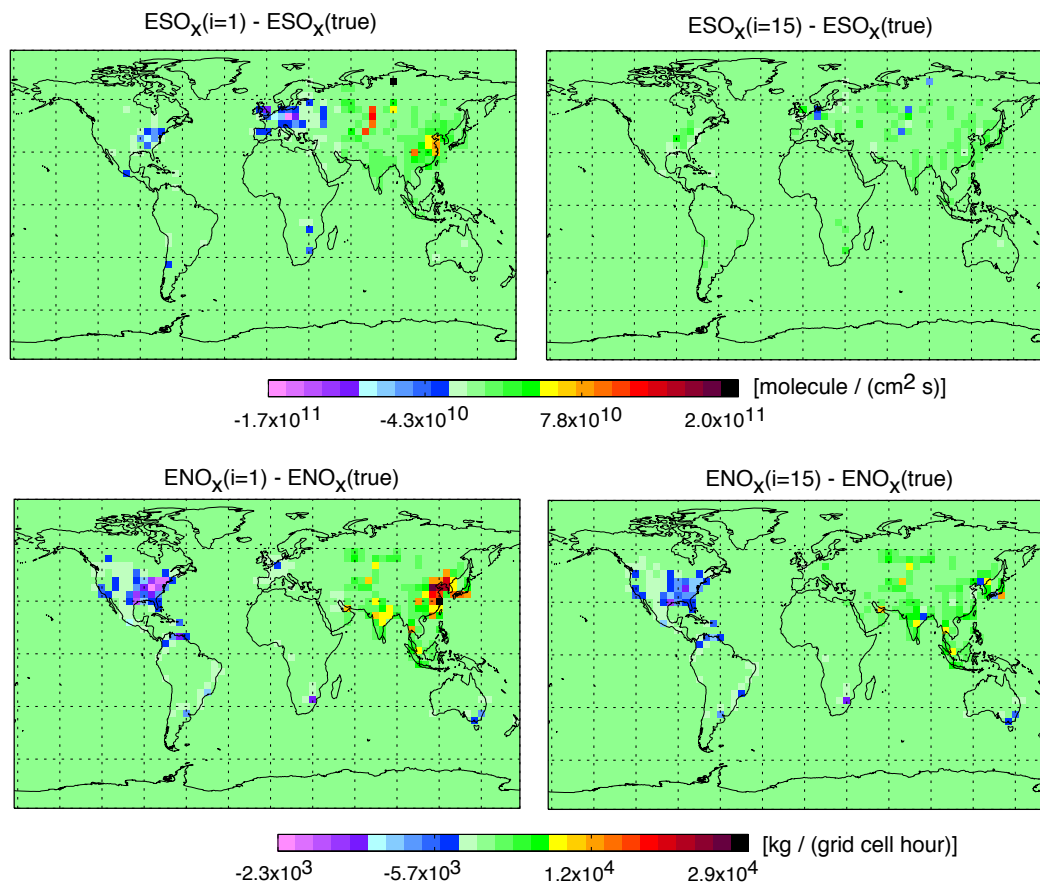


Fig. 14. Emissions inventory estimates for test DA3. Difference between the estimated emission inventory at iteration i and the “true” inventory, which was used to generate the pseudo-observations. Results are shown for the initial estimate (left column) and after 15 iterations (right column).

Appendix A

Implementation of a Rosenbrock solver and comparison to SMVGEARII

Solving large systems of chemical rate equations in CTMs requires the use of special numerical tools, or solvers, that are specifically designed for this purpose. Taking the adjoint of such solvers manually, or using generic automatic differentiation tools, can be an onerous task. We desire to create the adjoint of the full chemical mechanism in GEOS-Chem using the KPP software library (Sandu et al., 2003; Daescu et al., 2003; Damian et al., 2002), which is a set of tools specifically built for automatic differentiation of chemical mechanisms and the numerical algorithms used to solve these systems. In order to make use of these tools, we must first implement the KPP generated numerical integration routines in the forward model. We investigate the feasibility and ramifications of replacing the current solver in GEOS-Chem, SMVGEARII (Jacobson, 1995), with a KPP generated Rosenbrock solver. We consider the amount of work required to make such a

switch, the efficiency of the Rosenbrock solver compared to the SMVGEARII solver, and the overall effect that such a switch has on the model predictions after a week-long simulation.

