



Development of a variational flux inversion system (INVICAT v1.0) using the TOMCAT chemical transport model

C. Wilson^{1,2}, M. P. Chipperfield², M. Gloor¹, and F. Chevallier³

¹School of Geography, University of Leeds, Leeds, UK

²School of Earth and Environment, University of Leeds, Leeds, UK

³Laboratoire des Sciences du Climat et l'Environnement, CEA-CNRS-UVSQ, UMR8212, IPSL, Gif-sur-Yvette, France

Correspondence to: C. Wilson (c.wilson@leeds.ac.uk) and M. P. Chipperfield (martyn@env.leeds.ac.uk)

Received: 27 September 2013 – Published in Geosci. Model Dev. Discuss.: 23 December 2013

Revised: 8 September 2014 – Accepted: 18 September 2014 – Published: 24 October 2014

Abstract. We present a new variational inverse transport model, named INVICAT (v1.0), which is based on the global chemical transport model TOMCAT, and a new corresponding adjoint transport model, ATOMCAT. The adjoint model is constructed through manually derived discrete adjoint algorithms, and includes subroutines governing advection, convection and boundary layer mixing, all of which are linear in the TOMCAT model. We present extensive testing of the adjoint and inverse models, and also thoroughly assess the accuracy of the TOMCAT forward model's representation of atmospheric transport through comparison with observations of the atmospheric trace gas SF₆. The forward model is shown to perform well in comparison with these observations, capturing the latitudinal gradient and seasonal cycle of SF₆ to within acceptable tolerances. The adjoint model is shown, through numerical identity tests and novel transport reciprocity tests, to be extremely accurate in comparison with the forward model, with no error shown at the level of accuracy possible with our machines. The potential for the variational system as a tool for inverse modelling is investigated through an idealised test using simulated observations, and the system demonstrates an ability to retrieve known fluxes from a perturbed state accurately. Using basic off-line chemistry schemes, the inverse model is ready and available to perform inversions of trace gases with relatively simple chemical interactions, including CH₄, CO₂ and CO.

1 Introduction

Chemical transport models (CTMs) are powerful tools with which we can describe transport and chemical processes in the Earth's atmosphere. CTMs provide global, three-dimensional (3-D) concentration fields of atmospheric trace gases and, through modification of model parameters and boundary conditions, they allow us to investigate the sensitivity of the state of the atmosphere to both anthropogenic and natural variations of these conditions. In order to accurately model the chemical composition of the atmosphere, it is important that a model's input parameters, such as chemical reaction rates, meteorological conditions and the magnitude and location of emissions of atmospheric species, are realistic. However, uncertainties still surround the spatial and temporal variation of the surface flux of some trace gases including, for example, carbon dioxide (CO₂), methane (CH₄) and carbon monoxide (CO) (e.g. Gurney et al., 2002; Frankenberg et al., 2008; Le Quéré et al., 2009; Kopacz et al., 2010). As well as enhancing the performance of atmospheric models, improving our estimates of the surface fluxes of atmospheric species helps to broaden our knowledge of the processes through which trace gases are emitted, and therefore improve our understanding of the interactions between anthropogenic activity, the biosphere, atmospheric composition and climate.

There are a number of methods available for estimating the surface flux of atmospheric species. These can broadly be divided into two method types: bottom-up and top-down. Bottom-up methods are those that attempt to estimate fluxes either through direct measurements or by modelling the

processes which lead to flux of the species into the atmosphere. However, direct measurements of emission rates over relatively small regions are difficult to make and are subject to significant errors when extrapolated to global scales (Jung et al., 2011). Process modelling, meanwhile, requires considerable understanding of the complex procedures which lead to the emissions, and may also be subject to extrapolation errors similar to those which affect the direct flux measurements. Top-down methods, in contrast, attempt to estimate emissions using information about the atmospheric distributions of the species and knowledge of atmospheric transport. This method has the benefit that the assimilation of observational data provides a constraint on the surface flux, assuming that the representation of atmospheric transport and chemistry is accurate. However, limitations of the top-down approach include insufficient global observational coverage and modelling inaccuracies (Dentener et al., 2003; Mikaloff Fletcher et al., 2004; Chen and Prinn, 2005). If measurements of the isotopic composition of a species are included in a top-down emission estimate, it may be possible to partition the distinct emission processes of the species. However, since we currently have relatively poor global coverage of these isotopic observations (Dlugokencky et al., 2011), it is necessary that bottom-up and top-down processes must be used in tandem in order to gain full understanding of trace gas emission budgets. Since an atmospheric model, such as a CTM, is generally used to characterise the atmospheric transport and chemistry in order to relate the concentration fields to the surface flux, top-down techniques are usually referred to as “inverse modelling”. This is in contrast to forward modelling, which relates surface flux estimates to atmospheric concentration fields. The increasing availability of satellite measurements of atmospheric constituents provides a powerful data set for use in data assimilation techniques such as inverse modelling. This, together with ongoing developments of available computational power, means that inverse techniques are increasingly achievable.

There exist a number of inverse modelling techniques available for constraining surface emissions of atmospheric species based on observations of atmospheric concentrations, many of which are detailed in Sandu and Chai (2011). The variational method used in this work is similar to the four-dimensional variational (4D-Var) method which has previously been used in numerical weather prediction (NWP) (e.g. Le Dimet and Talagrand, 1986; Fisher and Courtier, 1995) in order to optimise model variables under the strong constraint that the other sequences of the model state are obtained by prognostic equations. The variational inverse modelling method used in this study is not strictly identical to that of NWP 4D-Var data assimilation, since the inverse method optimises surface fluxes that are not updated as part of the prognostic equations of the forward model. It is true that the inverse method often includes the initial 3-D concentration field of a species within the state vector, but this is usually included as a side product of the inversion that ensures

consistency, and is not used after it has been computed. The term “4D-Var” has been used extensively in previous studies to describe inverse schemes similar to that presented here (e.g. Meirink et al., 2008a, b; Bergamaschi et al., 2010). However, in order to avoid confusion with the NWP community, the term “variational”, rather than “4D-Var”, will be used throughout this work. It should also be noted that we do not currently make use of the incremental variational method described by Courtier et al. (1994).

Variational techniques require the development of the adjoint version of a CTM, which evaluates the sensitivity of a scalar metric of atmospheric concentration fields to input parameters such as surface fluxes. Through data assimilation, the inverse variational technique minimises, in a least-squares sense, a cost function which measures the difference between model predictions and observations, whilst also limiting changes made to existing knowledge of the surface fluxes as much as is reasonable.

This variational technique, and the related adjoint technique, have both previously been applied in a number of studies relating to atmospheric science. Previous applications have included Lagrangian transport models (e.g. Elbern et al., 1997), air quality data assimilation (e.g. Elbern and Schmidt, 2001; Carmichael et al., 2008) and Eulerian CTMs with full atmospheric chemistry schemes (e.g. Henze et al., 2007). Since the turn of the century, the combination of increased computational ability supplemented by large, high-resolution observational data sets provided by remote sensing has allowed atmospheric modellers to significantly expand the potential of the inverse variational method. Previous studies to have used this method in order to quantify surface fluxes of atmospheric species include Chevallier et al. (2005), Pan et al. (2007), Meirink et al. (2008b), Bergamaschi et al. (2010) and Bousquet et al. (2011).

This paper details the development and testing of a new variational inversion model which uses the TOMCAT CTM (Chipperfield, 2006) as its basis. TOMCAT has been extensively used in the past for investigations into chemistry and tracer transport in the troposphere and stratosphere (e.g. Arnold et al., 2005; Monge-Sanz et al., 2007; Breider et al., 2010; Hossaini et al., 2010; Feng et al., 2011; Monks et al., 2012). The variational inverse modelling process initially required the development of adjoint versions of the transport routines in TOMCAT, and the development of this adjoint model is also documented here. We also evaluate TOMCAT’s representation of atmospheric transport, since accuracy in this respect is crucial for the formulation of accurate surface flux estimates. One purpose of this work, therefore, is to quantify how well the new TOMCAT variational system performs as a tool to estimate surface fluxes in future applications. The paper also includes a novel reciprocity test for the adjoint model, which is based upon the work of Hourdin and Talagrand (2006).

In Sect. 2 we describe the TOMCAT model, while in Sect. 3 we describe the variational inversion process. In

Sect. 4, we evaluate the representation of atmospheric transport in the TOMCAT model through comparisons with observations of the trace gas sulfur hexafluoride (SF_6), and we describe the development and testing of a new adjoint model in Sect. 5. Finally, in Sect. 6, we report the construction and testing of the new variational inverse model.

2 The TOMCAT CTM

The TOMCAT model is an Eulerian, grid point, off-line three-dimensional (3-D) CTM, described in Chipperfield et al. (1993), Stockwell and Chipperfield (1999) and Chipperfield (2006). The standard horizontal model grid in the TOMCAT model is made up of regular longitudes and irregular Gaussian latitudes, whilst the vertical grid uses combined σ - p coordinates. Whilst the model typically has a horizontal resolution of $2.8^\circ \times 2.8^\circ$, with 60 vertical levels up to 0.1 hPa, the high computational burden of the variational framework requires that a coarser resolution of $5.6^\circ \times 5.6^\circ$, with 31 vertical levels up to 10 hPa is used whilst testing the inverse model. Previous studies have attempted to alleviate the computational burden of the data assimilation or inversion process using reduced versions of non-linear models (Amsellem et al., 2013; Stefanescu et al., 2014), or with incremental optimisation using low-resolution linearised models (Courtier et al., 1994). Further study is necessary to explore the potential for these methods with the TOMCAT model. The model meteorology, including winds, temperature and pressure data, is read in from ERA-Interim analyses provided by the European Centre for Medium-Range Weather Forecasts (ECMWF, <http://www.ecmwf.int>) (Dee et al., 2011) and transformed onto the TOMCAT model grid. The model uses a process split method, in which separate routines representing the different transport processes are carried out in sequence. In the standard model set-up used in this study, the atmospheric transport consists of routines based upon the Eulerian conservation of second-order moments advection scheme developed by Prather (1986), a convection scheme based on that of Tiedtke (1989), and a boundary layer mixing scheme derived from that of Holtslag and Boville (1993). The advection routine is further split into three subroutines that each carry out tracer transport along one axis only (zonal, meridional or vertical). The model also contains the option of a full tropospheric chemistry scheme as detailed in Arnold et al. (2005), which may be replaced by a simpler “offline” scheme at the user’s discretion. The full chemistry scheme is not included in this work, however. A key characteristic of TOMCAT’s transport scheme is that it is linear in nature. This is beneficial, since non-linearities would complicate the development of the adjoint model.

