

DEVELOPMENT OF ADJOINT SENSITIVITY ANALYSIS CAPABILITIES FOR CMAQ

Amir Hakami* and John H. Seinfeld
California Institute of Technology

Adrian Sandu and Kumaresh Singh
Virginia Polytechnic Institute and State University

Daewon Byun, Peter P. Percell, and Violeta Coarfa
University of Houston

Qinbin Li
Jet Propulsion Laboratory, NASA

1. INTRODUCTION

Sensitivity analysis is receiving increasing attention among the air quality modeling community. It has proven to be a powerful tool in various applications ranging from development and evaluation of control strategies to inverse modeling or source apportionment. Traditionally, sensitivity calculations have been performed by one-at-a-time perturbation of input parameters. Brute-force sensitivity analysis is straightforward, but will become infeasible for systems with large numbers of input parameters. A number of formal sensitivity techniques have been developed and implemented in various air quality models including CMAQ [Cohan *et al.*, 2005]. These techniques have, so far, been limited to forward sensitivity calculations. For instance, the most widely used forward sensitivity technique in air quality models, the decoupled direct method (DDM) [Dunker, 1984], integrates a set of sensitivity equations alongside the concentrations for calculation of sensitivity coefficients. DDM propagates a perturbation in an input forward in time, and as a result is efficient in calculating sensitivities of all state variables (concentrations) with respect to few parameters (inputs).

Adjoint sensitivity analysis is a powerful complement to forward methods. While forward techniques are source-based, backward methods provide receptor-based sensitivity information. In Adjoint (backward) sensitivity analysis, a perturbation in an output variable or metric is propagated backward in time. Therefore, adjoint analysis is efficient in calculating sensitivities of

few state variables or metrics with respect to a large number of (input) parameters.

2. FORMULATION

The forward model integrates the atmospheric diffusion equation subject to appropriate initial and boundary conditions:

$$\frac{\partial C_i}{\partial t} = -\nabla \cdot (\mathbf{u} C_i) + \nabla \cdot (\mathbf{K} \nabla C_i) + R_i + E_i \quad (1)$$

where C_i is the concentration of species i , \mathbf{u} is the vector wind field, \mathbf{K} is the diffusivity tensor, the term E_i represents elevated emissions, and R_i is the net rate of production of species i by chemical reactions. A perturbation in a parameter can be propagated in the system of equations by using the following tangent linear model (TLM):

$$\frac{\partial \delta C_i}{\partial t} = -\nabla \cdot (\mathbf{u} \delta C_i) + \nabla \cdot (\mathbf{K} \nabla \delta C_i) + \mathbf{J}_i^T \delta C_i + \delta E_i \quad (2)$$

J_i is the i -th row of the Jacobian of the chemical reaction rates. Note that TLM equations are the same as those solved in DDM calculations, as both can be used to calculate local sensitivity coefficients. The adjoint system is an auxiliary set of equations to the TLM [Sandu *et al.*, 2005]:

$$-\frac{\partial \lambda_i}{\partial t} = \mathbf{u} \cdot \nabla \lambda_i + \nabla \cdot (\mathbf{K} \nabla \lambda_i) + \mathbf{J}_i^T \lambda + \varphi_i \quad (3)$$

*Corresponding author: Amir Hakami, M/C 210-41
California Institute of Technology, Pasadena, CA 91125;
e-mail: amir@cheme.caltech.edu

where λ is the vector of the adjoint variables, and \mathbf{J}_i is the i -th column of the Jacobian for the chemical reaction rates, and φ_i is the forcing term for the adjoint equations. The minus sign in the time derivative indicates that adjoint variables are integrated backward in time. At each time step adjoint variables represent sensitivities of a predefined, receptor-based cost function with respect to the concentrations of individual species. If the cost function is defined as the concentration of a species at a location and time, the four-dimensional adjoint field for each species will simply represent the sensitivities with respect to concentrations of that species at various locations/times. Therefore, integration of adjoints equation back to $t = 0$ results in sensitivities with respect to initial conditions. Sensitivities with respect to parameters other than initial conditions can be integrated during the adjoint calculations using the chain rule. The parametric derivative term is evaluated for the operator in the model where emission processing is carried out (vertical diffusion or chemistry in CMAQ).

3. IMPLEMENTATION

Adjoint calculations can be carried out by numerically integrating equation (3). This method is referred to as the continuous adjoint, and is the focus of this paper. Alternatively, one may develop an adjoint for the discretized form of the forward model, i.e. discrete adjoint. The two methods are not equivalent, as the adjoint and discretization operators are not commutable. Development of the discrete adjoint CMAQ is presented in a companion paper.