After manually translating the SMVGEARII mechanism input files to KPP input files, the KPP tools easily generate a set of Fortran code that solves the given system for a variety of supported Rosenbrock type integrators in a box model setting. Minimal manual adjustment to this code was required to interface with the 3-D GEOS-Chem model and to allow support for OpenMP parallelization. Some amount of modifications to the KPP code itself will be required to fully automate this process.

Next we consider the efficiency of the Rosenbrock solver and the SMVGEARII solver in a global simulation with only chemistry. For each species, in every cell, we compare the concentrations from benchmark solutions at the end of a day-long simulation to concentrations from a reference solution for each solver. The benchmark calculations span a set of tolerance levels $\{10^{-1} \leq \text{RTOL} \leq 10^{-5}, 10^6 \text{ molecules cm}^{-3} \geq \text{ATOL} \geq 10^{-2}$

molecules cm^{-3} while the reference solutions were computed using tight tolerances ($\text{RTOL}=10^{-8}$, $\text{ATOL}=10^2$ molecules cm^{-3}). RTOL and ATOL are the relative and absolute error tolerance levels, respectively. Looser tolerance levels result in repeated failure to converge in numerous grid cells.

To assess the accuracy of the two methods, following Sandu et al. (1997) we define the significant digits of accuracy (SDA) as

$$\text{SDA} = -\log_{10}(\max_k ER_k)$$

where ER_k is a spatially modified root mean square norm of the relative error of the benchmark solution ($\hat{c}_{k,j}$) with respect to a reference solution ($c_{k,j}$) for species k in grid cell j ,

$$ER_k = \sqrt{\frac{1}{|\theta_k|} \cdot \sum_{j \in \theta_k} \left| \frac{c_{k,j} - \hat{c}_{k,j}}{c_{k,j}} \right|^2}$$

For Θ total grid cells, θ_k is the set of all locations of significant concentrations of species k , $\{0 \leq \theta \leq \Theta : c_{k,j} \geq a\}$. A threshold value of $a=10^6$ molecules cm^{-3} is chosen to avoid inclusion of errors from locations where concentrations of a given species are less than chemically meaningful values.

We present the results in the form of a work – precision diagram, wherein the value of SDA for each test is plotted versus the average computational expense for the solver to integrate the chemical mechanism for one hour. When calculating this average, we do not consider the time required during the initial six hours of the simulation, as each solver requires a bit of “spin up” time in order to adjust internal time steps to values more appropriate than the default starting step size according to the stiffness of the local system. Such spin up time is negligible with respect to the total computational cost of any simulation longer than a few days.

Figure A1 shows the work-precision diagram for the global benchmark simulations. The Rosenbrock solver is nearly twice as efficient as the SMVGEARII solver during these tests. Based on this analysis, we choose to run the Rosenbrock solver at tolerance levels that yield an SDA of ~ 1.0 as the standard setting for this work.

For practical applications, we are interested in the difference in the total model predictions, including all model processes, incurred by switching to the Rosenbrock solver. We compare the daily average concentrations after a week-long simulation, including all model processes, calculated using the new standard Rosenbrock settings versus the standard SMVGEARII settings. Figure A2 shows the values of ER_k for each species k using the Rosenbrock solver to generate the test solution and SMVGEARII for the reference solution. This figure shows that after switching to this Rosenbrock solver, the solution is changed by less than 10% for most species. The difference is larger, between 10 and 15%, for HNO_2 , HNO_4 , IAP,