3 Variational inverse modelling

The theory of the variational inversion technique with regards to estimating surface flux of atmospheric species has been well documented in studies such as Chevallier et al. (2005), Henze et al. (2007) and Meirink et al. (2008a). We summarise it here, as a full understanding of the theory is important to the analysis of the inverse model results. All notation follows that of Ide et al. (1997). The objective of variational inverse modelling is to optimise the value of a state vector, \mathbf{x} , with n elements, in order to improve the prediction of the model in comparison with a set of m observations, \mathbf{y} . The state vector in this case is a vector including the set of surface fluxes of atmospheric species, together with the initial 3-D atmospheric distribution of the species on the model grid resolution. Currently in TOMCAT, the surface flux has a monthly temporal resolution and spatial resolution dependent on the model grid. It is important to include the initial 3-D field within the state vector for long-lived species such as CH_4 and CO_2 . The optimisation is defined via a cost function, $J(\mathbf{x})$, which accounts both for the accuracy of the model prediction compared with the observations and for departures from an a priori estimate of the state vector, \mathbf{x}^b . $J(\mathbf{x})$ is defined as follows:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \frac{1}{2}(\mathbf{y} - \mathbf{H}\mathbf{x})^T \mathbf{R}^{-1}(\mathbf{y} - \mathbf{H}\mathbf{x}), \quad (1)$$

where the superscript T denotes the transpose of a matrix and $^{-1}$ denotes its inverse. The matrix \mathbf{H} represents the forward model simulation that maps from the state vector \mathbf{x} onto the location and time of the observations in \mathbf{y} . The $n \times n$ matrix \mathbf{B} is the error covariance matrix for the a priori state vector, \mathbf{x}^b , while \mathbf{R} is the error covariance matrix for the observations and has size $m \times m$. \mathbf{R} includes instrument error, representation error and model error. For Bayesian theory to hold with Eq. (1), all errors must be assumed to be Gaussian and unbiased. The a priori state vector, \mathbf{x}^b , is in this case our current “best estimate” of the state vector, while the error covariance matrix \mathbf{B} represents the uncertainty of this estimate. The cost function may be described in terms of the “background term” and the “observation term”, which refer to the contribution made by the departures made to the state vector and the model–observation differences, respectively. Finding the minimum value of the cost function is equivalent to finding the optimum value of \mathbf{x} . There are different methods available for solving this problem, and variational schemes use iterative methods in order to find the point at which the gradient of the cost function with respect to the state vector, notated $\nabla_{\mathbf{x}} J(\mathbf{x})$, is zero. This can only occur at a stationary point (minimum or maximum) of J . That is, when

$$\nabla_{\mathbf{x}} J(\mathbf{x}) = \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \mathbf{H}^T [\mathbf{R}^{-1}(\mathbf{y} - \mathbf{H}\mathbf{x})] = 0. \quad (2)$$

Numerically, difficulties arise in the calculation of \mathbf{H}^T , which is the transpose of the matrix \mathbf{H} , since it is impractical to find either of these matrices explicitly on the model grid resolution. Instead we use a forward model, M – in our case TOMCAT – together with a linear interpolation from the model grid onto the observation space in place of the matrix \mathbf{H} in Eq. (1). Through linear interpolation from the model grid onto the latitude, longitude and altitude of each observation in \mathbf{y} at the time that the observation was made, we thus obtain a mapping from the state vector onto the observation space. Since the transport in TOMCAT is linear, the use of this model together with the linear interpolation onto \mathbf{y} is exactly equivalent to the matrix \mathbf{H} .

The forward model is, in reality, a set of discrete mathematical operations and conditional statements, rather than a matrix operator, and we can therefore not explicitly find \mathbf{H}^T using the forward model only. Instead, we use an adjoint model, M^* , as shown by Talagrand and Courtier (1987). Together with an adjoint linear interpolation from the observation space onto the model grid, the adjoint model is used instead of explicitly finding \mathbf{H}^T . The theory of adjoint modelling is expanded upon in the next section. Again, due to the linearity of TOMCAT transport, this adjoint mapping from the observations onto the state vector is equivalent to \mathbf{H}^T . The adjoint version of the TOMCAT model was written for use with the new inverse model. $\nabla_{\mathbf{x}} J(\mathbf{x})$ can therefore be found as in Eq. (2) using the forward and adjoint models in place of \mathbf{H} and \mathbf{H}^T . Once $\nabla_{\mathbf{x}} J(\mathbf{x})$ has been found, it can be used in order to choose an appropriate descent direction along which to minimise $J(\mathbf{x})$. This process of finding the Jacobian of $J(\mathbf{x})$ and using it to minimise the cost function is repeated iteratively until some convergence criterion is met. The iterative process of minimisation will be discussed further in Sect. 6.

3.1 Adjoint modelling

The forward model M can be defined such that, for a concentration field \mathbf{c} at time t_i ,

$$\mathbf{c}(t_{i+1}) = M[\mathbf{c}(t_i)], \quad (3)$$

where time t_{i+1} denotes one model time step after time t_i . In practice, the model consists of parameterisations of various transport and chemical processes, and each of these processes is made up of a finite number of mathematical operations, M_j :

$$M[\mathbf{c}(t_i)] = \prod_j M_j[\mathbf{c}(t_i)]. \quad (4)$$

In our case, each operation M_j is linear and differentiable. However, this is not true for all models, many of which contain non-linear operators. Assuming however that each model operator M is differentiable, a first-order approximation of its first derivative, or Jacobian, can be represented by a tangent linear model (TLM), M' . The TLM simulates the propagation of perturbations forward in time and is dependent upon

the model state at which the linearisation takes place. The TLM is therefore defined such that

$$\delta\mathbf{c}(t_{i+1}) = M'[\mathbf{c}(t_i)]\delta\mathbf{c}(t_i) = \frac{\partial M[\mathbf{c}(t_i)]}{\partial \mathbf{c}}\delta\mathbf{c}(t_i). \quad (5)$$

Note that the model operator is differentiated with respect to the concentration field \mathbf{c} , and not the perturbation $\delta\mathbf{c}$. In practice, this means that elements of the forward model state must be made available in order to run the TLM. This process is known as a ‘‘checkpointing’’. Each mathematical operation M_j must be individually differentiated to create M'_j :

$$\delta\mathbf{c}(t_{i+1}) = \prod_j M'_j[\mathbf{c}(t_i)]\delta\mathbf{c}(t_i) = \prod_j \frac{\partial M_j[\mathbf{c}(t_i)]}{\partial \mathbf{c}}\delta\mathbf{c}(t_i). \quad (6)$$

As mentioned, if the model M is already linear, as with TOMCAT, then the TLM is identical to the forward model. TLM development and checkpointing are then unnecessary. From the TLM, the adjoint model (ADM), M^* , can be developed. The adjoint model propagates variables backwards through time in order to give the sensitivity of a scalar metric of \mathbf{c} (such as J) to the model input parameters. M^* is defined such that, for an inner product $\langle \cdot, \cdot \rangle$ and for vectors \mathbf{u} and \mathbf{v} , it holds that

$$\forall \mathbf{u}, \forall \mathbf{v} \quad \langle M'\mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{u}, M^*\mathbf{v} \rangle. \quad (7)$$

When creating the adjoint model, it is necessary to choose whether to use the adjoint of the transport equations on which the forward model is based, known as a continuous adjoint, or to find the adjoint from the forward model code directly, known as a discrete adjoint. Sirkes and Tziperman (1997) showed that, when using a ‘‘leap-frog’’ advection scheme, their continuous adjoint differed from the actual numerical gradient of $J(\mathbf{x})$, which slowed down the minimisation, but that it did not introduce non-physical behaviour, such as a two-timestep leapfrog computational mode. Gou and Sandu (2011) carried out a comparison between the continuous and discrete adjoints for a piece-wise parabolic advection scheme. They found that the discrete adjoint in this case led to worse performance than the continuous adjoint for variational applications in experimental settings, but that there was little difference in the performance of the two schemes in an inversion using real observational data. Recently, Haines et al. (2014) used an Eulerian backtracking method in order to combine the advantages of both the discrete and continuous adjoint formulations into a single adjoint model. We decided that we would use the discrete adjoint of TOMCAT for this work, in order to make it consistent with the forward model. However, further tests in the future should be carried out in order to examine the performance of the continuous and discrete adjoints of the TOMCAT advection scheme.

In practice, the forward model consists of several thousand lines of computer code, each carrying out a mathematical operation or creating a conditional statement or loop. TLM

and ADM codes can therefore be created from the forward code by hand or through the use of automatic code generators. A variety of these are available, such as TAMC (<http://www.autodiff.com/tamc>) and TAPENADE (www-sop.inria.fr/tropics/tapenade.html), which create a TLM or ADM from a supplied forward model without the need for the time-consuming hand-coding process. Whether coded by hand or automatically, the creation process may produce errors or inconsistencies in the adjoint code, due to human error or through bugs introduced in the automatic coding process (Nehrkorn et al., 2006), and it can often be more efficient to code by hand in the first place. A major problem with automatic procedures is that they may also reduce the potential for optimisation of the adjoint code, such as parallelisation. Since the variational approach described here relies on a number of iterations of forward and adjoint model simulations being carried out, parallelisation of the model code may be necessary in order to carry out inversions with long inversion time-periods. It was therefore decided to develop the adjoint version of the TOMCAT model by hand, rather than through the use of an online tool, as this would allow a greater level of control over the format of the code, and a better understanding of the details of the development process. Other groups have previously detailed the development of adjoint versions of atmospheric models using continuous adjoints and automatic code generators (e.g. Meirink et al., 2006; Hakami et al., 2007; Henze et al., 2007; Huang et al., 2009).

For this work it was decided that only the TOMCAT routines concerning atmospheric transport would be included within the inversion. The model's full chemistry scheme was not included as this would significantly increase the running time of the inverse program, which we initially intend to apply to species such as CO₂, which is treated as a passive tracer, and CH₄, which can be simulated using specified chemical destruction fields, and so do not require the full model chemistry scheme. It is currently intended that the full chemical adjoint scheme will be included at a later date. It was therefore necessary to produce adjoint versions of the TOMCAT model's advection, convection and planetary boundary layer transport schemes, along with other sub-routines concerning the model's calendar and initialisation.

Since each operation within every section of the model's transport scheme was already linear, there was no need to produce a TLM. This meant that we could proceed directly to the production of the model's adjoint. With the adjoint model completed, it was important to thoroughly validate each sub-routine against its equivalent in the forward model in order to ensure its accuracy. The details of this testing will be discussed in Sect. 5.

4 Assessment of tropospheric transport in TOMCAT

In order that the TOMCAT model may be used as an inverse modelling tool, it is important to assess the accuracy of the model's representation of atmospheric transport. The variational inversion process assumes that all model errors are Gaussian and unbiased, and therefore significant model biases would propagate through the model and violate these basic assumptions.

Whilst the observation error covariance matrix **R** should theoretically take account of model transport error, in a variational inverse simulation, in practice this is difficult to quantify and so it is usually assumed that model transport is accurate and unbiased when performing an inversion. Here we use observations of the atmospheric trace gas sulfur hexafluoride (SF₆), together with output from multiple model simulations at different grid resolutions in order to ascertain that the accuracy of the forward model transport is good enough to be used for inverse modelling. SF₆ is useful for investigating aspects of atmospheric transport, and has been used previously in a number of such studies (e.g. Gloor et al., 2007; Patra et al., 2009). It is especially suited to examining interhemispheric transport, since its sources are almost exclusively located in the Northern Hemisphere (NH).

