**Atmospheric
Chemistry
and Physics**

An improved Kalman Smoother for atmospheric inversions

L. M. P. Bruhwiler[1], A. M. Michalak[2], W. Peters[3], D. F. Baker[4], and P. Tans[1]

[1]NOAA Climate Monitoring and Diagnostics Laboratory, Boulder, Colorado, USA
[2]Department of Civil and Environmental Engineering, University of Michigan, Ann Arbor, Michigan, USA
[3]Cooperative Institute for Research in Environmental Sciences, University of Colorado, Boulder, Colorado, USA
[4]National Center for Atmospheric Research, Boulder, Colorado, USA

Abstract. We explore the use of a fixed-lag Kalman smoother for sequential estimation of atmospheric carbon dioxide fluxes. This technique takes advantage of the fact that most of the information about the spatial distribution of sources and sinks is observable within a few months to half of a year of emission. After this period, the spatial structure of sources is diluted by transport and cannot significantly constrain flux estimates. We therefore describe an estimation technique that steps through the observations sequentially, using only the subset of observations and modeled transport fields that most strongly constrain the fluxes at a particular time step. Estimates of each set of fluxes are sequentially updated multiple times, using measurements taken at different times, and the estimates and their uncertainties are shown to quickly converge. Final flux estimates are incorporated into the background state of CO_2 and transported forward in time, and the final flux uncertainties and covariances are taken into account when estimating the covariances of the fluxes still being estimated. The computational demands of this technique are greatly reduced in comparison to the standard Bayesian synthesis technique where all observations are used at once with transport fields spanning the entire period of the observations. It therefore becomes possible to solve larger inverse problems with more observations and for fluxes discretized at finer spatial scales. We also discuss the differences between running the inversion simultaneously with the transport model and running it entirely off-line with pre-calculated transport fields. We find that the latter can be done with minimal error if time series of transport fields of adequate length are pre-calculated.

1 Introduction

Understanding the interannual variability of the sources and sinks of atmospheric carbon dioxide is critical to the success of managing carbon reservoirs and emissions. One approach that has been used over recent years to quantify the atmospheric carbon budget involves estimating net fluxes from oceans and terrestrial regions using inverse techniques. These techniques involve comparing predictions from atmospheric transport models and measurements of atmospheric carbon abundances at observation sites distributed over the regions of interest. The spatial pattern of the observed and predicted differences is used to infer the spatial distribution of sources and sinks of carbon dioxide by seeking a distribution of fluxes that in a least squares sense minimizes the difference between the model predictions and observations, as well as any prior information used to constrain the problem.

The technique that has been most commonly employed to estimate carbon fluxes thus far is the Bayesian synthesis inversion (e.g. Enting et al., 1995). For this method, a cost function is formulated that has two terms; one involving the observations and one involving a prior estimate of the fluxes. The resulting flux estimates are therefore constrained both by the observations and a prior guess of the solution. Prior information is needed for the case of atmospheric inversions because the observational network is generally too sparse to permit estimation of fluxes on the scales of interest. More specifically, the problem tends to be underdetermined in regions where observations are sparse, and possibly overdetermined in regions where there are many observations, depending on the spatial scale of the fluxes to be estimated.

Most studies employing the Bayesian synthesis inversion technique have been executed in what has been referred to as "batch" mode (Gelb, 1974) wherein fluxes for all source regions are estimated at all times simultaneously using all of the observations. In the case of an annual mean inversion,

Correspondence to: L. M. P. Bruhwiler
(lori.bruhwiler@noaa.gov)

the vector containing the sources and sinks is the size of the number of fluxes being estimated, and annual averages of observations are used. Examples of annual mean studies of this type are described by Fan et al. (1998), Gurney et al. (2002) and Bousquet et al. (1999). The latter authors solved for annual mean fluxes using monthly observations and thus scaled prior seasonal fluxes by an estimated annual coefficient. Annual mean inversions are comparatively simple and do not require much computational expense, however, they do not result in information about seasonal cycles.

An approach that yields average seasonal cycles of sources and sinks is the cyclo-stationary approach. Multi-year monthly average observations and a state vector that includes monthly values for each source region are used to estimate monthly fluxes (e.g. Kaminski et al., 1999 and Gurney et al., 2004). This approach does not give information about the variability of carbon sources and sinks from year to year. In addition, since monthly average transport information rather than annual average transport information must be used, the size of the problem is significantly larger than for the annual case. The state and observation vectors are a factor of 12 larger, while the matrix of response functions (giving the response at each observation site from pulses emitted and transported from each source region) is a factor of 12 squared larger. Thus more computational resources are required for the cyclo-stationary problem.

The variability of carbon sources and sinks from month to month and year to year is of considerable interest. In particular, the behavior of carbon fluxes over the recent decades at contintental and ocean basin spatial scales may be estimated using existing observational network data and analyzed for changes over time. The first such study was described by Rayner and Law (1999) and a more recent studies are that of Bousquet et al. (2000) and Rodenbeck et al. (2003). The latter authors estimated fluxes at the spatial scale of the transport model grid used.

Estimation of fluxes at very high spatial resolution is a way to hopefully minimize biases and errors arising from using a few sites to constrain large regions with heterogeneous sources and sinks. In addition, when estimating fluxes for continental-scale regions, it is common practice to use prior flux estimates of the oceans and biosphere to obtain spatial distributions of emissions that are used as subregional emission patterns for calculation of the basis functions. If the spatial distribution of the prior fluxes has errors or biases, then the estimated fluxes will be biased also. The strategy of resolving fluxes at grid scales therefore involves estimating large numbers of parameters, and using prior flux and covariance information to constrain the solution where the data have minimal influence (e.g. Kaminski et al., 2001 and Engelen et al., 2002). Estimates for grid-scale fluxes may then be combined in order to obtain estimates for continental and ocean basin-scale fluxes and uncertainties that hopefully do not suffer from aggregation error (although they may indeed be biased by the prior flux estimates themselves). For

these types of problems, the batch Bayesian synthesis grows very computationally demanding since the size of the matrices in the estimation equation gets very large, and since the amount of computation involved in generating basis functions becomes prohibitive, even if an adjoint transport model is used.

In this study we propose a technique wherein the meaurements are used to sequentially estimate fluxes. This technique relies on the observation that, at a particular time, current measurements no longer constrain fluxes from sufficiently far into the past very well due to the tendency of atmospheric mixing to smooth out spatial gradients over time. We find that observations and transport information only from the most recent 6 to 9 months needs to be retained in order to achieve very good agreement with flux estimates from the batch method. Stepping through the observations and keeping only a subset of transport information results in considerable computational savings, and it becomes possible to efficiently estimate fluxes over multiple decades without significant loss of information contained in the observations. This numerical efficiency comes at the cost of a slightly increased estimated flux uncertainty as we will show.

We refer to our technique as a Kalman smoother because it produces estimates of fluxes at a particular time using observations from that time step as well as observations from subsequent times. A filter, on the other hand, would use only past observations to estimate fluxes at a particular time step. Our technique is built upon the optimal filtering technique originally developed by Kalman (1960) for estimating the state of a system given imperfect prior information and observations. Among the first applications for Kalman filtering was to navigation systems.

The application of the Kalman filter to atmospheric problems was pioneered by Hartley (1992) and Hartley and Prinn (1993) for the estimation of regional emissions of the chloroflourocarbon CFC11. The problem of inverting for time-varying fluxes was explored by HaasLaursen et al. (1996) using an idealized transport model. They proposed an adaptive-iterative Kalman filter wherein fluxes at a particular time step are repeatedly estimated until convergence, and the covariance is reinitialized if the posterior model-observation difference grows too large. The technique we describe here is different from this work in several important ways. We estimate fluxes at each time step multiple times using a different set of observations each time. By using observations from multiple time steps into the future to constrain fluxes at a particular time, we are able to allow information from source regions to propagate for as long as necessary. The time period over which transport information is retained is the "fixed-lag".