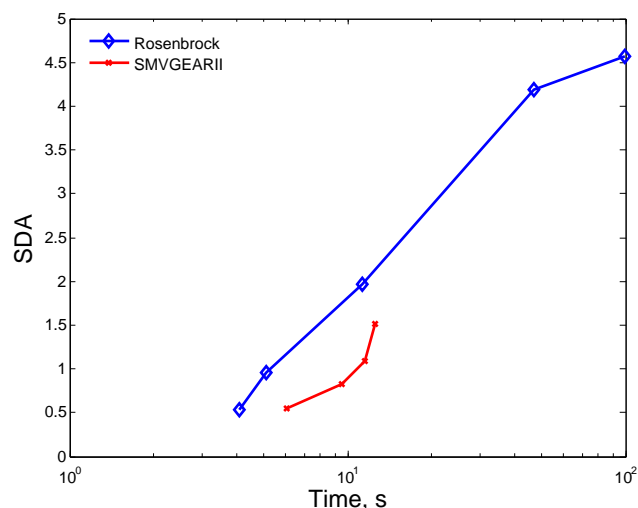


Fig. A1. Work-precision diagram for the Rosenbrock (blue circles), and SMVGEARII (red crosses) chemical solvers. Each solver is implemented in the 3-D model and run for one day using a 1 h external chemical time step. Plot shows the average time taken per external chemical time step versus the significant digits of accuracy (SDA) achieved. Tests performed using dual 1.5 GHz Itanium processors.

INO_2 , ISOP, N_2O_5 , NO, NO_2 , PP, and RIP (for full definition of species, see <http://www.env.leeds.ac.uk/~mat/GEOS-CHEM/GEOS-CHEM.Chemistry.htm>). Determining whether or not this is an actual improvement in the accuracy of the forward model itself would require further comparison to observations. At the very least, the switch results in an improvement in the numerical solution of the forward model equations for slightly less computational cost.

Overall, while a more detailed analysis (requiring optimization of specific species tolerance levels and the parameters that control internal step size expansion and contraction) is necessary to determine unequivocally which method is more efficient, in our experience, not only is the Rosenbrock method desirable because of its differentiability, but it also appears to improve forward model performance by providing more accurate solutions to the model’s chemical mechanism than the SMVGEARII solver for less computational expense. We have reported only the results using the Rodas-3 set of Rosenbrock coefficients; however, additional tests were performed using the other available sets (Ros-2, Ros-3, Ros-4, Rodas-4), and the trends were similar. It must also be emphasized that these comparisons should not be generalized to other platforms or CTMs; the SMVGEARII algorithm is designed to perform most efficiently on vector platforms by re-ordering the grid cells every external chemistry time step, an operation which serves only to increase the cost of this method by $\sim 5\%$ on non-vector machines such as those used in this study, and most other GEOS-Chem studies.

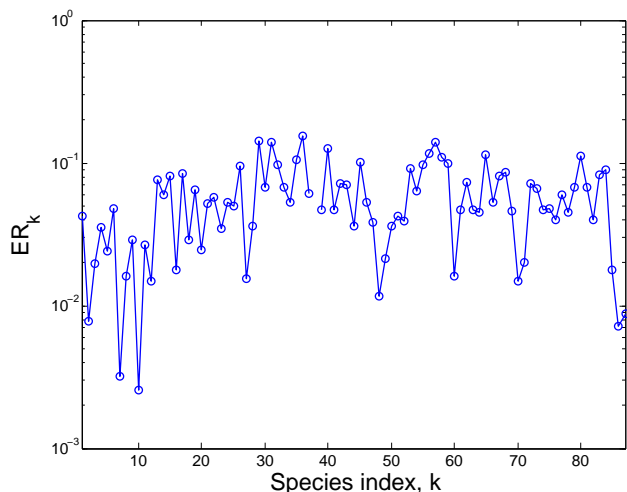


Fig. A2. Difference between the new standard GEOS-Chem simulation using the Rosenbrock solver with respect to the original GEOS-Chem solution using SMVGEARII after a week-long run. The effect of switching solvers is a ~ 5 – 10% change in species concentrations.

Appendix B

Discrete adjoint derivatives with respect to reaction rate constants

We desire to calculate the gradient of the cost function with respect to NO_x emissions. In GEOS-Chem, the emission and dry deposition of many species, such as NO_x , are incorporated into reactions in the tropospheric chemical mechanism as,

$$\frac{dc_{\text{NO}}}{dt} = \text{ENO}_x + \dots \quad (\text{B1})$$

where ENO_x is the NO_x emission rate, emitted as NO . The strong influence of NO_x on the overall chemistry precludes using the continuous adjoint equation of the above equation,

$$\lambda_{\text{ENO}_x} = \int \lambda_{\text{NO}} dt \quad (\text{B2})$$

Hence, we must calculate the sensitivity of the discrete chemical solver itself with respect to the reaction rate coefficients. We present a derivation of these equations here, as they have not yet been presented elsewhere, and they are necessary for accurate calculation of the desired adjoint sensitivities.