SF₆, which is a potent greenhouse gas, is inert in the troposphere and stratosphere, giving it an extremely long atmospheric lifetime which has been estimated to be between 800 and 3200 yr (Ravishankara et al., 1993; Morris et al., 1995). The only atmospheric sinks of SF₆ are a relatively slow photochemical destruction process and electron capture, both of which only occur in the atmosphere above 60 km, therefore having only a small impact on its atmospheric concentration (Reddmann et al., 2001). Hall and Waugh (1998) found that ignoring the effect of mesospheric destruction when simulating SF₆ may lead to over-estimation of SF₆ concentration in the high-latitude middle stratosphere (above 30 km), but only has a small effect elsewhere.

This property is one of many which make SF₆ a good tracer for testing the simulated long-term atmospheric transport in CTMs. The fact that SF₆ is inert in the troposphere and stratosphere means that there is no need to include chemical processes in the model. Also, the release of SF₆ into the atmosphere is almost entirely anthropogenic in nature. This means both that emissions are fairly constant in time (Levin et al., 2010), with a negligible seasonal cycle (Olivier and Berdowski, 2001), and that we can produce spatially accurate surface emission estimates by distributing sales numbers within each nation according to electrical energy use (Olivier, 2002).

The TOMCAT model previously submitted the results of long-term SF₆ simulations to the TransCom CH₄ intercomparison project (Patra et al., 2011), where it captured the seasonal cycle of SF₆ at three ground-based observation sites to a high level of statistical significance, and reproduced the interhemispheric gradient of SF₆ to within 0.05 parts per

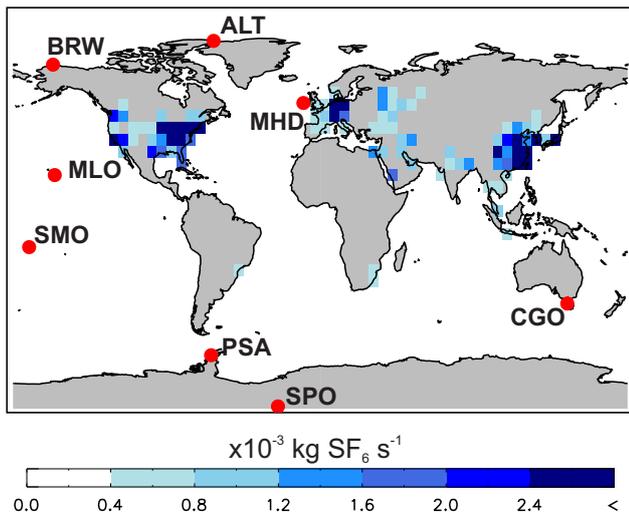


Figure 1. Emissions of SF_6 ($\times 10^{-3} \text{ kg grid cell}^{-1} \text{ s}^{-1}$) on the $5.6^\circ \times 5.6^\circ$ TOMCAT grid for the year 2008. Locations of NOAA surface stations used in this study are also shown.

trillion (ppt). However, those simulations were carried out using the standard TOMCAT model grid resolution ($2.8^\circ \times 2.8^\circ$ with 60 vertical levels up to 0.1 hPa), whilst the variational inverse model will initially be run using a coarser resolution ($5.6^\circ \times 5.6^\circ$, 31 vertical levels up to 10 hPa) in order to reduce simulation times and memory requirements as much as possible. Therefore, new SF_6 simulations were carried out with the TOMCAT model at this coarse resolution in order to assess the effect that reducing the model's resolution has on its performance. For these simulations, the 3-D SF_6 field was initialised on 1 January 1988, with initial values provided by TransCom, and ran up until 31 December 2010, with full 3-D output every 3.75 days. The model timestep was 60 min. We linearly interpolated the TransCom initial concentration field from a $2.8^\circ \times 2.8^\circ$ horizontal grid with 60 vertical levels to the coarser $5.6^\circ \times 5.6^\circ$, 31 level TOMCAT grid. Due to the long spin-up time for this simulation, the initial field should not significantly influence the results in the more recent years. Emissions were also supplied by TransCom, and were originally taken from the Emission Database for Global Atmospheric Research (EDGAR), Version 4.0 (Olivier and Berdowski, 2001), and scaled as in Reddman et al. (2001). Figure 1 shows annual mean SF_6 emissions for the year 2008 on the $5.6^\circ \times 5.6^\circ$ TOMCAT model grid, showing that the majority of SF_6 emissions are from NH industrialised countries. Model output was compared with flask measurements of surface SF_6 from remote station sites in the National Oceanic and Atmospheric Administration (NOAA) network, who have taken weekly measurements of SF_6 at a number of stations since 1995. The locations of the stations used for comparison with the model are also shown in Fig. 1. Measurements have an accuracy of approximately 0.04 ppt.

Table 1. Pearson's correlation value ($r_{5.6}$ and $r_{2.8}$) and root-mean-square error (RMSE_{5.6} and RMSE_{2.8}) for modelled and observed seasonal cycle of SF_6 at eight surface stations, shown in Fig. 2, for two different model grid resolutions.

Station	$r_{5.6}$	$r_{2.8}$	RMSE _{5.6} (ppt)	RMSE _{2.8} (ppt)
ALT	0.29	0.58	0.018	0.015
BRW	0.52	0.64	0.012	0.011
MHD	0.60	0.67	0.020	0.015
MLO	0.36	0.32	0.017	0.018
SMO	0.85	0.89	0.014	0.009
CGO	0.13	0.53	0.011	0.009
PSA	0.22	0.42	0.012	0.011
SPO	0.49	0.75	0.012	0.010

Figure 2 shows the modelled and observed mean detrended seasonal anomalies of SF_6 at each station at both the $2.8^\circ \times 2.8^\circ$ and $5.6^\circ \times 5.6^\circ$ grid resolutions. Table 1 shows the Pearson's correlation value (r) and the root-mean-square error (RMSE) between observed and modelled SF_6 anomalies at each station for both model resolutions. For both modelled and observed SF_6 , in order to display only the seasonal cycle due to transport, the linear trend displayed by SF_6 at the South Pole station (SPO), the site furthest from the source regions, was removed from all data. The modelled and observed SF_6 was averaged over the years 2005 to 2010, and the mean value at each station over this time period was subtracted. This figure shows that switching to the lower resolution does not have a significant impact upon the model's representation of the seasonal cycle, especially in the SH. Prather (1986) showed that the advection scheme used in the model performs well at low resolutions and is relatively non-diffusive. In the NH, however, the proximity of the majority of SF_6 emissions mean that the larger grid boxes produce a diffusive effect as emissions are more rapidly mixed across grid boxes, which slightly alters the model concentrations. Differences between the two resolutions are never greater than 0.03 ppt, however. The greatest difference between the two resolutions is at the MHD station, due to the fact that this particular station, located on the west coast of Ireland, is subject to numerical diffusion of high UK and Irish emissions through the model grid box to different extents depending on the grid box size. Due to the effect of these local emissions, MHD has the largest RMSE of any of the stations, but the correlation is relatively high (≥ 0.60), since the timing of the variations is captured well in the model. At the two Arctic stations, ALT and BRW, there appears to be a systematic underestimation of the negative seasonal anomaly during September and October. This may indicate strong model transport into the Arctic during these months, or weak transport away from the region.

SH sites such as CGO, PSA and SPO have relatively weak seasonal cycles, and the model shows extremely little variation around the mean in the SH. Some SH seasonal

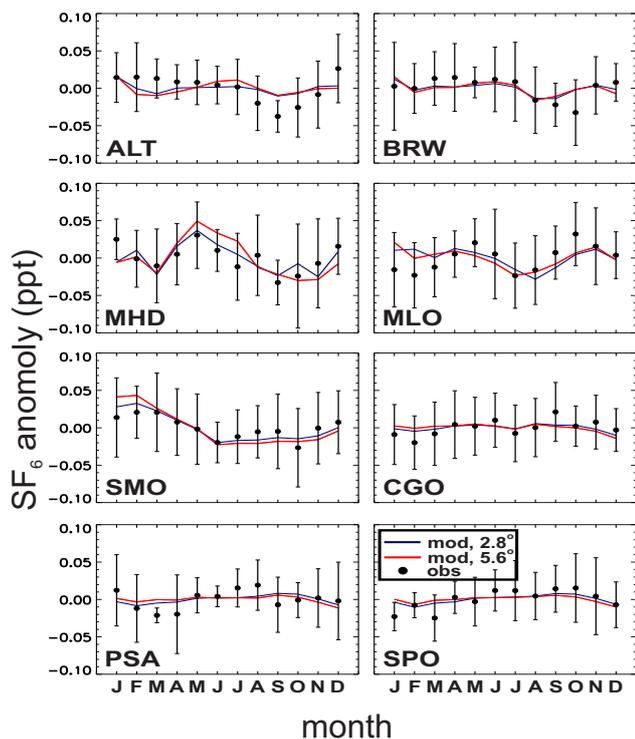


Figure 2. Modelled and observed monthly mean SF₆ anomaly (ppt) averaged over the period 2005–2010. The blue line represents modelled SF₆ using the 2.8° × 2.8° TOMCAT model grid, while the red line shows results for the 5.6° × 5.6° model grid. Black dots represent flask observations from NOAA surface station sites, and error bars show one standard deviation of the monthly mean of the observations over the 6-year period.

variation may be missing in the model, as the observations display larger anomalies than those produced by the simulation. This may indicate that there is too much homogeneity in the modelled SH troposphere. The decrease in resolution produces very little impact on the modelled seasonal cycle at these stations. Due to the weak seasonal cycle, small model–observation correlations are not necessarily representative of poor model performance, and the fact that the RMSE is less than 0.012 ppt at each of these stations indicates that the model is representing transport to the SH well. SMO displays positive anomalies in December through to March and negative anomalies for the rest of the year due to its position relative to the Intertropical Convergence Zone (ITCZ), which is the meteorological (rather than the notional 0°) boundary between the NH and SH. As the position of the ITCZ varies throughout the year due to the changing location of the sun’s zenith point, SMO alternates between sampling NH and SH air. This oscillation is reproduced in both of the model simulations, with correlations greater than 0.8 produced by each. MLO displays a biannual seasonal cycle, due to the increased influence of SH air at MLO during the NH summer and winter (Lintner et al., 2006), and the model

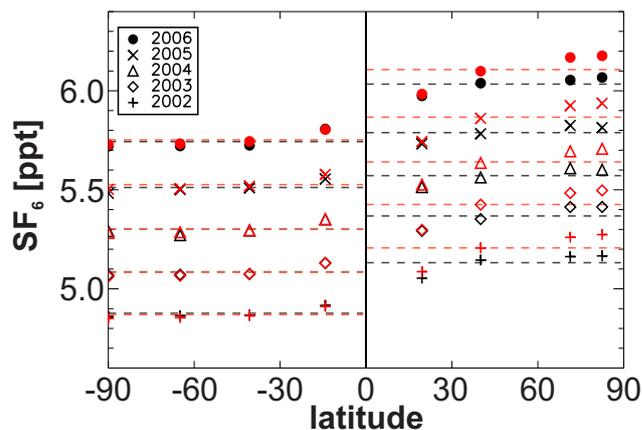


Figure 3. Modelled and observed annual mean SF₆ concentrations (ppt) at different latitudes for the period 2002–2006. The black symbols represent observations at eight surface stations, while the red symbols show the equivalent modelled concentrations using the 5.6° × 5.6° model grid. Different years are represented by different symbols, as shown in the legend. The dotted lines represent the observed (black) and modelled (red) NH and SH mean concentration for each year.

produces the same semi-annual variation, albeit with smaller negative SF₆ anomalies during January and February. The correlations here are relatively low (0.32–0.36) due to this fact.