As has been pointed out by Enting (2002), a shortcoming of a sequential estimation technique such as that proposed here, is that the covariance is not necessarily propagated from states (fluxes) that are no longer being updated by observations to the parts of the state vector that are still being updated. This effect is small if observations and transport fields

from many time steps are used. If shorter periods are used in the sequential inversion, then the estimated flux uncertainty may be significantly underestimated. In this study, we introduce an algorithm that corrects this deficiency by correctly propagating covariance from non-active to active parts of the state vector.

The next section briefly describes the transport model used in this study. This is followed by Sect. 3, which describes the formulation of the fixed-lag Kalman smoother. We then show how the flux estimates produced using our fixed-lag Kalman smoother compare with those calculated using the standard batch Bayesian synthesis inversion technique. We introduce a method for propagating the covariance between time steps no longer being estimated to time steps still in the state. In the last section, we discuss the issue of running the fixed-lag Kalman smoother concurrently with the transport model as opposed to using it with pre-calculated transport fields "off-line" from the transport model.

2 The transport model

The transport model used for this study is the coarse grid Tracer Model version 3 (TM3). The horizontal resolution is roughly $7.5° \times 10°$, with 9 vertical levels spanning the surface to 10 hPa. The TM3 global transport model may be driven by either analyzed meteorological fields or those calculated by a general circulation model. For a detailed description of the model and its physical parameterizations, see Heimann and Koerner (2003). TM3 integrates the tracer continuity equation for an arbitrary number of tracers using the slopes advection scheme of Russell and Lerner (1981). Also included are stability-dependent vertical diffusion using the parameterization of Louis (1979), and a detailed convective mass transport scheme by Tiedke (1989). The integration time step for TM3 is 3 h. As noted by Denning et al. (1999) in a model intercomparison study of an inert atmospheric tracer (SF_6), TM3 lies in the group of models that tends to have weak vertical mixing.

Mass fluxes used by the TM3 model to transport atmospheric trace species must be pre-calculated by post-processing assimilated meteorological data. This is a large computational task, and for this study we used coarse resolution fields because we did not have access to higher resolution fields that covered the entire period spanning 1980 to 2001. Furthermore, in the work presented here, we are interested in estimating monthly average fluxes using monthly average observations and response functions. Use of higher resolution transport fields are therefore not likely to impact the results discussed here significantly, except that some observation sites may be excluded from our inversions because the transport at these sites cannot be adequately simulated with a coarse resolution model. Furthermore, conclusions concerning the performance of the fixed-lag Kalman smoother technique itself are unlikely to be dependent on the resolution

of the transport used. The windfields used to calculate mass fluxes used in this study are European Centre for Medium Range Weather Forecasting (ECMWF) reanalysis products from 1979 through 1992 or the National Center for Environmental Prediction reanalysis products from 1983 through 2001.

3 The Kalman Smoother

In this section we discuss the development of the Fixed-Lag Kalman Smoother starting with basic consideration of Bayesian estimation. We also discuss a covariance propagation scheme and show how results obtained using the new technique compare to those obtained with the more conventional Bayesian Synthesis inversion.

3.1 Background

Neglecting for the moment the use of prior information, the linear estimation problem may be characterized by solution of

$$z = \mathbf{H}s + v \tag{1}$$

where z is a vector of observations, and N is the total number of observations. Alternatively, z may also be a vector of differences between model predictions and observations. In this case, the estimated fluxes are interpreted as adjustments to the fluxes used in the model predictions (the "priors"). The matrix of basis functions, $\mathbf{H}$, is dimensioned N by M, where M is the number of fluxes to be estimated. The elements of $\mathbf{H}$ are calculated using an atmospheric transport model, and give the response at each measurement site for each time due to emissions originating from each source region at each time. s and v are random vectors where s is the vector of source strengths to be estimated and v is the "data uncertainty", which actually represents the inability of the coarse-resolution grid model to simulate near-field influences on the observations as well as long-range transport errors. In this study, we refer to this as the model-data mismatch error.

For the problem under consideration, v arises chiefly from misrepresentation of transport and near-field sources at measurement sites. The uncertainty of the CO_2 measurements is about %0.05 while the transport and misrepresentation errors are typically an order of magnitude larger. In general, model-data mismatch errors occur because the transport model calculates transport for grid boxes several hundred kilometers in extent for comparison to what is essentially a point measurement. Small-scale processes that dominate the observed signal are not likely to be represented well by the transport model. Examples of these are circulations associated with coastlines and mountains, sporadic transport of plumes from urban areas, and small-scale mixing processes associated with the planetary boundary layer. In addition, there may also be errors in the wind fields used.

Typically, Gaussian statistics are assumed and the probability density function (pdf), p, of observation vector, z, given source strength vector, s, is

$$p(z|s) = \frac{1}{\sqrt{2\pi |\mathbf{R}|}} e^{-\frac{1}{2}(z-\mathbf{H}s)^T \mathbf{R}^{-1}(z-\mathbf{H}s)} \tag{2}$$

where $\mathbf{R}$ is the model-data mismatch error covariance matrix obtained by noting that the expectation values $<z-\mathbf{H}s>=<v>=0$ so that

$$\mathbf{R} = <(z - \mathbf{H}s - <z - \mathbf{H}s>)(z - \mathbf{H}s - <z - \mathbf{H}s>)^T> = <vv^T> \tag{3}$$

Likewise, s is also assumed to be normally distributed about prior values $s_\mathbf{p}$ so that its prior probability density function is given by

$$p(s) = \frac{1}{\sqrt{2\pi |\mathbf{Q}|}} e^{-\frac{1}{2}(s-s_\mathbf{p})^T \mathbf{Q}^{-1}(s-s_\mathbf{p})} \tag{4}$$

where $\mathbf{Q}$ is the covariance matrix specified for deviations from the prior flux estimates, $s_\mathbf{p}$. $\mathbf{R}$ and $\mathbf{Q}$ are typically prescribed as diagonal matrices, the values of which specify the relative confidence in observations and information about prior fluxes.

It is important to note that assuming that $<z-\mathbf{H}s>=0$ and $<s-s_\mathbf{p}>=0$ implies no bias in either the transport or prior flux estimates. In either case, this assumption is unlikely to be correct. Biases may, of course, be modeled with arbitrary magnitude. In the case of transport, it is difficult to accurately characterize potential biases. Possible biases in the prior flux estimates are also of concern, especially since the relative sparseness of current observational networks used for atmospheric inversions ensures that the prior flux estimates will dominate the estimated fluxes in some regions.

Invoking Bayes' Theorem,

$$p'(s|z) = \frac{p(z|s)\,p(s)}{\int p(z|s)\,p(s)ds} \tag{5}$$

which states that the probability of a particular source vector given the observational data is equal to the probability of the data given that source vector times the prior probability of that source vector normalized by the total probability of the data for all source vectors. Considering the above pdfs, it follows that to maximize the probability of the sources given the observations, the objective function that must be minimized is:

$$L_S = (z - \mathbf{H}s)^T \mathbf{R}^{-1}(z - \mathbf{H}s) + (s - s_\mathbf{p})^T \mathbf{Q}^{-1}(s - s_\mathbf{p}) \tag{6}$$

At the minimum, the derivative of L_S with respect to s must be zero:

$$\left.\frac{\partial L_S}{\partial s}\right|_{s'} = -(z - \mathbf{H}s')^T \mathbf{R}^{-1}\mathbf{H} + (s' - s_\mathbf{p})^T \mathbf{Q}^{-1} = 0 \tag{7}$$

where the s' indicates the "posteriori" source strength vector.