Forward and backward sensitivity analysis capabilities for the Rosenbrock family of chemical solvers have been already implemented in the kinetic pre-processor (KPP) model (Sandu et al., 2003). In order to facilitate DDM (TLM) and adjoint implementations in CMAQ, KPP is first integrated into CMAQ. In doing so, CMAQ acquires significant flexibility in its chemistry processes. KPP has proven to be an efficient and easy-to-use chemical mechanism pre-processor, which can now be used to implement new and/or modified mechanisms in CMAQ. KPP offers a choice between 5 different Rosenbrock solvers, while the new version of KPP adds two Runge-Kutta solvers, together with their corresponding TLM and adjoint integrators. KPP solvers show better accuracy (with the same accuracy settings in

terms of the absolute and/or relative tolerances) than EBI and similar to SMVGEAR. KPP solvers are 30-80% slower than EBI, while they are 2-3 times faster than SMVGEAR (in single-processor simulations). The adjoint sensitivities show very good agreement with brute-force for chemistry integrations. An example can be seen in Figure 1 of the sensitivities of ozone concentrations at each grid cell with respect to the initial ozone concentrations, when chemistry is the only process included (no transport), as calculated by the adjoint and brute-force approaches.

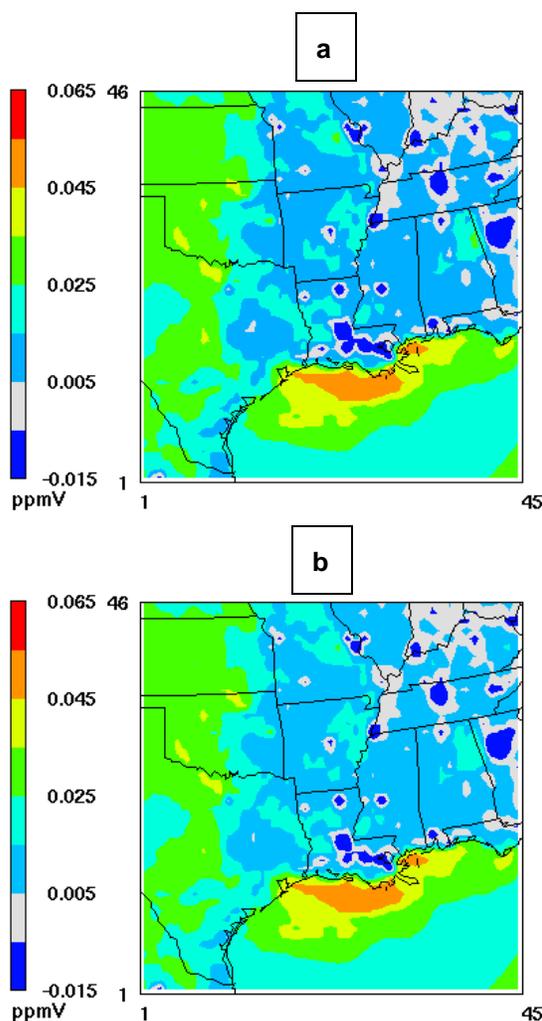


Fig. 1. Sensitivity coefficients calculated by a) brute-force method, and b) adjoint method. The calculations only include chemistry.

In the presence of nonlinearity, integration of the adjoints in equation (3) requires knowledge of concentrations during the backward simulation. Therefore, concentrations of all species are saved during the forward integration to the hard disk. As the internal chemistry time steps can be very small, a two-level checkpointing scheme is employed [Sandu *et al.*, 2005]. In the continuous approach, only the chemistry module requires checkpointing, as it is the sole nonlinear process. In reality, advection schemes such as PPM or Bott, also introduce nonlinearity into the advection process by employing flux-adjusting limiters. As a result, for the discrete adjoint implementation, the (second-level) checkpointing scheme will also need to resolve concentrations for individual chemistry and advection operators. In our implementation of continuous adjoint, we use the synchronization/chemistry time step as the checkpointing interval. For the current implementation a fixed checkpointing frequency is imposed for the duration of the simulation, and for all layers. This approach, however, will be replaced by a variable-time and layer-dependent checkpointing scheme in future.

The horizontal diffusion operator is symmetrical and therefore can be used without change in the adjoint. However, the horizontal diffusion is moved outside the couple/decouple loop, and internal decoupling with density is required. Vertical diffusion in CMAQ is solved using a tridiagonal system, which can also be employed for both continuous and discrete formulation (in transpose form). If sensitivities with respect to emissions are desired, they can be integrated during the adjoint integrations and over the internal vertical diffusion time steps. Integration of sensitivities with respect to the emissions of each species requires one back-substitution in each call to the tridiagonal solver, and therefore, the cost associated with emission sensitivity calculations is negligible compared to the overall cost of adjoint calculations. Vertical diffusion of the adjoints is accurately captured in our simulations, as shown in Figure 2, where sensitivity of the concentrations of ozone in the first-layer cells with respect to emissions of NO in the first layer is calculated using both the adjoint and brute-force methods when chemistry and vertical diffusion are the only included processes.