Figure 3 shows the annual mean modelled and observed SF₆ at eight surface stations located at different latitudes for the years 2002 to 2006. The modelled SF₆ is taken from the simulation which used the 5.6° × 5.6° grid resolution, and in order to remove any bias caused by the model initialisation, the mean model bias at SPO for January 2000 was removed uniformly from all modelled SF₆ concentrations at all times. The model captures the annual increase of SF₆ at the surface well, but the interhemispheric difference (IHD) is too large in the model compared with the observations. The mean observed IHD is approximately 0.28 ppt for this period, while the mean modelled IHD is 0.34 ppt, which is approximately 18 % too high. This shows that interhemispheric transport in TOMCAT is likely too slow.

Overall, the tracer transport in the TOMCAT model performs well in comparisons with observed SF₆. Comparisons with flask samples at station sites provide a validation of the large-scale transport in the model such as interhemispheric and zonal transport and representation of seasonal large-scale atmospheric variations such as the ITCZ. The simulations reproduce the phase and amplitude of the seasonal cycle at most stations. Some SH seasonal transport variations are not reproduced in the model, and model transport in the Arctic may not be strong enough during the NH autumn. However, the timing and magnitude of the effect of the ITCZ is captured well. The observation accuracy of 0.04 ppt at these stations is close to the absolute seasonal variation in many

places, and the model is always within this level of accuracy. Work is always ongoing to improve the TOMCAT model's representation of physical processes, and the adjoint model will be similarly maintained in the future.

5 Construction and validation of the adjoint model

As discussed in Sect. 3, in order to use the TOMCAT model in a variational framework, it is first necessary to produce an adjoint version of the model. In order to reduce the running time of the adjoint model, the forward TOMCAT model was altered so it saved any variables at each model timestep that are also needed by the adjoint model. This meant that these variables did not have to be recalculated for use in the adjoint model. Adjoint versions of the advection, convection and boundary layer transport schemes were produced, and each of the new forward and adjoint routines were coded by hand. Each subroutine was individually and thoroughly tested to confirm its accuracy in relation to the original forward version of the routine. They were then combined to produce the full adjoint version of the TOMCAT model, known as ATOMCAT. The accuracy of the full adjoint model was then also tested.

5.1 Adjoint model tests

This section describes the tests which were used in order to assess the accuracy of ATOMCAT. If the transport in the adjoint model is not exactly accurate, then at best, it will introduce errors into the a posteriori estimate of the state vector. In the worst case, the cost function may not converge at all. Three different tests were carried out, the first of which confirmed that the adjoint identity equation, as defined in Eq. (7), held for ATOMCAT. This first test should in fact be sufficient to be assured of the accuracy of the adjoint model, but a second test was also performed which ensured that tracer transport in ATOMCAT is reciprocal to that of the TOMCAT model, a property that should hold for adjoint models (Hourdin and Talagrand, 2006). This test is an extension of the adjoint identity test, and provides further validation of the adjoint transport over longer time periods. Finally, a test we refer to as the “alpha test”, described in Navon et al. (1992), was carried out. This tested the accuracy of the adjoint model using a Taylor expansion of the cost function. While none of these tests are exhaustive, in the sense that they cannot possibly be carried out using all possible input values, they do provide a strong endorsement of the accuracy of the adjoint model.

For each individual subroutine, it was checked that the identity shown in Eq. (7) held. In practice, it was checked that the following identity held up to the level of accuracy possible due to the rounding error introduced on the machine used to perform the simulation, known as the machine epsilon:

$$\forall \mathbf{u}, \forall \mathbf{v} \quad \frac{\langle M' \mathbf{u}, \mathbf{v} \rangle}{\langle \mathbf{u}, M^* \mathbf{v} \rangle} = 1. \quad (8)$$

Equation (8) was tested for each of the three subroutines representing advection in each dimension, and also for the convection and boundary layer mixing schemes, as well as for the full ATOMCAT model. The input variable for the forward subroutine, \mathbf{u} , was defined as a normally distributed random variable, and in the input variables for the adjoint subroutines were defined to be equal to the output from the corresponding forward subroutine, $M'(\mathbf{u})$. The identity in Eq. (8) is then

$$\frac{\|M'(\mathbf{u})\|^2}{\langle \mathbf{u}, M^*(M'(\mathbf{u})) \rangle} = 1. \quad (9)$$

The ADM was tested based on Eq. (9) using 10 different random initialisations for 480 iterations (equal to 20 model days) of each subroutine. We found that for each subroutine and each initialisation, the identity given in Eq. (9) holds up to machine epsilon, strongly indicating that the adjoint model has been accurately coded from the forward model. The level of accuracy of the results of this test implies that the adjoint model is likely to be correct.

5.2 Reciprocity of atmospheric transport

In order to further test the accuracy of the adjoint transport in ATOMCAT, the property of reciprocity of model transport was investigated. It has been previously discussed that for a linear model, transport in the adjoint model is reciprocal to transport in the forward model (e.g. Hourdin and Talagrand, 2006; Hourdin et al., 2006). This result is equivalent to Eq. (7) and implies that the accuracy of ATOMCAT may be tested by examining the reciprocity of its transport. The tests in this section are therefore extensions of those performed in Sect. 5.1, and examine the accuracy of the adjoint model over multiple model time steps. The reciprocity test posits that if the adjoint model is initialised with a mass, m , of tracer in any given model grid box, D , and integrated backwards through time from time t_n to time t_0 , then the mass in any other specified grid box S at t_0 is equal to that which is found in D if the forward model is integrated from t_0 to t_n after being initialised with mass m of tracer in grid box S . Figure 4 shows a schematic of this theory. Due to the high computational burden of adjoint modelling and the increased simulation time required to carry out both forward and adjoint simulations, the reciprocity of tracer transport in the ATOMCAT model was examined on two different timescales. The adjoint transport over 1 day was examined from every surface grid box, while longer simulations were carried out which investigated adjoint transport from selected grid boxes only.

In order to test the short-term reciprocity of the ADM transport, separate forward simulations were carried out in

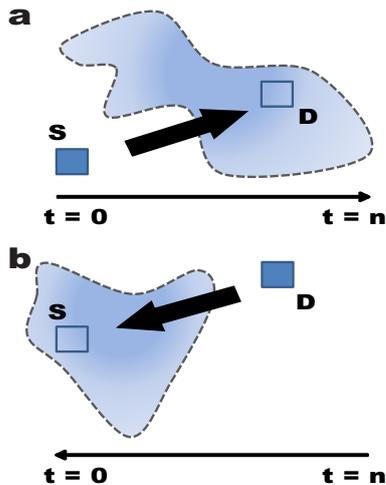


Figure 4. Diagram showing the principle of reciprocity of atmospheric transport using an adjoint model: (a) shows the dispersion of a mass m of tracer in the forward model, released from box “S” at time $t = 0$ after n model time steps; concentration is indicated by colour intensity; (b) shows the equivalent adjoint transport for a mass m of tracer, released from grid box “D” at time n , integrated backwards to time 0. The mass of tracer in box “D” in (a) at time n and in box “S” in (b) at time 0 is identical.

which one surface grid box, S , was initialised with an (arbitrary) concentration of tracer mass of 100 kg, with zero mass elsewhere. One separate simulation was performed for each surface grid box. After the simulation period of 1 day (1 July 2008) was complete, the location D and value m_D of the maximum tracer mass in each simulation was noted. Following this, separate adjoint simulations were carried out in which a pulse of 100 kg was placed into each box D and the ADM was integrated backwards over the same day. In an accurate adjoint model, the total mass, m_S contained in grid box S at the end of the adjoint simulation should be equal to m_D . All forward and adjoint simulations included all of the transport processes available in the model. This was repeated for every surface grid box, using the $5.6^\circ \times 5.6^\circ$ resolution (giving $64 \times 32 = 2048$ simulations). For each simulation, the values of m_D and m_S were exactly equal at machine epsilon, indicating that the adjoint transport in ATOMCAT is correct over short timescales.

In order to test the reciprocity property of adjoint transport in the ATOMCAT model over a longer time period, simulations were carried out in which the reciprocity experiment described above was repeated over a time period of 1 month (July 2008), but for 10 surface grid boxes S_n , $1 \leq n \leq 10$, only. This test was carried out only at certain locations due to the computational burden and time that a more large-scale test would necessitate. Again, at the end of the forward simulation the grid box D_n with the largest mass of tracer m_D^n was found and chosen to be the initial grid box for an adjoint simulation over the same month. Figure 5 shows the results

of this experiment for one of the sites chosen for the emission pulse’s starting point, with the location of the other sites marked in the uppermost panel. Figure 5b shows the tracer mass distribution at the seventh vertical model level from the surface on 31 July, 1 month after release from the grid box marked “S”, located at 84.4° W, 30.5° N. The grid box containing the maximum tracer mass at this time is marked “D”. Figure 5c meanwhile, displays the surface level distribution of tracer mass on 1 July, at the end of an adjoint simulation initialised on 31 July with a tracer mass of 100 kg released from grid box “D”. In this case, and in each of the other cases, the tracer masses m_S^n and m_D^n are exactly equal at machine epsilon. Again, this implies that the tracer transport in ATOMCAT is consistent with that of the forward model, and suggests that over a longer time period the adjoint model is representative of the forward model transport to a high level of accuracy.

Finally, we carried out the accuracy test described by Navon et al. (1992), in which the Taylor series of the cost function is evaluated. Taylor’s theorem states that

$$J(\mathbf{x} + \alpha \delta \mathbf{x}) = J(\mathbf{x}) + \alpha \delta \mathbf{x}^T \nabla_{\mathbf{x}} J(\mathbf{x}) + O(\delta \mathbf{x}^2), \quad (10)$$

where α is a small scalar and $\delta \mathbf{x}$ is a vector of the size of \mathbf{x} . Rewriting this equation, we have

$$\phi(\alpha) = \frac{J(\mathbf{x} + \alpha \delta \mathbf{x}) - J(\mathbf{x})}{\alpha \delta \mathbf{x}^T \nabla_{\mathbf{x}} J(\mathbf{x})} = 1 + O(\delta \mathbf{x}). \quad (11)$$

Therefore, for values of α that are small, but not too close to the order of accuracy of the machine, $\phi(\alpha)$ should be close to 1. We tested this identity by letting the state vector, \mathbf{x} , be the emissions of SF_6 on the TOMCAT model grid, shown in Fig. 1. Note that, for this test, we did not include the initial atmospheric concentrations in the state vector. We also discarded the background term of the cost function, as it would not affect the result, and generated a set of pseudo-observations, \mathbf{y} , by running the forward model for 7 days from 1 July 2008 and multiplying the resultant 3-D field by a factor of 1.5. We defined $\delta \mathbf{x}$ to equal $\mathbf{x}/10$, and let α vary between 1×10^{-14} and 1. Figure 6 shows the results of this test. It is clearly seen that for values of α between 1×10^{-11} and 1×10^{-2} , a unit value of $\phi(\alpha)$ is obtained. As previously explained, whilst these tests do not completely validate the accuracy of the adjoint model, the perfect level of accuracy attained very strongly suggests that the adjoint transport in ATOMCAT is correct.

