Atmos. Chem. Phys. Discuss., 6, S6637–S6646, 2007
www.atmos-chem-phys-discuss.net/6/S6637/2007/

Atmospheric Chemistry and Physics Discussions

ACPD

6, S6637–S6646, 2007

Interactive Comment

Interactive comment on "Development of the adjoint of GEOS-Chem" *by* D. K. Henze and J. H. Seinfeld

D. K. Henze and J. H. Seinfeld

Received and published: 4 February 2007

Response to reviewer #2

1. *The authors have apparently not found a suitable adjoint for the advection scheme. While they claim this is not the focus of their work, advection is such a fundamental process that it casts doubt on the adjoint model as a whole. The adjoint may give suitable answers when global sensitivities as a whole are evaluated over short time periods (less than a day). However, there appear to be some real problems with it.*

   The revised manuscript further demonstrates that the continuous approach is indeed suitable. That there is a difference between discrete and continuous adjoints of advection is well known. The reviewer's stance favors the former

EGU

approach. Indeed, our original submission was written in a manner that likely reinforced this viewpoint. However, 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 confirms that the continuous approach is adequate, if not preferable. This is not surprising, as it is well documented that discrete adjoints of sign preserving and monotonic (i.e. nonlinear and discontinuous) advection schemes are not well behaved and can contain undesirable numerical artifacts. Thuburn and Haine (2001) demonstrated that discrete adjoint (or finite difference) sensitivities of higher order nonoscillatory or sign-preserving schemes are problematic, while Vukićević et al. (2001) concluded that solving the continuous adjoint equations using accurate methods (i.e. that from the forward model) was preferable to implementing the discrete adjoint of such routines for inverse modeling. This was recommended over using simplified advection schemes with well behaved discrete adjoints. Liu and Sandu (2006[1]) clearly explain the source of such discrepancies for several types of linear and nonlinear advection routines. Recent works by Hakami et al. (2006[2]) and Singh et al. (2007[3]) have compared continuous adjoint, discrete adjoint, and finite difference sensitivities for horizontal advection in the regional model CMAQ (also using a PPM advection scheme), similarly concluding that continuous adjoint gradients of this type of advection scheme are preferable for sensitivity analysis and inverse modeling.

The reviewer is correct in that transport is of fundamental importance to applications of the GEOS-Chem adjoint. Further discussion and demonstration of the

[1] Liu, Z. and Sandu, A.: Analysis of discrete adjoints of numerical methods for the advection equation, Int. J. Numer. Meth. Fl., submitted, 2006.

[2] Hakami, A., Seinfeld, J. H., Sandu, A., Singh, K., Byun, D., Percell, P. P.,Coarfa, V., Li, Q.: Development of adjoint sensitivity capabilities for CMAQ, extended abstract, CMAS Conference, 2006.

[3] Singh, K. Sandu, A., and Hakami, A.: 4D-Var data assimilation for EPA's CMAQ chemistry and transport model, submitted abstract to International Conference on Computational Science, 2007.

ACPD
6, S6637–S6646, 2007

issues mentioned above are clearly required (rather than simply referring to the literature); hence, the revised work more explicitly addresses the significance of the differences between the adjoint and finite difference sensitivities of horizontal advection. 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 comparison 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 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). Adjoint sensitivities of the discrete advection algorithm would contain similar features. Overall, these discrepancies should not be viewed as errors in the continuous adjoint solution. It has been clarified that comparisons of the results from these two methods are only tests of accuracy for components of the adjoint derived from the discrete forward model. That being said, we will try to address the reviewer's specific concerns.

(a) Evaluating the cost function regionally leads to large discrepancies. This seems to severely limit the type of data which can be used: it precludes using regional data and it necessitates long data assimilation windows when using satellite data. As the adjoint solution deteriorates over long integration times it is not clear under what circumstances one can use this adjoint in realistic data assimilation problems. (b) The authors claim that this could be an impediment if only sparse or infrequent measurements are available. There is a large gap between "sparse and infrequent measurements" and global sensitivities. At what spatial and temporal scales will measurements be of value? (c) The solution deteriorates dramatically with long assimilation windows (even 2 days). Yet most satellites do not achieve global coverage in even 2 days. What global measurement system

ACPD

6, S6637–S6646, 2007

do they envisage using with rapid global coverage? (d) The authors state that this deterioration of the adjoint will be more critical for longer lived species whose distributions are chiefly determined by transport. Certainly this applies to aerosols which are transported over thousands of kilometers. What type of species do they expect to be able to use in realistic analysis?