For completeness, we first present the equations for the Rosenbrock method, which advances the forward model solution (c^n) from one step to the next using the following formulas,

$$c^{n+1} = c^n + \sum_{i=1}^s m_i k_i, \quad \text{Err}^{n+1} = \sum_{i=1}^s e_i k_i \quad (\text{B3})$$

$$T_i = t^n + \alpha_i h, \quad C_i = c^n + \sum_{j=1}^{i-1} a_{ij} k_j \quad (\text{B4})$$

$$A = \left[\frac{1}{h\gamma} - J^T(t^n, c^n) \right] \quad (\text{B5})$$

$$A \cdot k_i = f(T_i, C_i) + \sum_{j=1}^{i-1} \frac{b_{ij}}{h} k_j + h\gamma_i f_i(t^n, c^n) \quad (\text{B6})$$

where s is the number of stages, $\alpha_i = \sum_j \alpha_{ij}$, $\gamma_j = \sum_j \gamma_{ij}$, m_i , α_{ij} , $a_{i,j}$, b_{ij} , γ_{ij} , and e_i are method coefficients, $f(\cdot, \cdot)$ is the ODE derivative function: $c' = f(t, c)$, $f_t(\cdot, \cdot)$ is the partial time derivative: $f_t(t, c) = \partial f(t, c) / \partial t$, $J(\cdot, \cdot)$ is the Jacobian: $J(t, c) = \partial f(t, c) / \partial c$, $J_t(\cdot, \cdot)$ is the partial time derivative of the Jacobian: $J_t(t, c) = \partial J(t, c) / \partial t$, and $H(\cdot, \cdot)$ is the Hessian: $H(t, c) = \partial^2 f(t, c) / \partial c^2$. A is the system matrix, and C_i , T_i , k_i are internal stage quantities defined by the method. The J and γ used here are not likely to be confused with the use of these notations in the cost function definitions, and allow us to maintain consistent notation with the KPP documentation, which should be consulted for further explanations and values of the method coefficients.

The equation for the adjoint of the concentrations, λ_c , is obtained by differentiating the method with respect to c_n , see Eq. (11).

$$A \cdot u_i = m_i \lambda_c^{n+1} + \sum_{j=i+1}^s \left(a_{ji} v_j + \frac{b_{ji}}{h} u_j \right) \quad (\text{B7})$$

$$v_i = J^T(T_i, Y_i) \cdot u_i, \quad i = s, s-1, \dots, 1 \quad (\text{B8})$$

$$\lambda_c^n = \lambda_c^{n+1} + \sum_{i=1}^s (H(t^n, c^n) \times k_i)^T \cdot u_i \quad (\text{B9})$$

$$+ h J_t^T(t^n, c^n) \cdot \sum_{i=1}^s \gamma_i u_i + \sum_{i=1}^s v_i$$

where v_i and u_i are internal stage vectors defined by the method. For GEOS-Chem, the reaction rates are constant over the internal time steps, hence we use the reduced form of this equation for autonomous systems,

$$\lambda_c^n = \lambda_c^{n+1} + \sum_{i=1}^s (H(t^n, c^n) \times k_i)^T \cdot u_i \quad (\text{B10})$$

$$+ \sum_{i=1}^s J^T(T_i, C_i) \cdot u_i.$$

Taking the derivative of the Rosenbrock method with respect to the reaction rate parameters, and applying Eq. (14), gives the following equation, again for autonomous systems,

$$\lambda_p^n = \lambda_p^{n+1} + \sum_{i=1}^s (J_p(t^n, c^n) \times k_i)^T \cdot u_i \quad (\text{B11})$$