6 The TOMCAT variational inverse model

Once ATOMCAT had been completed and fully tested, it could be included in the new variational inverse version of the TOMCAT model, named INVICAT. The variational scheme used was based upon the system developed by Chevallier et al. (2005), which makes use (non-exclusively)

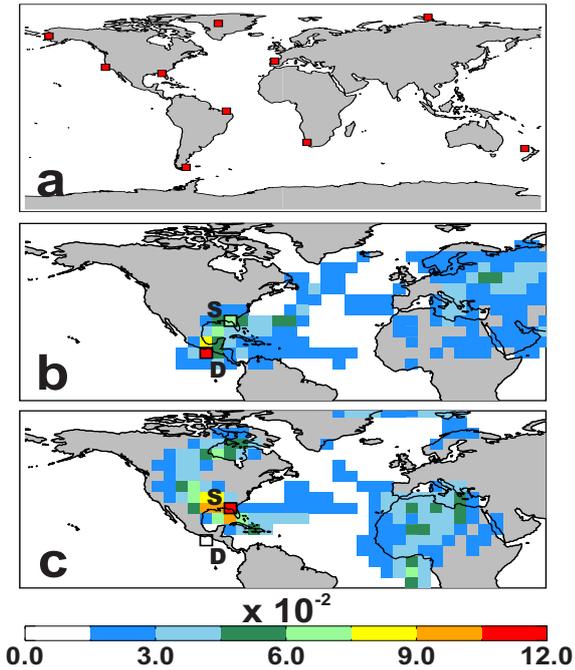


Figure 5. (a) Locations of grid cells from which tracer is released for 10 individual 1-month transport reciprocity tests described in Sect. 5. (b) Tracer mass (kg) at the seventh model level from the surface (approximately 2.6 km) at the end of a 30-day forward simulation initialised with a mass of 100 kg at the surface below the grid cell labelled “S” on 1 July 2008. The mass of tracer in the grid cell labelled “D” is 10.807×10^{-2} kg. (c) Surface tracer mass (kg) at the end of a 30-day adjoint simulation initialised with a mass of 100 kg above grid cell labelled “D” on 31 July 2008. The tracer mass in “S” at the surface is also equal to 10.807×10^{-2} kg.

of the M1QN3 minimisation program, described by Gilbert and Lemarechal (1989), in order to minimise the cost function $J(\mathbf{x})$. This program uses the Limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm which reduces the amount of computer memory necessary to find the optimal solution of a problem (Perry, 1977; Shanno, 1978; Nocedal, 1980).

The minimisation program is first called once the initial value and gradient of the cost function have been found. It is then repeated iteratively until the cost function or its gradient have met some pre-defined convergence criterion, when the program returns the a posteriori state vector. As in Chevallier et al. (2005), a preconditioning transformation is applied to the state vector in order to optimise the speed of the minimisation. Instead of minimising \mathbf{x} directly, the variable \mathbf{z} is defined such that $\mathbf{z} = \mathbf{B}^{-1/2}(\mathbf{x} - \mathbf{x}^b)$, and this variable is minimised instead. This increases the efficiency of the minimisation by reducing the ratio of its largest and smallest eigenvalues (known as the condition number) of the Hessian of the cost function ($\nabla^2 J(\mathbf{x})$) (e.g. Andersson et al., 2000).

At each iteration k , $k \geq 1$, the program determines an appropriate descent direction, \mathbf{d}_k , of $J(\mathbf{x})$ at \mathbf{x}_k , where \mathbf{x}_k is

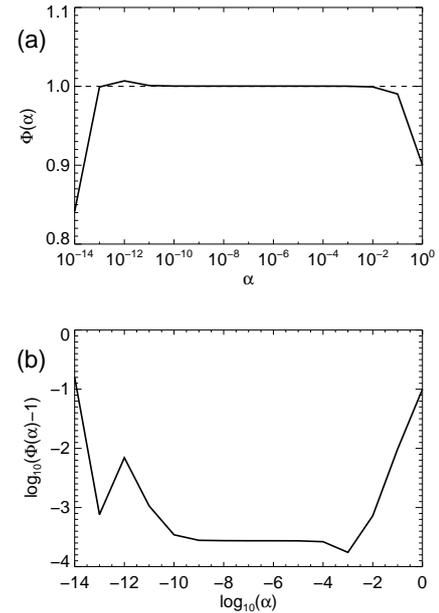


Figure 6. Verification of the alpha test. (a) variation of $\phi(\alpha)$ with respect to α ; (b) variation of $\log_{10}(\phi(\alpha) - 1)$ with respect to $\log_{10}(\alpha)$.

the updated state vector at iteration k . A quasi-Newtonian (QN) method is used in order to choose an appropriate descent direction. This has the advantage over other methods of not needing an exact line search in order to find the minimum along \mathbf{d}_k , although it requires a relatively large amount of computer memory when compared to conjugate gradient methods (Gilbert and Lemarechal, 1989). INVICAT’s minimisation program chooses the descent direction with the value $\mathbf{d}_k = -W_k \mathbf{g}_k$, where W_k approximates the inverse Hessian of $J(\mathbf{x}_k)$, and \mathbf{g}_k is the gradient of the cost function at \mathbf{x}_k . Once this descent direction is chosen, the step size, α_k , to be taken along this direction is determined by the line search procedure. At the next iteration the state vector therefore has the form $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{d}_k$.

Initially, α_k is set to equal $\frac{2J(\mathbf{x}_k)}{\langle \mathbf{g}_k, \mathbf{g}_k \rangle}$, before iteratively being quickly reduced to an appropriate length to find the minimum along \mathbf{d}_k by testing the Wolfe conditions, which are as follows:

$$J(\mathbf{x}_k + \alpha_k \mathbf{d}_k) \leq J(\mathbf{x}_k) + \omega_1 \langle \mathbf{g}_k, \mathbf{d}_k \rangle, \quad (12)$$

$$\langle \mathbf{g}_{k+1}, \mathbf{d}_k \rangle \geq \omega_2 \langle \mathbf{g}_k, \mathbf{d}_k \rangle, \quad (13)$$

where it is necessary to have $0 < \omega_1 < \frac{1}{2}$ and $\omega_1 < \omega_2 < 1$. For this study, values of $\omega_1 = 0.0001$ and $\omega_2 = 0.9$ were chosen. The line-search algorithm iteratively reduces the value of α_k until Eqs. (12) and (13) both hold.

Figure 7 shows a flowchart representing the steps undertaken by INVICAT model in order to find the a posteriori flux estimate. As mentioned, since the adjoint model needs to read some forward model data at every time step, these

data are saved to output files during each forward simulation, and is later read by the adjoint model. It was decided that the data must be written to files rather than held in the machine’s memory due to the large memory requirements necessary, especially for long simulations. For example, a 1 year simulation requires approximately 90 Gb of available storage in order to run. This amount of data storage is currently readily available, and it is not feasible for the machine to hold all of the necessary data in the internal memory during such an inversion.

6.1 Validation of INVICAT

In order to examine the potential of INVICAT to retrieve surface fluxes of an atmospheric species, an inversion was carried out in which pseudo-observations of SF₆ were created using the TOMCAT model. In this experiment, a 2-D field flux array was constructed with the same distribution as the SF₆ emissions for the year 2008, shown in Fig. 1, but multiplied by a factor of 1000 so that the fluxes alter the atmospheric concentration significantly over a short time span. This field was considered for this experiment to be the “true” emissions, x^{tr} , and were used to produce a set of atmospheric SF₆ concentrations y^{tr} . Note that, for this test, the initial 3-D SF₆ distribution was not included in the state vector. A simulation was performed which was initialised at midnight on 1 July 2008 and ran for 7 days. Each grid cell of the 3-D model field was initialised with a mixing ratio of 5 ppt. y^{tr} was defined to be the modelled surface layer SF₆ field at the end of each 24 h period over the course of the simulation, giving y^{tr} a dimension of $64 \times 32 \times 7 = 14336$. These are then used as pseudo-observations in INVICAT in order to attempt to reproduce x^{tr} from a perturbed a priori flux estimate. The a priori was defined to have the same spatial distribution as the “true” fluxes, but with random perturbations which are consistent with the background error covariance matrix \mathbf{B} , i.e.

$$x^b = x^{tr} + qw^{1/2}, \tag{14}$$

where q is an array of random numbers with the same dimension as x and a standard normal distribution, whilst w is an array containing the eigenvalues of \mathbf{B} . \mathbf{B} was defined to be diagonal, with the a priori fluxes being given errors of 20 % if the “true” emissions for that grid cell were non-zero, or 0.001 kg s⁻¹ otherwise (to avoid dividing by zero during the initialisation of the state vector). Meanwhile the observation error covariance matrix \mathbf{R} was also defined to be diagonal, with all observations having a relatively small error of 0.1 ppt. The simulation required approximately 70 min to complete the chosen 10 minimisation iterations (although 11 forward and adjoint simulations were carried out in total, along with two extra forward simulations) on a machine with eight cores running at 2.50 GHz. We also performed the same experiment with 15, 20 and 25 minimisation iterations. However, whilst these extra iterations required longer simulation

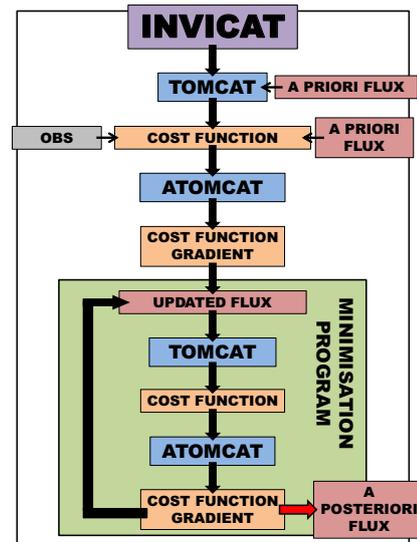


Figure 7. Flowchart depicting the individual steps taken by the INVICAT 4D-Var inverse model. TOMCAT and ATOMCAT are shown in blue, cost function evaluation programs are orange, flux estimates are red and observations are grey. The green section is repeated iteratively until some convergence criterion is met, when the a posteriori flux estimate is evaluated.

times, they brought only slight improvements to the result of the inversion, and therefore these results are not shown here.