The posterior covariance estimate is found from the inverse of the Hessian of the objective function:

$$\mathbf{Q}' = \left(\frac{\partial^2 L_s}{\partial s^2}\right)^{-1} = \left(\mathbf{H}^T \mathbf{R}^{-1}\mathbf{H} + \mathbf{Q}^{-1}\right)^{-1}. \tag{8}$$

Equations (7) and (8) may be re-arranged to produce

$$s' = s_\mathbf{p} + \mathbf{Q}\mathbf{H}^T\left(\mathbf{R} + \mathbf{H}\mathbf{Q}\mathbf{H}^T\right)^{-1}\left(z - \mathbf{H}s_\mathbf{p}\right) \tag{9}$$

$$\mathbf{Q}' = \mathbf{Q} - \mathbf{Q}\mathbf{H}^T\left(\mathbf{R} + \mathbf{H}\mathbf{Q}\mathbf{H}^T\right)^{-1}\mathbf{H}\mathbf{Q} \tag{10}$$

Equations (9) and (10) are the discrete Kalman filter update equations (Kalman, 1960; Gelb, 1974) where the Kalman gain matrix, $\mathbf{K}$, may be identified as

$$\mathbf{K} = \mathbf{Q}\mathbf{H}^T\left(\mathbf{R} + \mathbf{H}\mathbf{Q}\mathbf{H}^T\right)^{-1} \tag{11}$$

In the above equations, the Kalman gain matrix functions as a weighting factor between the prior values of the source vector and flux error covariance matrix ($s_\mathbf{p}$ and $\mathbf{Q}$) and new information from additional observations. The prior values may be those specified externally, or a previous estimate. As the misrepresentation error, $\mathbf{R}$, approaches zero, the Kalman gain matrix approaches $\mathbf{H}^\#$ (the pseudo-inverse of $\mathbf{H}$, since $\mathbf{H}$ is a non-square matrix) and it may be shown that $s' \rightarrow (\mathbf{H^T H})^{-1}\mathbf{H^T}z$. In this limiting case, only the observations and available transport information are used to estimate s'. If the model-data mismatch error is much larger than the prior flux error ($\mathbf{R} \gg \mathbf{Q}$; i.e. $\mathbf{Q} \approx 0$), the Kalman gain matrix goes to zero and $s' \approx s_\mathbf{p}$ so that the observations are de-emphasized in the inversion.

Equations (9) and (10) may be manipulated to give an alternate form of the Kalman update equations:

$$s' = s_\mathbf{p} + \left[\mathbf{H}^T \mathbf{R}^{-1}\mathbf{H} + \mathbf{Q}^{-1}\right]^{-1}\mathbf{H}^T \mathbf{R}^{-1}\left(z - \mathbf{H}s_\mathbf{p}\right) \tag{12}$$

$$\mathbf{Q}' = \left[\mathbf{H}^T \mathbf{R}^{-1}\mathbf{H} + \mathbf{Q}^{-1}\right]^{-1}. \tag{13}$$

Note that this form for $\mathbf{Q}$ is the same as Eq. (8).

The dimension of the matrix to be inverted in the above equations is $M \times M$, whereas the corresponding term is of dimension $N \times N$ in Eqs. (9) and (10). If the number of fluxes to be estimated is significantly smaller than the number of observations, it is more computationally efficient to use the latter form even though $\mathbf{R}^{-1}$ and $\mathbf{Q}^{-1}$ must be calculated. $\mathbf{R}$ is generally assumed to be diagonal, and $\mathbf{Q}$ is relatively small in dimension if the number of source regions is not large.

3.2 Time-stepping

Assume that J months of observations are available at n observation sites with monthly observations and that we wish

to estimate fluxes at with monthly resolution. Equation (1) may then be expressed (omitting $\boldsymbol{v}$ for the moment) as

$$
\begin{bmatrix} z_J \\ z_{J-1} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ z_1 \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{J,J} & \mathbf{H}_{J,J-1} & \cdot & \cdot & \cdot & \mathbf{H}_{J,1} \\ 0 & \mathbf{H}_{J-1,J-1} & \mathbf{H}_{J-1,J-2} & \cdot & \cdot \mathbf{H}_{J-1,1} \\ 0 & 0 & \mathbf{H}_{J-2,J-2} & \mathbf{H}_{J-2,J-3} & \cdot \mathbf{H}_{J-2,1} \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & 0 & \cdot & \cdot \mathbf{H}_{1,1} \end{bmatrix} \begin{bmatrix} s_J \\ s_{J-1} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ s_1 \end{bmatrix}
$$

(14)

where each element of the vectors and matrix is itself a vector or matrix. Each vector z_J is a vector of length equal to the number of observation sites, N, each vector s_J is of length equal to the number of source regions, M, and each matrix $\mathbf{H}_{J,J}$ in Eq. (14) is a matrix dimensioned N, by M. Specifically,

$$
\begin{bmatrix} z_1 \\ z_2 \\ \cdot \\ \cdot \\ \cdot \\ z_n \end{bmatrix}, \begin{bmatrix} h_{1,1} & h_{2,1} & \cdots & h_{1,M} \\ h_{2,1} & h_{2,2} & \cdots & h_{2,M} \\ \cdot & \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdot & \cdots & \cdot \\ h_{n,1} & \cdot & \cdots & h_{n,M} \end{bmatrix}, \begin{bmatrix} s_1 s_2 \\ \cdot \\ \cdot \\ \cdot \\ s_M \end{bmatrix}.
$$

(15)

Note that in Eq. (14) the subscript, J, denotes the set of most recent observations and flux estimates, and the subscript 1 represents the earliest observations.

The size of the large matrix of basis functions in Eq. (14) (which we will refer to as $\mathbf{H}^*$) grows rapidly as the number of observations and source regions increases. For example, for 20 source regions and a modest observational network of 50 sites with 20 years of monthly average data $\mathbf{H}^*$ is 12 000 by 4800. Consider a model with relatively coarse horizontal resolution of 8° latitude by 10° longitude, or 864 surface grid cells. An inversion for which every grid cell is a source region would have $\mathbf{H}^*$ dimensioned 12 000 by 207 360. As the size of the matrices in the problem grows, the computational cost becomes an important issue.

The atmospheric inversion problem becomes more efficient by observing that the basis functions, which are calculated by transporting pulses forward from each source region or backwards in time from each observation site using an adjoint model, are mixed throughout the troposphere until constant values are reached. Figure 1 illustrates this for two sampling locations, and shows that most of the signal from each source region occurs during the first 4–6 months. In the tropics, pulses from source regions decay more rapidly than at high latitudes due to rapid vertical mixing throughout a relatively deep tropospheric column. Note also that signals from adjacent regions dominate the responses at observation sites. This implies that a pulse traveling from Australia to Mace Head, Ireland would contribute relatively little signal compared to recent pulses from nearby source regions. In other words, Mace Head does not constrain Australian sources well. We therefore keep only a subset of the
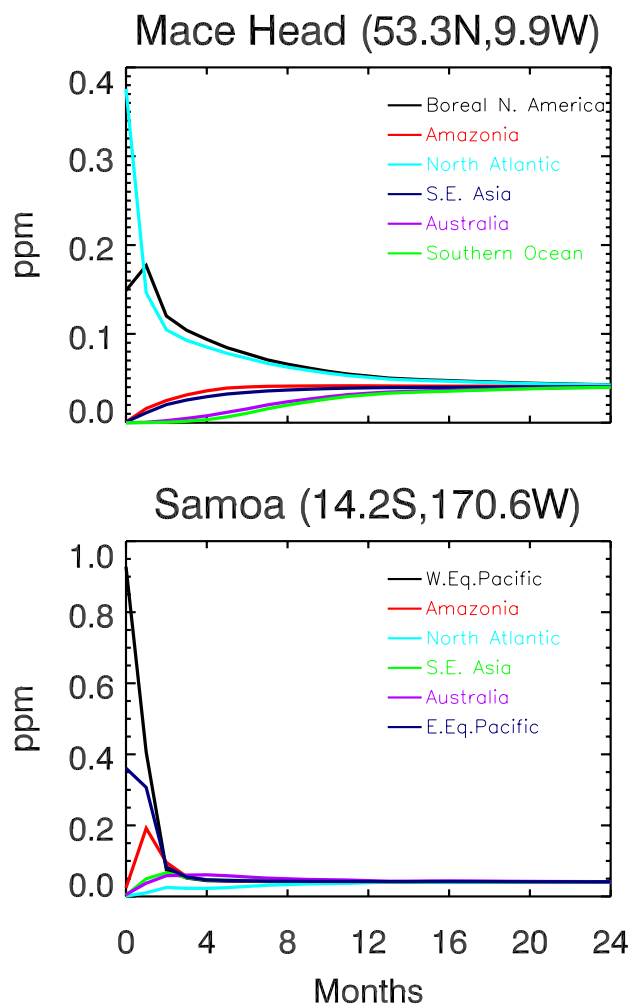