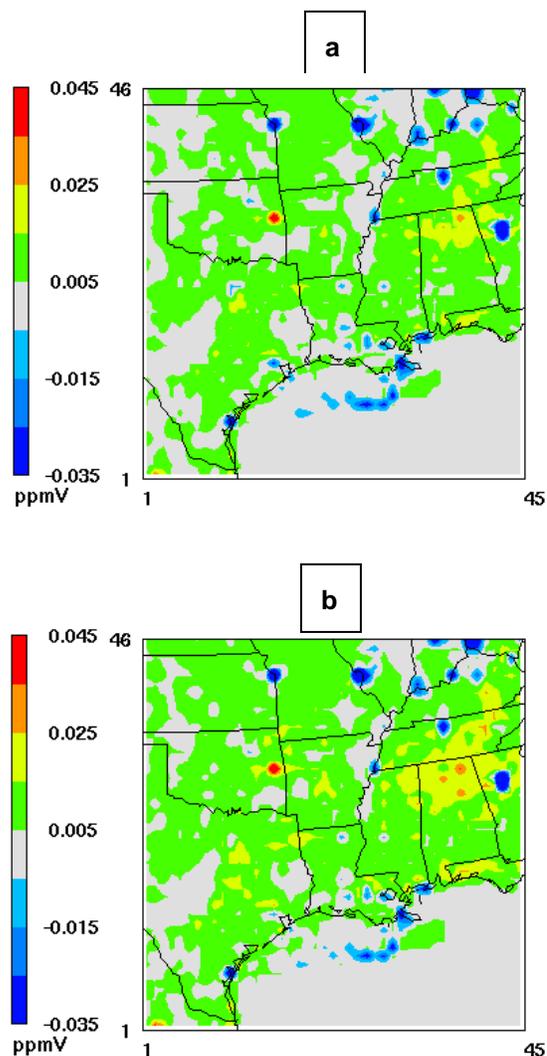


Fig. 2. Sensitivity coefficients calculated by a) brute-force method, and b) adjoint method. The calculations only include chemistry and vertical diffusion.

In the continuous implementation, the adjoint advection process is treated as the reverse-wind advection. CMAQ (in PPM) integrates the flux form of the advection process, which in 1-D is written as:

$$\frac{\partial(\rho C)}{\partial t} = -\frac{\partial(\rho u C)}{\partial t} \quad (4)$$

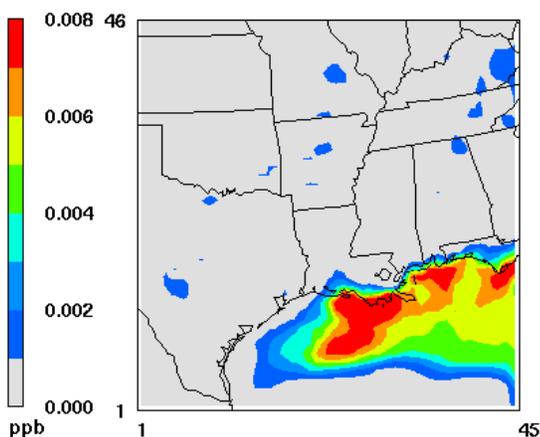


Fig. 3. An adjoint field representing the sensitivity of domain-wide average ozone at the end of the day (June 20th) with respect to initial ozone concentrations at each location in the beginning of the day.

The adjoint of the rearranged equation can be written as:

$$-\frac{\partial \lambda}{\partial t} = \frac{\partial(u\lambda)}{\partial t} \quad (5)$$

Therefore, the same numerical scheme used for concentrations can be used for the adjoints, without the need for “coupling” and “decoupling” of the adjoints. In this implementation we have reverted to ADJCON for mass conservation, but adjustments of vertical wind profiles as implemented in CMAQ 4.5 will be available as an option in future. For use in mass conservation (in either approaches), advected air densities are also checkpointed in the forward simulation. An example of an adjoint field with all processes included is shown in Figure 3, where the adjoints represent sensitivity of the average domainwide ozone at the end of the day with respect to initial ozone concentrations at each point. Note that the lack of horizontal transport processes in Figures 1 and 2 allows for a domainwide comparison of the brute-force (a forward method) and adjoint fields. In general this is not applicable, and only one point of comparison is available from each pair of forward/backward sensitivity calculations. As a result, a full and direct comparison between forward and adjoint sensitivities in presence of all processes is usually infeasible. In our work we will

present a thorough comparison with forward sensitivities, and evaluate the accuracy of the adjoint calculations.

4. FUTURE WORK

This work (together with its companion paper) provides a set of tools for backward (Adjoint) and forward (DDM/TLM) sensitivity calculations in the gas-phase. We are working to expand these capabilities to other modules in CMAQ such as aqueous chemistry, cloud processes, and aerosols. As the final end product, we hope to provide CMAQ users with full continuous/discrete, forward/backward sensitivity capabilities.

5. REFERENCES

Cohan, D. S. et al. 2005. Nonlinear response of ozone to emissions: Source apportionment and sensitivity analysis, *Environ. Sci. Technol.*, 39, 6739-6748.

Dunker, A.M. 1984. The decoupled direct method for calculating sensitivity coefficients in chemical kinetics, *J. Chem. Phys.* 81, 2385-2393.

Sandu, A. et al., 2003. Direct and adjoint sensitivity analysis of chemical kinetic systems with KPP : Part I : theory and software tools, *Atmos Environ.*, 37, 5083-5096.

Sandu, A., 2005. Adjoint sensitivity analysis of regional air quality models, *J. Comput. Physics*, 204, 222-252.