Figure 8a shows the reduction of the cost function $J(x)$ as INVICAT performs 10 minimisation iterations relative to its initial value J_1 . Figure 8b displays the development of the normalised cost function gradient norm, $|\nabla J(x)|/|\nabla J_1|$, where $\nabla_x J_1$ is the initial value of the gradient of the cost function. $J(x)$ decreases steadily throughout the run, reaching a value almost 2 orders of magnitude less than its initial value by the end of the minimisation, whilst the cost function gradient norm is 8 orders of magnitude lower than its initial value after 10 iterations. The contribution of the background term to the total cost function is negligible (not shown), and therefore the value of the observational term decreases by around 99 % during the inversion. The fact that the cost function is reduced so quickly indicates that the minimisation program is, at least in this idealised case, working efficiently.

Figure 9a shows the difference between the a priori fluxes, x^b , and the “true” fluxes x^{tr} , while Fig. 9b shows the difference between the updated a posteriori fluxes after 10 minimisation iterations, x^{out} , and x^{tr} . The true fluxes have been almost completely retrieved in all grid cells, with only two grid cells still having errors larger than 0.25 kg SF₆ s⁻¹. The RMSE of a flux vector x , $RMSE_x$, is defined as

$$RMSE_x = \sqrt{(x - x^{tr})^2}. \tag{15}$$

The RMSE of the a priori emissions, $RMSE_{x^b}$, is equal to 0.1 kg s⁻¹, while $RMSE_{x^{out}}$ is equal to 0.02 kg s⁻¹, meaning that approximately 80 % of the total error in the a priori

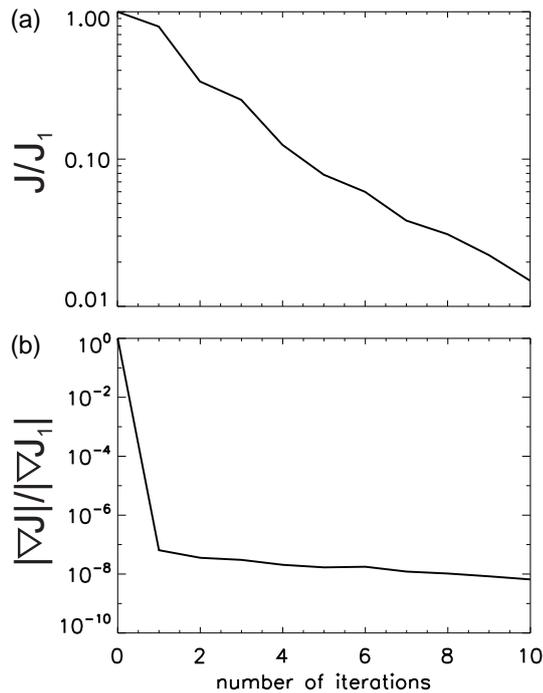


Figure 8. (a) Normalised cost function reduction for INVICAT retrieval using pseudo SF₆ observations. (b) Normalised cost function gradient development for the same experiment.

emissions have been corrected by INVICAT. It should be noted that the test described here does not include any random error term being ascribed to the observations, which would be a significant issue for future inversions using real data. This test should therefore not be seen to be representative of the performance of INVICAT in a real application, but as an assessment of the capability of ATOMCAT and the minimisation program to retrieve a known set of emissions.

An important step in the development of our variational system is finding a way to measure the error reduction achieved by an inversion. The variational inverse method does not allow for explicit output of the a posteriori error covariance matrix, and therefore must be approximated from the variables which can be produced during the inversion. For experiments such as the one described in this section, which use the forward model to produce pseudo-observations that are consistent with the error covariance matrices, an ensemble of observations can be carried out in order to measure the robustness of the result, as in Chevallier et al. (2007). For inversions which assimilate genuine observations of trace gases, meanwhile, it is possible to output the leading eigenvectors and eigenvalues of the Hessian of $J(\mathbf{x})$ as a by-product of the inversion, which can be used to approximate the posterior error covariance matrix as shown in Chevallier et al. (2005) and Meirink et al. (2008b). It is our aim to develop this technique within INVICAT, allowing us to quantify the error reduction of a given inversion.

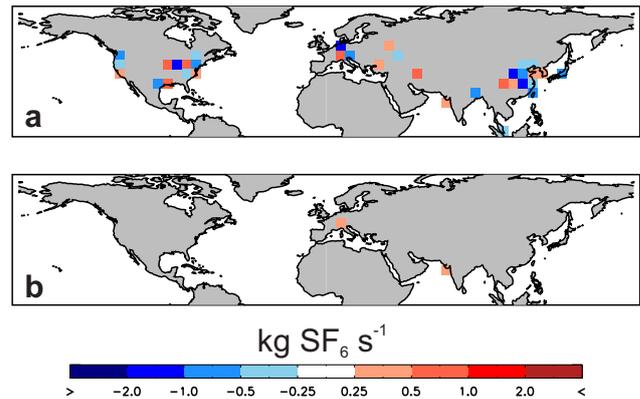


Figure 9. (a) A priori flux error ($\text{kg grid cell}^{-1} \text{s}^{-1}$) for INVICAT experiment using pseudo-observations, defined as $\mathbf{x}^b - \mathbf{x}^{\text{tr}}$, where \mathbf{x}^b is the a priori flux estimate and is created by randomly perturbing the “true” fluxes, \mathbf{x}^{tr} , consistently with the error statistics of the a priori. (b) A posteriori flux error for the same experiment after 10 minimisation iterations, defined as $\mathbf{x}^{\text{out}} - \mathbf{x}^{\text{tr}}$.

The scenario described in this experiment is clearly idealised, as it serves to test the capabilities of the inverse model in the most optimal conditions. The cost function is converging towards a minimum, with a relatively small number of iterations, and is reproducing \mathbf{x}^{tr} with high level of accuracy. This indicates that the performance of INVICAT is currently robust enough to allow us to carry out inversions with genuine observational data.

7 Summary

We have presented thorough details of the development and testing of an adjoint version of the transport section of the TOMCAT CTM and a variational inverse model for the purpose of updating surface fluxes through data assimilation. These models are named ATOMCAT and INVICAT, respectively, and are initially intended to analyse the carbon and methane cycles using observed atmospheric concentration data from in situ measurements and remote sensing. The adjoint model was coded by hand, without the use of automatic differentiation tools, for transport subroutines representing advection in three separate directions, convection and boundary layer mixing. In each case, the discrete adjoint of each subroutine was created directly from the original TOMCAT code, rather than developing a continuous adjoint from the equations governing the transport in the forward model. Meanwhile, since the adjoint model depends upon the state of the forward model at each time step, the forward TOMCAT model was updated in order to save the necessary information so that they may be read by ATOMCAT.

We investigated the accuracy of the transport scheme included in the TOMCAT model using the atmospheric trace gas SF₆. TOMCAT had previously been included in the

TransCom CH₄ model intercomparison, where it had performed accurately in comparison both with observations and with other transport models for SF₆. Further tests showed that TOMCAT captures the seasonality of atmospheric transport well at surface stations, and that reducing the resolution of the model grid by approximately 50 % did not significantly impact the model transport representation. Initially, it is therefore likely that INVICAT will initially be run using the lower model grid resolution in order to maximise running speed and minimise data storage.

Each individual transport routine contained within the adjoint model was tested to ensure that they satisfied the adjoint identity equation, before being combined to form the complete adjoint transport model. This too was tested thoroughly, using the adjoint identity equation, the property of reciprocity that must hold for adjoint transport, and finally the Taylor expansion of the cost function. The reciprocity condition held up to machine epsilon, for tests involving every surface grid cell over a time period of 1 day, and also for selected grid cells over a longer time period of 1 month. This level of accuracy, exact up to the accuracy of the machine used to carry out the tests, strongly implies that the adjoint transport model is indeed analogous to the transport in TOMCAT.

Concurrently with the development of INVICAT and ATOMCAT, a group at the University College London have developed an alternative adjoint version of the TOMCAT model (RETRO-TOM) (Haines et al., 2014). This takes advantage of the property of time symmetry held by the Prather advection scheme by using an Eulerian backtracking method to find adjoint sensitivities. Due to the differing philosophies behind RETRO-TOM and ATOMCAT, each adjoint model has its own distinct advantages. RETRO-TOM does not require output from a forward model simulation, and so is currently likely to be better suited as an alternative to Lagrangian backtracking (e.g. Haines and Esler, 2014). This also means that RETRO-TOM can easily be run using different model grid resolutions. However, ATOMCAT has the advantage that it is not reliant upon its component transport routines being time symmetric, which allows greater freedom of choice than for RETRO-TOM. RETRO-TOM is not currently coupled to a data assimilation of inverse modelling framework, but it would be an interesting avenue of further study to explore the possibility of including RETRO-TOM as an option for adjoint transport within INVICAT. Examining differences in the performances of the two models would allow us to produce the most accurate and consistent inversion results.

Finally, the ability of the variational system INVICAT, which incorporates ATOMCAT, to update surface fluxes through data assimilation was investigated using pseudo “observations” produced using TOMCAT. The model was able to reproduce a set of surface fluxes from a perturbed a priori very closely, reducing the cost function by approximately two orders of magnitude within ten minimisation iterations. The model will now be applied to studies of CH₄ and CO₂,

assimilating real observations from both in situ measurements and remote sensing instruments. However, real inverse modelling studies would require investigation into the optimal values of the error covariance matrices **R** and **B**, which is discussed further in Singh et al. (2011) and Berchet et al. (2013).

Code availability

TOMCAT/SLIMCAT (www.see.leeds.ac.uk/tomcat) is a UK community model. It is available to UK (or NERC-funded) researchers who normally access the model on common facilities or who are helped to install it on their local machines. As it is a complex research tool, new users will need help to use the model optimally. We do not have the resources to release and support the model in an open way. Any potential user interested in the model should contact Martyn Chipperfield.

The model updates described in this paper (INVICAT and ATOMCAT) will be included in the standard model library in the future and therefore will be similarly available only to those who are able to support TOMCAT. The minimisation code M1QN3, meanwhile, is protected by copyright and cannot be distributed except with the permission of its authors. For the review process a limited version of the INVICAT code was made available to the editors, which included sections of the ATOMCAT and TOMCAT code. This code formed the basis of that used to carry out the accuracy experiments described in Sect. 6.1. Inquiries into the availability of the ATOMCAT/INVICAT code can be addressed to the authors.

Acknowledgements. This work was supported by funding from NERC and NCEO. The TOMCAT model is also supported by NCAS. Manuel Gloor was financially supported by the NERC Consortium grant AMAZONICA (NE/F005806/1) and the EU 7th Framework GEOCARBON project (grant no. agreement 283080), Chris Wilson by a NERC DARC studentship and the NERC consortium grant GAUGE. The authors are grateful for data obtained from the World Data Centre of Greenhouse Gases, provided by NOAA/ESRL. The minimisation program M1QN3 is part of the MODULOPT library at INRIA (<https://who.rocq.inria.fr/Jean-Charles.Gilbert/modulopt/>), and the authors thank the company for its permission to use the minimiser. We also thank Wuhu Feng (NCAS).