Fig. 1. Monthly average basis functions calculated as one-month pulses transported forward from each source region and sampled at a particular observation site (Mace Head and Samoa are shown here). The rate at which carbon was emitted from each region was 1 GtC/yr and emissions over each region were distributed using scaling derived from the CASA model net primary productivity (Randerson et al., 1997; Takahashi et al., 1999) and ocean flux estimates based on observed partial pressure of CO_2 in seawater.

transport information so that Eq. (14) effectively becomes:

$$
\begin{bmatrix} z_j \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{j,j} & \mathbf{H}_{j,j-1} & \cdot \cdot & \mathbf{H}_{j,j-P} \end{bmatrix} \begin{bmatrix} s_j \\ s_{j-1} \\ \cdot \\ \cdot \\ s_{j-P} \end{bmatrix}
$$

(16)

where j denotes the current time step and P is the number of months of transport information kept at each time step. As Eq. (16) implies, the source vector at each time step will be estimated P times, each time being compared to a different data vector. The $\mathbf{H}^*$ matrix gives the response to pulses emitted a time step at each of the observation sites, where $\mathbf{H}_{j,j}$
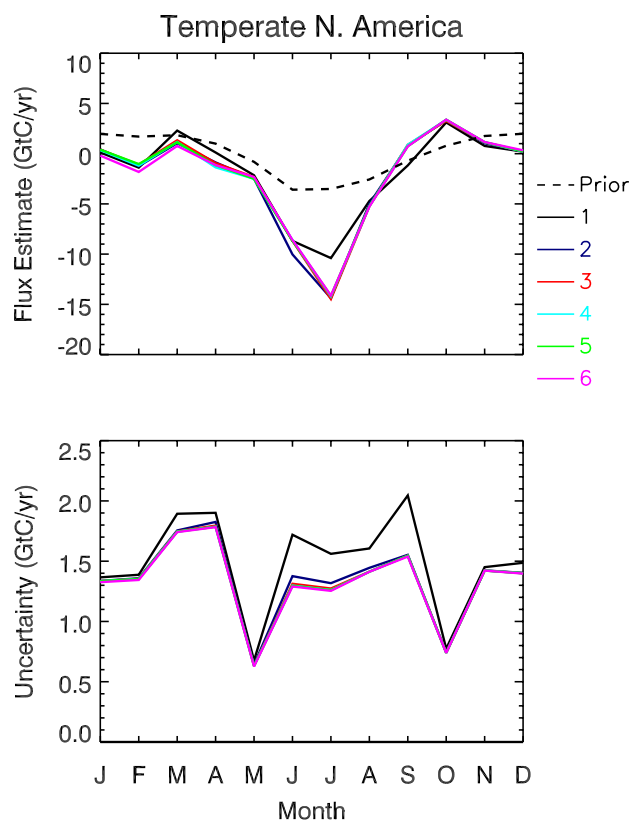
Fig. 2. Flux estimates (top) and uncertainties (bottom) for Temperate North America for one year out of a 5 year test inversion with 22 source regions. The legend refers to cases where 1 to 6 months of transport were retained in the state. The test inversion used a subset of the GLOBALVIEW data product, prior flux estimates from CASA model (Randerson et al., 1997; Takahashi et al., 1999), and cyclic meteorology from ECMWF (the European Centre for Medium-Range Weather Forecasting). 88 observation sites were used and the standard deviation of the model-data mismatch error was assumed to be 0.3 ppm for marine boundary layer sites and 1.5 ppm for continental sites. The uncertainties of the prior flux estimates were assumed to be 10 GtC/yr for the terrestrial biosphere, and 1 GtC/yr for the oceans.

are pulses emitted from the most recent month, and $\mathbf{H}_{j, j-P}$ are responses due to pulses emitted P months in the past. At each subsequent time step, a particular source vector is shifted downward, while the basis functions for each month are shifted to the right. A source vector and basis function matrix for a new month are added to the right hand side of Eq. (16). Once a source vector has been estimated P times it is dropped from the part of the state still being estimated. As shown in Fig. 2, both the estimate and its uncertainty converge with repeated iterations, and the last estimate is taken as the best estimate. Note that the largest changes between iterations occur after the first iteration. The difference between the second and third iteration is much smaller, and it is difficult to tell the difference between subsequent iterations.
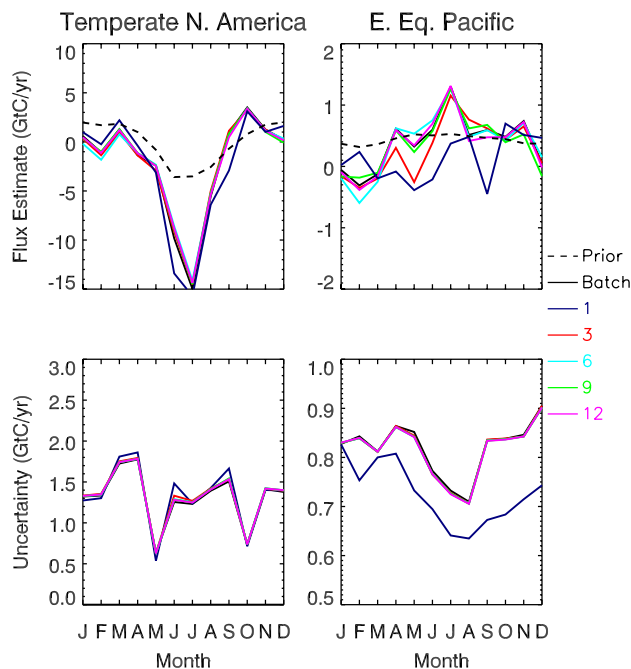

Fig. 3. Flux estimates (top) and uncertainties (bottom) for Temperate North America and the Eastern Equatorial Pacific for one year out of a 5 year test inversion. The prior flux estimates are shown as a black dashed line, estimates produced using the batch technique are represented by the solid black line, and the colored lines show estimates where varying numbers of months of transport were retained in the state. Details of the calculation are described in the text and in the Fig. 2 caption. The behavior of these two source regions is representative of all 22 source regions. Note that agreement between the Batch and the Kalman Smoother is very good for both regions if at least 6 months of transport is retained in the state.

As expressed in Eq. (16), observations from only one time level are used to estimate fluxes at P time steps, whereas in the Batch calculation all observations are used at once to estimates fluxes at all times. For the Kalman Smoother, the effect of stepping through the observations in this manner implies that for a particular time step, the use of observations from previous times is discontinued at a point where the observations no longer significantly constrain the current time step. For this reason, no significant amount of information is lost by stepping through the observations, although higher posterior uncertainties are expected since less information is used to constrain each set of fluxes. Even these differences will be shown to be small in the next section.

4 Comparison to the Batch Bayesian Synthesis Inversion

Solving Eqs. (9) and (10) (or Eqs. 12 and 13) using observations at all available times for flux estimates and uncertainties is referred to as the "Batch" technique. For this case,

each estimate is constrained by observations at all subsequent times so that all of the available data is used at once to determine the entire time series of fluxes and uncertainty estimates. When a subset of the observations is used, it is reasonable to expect that the decreased computing cost will come at the expense of larger estimated uncertainty, since each estimate is now constrained by less data. Since recent emissions produce the largest signals at measurement sites and the largest spatial gradients, the increases in estimated uncertainty are generally acceptably small provided that at least several months of transport are used.

Figure 3 shows results from a 5 year test inversion for fluxes from 22 source regions corresponding to those used in the TransCom 3 model intercomparison (Gurney et al., 2002, 2003; Law et al., 2003). For these calculations, the transport model was run with repeating meteorology (ECMWF assimilated winds from 1980). The prior fluxes were specified from the same sources used for TransCom 3 (Takahashi et al., 1999, CASA model), and were aggregated onto the Tracer Model 3 (TM3; Heimann and Koerner, 2003) 8×10 horizontal grid using the procedure described by Gurney et al. (2000). Representative values for the prior uncertainties and model-data mismatch errors were used for all calculations in this section (10 and 1 GtC/yr for the land biosphere and ocean prior fluxes, respectively, and 0.3 and 1.5 ppm for marine boundary layer and continental sites, respectively). Monthly average CO_2 abundances from the GLOBALVIEW data product for 1980–1985 were used at a subset of 88 measurement sites. The measurement sites were selected based on whether they represented marine boundary layer samples, were located far from large sources of emissions, or were generally fairly well-simulated by the transport model. Vertical profiles from tall towers and aircraft platforms were not used in these test inversions.