$$+ \sum_{i=1}^s f_p^T(T_i, C_i) \cdot u_i.$$

Though Eq. (B11) is not implemented in the KPP generated adjoint code, KPP does generate the necessary routines for calculation of f_p ($dFun.dRcoeff$) and $(J_p(t^n, c^n) \times k_i)^T$ ($dJac.dRcoeff$). For emissions, the function derivative is simply the identity matrix, and the Jacobian derivative is zero as the emission ODE is independent of any other species concentrations, so the discrete adjoint of the emission rates is

$$\lambda_E^n = \lambda_E^{n+1} + \sum_{i=1}^s I \cdot u_i. \quad (\text{B12})$$

Acknowledgements. This work was supported by U.S. Environmental Protection Agency, grant R832158, and the National Science Foundation, grant NSF ITR AP&IM 0205198, which provided access to the TeraGrid resources at the National Center for Supercomputing Applications. We wish to thank M. Kopacz and A. Sandu for valuable insight and technical support.

Edited by: A. Nenes

References

- Allen, D. J., Kasibhatla, P., Thompson, A. M., Rood, R. B., Doddridge, B. G., Pickering, K. E., Hudson, R. D., and Lin, S. J.: Transport-induced interannual variability of carbon monoxide determined using a chemistry and transport model, *J. Geophys. Res.*, 101, 28 655–28 669, 1996.
- Binkowski, F. S. and Roselle, S. J.: Models-3 community multiscale air quality (CMAQ) model aerosol component – 1. Model description, *J. Geophys. Res.*, 108, 4183, doi:10.1029/2001JD001409, 2003.
- Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C.: A limited memory algorithm for bound constrained optimization, *Scientific Computing*, 16, 1190–1208, 1995.
- Chai, T. F., Carmichael, G. R., Sandu, A., Tang, Y. H., and Daescu, D. N.: Chemical data assimilation of Transport and Chemical Evolution over the Pacific (TRACE-P) aircraft measurements, *J. Geophys. Res.*, 111, D02301, doi:10.1029/2005JD005883, 2006.
- Chin, M., Rood, R. B., Lin, S. J., Muller, J. F., and Thompson, A. M.: Atmospheric sulfur cycle simulated in the global model GOCART: Model description and global properties, *J. Geophys. Res.*, 105, 24 671–24 687, 2000.
- Colella, P. and Woodward, P. R.: The Piecewise Parabolic Method (PPM) for Gas-Dynamical Simulations, *J. Comp. Phys.*, 54, 174–201, 1984.
- Collins, W., Rasch, P., Eaton, B., Khattatov, B., Lamarque, J.-F., and Zender, C.: Simulating aerosols using a chemical transport model with assimilation of satellite aerosol retrievals: Methodology for INDOEX, *J. Geophys. Res.*, 106, 7313–7336, 2001.
- Daescu, D., Sandu, A., and Carmichael, G. R.: Direct and adjoint sensitivity analysis of chemical kinetic systems with KPP: II – Validation and Numerical Experiments, *Atmos. Environ.*, 37, 5097–5114, 2003.
- Damian, V., Sandu, A., Potra, F., and Carmichael, G. R.: The Kinetic PreProcessor KPP – A software environment for solving chemical kinetics, *Comput. Chem. Eng.*, 26, 1567–1579, 2002.
- Dubovik, O., Lapyonok, T., Kaufman, Y. J., Chin, M., Ginoux, P., Remer, L. A., and Holben, B. N.: Inversion of global distribution of aerosol sources using MODIS and AERONET data, *Optica Pura y Aplicada*, 37, 3349–3358, 2004.
- Elbern, H. and Schmidt, H.: A four-dimensional variational chemistry data assimilations scheme for Eulerian chemistry transport modeling, *J. Geophys. Res.*, 104, 18 583–18 598, 1999.
- Elbern, H., Schmidt, H., and Ebel, A.: Variational data assimilation for tropospheric chemistry modeling, *J. Geophys. Res.*, 102, 15 967–15 985, 1997.
- Errera, Q. and Fonteyn, D.: Four-dimensional variational chemical assimilation of stratospheric measurements, *J. Geophys. Res.*, 106, 12 253–12 265, 2001.
- Errico, R. M. and Vukicevic, T.: Sensitivity Analysis Using an Adjoint of the PSU-NCAR Mesoscale Model, *Mon. Wea. Rev.*, 120, 1644–1660, 1992.
- Fisher, M. and Lary, D. J.: Lagrangian four-dimensional variational data assimilation of chemical species, *Quart. J. Roy. Meteorol. Soc.*, 121, 1681–1704, 1995.
- Giering, R. and Kaminski, T.: Recipes for Adjoint code Construction, *ACM Trans. Math. Softw.*, 24, 437–474, 1998.
- Giering, R., Kaminski, T., Todling, R., Errico, R., Gelaro, R., and Winslow, N.: Generating tangent linear and adjoint versions of NASA/GMAO's Fortran-90 global weather forecast model, in: *Automatic Differentiation: Applications, Theory, and Implementations*, edited by: Bücker, H. M., Corliss, G., Hovland, P., Naumann, U., and Norris, B., volume 50 of “Lecture Notes in Computational Science and Engineering”, pages 275–284, Springer, New York, NY, 2005.
- Giles, M. B. and Pierce, N. A.: An Introduction to the Adjoint Approach to Design, *Flow, Turbulence and Combustion*, 65, 393–415, 2000.
- Gilliland, A. and Abbitt, P. J.: A sensitivity study of the discrete Kalman filter (DKF) to initial condition discrepancies, *J. Geophys. Res.*, 106, 17 939–17 952, 2001.
- Gilliland, A. B., Dennis, R. L., Roselle, S. J., and Pierce, T. E.: Seasonal NH₃ emission estimates for the eastern United States based on ammonium wet concentrations and an inverse modeling method, *J. Geophys. Res.*, 108, 4477, doi:10.1029/2002JD003063, 2003.
- Gilliland, A. B., Appel, K. W., Pinder, R. W., and Dennis, R. L.: Seasonal NH₃ emissions for continental United States: Inverse model estimation and evaluation, *Atmos. Environ.*, 40(26), 4986–4998, 2006.
- Griewank, A. and Walther, A.: Algorithm 799: Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation, *ACM Trans. Math. Softw.*, 26, 19–45, 2000.
- Hakami, A., Henze, D. K., Seinfeld, J. H., Chai, T., Tang, Y., Carmichael, G. R., and Sandu, A.: Adjoint inverse modeling of black carbon during the Asian Pacific Regional Aerosol Characterization Experiment, *J. Geophys. Res.*, 110, D14301, doi:10.1029/2004JD005671, 2005.
- Hakami, A., Seinfeld, J. H., Chai, T. F., Tang, Y. H., Carmichael, G. R., and Sandu, A.: Adjoint sensitivity analysis of ozone nonattainment over the continental United States, *Environ. Sci. Technol.*, 40, 3855–3864, 2006.
- Henze, D. K., Seinfeld, J. H., Liao, W., Sandu, A., and Carmichael, G. R.: Inverse modeling of aerosol dynamics: Condensational growth, *J. Geophys. Res.*, 109, D14201, doi:10.1029/2004JD004593, 2004.