Edited by: A. Sandu

References

- Amsallem, D., Zahr, M., Choi, Y., and Farhat, C.: Design Optimization Using Hyper-Reduced-Order Models, Technical report, Stanford University, 2013.
- Andersson, E., Fisher, M., Munro, R., and McNally, A.: Diagnosis of background errors for radiances and other observable quantities in a variational data assimilation scheme, and the explana-

- tion of a case of poor convergence, *Q. J. Roy. Meteor. Soc.*, 126, 1455–1472, 2000.
- Arnold, S. R., Chipperfield, M. P., and Blitz, M. A.: A three-dimensional model study of the effect of new temperature-dependent quantum yields for acetone photolysis, *J. Geophys. Res.*, 110, D22305, doi:10.1029/2005JD005998, 2005.
- Berchet, A., Pison, I., Chevallier, F., Bousquet, P., Conil, S., Geever, M., Laurila, T., Lavrič, J., Lopez, M., Moncrieff, J., Necki, J., Ramonet, M., Schmidt, M., Steinbacher, M., and Tarniewicz, J.: Towards better error statistics for atmospheric inversions of methane surface fluxes, *Atmos. Chem. Phys.*, 13, 7115–7132, doi:10.5194/acp-13-7115-2013, 2013.
- Bergamaschi, P., Krol, M., Meirink, J. F., Dentener, F., Segers, A., van Aardenne, J., Monni, S., Vermeulen, A. T., Schmidt, M., Ramonet, M., Yver, C., Meinhardt, F., Nisbet, E. G., Fisher, R. E., O'Doherty, S., and Dlugokencky, E. J.: Inverse modeling of European CH₄ emissions 2001–2006, *J. Geophys. Res.*, 115, D22309, doi:10.1029/2010JD014180, 2010.
- Bousquet, P., Ringeval, B., Pison, I., Dlugokencky, E. J., Brunke, E.-G., Carouge, C., Chevallier, F., Fortems-Cheiney, A., Frankenberg, C., Hauglustaine, D. A., Krummel, P. B., Langenfelds, R. L., Ramonet, M., Schmidt, M., Steele, L. P., Szopa, S., Yver, C., Viovy, N., and Ciais, P.: Source attribution of the changes in atmospheric methane for 2006–2008, *Atmos. Chem. Phys.*, 11, 3689–3700, doi:10.5194/acp-11-3689-2011, 2011.
- Breider, T. J., Chipperfield, M. P., Richards, N. A. D., Carslaw, K. S., Mann, G. W., and Spracklen, D. V.: Impact of BrO on dimethylsulfide in the remote marine boundary layer, *Geophys. Res. Lett.*, 37, L02807, doi:10.1029/2009GL040868, 2010.
- Carmichael, G. R., Sandu, A., Chai, T., Daescu, D. N., Constantinescu, E. M., and Tang, Y.: Predicting air quality: improvements to integrate models and measurements, *J. Comput. Phys.*, 227, 3540–3571, 2008.
- Chen, Y. H. and Prinn, R. G.: Atmospheric modeling of high- and low-frequency methane observations: importance of inter-annually varying transport, *J. Geophys. Res.*, 110, D10303, doi:10.1029/2004JD005542, 2005.
- Chevallier, F., Fisher, M., Peylin, P., Serrar, S., Bousquet, P., Breon, F. M., Chedin, A., and Ciais, P.: Inferring CO₂ sources and sinks from satellite observations: method and application to TOVS data, *J. Geophys. Res.*, 110, D24309, doi:10.1029/2005JD006390, 2005.
- Chevallier, F., Bréon, F.-M., and Rayner, P. J.: Contribution of the Orbiting Carbon Observatory to the estimation of CO₂ sources and sinks: theoretical study in a variational data assimilation framework, *J. Geophys. Res.*, 112, D09307, doi:10.1029/2006JD007375, 2007.
- Chipperfield, M. P.: New version of the TOMCAT/SLIMCAT offline chemical transport model: Intercomparison of stratospheric tracer experiments, *Q. J. Roy. Meteor. Soc.*, 132, 1179–1203, doi:10.1256/qj.05.51, 2006.
- Chipperfield, M., Cariolle, D., Simon, P., Ramarosan, R., and Lary, D.: A 3-dimensional modeling study of trace species in the Arctic lower stratosphere during winter 1989–1990, *J. Geophys. Res.*, 98, 7199–7218, doi:10.1029/92JD02977, 1993.
- Courtier, P., Thepaut, J.-N., and Hollingsworth, A.: A strategy for operational implementation of 4D-Var, using an incremental approach, *Q. J. Roy. Meteor. Soc.*, 120, 1367–1387, 1994.
- Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hólm, E. V., Isaksen, I., Kållberg, P., Köhler, M., Matricardi, M., McNally, A. P., Monge-Sanz, B. M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J.-N., and Vitart, F.: The ERA-Interim reanalysis: configuration and performance of the data assimilation system, *Q. J. Roy. Meteor. Soc.*, 137, 553–597, doi:10.1002/qj.828, 2011.
- Dentener, F., Peters, W., Krol, M., v. Weele, M., Bergamaschi, P., and Lelieveld, J.: Interannual variability and trend of CH₄ lifetime as a measure for OH changes in the 1979–1993 time period, *J. Geophys. Res.*, 108, 4442, doi:10.1029/2002JD002916, 2003.
- Dlugokencky, E. J., Nisbet, E. G., Fisher, R., and Lowry, D.: Global atmospheric methane: budget, changes and dangers, *Philos. T. Roy. Soc.*, 369, 2058–2072, doi:10.1098/rsta.2010.0341, 2011.
- Elbern, H. and Schmidt, H.: Ozone episode analysis by four-dimensional variational data assimilation, *J. Geophys. Res.*, 106, 3569–3590, 2001.
- Elbern, H., Schmidt, H., and Ebe, A.: Variational data assimilation for tropospheric chemistry modeling, *J. Geophys. Res.*, 102, 15967–15985, 1997.
- Feng, W., Chipperfield, M. P., Dhomse, S., Monge-Sanz, B. M., Yang, X., Zhang, K., and Ramonet, M.: Evaluation of cloud convection and tracer transport in a three-dimensional chemical transport model, *Atmos. Chem. Phys.*, 11, 5783–5803, doi:10.5194/acp-11-5783-2011, 2011.
- Fisher, M. and Courtier, P.: Estimating the covariance matrices of analysis and forecast error in variational data assimilation, ECMWF Tech. Memo 220, Eur. Cent. for Medium-Range Weather Forecasts, Reading, UK, 35 pp., available at: http://www.ecmwf.int/publications/library/ecpublications/_pdf/tm/001-300/tm220.pdf (last access: 20 December 2013), 1995.
- Frankenberg, C., Bergamaschi, P., Butz, A., Houweling, S., Meirink, J. F., Notholt, J., Petersen, A. K., Schrijver, H., Warneke, T., and Aben, I.: Tropical methane emissions: a revised view from SCIAMACHY onboard ENVISAT, *Geophys. Res. Lett.*, 35, L15811, doi:10.1029/2008GL034300, 2008.
- Gilbert, J. and Lemarechal, C.: Some numerical experiments with variable-storage quasi-Newton algorithms, *Math. Program.*, 45, 407–435, doi:10.1007/BF01589113, 1989.
- Gloor, M., Dlugokencky, E., Brenninkmeijer, C., Horowitz, L., Hurst, D. F., Dutton, G., Crevoisier, C., Machida, T., and Tans, P.: Three-dimensional SF₆ data and tropospheric transport simulations: signals, modeling accuracy, and implications for inverse modeling, *J. Geophys. Res.*, 112, D15112, doi:10.1029/2006JD007973, 2007.
- Gou, T. and Sandu, A.: Continuous versus discrete advection adjoints in chemical data assimilation with CMAQ, *Atmos. Environ.*, 45, 4868–4881, 2011.
- Gurney, K. R., Law, R. M., Denning, A. S., Rayner, P. J., Baker, D., Bousquet, P., Bruhwiler, L., Chen, Y.-H., Ciais, P., Fan, S., Fung, I. Y., Gloor, M., Heimann, M., Higuchi, K., John, J., Maki, T., Maksyutov, S., Masarie, K., Peylin, P., Prather, M., Pak, B. C., Randerson, J., Sarmiento, J., Taguchi, S., Takahashi, T., and Yuen, C.-W.: Towards robust regional estimates