As shown in Fig. 3, the central value of the flux estimate probability distribution calculated with the fixed-lag Kalman smoother agrees well with the batch calculation, especially if the basis functions are transported for six months or more (i.e. each month of fluxes is estimated using at least six months of subsequent observations). The differences between the batch and Kalman smoother are largest for the case where only one month of transport is retained, and this is not recommended since considerable information remains in pulses well after the first month after emission. The length of time that basis functions need to be transported is determined by two opposing factors; the time taken to transport pulses from source regions to observation sites that constrain the fluxes, and the diffusion of pulses by atmospheric mixing which flattens spatial gradients.

Differences between the Kalman smoother and batch flux estimates are shown in Figs. 4 and 5 for a subset of flux regions. Although results are shown for only 4 of 22 source regions, they are representative of all source regions. As shown in the top panels of Figs. 4 and 5, the central value of the posterior probability distribution function agrees well for both techniques for North America with the exception of the case for which the basis functions are transported for only one month. The Kalman smoother estimates approach the batch estimates as the number of months of transport increases, as expected. For the Eastern Equatorial Pacific, the agreement is best if the basis functions are transported for at least nine months. For Amazonia, the differences between the Kalman smoother and the batch technique are approximately as large as the prior flux estimate for some months. It is interesting to note that the solutions closest to the batch are those for which the basis functions are transported for only three months. A possible explanation is that the vertical mixing in the transport model is too weak, allowing tropical continental sites to "see" pulses which have been transported for longer periods of time.

Figures 4 and 5 also show that the differences between the smoother and batch uncertainty estimates tend to be fairly small except for the case where the basis functions are transported for only one month. It is noteworthy that the estimated uncertainty for the Kalman smoother is sometimes smaller than that for the batch case. As noted above, since the batch technique uses all of the available data rather than a subset, it should give the lowest uncertainties. In addition, the Kalman smoother uncertainty estimates are often lowest for cases where the basis functions have been transported for the least amount of time. For example, see the uncertainty differences for Amazonia in Fig. 5. Here the estimate uncertainties for 3 month transport of basis functions are often smaller than those obtained using 6 or more months of transport. Uncertainties estimated using only one month of transport are the most questionable, since these estimates falsely appear to be much less uncertain than the batch estimates.

The explanation for this behavior lies in the fact that the Kalman smoother does not take into account the covariance between monthly fluxes that are no longer being estimated and those still being estimated. In addition, once a final estimate is made of a set of monthly fluxes, these fluxes are incorporated into the background state that is propagated forward in time by using the calculated transport fields (e.g. the basis functions). In effect, they are treated as known quantities, whereas they actually have associated uncertainties and correlations with fluxes still being estimated. The flux is therefore essentially assumed to be perfectly well-known and error correlations with time steps currently being estimated are neglected. Since time-dependent flux estimates tend to over- and undershoot the solution at successive time steps, the total uncertainty aggregated over time is likely to be smaller due to these temporal correlations. On the other hand, the incorporation of estimated fluxes into the background state without accounting for estimated uncertainty implies a certain level of uncertainty underestimation for subsequent time steps. In the next subsection, we propose a technique to correct these shortcomings.
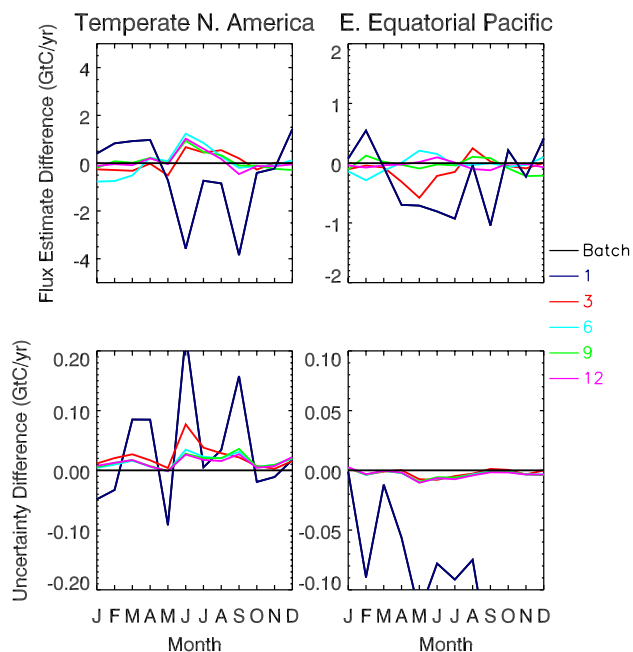
Fig. 4. Flux estimate differences (top, Kalman smoother – Batch) and uncertainty differences (bottom) for Temperate North America and the Eastern Equatorial Pacific for one year out of a 5 year test inversion. The legend refers to estimates produced with the batch technique, or estimates where varying numbers of months of transport were retained in the state. Details of the test inversion are described in the text and the Fig. 2 caption. Note that the differences are largest when only one month of transport is retained in the the state. Note also that the Kalman Smoother can produce smaller uncertainties than the Batch technique.
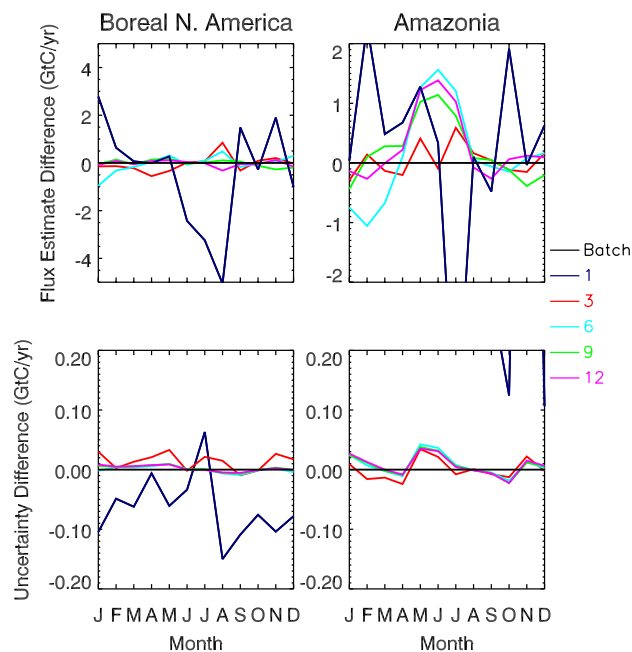
Fig. 5. Difference (Kalman smoother – Batch) in flux estimates (top) and uncertainties (bottom) for Boreal North America and Amazonia for one year out of a 5 year test inversion. The legend refers to estimates produced with the batch technique, or estimates where varying numbers of months of transport were retained in the state. Details of the test inversion are described in the text and the Fig. 2 caption. Note that the differences are quite large for the case where only one month of transport is retained in the state. The differences are also significant for Amazonia, a region that is not well-constrained by observations. The Kalman Smoother also produces uncertainties that may be larger than those using the Batch method.