The reviewer's concerns are warranted in that discrepancies between adjoint and finite difference sensitivities are larger when the cost function is evaluated regionally and infrequently, and can accumulate over time. However, the impediment is not that sparse data, or observations of long lived species, are of no use. Rather, checking the validity of adjoint code for practical setups by comparing adjoint to finite difference sensitivities is hindered by the fact that such sparse data exacerbate the discrepancies between these approaches. Explaining why, and by how much, such discrepancies become inflated under practical conditions is the motivation for the additional tests shown in panels (b) – (d) of Fig. 7 (previously Figs. 5 and 6.).

Additionally, we now consider a more realistic example involving sparse, regional data and long range transport. The model is compared to measurements of aerosol nitrate from the IMPROVE network of monitoring stations. 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 Jan 1, 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.

ACPD

6, S6637–S6646, 2007

Interactive Comment

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

EGU

ACPD

6, S6637–S6646, 2007

Interactive Comment

(e) The discrepancy seems to get worse at high sensitivities. Aren't these exactly the points which yield the most information about the solution?

Indeed, as shown in Liu and Sandu (2006[1]), and also Fig. 4, discrepancies are often pronounced near peaks. Since such peaks are of maximal importance for analysis, using the continuous adjoint approach to avoid spurious enhancement or dampening of such peaks seems preferable.

...While I can understand the authors' reluctance to further explore the adjoint of the advection scheme, and while I realize the advection scheme is not the authors' primary interest, transport is intrinsic to the problems they are addressing. Did they authors consider implementing another advective scheme more amendable to adjoint solutions?

Yes, implementing a more approximate transport scheme with a better behaved discrete adjoint was considered, but rejected for two reasons: (1) that transport is crucial for aerosol related applications requires an accurate transport algorithm in the forward model and (2) it is preferable to stick with the method used to generate the GEOS-3 meteorology for the sake of consistency.

...The authors have failed to show that their system is adequate for realistic data assimilation and inverse problems. Slope biases of 0.8 to 1.3 seem very significant for such an idealized test. Furthermore, the points listed above need to be adequately addressed. I think it is incumbent on the authors' to show they have produced an adjoint which is adequate for use under realistic conditions.

We recognize that the justification for our approach given in the original manuscript was not sufficient. In our revised work, we have more explicitly shown that observed differences between adjoint and finite difference sensitivities of advection are consistent with previous works showing differences between continuous and discrete adjoints, and have elaborated on why the continuous approach is deemed preferable. We have also more clearly explained how such discrepancies can hinder adjoint code validation, and yet how simple tests can still be indicative of model performance in realistic scenarios.

2. *This paper seems to be an application paper instead of a numerical methods paper. I believe many of the numerical methods used have been published elsewhere. While I commend the authors for their careful checking of the adjoint solution, and while the tests they conducted are necessary checks on their modeling system, these checks should not be the focal point of the paper. It would be sufficient to state that tests are conducted and the adjoint solution is accurate to x% (with perhaps one figure) and then to go on and apply the system to a real scientific problem. In other words the checks on the adjoint, which are certainly necessary, do not make a paper. (I suspect one reason separate checks were performed on different modeling components is to identify the advective adjoint as a problem. This would not have been necessary if a more accurate adjoint of the advective scheme had been found.).*

   The reviewer is correct in that the hybrid approach adopted here necessitated component-wise inspection of the adjoint model performance. This point, previously only mentioned in the conclusion, has now been stated more clearly in the second paragraph of Sect. 3. However, demonstration of the adjoint model performance on a component-wise basis is warranted for several additional reasons. First, the sensitivity with respect to aerosol thermodynamics deserves attention given the originality and difficulty of such calculations. The performance of the chemistry adjoint must also be demonstrated as the equations used for these sensitivities were derived for the first time in this work, see Appendix B. Furthermore, as GEOS-Chem has many routines common to other models, it behooves us to consider the adjoint of these routines separately. Component-wise characterization of model performance is necessary for future community-based model development.

   ...The “science” in the paper seems rather haphazard – more in the vein of showing the power of the adjoint, instead of investigating a scientific problem in depth.

ACPD

6, S6637–S6646, 2007

Interactive Comment

The paper needs to be science driven.

Publishing a description of the adjoint model's development and validation thus far is important for establishing the capabilities of this tool for scientific interests among a rapidly growing body of users. It was reasoned that the results from further application of this tool are enough to comprise a separate work.