- Jacob, D. J., Liu, H., Mari, C., and Yantosca, B. M.: Harvard wet deposition scheme for GMI, http://gmi.gsfc.nasa.gov/models/jacob_wetdep.pdf, 2000.
- Jacobson, M. Z.: Computation of Global Photochemistry with SMVGEAR-II, *Atmos. Environ.*, 29, 2541–2546, 1995.
- Kahn, R. A., Ogren, J. A., Ackerman, T. P., Bosenberg, J., Charlson, R. J., Diner, D. J., Holben, B. N., Menzies, R. T., Miller, M. A., and Seinfeld, J. H.: Aerosol data sources and their roles within PARAGON, *Bull. Amer. Meteorol. Soc.*, 85, 2004.
- Kim, Y. P., Seinfeld, J. H., and Saxena, P.: Atmospheric Gas-Aerosol Equilibrium – I. Thermodynamic Model, *Aerosol Sci. Technol.*, 19, 157–181, 1993.
- Knipping, E. M., Kumar, N., Pun, B. K., Seigneur, C., Wu, S. Y., and Schichtel, B. A.: Modeling regional haze during the BRAVO study using CMAQ-MADRID: 2. Source region attribution of particulate sulfate compounds, *J. Geophys. Res.*, 111, D06303, doi:10.1029/2004JD005609, 2006.
- Lamb, R. G., Chen, W. H., and Seinfeld, J. H.: Numerico-empirical analysis of atmospheric diffusion theories, *J. Atmos. Sci.*, 32, 1794–1807, 1975.
- Liao, H., Yung, Y. L., and Seinfeld, J. H.: Effects of aerosols on tropospheric photolysis rates in clear and cloudy atmospheres, *J. Geophys. Res.*, 104, 23 697–23 707, 1999.
- Lin, S. J. and Rood, R. B.: Multidimensional flux-form semi-Lagrangian transport schemes, *Mon. Wea. Rev.*, 124, 2046–2070, 1996.
- Mallet, V. and Sportisse, B.: 3-D chemistry-transport model Polair: numerical issues, validation and automatic-differentiation strategy, *Atmos. Chem. Phys. Discuss.*, 4, 1371–1392, 2004, <http://www.atmos-chem-phys-discuss.net/4/1371/2004/>.
- Mallet, V. and Sportisse, B.: Uncertainty in a chemistry-transport model due to physical parameterizations and numerical approximations: An ensemble approach applied to ozone modeling, *J. Geophys. Res.*, 111, D01302, doi:10.1029/2005JD006149, 2006.
- Marchuk, G.: Numerical solution of the problems of the dynamics of the atmosphere and the ocean (In Russian), *Gidrometeoizdat*, 1974.
- Martien, P. T. and Harley, R. A.: Adjoint sensitivity analysis for a three-dimensional photochemical model: Application to Southern California, *Environ. Sci. Technol.*, 40, 4200–4210, 2006.
- Martien, P. T., Harley, R. A., and Cacuci, D. G.: Adjoint Sensitivity Analysis for a three-dimensional photochemical model: implementation and method comparison, *Environ. Sci. Technol.*, 40(8), 2663–2670, doi:10.1021/es0510257, 2006.
- Martin, R. V., Jacob, D. J., Yantosca, R. M., Chin, M., and Ginoux, P.: Global and regional decreases in tropospheric oxidants from photochemical effects of aerosols, *J. Geophys. Res.*, 108, 4097, doi:10.1029/2002JD002622, 2003.
- Meirink, J. F., Eskes, H. J., and Goede, A. P. H.: Sensitivity analysis of methane emissions derived from SCIAMACHY observations through inverse modelling, *Atmos. Chem. Phys.*, 6, 1275–1292, 2006, <http://www.atmos-chem-phys.net/6/1275/2006/>.
- Mendoza-Dominguez, A. and Russell, A. G.: Iterative inverse modeling and direct sensitivity analysis of a photochemical air duality model, *Environ. Sci. Technol.*, 34, 4974–4981, 2000.