5 Propagation of covariance

Equation (6) may be rewritten in the form

$$L_{s_u,s_v} = \left(z - \left[\mathbf{H}_u\mathbf{H}_v\right]\begin{bmatrix}s_u \\ s_v\end{bmatrix}\right)^T \mathbf{R}^{-1}\left(z - \left[\mathbf{H}_u\mathbf{H}_v\right]\begin{bmatrix}s_u \\ s_v\end{bmatrix}\right)$$
$$+ \left(\begin{bmatrix}s_u \\ s_v\end{bmatrix} - \begin{bmatrix}s_{\mathbf{p},u} \\ s_{\mathbf{p},v}\end{bmatrix}\right)^T \begin{bmatrix}\mathbf{Q}_{uu}\mathbf{Q}_{uv} \\ \mathbf{Q}_{vu}\mathbf{Q}_{vv}\end{bmatrix}^{-1}\left(\begin{bmatrix}s_u \\ s_v\end{bmatrix} - \begin{bmatrix}s_{\mathbf{p},u} \\ s_{\mathbf{p},v}\end{bmatrix}\right)$$
$$\tag{17}$$

where the subscript, u, pertains to the part of the state still being estimated and v pertains to the part of the state that is no longer being estimated. For example, if six months of transport information is used in the estimation, then $\mathbf{H}_u$ is dimensioned the number of measurement sites, n, by the number of fluxes to be estimated (six times the number of source regions for six months of transport). $\mathbf{H}_v$ is transport information from some number of time steps farther back than six months ago, for which estimates are no longer being made. It is dimensioned n by the number of months for which we intend to consider correlations times the number of

source regions. Likewise, the terms $\mathbf{Q}_{uv}$ and $\mathbf{Q}_{vu}$ represent the covariance between states still being estimated and those no longer being estimated.

Taking the derivative of Eq. (17) with respect to s_u gives

$$\frac{\partial L_{s_u,s_v}}{\partial s_u} = -(z - \mathbf{H}_u s_u - \mathbf{H}_v s_v)^T \mathbf{R}^{-1}\mathbf{H}_u + (s_u - s_{\mathbf{p},u})^T \mathbf{Q}_{aa}^{-1}$$
$$+ (s_v - s_{\mathbf{p},v})^T \mathbf{Q}_{ba}^{-1} \tag{18}$$

where

$$\mathbf{Q}^{-1} = \begin{bmatrix}\mathbf{Q}_{uu} & \mathbf{Q}_{uv} \\ \mathbf{Q}_{vu} & \mathbf{Q}_{vv}\end{bmatrix}^{-1} = \begin{bmatrix}\mathbf{Q}_{aa}^{-1} & \mathbf{Q}_{ab}^{-1} \\ \mathbf{Q}_{ba}^{-1} & \mathbf{Q}_{bb}^{-1}\end{bmatrix} \tag{19}$$

We then set Eq. (18) to zero. Given that we are no longer updating our estimate of s_v, then $s'_v = s_{\mathbf{p},v}$, and the last term on the right-hand side of Eq. (18) drops outs. Solving for s_u gives

$$s'_u = s_{\mathbf{p},u} + \mathbf{Q}_{aa}\mathbf{H}_u^T\left[\mathbf{R} + \mathbf{H}_u\mathbf{Q}_{aa}\mathbf{H}_u^T\right]^{-1}\left(z - \mathbf{H}_u s_{\mathbf{p},u} - \mathbf{H}_v s'_v\right). \tag{20}$$

Finally,

$$\mathbf{Q}_{aa}^{-1} = (\mathbf{Q}_{uu} - \mathbf{Q}_{uv}\mathbf{Q}_{vv}^{-1}\mathbf{Q}_{vu})^{-1} \qquad (21)$$

$$\mathbf{Q}_{ab}^{-1} = (\mathbf{Q}_{ba}^{-1})^T = -(\mathbf{Q}_{uu} - \mathbf{Q}_{uv}\mathbf{Q}_{vv}^{-1}\mathbf{Q}_{vu})^{-1}\mathbf{Q}_{uv}\mathbf{Q}_{vv}^{-1} \quad (22)$$

$$\mathbf{Q}_{bb}^{-1} = (\mathbf{Q}_{vv} - \mathbf{Q}_{vu}\mathbf{Q}_{uu}^{-1}\mathbf{Q}_{uv})^{-1} \qquad (23)$$

by using a matrix partitioning identity. Note that Eq. (20) looks exactly like Eq. (9), the original Kalman filter update equation, except that $\mathbf{H}_v\mathbf{s}_v$ is subtracted from $\mathbf{z}$ and $\mathbf{Q}$ is replaced by $\mathbf{Q}_{aa}$. $\mathbf{H}_v\mathbf{s}'_v$ accounts for the final estimate of $\mathbf{s}_v$ in the model-observation difference vector, $\mathbf{z}$, while $\mathbf{Q}_{aa}$ takes into account any correlations of $\mathbf{s}_v$ with states still be estimated ($\mathbf{s}_u$). As suggested by Eq. (20), $\mathbf{Q}_{aa}$ is the result of a correction to the covariance matrix of states still being estimated, $\mathbf{Q}_{uu}$. This correction takes the form of correlations between states still being estimated and the set of final estimates, normalized by the covariance matrix of the final estimates. This implies that the correction will be largest when the correlations are large and the uncertainty of $\mathbf{s}_v$ is small. The correction will be small when the correlations are small or the uncertainty of the final estimate is large. The correlation between successive time steps is expected to be the most significant; however, it is possible to apply the correction for up to one less than the number months of states which are still being estimated.

An expression for the covariance is found by calculating the inverse of the Hessian of the revised objective function (Eq. 17). This leads to

$$\mathcal{H}^{-1} = \begin{bmatrix} \mathbf{H}_u^T\mathbf{R}^{-1}\mathbf{H}_u + \mathbf{Q}_{aa}^{-1} & \mathbf{H}_u^T\mathbf{R}^{-1}\mathbf{H}_v + \mathbf{Q}_{ab}^{-1} \\ \mathbf{H}_v^T\mathbf{R}^{-1}\mathbf{H}_u + \mathbf{Q}_{ba}^{-1} & \mathbf{H}_v^T\mathbf{R}^{-1}\mathbf{H}_v + \mathbf{Q}_{bb}^{-1} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{Q}'_{uu} & \mathbf{Q}'_{uv} \\ \mathbf{Q}'_{vu} & \mathbf{Q}'_{vv} \end{bmatrix}$$
(24)

The top left term is the covariance of the part of the state that is still being estimated; the cross terms relate to the covariance of the part of the state that is no longer being estimated with the part of the state that is still being estimated. These are incorporated into the new expression for the covariance:

$$\mathbf{Q}' = \Big[(\mathbf{H}_u^T\mathbf{R}^{-1}\mathbf{H}_u + \mathbf{Q}_{aa}^{-1})$$
$$- (\mathbf{H}_u^T\mathbf{R}^{-1}\mathbf{H}_v + \mathbf{Q}_{ab}^{-1})(\mathbf{H}_v^T\mathbf{R}^{-1}\mathbf{H}_v + \mathbf{Q}_{bb}^{-1})^{-1}$$
$$(\mathbf{H}_v^T\mathbf{R}^{-1}\mathbf{H}_u + \mathbf{Q}_{ba}^{-1}) \Big]^{-1} \qquad (25)$$

where $\mathbf{Q}_{aa}^{-1}$, $\mathbf{Q}_{ab}^{-1}$, $\mathbf{Q}_{ba}^{-1}$ and $\mathbf{Q}_{bb}^{-1}$ are defined by Eqs. (21–23s). Note that $(\mathbf{H}_u^T\mathbf{R}^{-1}\mathbf{H}_u + \mathbf{Q}_{aa}^{-1})^{-1} = \mathbf{Q}'_{uu}$

Equation (23) is similar in form to Eq. (13) with $\mathbf{Q}_{aa}$ replacing $\mathbf{Q}$ and subtraction of a term describing correlations in covariance between states still being estimated and states for which final estimates have been calculated. Note that the maximum months of that may be retained in the correlation propagation scheme is the number of months of transport retained less one. As shown in Fig. 6, the differences between
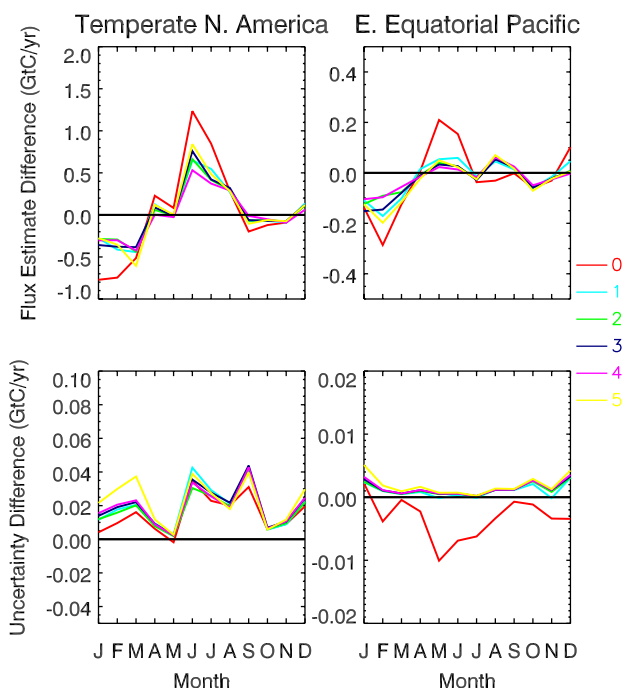