- of CO₂ sources and sinks using atmospheric transport models, *Nature*, 415, 626–630, 2002.
- Haines, P. E. and Esler, J. G.: Determination of the source regions for surface to stratosphere transport: an Eulerian backtracking approach, *Geophys. Res. Lett.*, 41, 1343–1349, 2014.
- Haines, P. E., Esler, J. G., and Carver, G. D.: Technical Note: Adjoint formulation of the TOMCAT atmospheric scheme in the Eulerian backtracking framework (RETRO-TOM), *Atmos. Chem. Phys.*, 14, 5477–5493, 2014, <http://www.atmos-chem-phys.net/14/5477/2014/>.
- Hakami, A., Henze, D. K., Seinfeld, J. H., Singh, K., Sandu, A., Kim, S., Byun, D., and Li, Q.: The adjoint of CMAQ, *Environ. Sci. Technol.*, 41, 7807–7817, 2007.
- Hall, T. M. and Waugh, D. W.: Influence of nonlocal chemistry on tracer distributions: inferring the mean age of air from SF₆, *J. Geophys. Res.*, 103, 13327–13336, doi:10.1029/98JD00170, 1998.
- Henze, D. K., Hakami, A., and Seinfeld, J. H.: Development of the adjoint of GEOS-Chem, *Atmos. Chem. Phys.*, 7, 2413–2433, doi:10.5194/acp-7-2413-2007, 2007.
- Holtstlag, A. and Boville, B.: Local vs. nonlocal boundary-layer diffusion in a global climate model, *J. Climate*, 6, 1825–1842, 1993.
- Hossaini, R., Chipperfield, M. P., Monge-Sanz, B. M., Richards, N. A. D., Atlas, E., and Blake, D. R.: Bromoform and dibromomethane in the tropics: a 3-D model study of chemistry and transport, *Atmos. Chem. Phys.*, 10, 719–735, doi:10.5194/acp-10-719-2010, 2010.
- Hourdin, F. and Talagrand, O.: Eulerian backtracking of atmospheric tracers, I: Adjoint derivation and parametrization of subgrid-scale transport, *Q. J. Roy. Meteor. Soc.*, 132, 567–583, 2006.
- Hourdin, F., Talagrand, O., and Idelkadi, A.: Eulerian backtracking of atmospheric tracers, II: Numerical aspects, *Q. J. Roy. Meteor. Soc.*, 132, 585–603, 2006.
- Huang, X.-Y., Xiao, Q., Barker, D. M., Zhang, X., Michalakes, J., Huang, W., Henderson, T., Bray, J., Chen, Y., Ma, Z., Dudhia, J., Guo, Y., Zhang, X., Won, D.-J., Lin, H.-C., and Kuo, Y.-H.: Four-dimensional variational data assimilation for WRF: formulation and preliminary results, *Mon. Weather Rev.*, 137, 299–314, 2009.
- Ide, K., Courtier, P., Ghil, M., and Lorenc, A. C.: Unified notation for data assimilation: operational, sequential and variational, *J. Meteorol. Soc. Jpn.*, 75, 181–189, 1997.
- Jung, M., Reichstein, M., Margolis, H. A., Cescatti, A., Richardson, A. D., Arain, M. A., Arneth, A., Bernhofer, C., Bonal, D., Chen, J., Gianelle, D., Gobron, N., Kiely, G., Kutsch, W., Lasslop, G., Law, B. E., Lindroth, A., Merbold, L., Montagnani, L., Moors, E. J., Papale, D., Sottocornola, M., Vaccari, F., and Williams, C.: Global patterns of land-atmosphere fluxes of carbon dioxide, latent heat, and sensible heat derived from eddy covariance, satellite, and meteorological observations, *J. Geophys. Res.*, 116, G00J07, doi:10.1029/2010JG001566, 2011.
- Kopacz, M., Jacob, D. J., Fisher, J. A., Logan, J. A., Zhang, L., Megretskaia, I. A., Yantosca, R. M., Singh, K., Henze, D. K., Burrows, J. P., Buchwitz, M., Khlystova, I., McMillan, W. W., Gille, J. C., Edwards, D. P., Eldering, A., Thouret, V., and Nedelec, P.: Global estimates of CO sources with high resolution by adjoint inversion of multiple satellite datasets (MOPITT, AIRS, SCIAMACHY, TES), *Atmos. Chem. Phys.*, 10, 855–876, doi:10.5194/acp-10-855-2010, 2010.
- Le Dimet, F.-X. and Talagrand, O.: Variational algorithms for analysis and assimilation of meteorological observations, *Tellus A*, 38, 97–100, 1986.
- Le Quéré, C., Raupach, M. R., Canadell, J. G., Marland, G., Bopp, L., Ciais, P., Conway, T. J., Doney, S. C., Feely, R. A., Foster, P., Friedlingstein, P., Gurney, K., Houghton, R. A., House, J. I., Huntingford, C., Levy, P. E., Lomas, M. R., Mankut, J., Metz, N., Ometto, J. P., Peters, G. P., Prentice, I. C., Randerson, J. T., Running, S. W., Sarmiento, J. L., Schuster, U., Sitch, S., Takahashi, T., Viovy, N., van der Werf, G. R., and Woodward, F. I.: Trends in the sources and sinks of carbon dioxide, *Nat. Geosci.*, 2, 831–836, 2009.
- Levin, I., Naegler, T., Heinz, R., Osusko, D., Cuevas, E., Engel, A., Ilmberger, J., Langenfelds, R. L., Neininger, B., Rohden, C. v., Steele, L. P., Weller, R., Worthy, D. E., and Zimov, S. A.: The global SF₆ source inferred from long-term high precision atmospheric measurements and its comparison with emission inventories, *Atmos. Chem. Phys.*, 10, 2655–2662, doi:10.5194/acp-10-2655-2010, 2010.
- Lintner, B. R., Buermann, W., Koven, C. D., and Fung, I. Y.: Seasonal circulation and Mauna Loa CO₂ variability, *J. Geophys. Res.*, 111, GB4004, doi:10.1029/2005JD006535, 2006.
- 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, doi:10.5194/acp-6-1275-2006, 2006.
- Meirink, J. F., Bergamaschi, P., Frankenberg, C., d’Amelio, M. T. S., Dlugokencky, E. J., Gatti, L. V., Houweling, S., Miller, J. B., Roeckmann, T., Villani, M. G., and Krol, M. C.: Four-dimensional variational data assimilation for inverse modeling of atmospheric methane emissions: Analysis of SCIAMACHY observations, *J. Geophys. Res.*, 113, doi:10.1029/2007JD009740, 2008a.
- Meirink, J. F., Bergamaschi, P., and Krol, M. C.: Four-dimensional variational data assimilation for inverse modelling of atmospheric methane emissions: method and comparison with synthesis inversion, *Atmos. Chem. Phys.*, 8, 6341–6353, doi:10.5194/acp-8-6341-2008, 2008b.
- Mikaloff Fletcher, S. E., Tans, P. P., Bruhwiler, L. M., Miller, J. B., and Heimann, M.: CH₄ sources estimated from atmospheric observations of CH₄ and its ¹³C/¹²C isotopic ratios: 2. Inverse modeling of CH₄ fluxes from geographical regions, *Global Biogeochem. Cy.*, 18, GB4004, doi:10.1029/2004GB002224, 2004.
- Monge-Sanz, B. M., Chipperfield, M. P., Simmons, A. J., and Upala, S. M.: Mean age of air and transport in a CTM: comparison of different ECMWF analyses, *Geophys. Res. Lett.*, 34, L04801, doi:10.1029/2006GL028515, 2007.
- Monks, S. A., Arnold, S. R., and Chipperfield, M. P.: Evidence for El Niño-Southern Oscillation (ENSO) influence on Arctic CO interannual variability through biomass burning emissions, *Geophys. Res. Lett.*, 39, L14804, doi:10.1029/2012GL052512, 2012.
- Morris, R. A., Miller, T. M., Viggiano, A. A., Paulson, J. F., Solomon, S., and Reid, G.: Effects of electron and ion reactions on atmospheric lifetimes of fully fluorinated compounds, *J. Geophys. Res.*, 100, 1287–1294, doi:10.1029/94JD02399, 1995.
- Navon, I. M., Zou, X., Derber, J., and Sela, J.: Variational Data Assimilation with an Adiabatic Version of the NMC Spectral Model, *Mon. Weather Rev.*, 120, 1433–1446, 1992.

- Nehrkorn, T., Modica, G. D., Cerniglia, M., Ruggiero, F. H., Michalakes, J. G., and Zou, X.: MM5 adjoint development using TAMC: experiences with an automatic code generator, National Center for Atmospheric Research, Boulder, CA, 2006.
- Nocedal, J.: Updating Quasi-Newton Matrices with Limited Storage, *Mathematics of Computation*, 35, 773–782, 1980.
- Olivier, J.: Part III: greenhouse gas emissions: 1. Shares and trends in greenhouse gas emissions; 2. Sources and methods; greenhouse gas emissions for 1990 and 1995, CO₂ emissions from fuel combustion 1971–2000, International Energy Agency (IEA), Paris, 2002 Edn., III.1–III.31, 2002.
- Olivier, J. and Berdowski, J.: Global emissions sources and sinks, in: *The Climate System*, edited by: Berdowski, J., Guicherit, R., and Heij, B. J., A. A. Balem Publishers/Swets & Zeitlinger Publishers, Lisse, the Netherlands, 33–78, 2001.
- Pan, L., Chai, T., Carmichael, G. R., Tang, Y., Streets, D., Woo, J., Friedli, H. R., and Radke, L. F.: Top-down estimate of mercury emissions in China using four-dimensional variational data assimilation, *Atmos. Environ.*, 41, 2804–2819, 2007.
- Patra, P. K., Takigawa, M., Dutton, G. S., Uhse, K., Ishijima, K., Lintner, B. R., Miyazaki, K., and Elkins, J. W.: Transport mechanisms for synoptic, seasonal and interannual SF₆ variations and “age” of air in troposphere, *Atmos. Chem. Phys.*, 9, 1209–1225, doi:10.5194/acp-9-1209-2009, 2009.
- Patra, P. K., Houweling, S., Krol, M., Bousquet, P., Belikov, D., Bergmann, D., Bian, H., Cameron-Smith, P., Chipperfield, M. P., Corbin, K., Fortems-Cheiney, A., Fraser, A., Gloor, E., Hess, P., Ito, A., Kawa, S. R., Law, R. M., Loh, Z., Maksyutov, S., Meng, L., Palmer, P. I., Prinn, R. G., Rigby, M., Saito, R., and Wilson, C.: TransCom model simulations of CH₄ and related species: linking transport, surface flux and chemical loss with CH₄ variability in the troposphere and lower stratosphere, *Atmos. Chem. Phys.*, 11, 12813–12837, doi:10.5194/acp-11-12813-2011, 2011.
- Perry, A.: A Class of Conjugate Gradient Algorithms with a Two-Step Variable Metric Memory. Discussion Papers 269, Northwestern University, Centre for Mathematical Studies in Economics and Management Science, available at: <http://ideas.repec.org/p/nwu/cmsems/269.html> (last access: 17 June 2014), 1977.
- Prather, M.: Numerical advection by conservation of 2nd-order moments, *J. Geophys. Res.*, 91, 6671–6681, doi:10.1029/JD091iD06p06671, 1986.
- Ravishankara, A. R., Solomon, S., Turnipseed, A. A., and Warren, R. F.: Atmospheric lifetimes of long-lived halogenated species, *Science*, 259, 194–199, doi:10.1126/science.259.5092.194, 1993.
- Reddmann, T., Ruhnke, R., and Kouker, W.: Three-dimensional model simulations of SF₆ with mesospheric chemistry, *J. Geophys. Res.*, 106, 14525–14537, doi:10.1029/2000JD900700, 2001.
- Sandu, A. and Chai, T. F.: Chemical data assimilation – an overview, *Atmosphere*, 2, 426–463, 2011.
- Shanno, D. F.: Conjugate gradient methods with inexact searches, *Math. Oper. Res.*, 3, 244–256, 1978.
- Singh, K., Jardak, M., Sandu, A., Bowman, K., Lee, M., and Jones, D.: Construction of non-diagonal background error covariance matrices for global chemical data assimilation, *Geosci. Model Dev.*, 4, 299–316, doi:10.5194/gmd-4-299-2011, 2011.
- Sirkes, Z. and Tziperman, E.: Finite difference of adjoint or adjoint of finite difference?, *Mon. Weather Rev.*, 125, 3373–3378, 1997.
- Stefanescu, R., Sandu, A., and Navon, I. M.: POD/DIEM Reduced-Order Strategies for Efficient Four Dimensional Variational Data Assimilation, Technical Report TR 3, Virginia Polytechnic Institute and State University, available at: <http://arxiv.org/abs/1402.5992> (last access: 17 July 2014), 2014.
- Stockwell, D. Z. and Chipperfield, M. P.: A tropospheric chemical-transport model: development and validation of the model transport schemes, *Q. J. Roy. Meteor. Soc.*, 125, 1747–1783, doi:10.1256/smsqj.55713, 1999.
- Talagrand, O. and Courtier, P.: Variational assimilation of meteorological observations with the adjoint vorticity equation, I: Theory, *Q. J. Roy. Meteor. Soc.*, 113, 1311–1328, doi:10.1002/qj.49711347812, 1987.
- Tiedtke, M.: A comprehensive mass flux scheme for cumulus parameterization in large-scale models, *Mon. Weather Rev.*, 117, 1779–1800, 1989.