- Mendoza-Dominguez, A. and Russell, A. G.: Emission strength validation using four-dimensional data assimilation: application to primary aerosol and precursors to ozone and secondary aerosol, *J. Air Waste Manage. Assoc.*, 51, 1538–1550, 2001.
- Menut, L.: Adjoint modeling for atmospheric pollution process sensitivity at regional scale, *J. Geophys. Res.*, 108, 8562, doi:10.1029/2002JD002549, 2003.
- Menut, L., Vautard, R., Beekmann, M., and Honore, C.: Sensitivity of photochemical pollution using the adjoint of a simplified chemistry-transport model, *J. Geophys. Res.*, 105, 15 379–15 402, 2000.
- Muller, J. F. and Stavrou, T.: Inversion of CO and NO_x emissions using the adjoint of the IMAGES model, *Atmos. Chem. Phys.*, 5, 1157–1186, 2005, <http://www.atmos-chem-phys.net/5/1157/2005/>.
- Napelenok, S. L., Cohan, D., Hu, Y., and Russel, A. G.: Decoupled direct 3D sensitivity analysis for particulate matter (DDM-3DPM), *Atmos. Environ.*, 40, 6112–6121, 2006.
- Nenes, A., Pandis, S. N., and Pilinis, C.: ISORROPIA: A new thermodynamic equilibrium model for multiphase multicomponent inorganic aerosols, *Aquat. Geochem.*, 4, 123–152, 1998.
- Nester, K. and Panitz, H. J.: Sensitivity analysis by the adjoint chemistry transport model DRAIS for an episode in the Berlin Ozone (BERLIOZ) experiment, *Atmos. Chem. Phys.*, 6, 2091–2106, 2006, <http://www.atmos-chem-phys.net/6/2091/2006/>.
- Park, R. J., Jacob, D. J., Chin, M., and Martin, R. V.: Sources of carbonaceous aerosols over the United States and implications for natural visibility, *J. Geophys. Res.*, 108, 4355, doi:10.1029/2002JD003190, 2003.
- Park, R. J., Jacob, D., Field, B. D., Yantosca, R., and Chin, M.: natural and transboundary pollution influences on sulfate-nitrate-ammonium aerosols in the United States: implications for policy, *J. Geophys. Res.*, 109, D15204, doi:10.1029/2003JD004473, 2004.
- Pudykiewicz, J. A.: Application of adjoint tracer transport equations for evaluating source parameters, *Atmos. Environ.*, 32, 3039–3050, 1998.
- Reuther, J. J., Jameson, A., Alonso, J. J., Rimlinger, M. J., and Saunders, D.: Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part 1, *J. Aircraft*, 36, 51–60, 1999.
- 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. 2. Rosenbrock solvers, *Atmos. Environ.*, 31, 3459–3472, 1997.
- Sandu, A., Daescu, D. N., and Carmichael, G. R.: Direct and adjoint sensitivity analysis of chemical kinetic systems with KPP: Part I – theory and software tools, *Atmos. Environ.*, 37, 5083–5096, 2003.
- Sandu, A., Daescu, D., Carmichael, G. R., and Chai, T.: Adjoint sensitivity analysis of regional air quality models, *J. Comput. Phys.*, 204, 222–252, 2005a.
- Sandu, A., Liao, W., Carmichael, G. R., Henze, D. K., and Seinfeld, J. H.: Inverse modeling of aerosol dynamics using adjoints: Theoretical and numerical considerations, *Aerosol Sci. Technol.*, 39, 677–694, 2005b.
- Schichtel, B. A., Malm, W. C., Gebhart, K. A., Barna, M. G., and Knipping, E. M.: A hybrid source apportionment model integrating measured data and air quality model results, *J. Geophys. Res.*, 111, D07301, doi:10.1029/2005JD006238, 2006.
- Schmidt, H. and Martin, D.: Adjoint sensitivity of episodic ozone in