Fig. 6. Difference (Kalman smoother – Batch) in flux estimates (top) and uncertainties (bottom) for Temperate North America and the Eastern Equatorial Pacific for one year out of a 5 year test inversion. Details of the test inversion are described in the text and the Fig. 2 caption. The numbers given in the legend refer to the number of months of correlations retained outside of the state currently being estimated. Note that the agreement between flux estimates calculated with Batch and Kalman Smoother improves if the covariance propagation scheme is used with the Kalman Smoother. Note also that the Kalman Smoother uncertainties are always larger than the Batch uncertainties if the covariance propagation scheme is included.

including 1 month of correlations and not propagating the covariance are large (the red curve, 0 months and the light blue curve, 1 month), but including additional months has a generally small effect. It is important to note that including the propagation of covariance improves agreement between the Kalman smoother and the Batch calculations significantly for the estimated fluxes. It is therefore recommended that the Fixed-Lag Kalman Smoother should also include the covariance propagation scheme described here. Although the differences are fairly small for the uncertainty estimates, the Kalman smoother uncertainty estimates are always greater than the Batch uncertainty estimates, as expected. Figure 7 shows the same comparison for Amazonia, where the agreement between the Kalman smoother and the Batch technique is also improved.

The relative effects of propagating the covariance on calculations with varying amounts of basis function transport are shown in Fig. 8. In contrast to Fig. 5, the calculation using only 3 months of transport generally has the highest
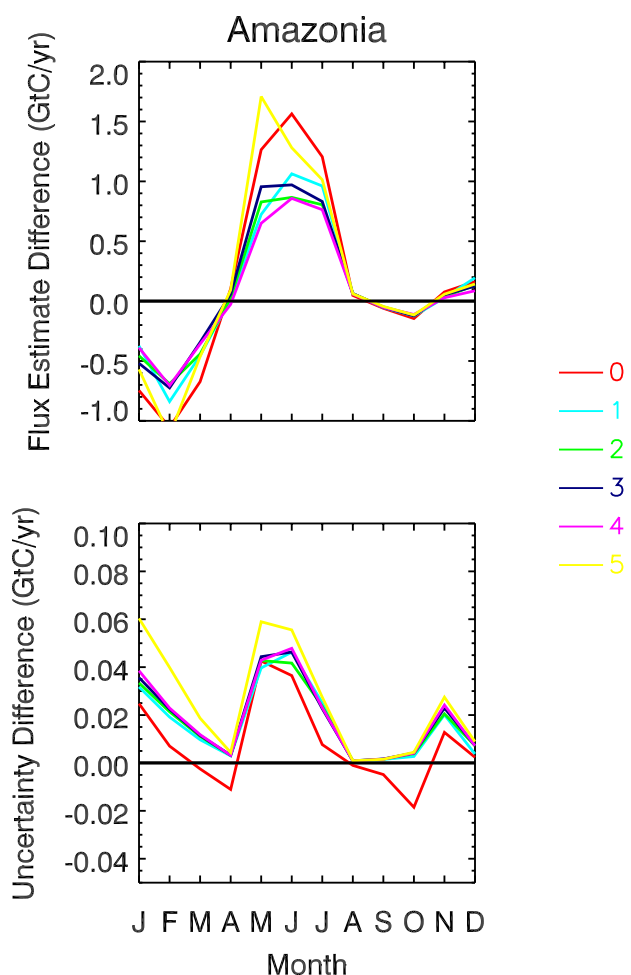
Fig. 7. Difference (Kalman smoother – Batch) in flux estimates (top) and uncertainties (bottom) for Amazonia for one year out of a 5 year test inversion. Details of the test inversion are described in the text and the Fig. 2 caption. The numbers given in the legend refer to the number of preceeding months of correlations retained outside of the state currently being estimated. 6 months of transport were retained in the state for this calculation. Note that the agreement of the Kalman Smoother flux estimates with the Batch estimates improves if the covariance propagation scheme is used. Note also that the Kalman Smoother uncertainty estimates are always larger than the Batch uncertainties for the case where the covariance propagation scheme is used.

estimated uncertainties, although the differences are fairly small. It is interesting to note that the estimates using 9 months of transport to have slightly higher uncertainty differences for December through March. This time period is the rainy season for the Amazon Basin, and it is likely that due to increased convection, the cases with fewer months of transport are not able to constrain this region as well as for the case with 9 months of transport. This is consistent with the expectation that estimates using the least amount of data should, in general, have the highest estimated uncertainties.
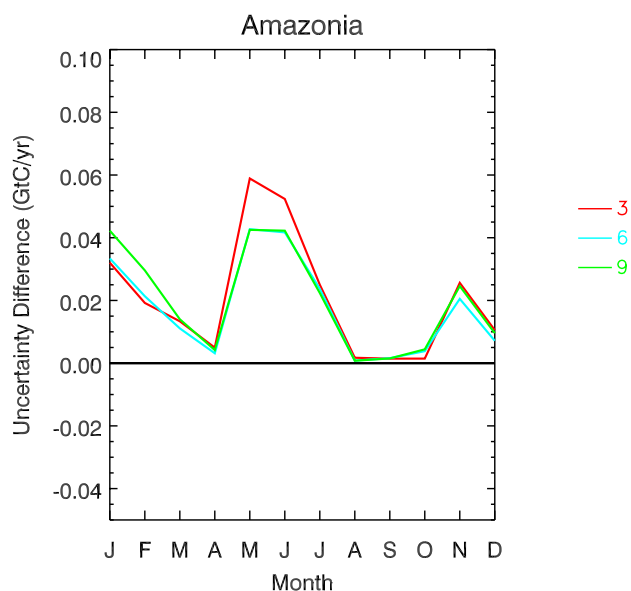

Fig. 8. Difference (Kalman smoother – Batch) in uncertainty estimates for Amazonia for one year out of a 5 year test inversion. Details of the test inversion are described in the text and the Fig. 2 caption. The numbers given in the legend refer to the number of months the basis functions are transported. One month of covariance was propagated from the part of the state no longer being estimated. Note that the uncertainties for the Kalman Smoother are always larger than those for the Batch technique, as expected.

6   Discussion

The method we have described up to this point assumes that the inversion is done "on-line". In other words, the transport model is run forward with prior flux estimates to produce a prediction of CO_2 abundances that are then compared with observed CO_2 at each site. The final flux estimates produced by the inversion are then incorporated into the background state of CO_2 using transport information in the form of the basis functions, and the transport model is run to the next inversion time step. The fixed-lag Kalman smoother may also be used "off-line" without actually running the transport model to make predictions of CO_2. The pre-calculated basis functions are used to reproduce the transport fields over the entire time period of interest, and are used to propagate the optimized fluxes forward in time, thus updating the background state. In principle, this gives exactly the same results as the "on-line" case, since the basis functions are a formal decomposition of the model's transport. Using the basis functions is more computationally efficient since only a few simple matrix multiplications are required to yield concentrations at each observations site used, compared to a full transport model run for the "on-line" fixed-lag Kalman smoother. If the basis functions are stored as three-dimensional arrays, then changes in the network configuration may be easily taken into account by appropriately sampling the fields. On

the other hand, if the basis functions are computed using the adjoint approach, then new adjoint runs must be initiated for each added observation site.