3. *I am puzzled by Equation (14). Why is it necessary to find the adjoint variable with respect to the parameters iteratively (page 10599). Usually this variable can be found through a direct application of the adjoint.*

   The reviewer's comment concerns the following equation,

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

   It is not necessary to solve this equation iteratively (from $n = N, \ldots, 1$), but it is much more efficient. Finding the adjoint parameter variable, λ_p, requires the value of the adjoint concentration variable, λ_c, and F_p at all times and grid cells within the boundary layer of the model run. Evaluation of this equation iteratively allows us to "discard" λ_c^n and F_p^n after updating λ_p^{n-1} from λ_p^n. Calculating λ_p in a non-iterative fashion would require saving the value of λ_c and F_p at each time step in each of these grid cells throughout the course of the adjoint model run. Storing these data alone would require several additional GB, while writing, reading, and reprocessing the data files would take considerable additional CPU time.

4. *The units on page 10604 are puzzling. The units given for emissions are valid for a source (i.e., emissions should molecules/(cm^2 s), not molecules/(cm^3 s)). Then the adjoint sensitivities of concentrations with respect to emissions should then be in units of sec/cm.*

   The reviewer's confusion is understandable. While emissions inventories themselves have units of fluxes, at this point in the model, emissions of NO_x are

ACPD

6, S6637–S6646, 2007

EGU

calculated as part of the chemical mechanism (see Eq. (30) of Appendix B); hence, they actually have units of zeroth order reaction rate constants, i.e. molecules/(cm^3 s). As pointed out by Reviewer #1, the correct units for the adjoint sensitivities should then be s. This table has been replaced by a figure showing the dimensionless ratio λ_{ENO_x} / Λ_{ENO_x}.

5. *The authors' state on page 10611 that the "diffusive nature of [first order, upstream, linear transport schemes] actually increases the bias". This is not at all clear to me. In fact linear (and hence diffusive) advection schemes are known to give accurate adjoints. Can the authors justify this statement?*

   While a discrete adjoint of this linear scheme would be accurate, continuous sensitivities are not necessarily consistent with discrete sensitivities even for a linear scheme. The reason, as shown in Liu and Sandu (2006[1]), is that discontinuities arise when the wind changes sign within a grid cell, leading to spurious source terms in the discrete adjoint equations. Tests similar to those shown in Fig. 4 were performed for the linear advection scheme, and negative values did occur, though they were small. Regardless, as this demonstration seemed somewhat extraneous given that the accuracy of this method is not worth considering, we shall leave this example for further study and have omitted it from present work, as it has yet to be satisfactorily explained.

6. *I am also puzzled by the derivation of the continuous adjoint (pages, 10607, 10608). One should not get equation (26) (an equation in flux form) from the advective form of the tracer continuity equation (equation 25) without additional assumptions.*

   The reviewer's question concerns the nonconservative advection equation,

$$\frac{\partial \mu}{\partial t} = -u\frac{\partial \mu}{\partial x}$$

ACPD
6, S6637–S6646, 2007

EGU

and its adjoint,

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

an advection equation in flux form. The adjoint equation comes from considering the variational equation

$$\int_{t_0}^{T}\int_0^L \frac{\partial\delta\mu}{\partial t}\lambda_\mu dxdt = \int_{t_0}^{T}\int_0^L -u\frac{\partial\delta\mu}{\partial x}\lambda_\mu dxdt$$

Integrating by parts and rearranging gives

$$\int_0^L [\delta\mu\lambda_\mu]_0^T dx + \int_{t_0}^{T} [\delta\mu(u\lambda_\mu)]_0^L dt = \int_{t_0}^{T}\int_0^L (\frac{\partial\lambda_\mu}{\partial t} + \frac{\partial(u\lambda_\mu)}{\partial x})\delta\mu dxdt$$

The right hand side determines the adjoint equation. The first term on the left side would relate the adjoint variable to the cost function, $\lambda_\mu = \frac{\partial J}{\partial\mu}$. The second term would provide boundary conditions for the adjoint and tangent linear equations, though specification as such is not an issue for global (i.e. periodic) systems.

7. *To be of value inverse problems need to involve many degrees of freedom as one attempts to localize the emissions. If I am not mistaken Figure 9 involves solving for a multiplicative factor for global emissions, and thus only involves a few degrees of freedom. While this "toy" problem is certainly important as a first test for checking ones solution, it is not sufficient to provide a rigorous test of the adjoint solution.*

   The scaling factors are independent for each grid cell, hence the sample problems actually have several thousand degrees of freedom. We have made this more transparent by amending the description of the numerical experiments to read:

ACPD
6, S6637–S6646, 2007

EGU

"An active subset of the parameters used to generate these observations is then perturbed using scaling factors, $p=\sigma p_a$, each of which is allowed to vary independently in every grid cell for each emitted species."

Interactive comment on Atmos. Chem. Phys. Discuss., 6, 10591, 2006.