- the Paris area to emissions on the continental scale, *J. Geophys. Res.*, 108, 8561–8577, doi:10.1029/2001D001583, 2003.
- Sirkes, Z. and Tziperman, E.: Finite difference of adjoint or adjoint of finite difference?, *Mon. Wea. Rev.*, 125, 3373–3378, 1997.
- Stavrakou, T. and Muller, J. F.: Grid-based versus big region approach for inverting CO emissions using Measurement of Pollution in the Troposphere (MOPITT) data, *J. Geophys. Res.*, 111, D15304, doi:10.1029/2005JD006896, 2006.
- Talagrand, O. and Courtier, P.: Variational assimilation of meteorological observations with the adjoint of the vorticity equations. Part I: Theory, *Quart. J. Roy. Meteorol. Soc.*, 113, 1311–1328, 1987.
- Thuburn, J. and Haine, T. W. N: Adjoints of nonoscillatory advection schemes, *J. Comp. Phys.*, 171, 616–631, 2001.
- Vautard, R., Beekmann, M., and Menut, L.: Applications of adjoint modelling in atmospheric chemistry: sensitivity and inverse modelling, *Environ. Modell. Softw.*, 15, 703–709, 2000.
- Vukićević, T. and Hess, P.: Analysis of tropospheric transport in the Pacific Basin using the adjoint technique, *J. Geophys. Res.*, 105, 7213–7230, 2000.
- Vukićević, T., Steyskal, M., and Hecht, M.: Properties of advection algorithms in the context of variational data assimilation, *Mon. Wea. Rev.*, 129, 1221–1231, 2001.
- Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J.: L-BFGS-B: a limited memory FORTRAN code for solving bound constrained optimization problems, Tech. rep., Northwestern University, 1994.