A disadvantage of the offline method is that a perfect decomposition of forward transport requires the basis functions to extend infinitely long in time, eliminating the advantage of needing only a few months of basefunctions at a time in solving the inverse problem with the Kalman smoother. For batch inversions, basis functions are often propagated for a period of a few years and are thereafter approximated by an asymptotic approach to the constant well-mixed value eventually achieved by atmospheric mixing. We have chosen to truncate the basefunctions after several months, since information about a given initial pulse is diminishing due to mixing in the atmosphere. The truncated part is replaced by a simple exponential decay to the asymptotic value achieved after a sufficiently long period of mixing by atmospheric transport. In this way we hope to preserve the numerical efficiency of the Kalman smoother for off-line calculations. Clearly, one must ensure that this truncation is not too short, and one must assess the impact on flux estimates.

We find that when at least six months of basis functions are used to recreate the background state, the difference in model predicted concentrations and estimated fluxes are acceptably small. The differences are random and not in the form of a bias (i.e. they have a zero mean), and always well within the regional and monthly flux uncertainty. Using the Kalman smoother in "off-line" mode therefore comes at a small price in increased uncertainty, with great savings in computational costs. Solving the inverse problem can be done off-line with a subset of the basis functions, and does not require further use of a transport model. Using the off-line Kalman smoother, the solution was obtained in approximately five minutes on a simple Macintosh PC and implemented using IDL. With this computational efficiency, many more experiments can be done quickly to explore the sensitivity of the flux estimates along several axes of uncertainty.

Another aspect of our fixed-lag Kalman smoother that we wish to highlight is its inherent flexibility in regard to the use of prior flux and uncertainty estimates. In this study we have used prior flux information from the (Takahashi et al., 1999, CASA model) as first guesses when the fluxes at a particular time step are estimated for the first time. Thereafter, the previous smoother estimate and estimated uncertainties are used as the first guess for each subsequent iteration. An alternative approach is to use results generated by the smoother as the first guess for each new time step (a persistence model). These could be in the form of estimates from the previous time step, or averages of previous estimates. Estimates produced this way would be determined only by observations (with the exception of the very first time step, for which a first guess must be used). The fact that estimation of time-varying fluxes occurs sequentially in our method makes a calculation such as this straightforward, and we are currently exploring the use of our smoother in this way.

7 Conclusions

The fixed-lag Kalman smoother technique introduced in this study offers a numerically efficient method for estimating fluxes and flux uncertainties of atmospheric trace species. This method relies on the fact that information about the spatial distribution of sources is preserved only for a limited amount of time before mixing by atmospheric transport dilutes signals from source regions. We have shown that excellent agreement with the standard Bayesian synthesis "batch" technique can be achieved by retaining transport information in the basis functions for as little as 6 months. Transporting basis functions for this shorter period of time, rather than the usual multiple years greatly reduces the computational expense of the flux estimation. In addition, the sizes of the matrices in the problem are also much smaller. With this technique, it therefore becomes possible to do inversions spanning multiple decades with relatively little computational expense. Calculations testing sensitivity to specified parameters are therefore feasible. If three dimensional basis functions are pre-calculated and stored, then it also becomes possible to evaluate issues related to the distribution of observing sites. Networks with large numbers of sites may be easily used to produce flux estimates with the fixed-lag Kalman smoother.

We have shown that the cost of the numerical efficiency gains is a relatively small increase in the estimated flux uncertainty. A problem that we encountered with the fixed-lag Kalman smoother in its original form was that the estimated uncertainty for the fixed-lag Kalman smoother was often smaller than for the Batch technique. The reason for this was that the covariance with fluxes no longer being estimated was not being propagated forward in time. Essentially, once a final estimate of the flux at a particular time step was computed, its effect was included into the background state with no uncertainty and no correlation with states still being estimated. We therefore developed a method to correct for this, resulting in flux uncertainty estimates that are consistent with the expectation that the uncertainties for the Batch technique are the smallest since the most data are used to constrain fluxes at a particular time, whereas for the Kalman smoother only a subset of observations are used for any one time-step.

The fixed-lag Kalman smoother introduced in this study is a promising technique for future atmospheric inversion problems because it can rather easily handle inversions over long periods of time, such as multiple decades. Furthermore, it can potentially be used to handle large observational networks. In the future, atmospheric inversions may well be used to estimate fluxes at ever smaller scales using large amounts of data collected at high time frequency. The fixed-lag Kalman smoother described in this study could potentially be used as the basis of an ensemble technique capable of handling the demanding flux estimation problems of the future (Peters, 2005).

Edited by: M. Heimann

References

Bousquet, P., Peylin, P., Ciais, Le Quere, C., Friedlingstein, P., and Tans, P. P.: Regional changes in carbon dioxide fluxes of land and oceans since 1980, Science, 290, 5495 1342–1346 2000.

Bousquet, P., Peylin, P., Ciais, P., Ramonet, M., and Monfray, P.: Inverse modeling of annual atmospheric CO2 sources and sinks, 2. Sensitivity study, J. Geophys. Res.-A, 104, 26 179–26 193, 1999.

Denning, A. S., Holzer, M., Gurney, K. R., Heimann, M., Law, R. M., Rayner, P. J., Fung, I. Y., Fan, S. M., Taguchi, S., Friedlingstein, P., Balkanski, Y., Taylor, J., Maiss, M., Levin, I.: Three-dimensional transport and concentration of SF6 – A model inter-comparison study (TransCom 2), Tellus B, 51, 2 266–297 1999.

Enting, I. G.: Inverse problems in atmospheric constituent transport, atmospheric and space science series, Cambridge University Press, Cambridge, 2002.

Enting, I. G., Trudinger, C. M., and Francey, R. J.: A Synthesis Inversion of the Concentration and Delta-C-13 of Atmospheric CO2, Tellus Series B-Chemical and Physical Meteorology, 47, 35–52, 1995.

Engelen, R. J., Denning, A. S., Gurney, K. R., and Modelers, TransCom 3.: On error estimation in atmospheric CO2 inversions, J. Geophys. Res.-A, 107, 22, 4635, doi:10.1029/2002JD002195,2002.

Fan, S., Gloor, M., Mahlman, J., Pacala, S., Sarmiento, J., Takahashi, T., and Tans, P.: A large terrestrial carbon sink in North America implied by atmospheric and oceanic carbon dioxide data and models, Science, 282, 442–446, 1998.

Gelb, A.: Applied Optimal Estimation, MIT Press, Cambridge, MA, 1974.

Gurney, K. R., Law, R., Rayner, P. J., and Denning, A. S.: TransCom 3 Experimental Protocol, Tech. Rep. Paper No. 707, Department of Atmospheric Science, Colorado State University, USA, 2000.

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.

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. M., Fung, I. Y., Gloor, M., Heimann, M., Higuchi, K., John, J., Kowalczyk, E., Maki, T., Maksyutov, S., Peylin, P., Prather, M., Pak, B. C., Sarmiento, J., Taguchi, S., Takahashi, T., and Yuen, C. W.: TransCom 3 CO2 inversion intercomparison: 1. Annual mean control results and sensitivity to transport and prior flux information, Tellus Series B-Chemical and Physical Meteorology, 55, 555–579, 2003.

Gurney, K. R., Law, R. M., Denning, A. S., Rayner, P. J., Pak, B. C., Baker, D., Bousquet, P., Bruhwiler, L., Chen, Y. H., Ciais, P., Fung, I., Heimann, M., John, J., Maki, T., Maksyutov, S., Peylin, P., Prather, M., and Taguchi, S.: Transcom 3 inversion intercomparison: model mean results for the estimation of sea-

sonal carbon sources and sinks, Global Biogeochem. Cycles, 18, doi:10.1029/2003GBC002108, 2004.

HaasLaursen, D. E., Hartley, D. E., and Prinn, R. G.: Optimizing an inverse method to deduce time-varying emissions of trace gases, J. Geophys. Res.-A, 101, 22 823–22 831, 1996.

Hartley, D. E.: Deducing trace gas emissions using an inverse method in three-dimensional chemical transport models, Tech. Rep. 17, MIT Center for Global Change Science, 1992.

Hartley, D. E. and Prinn, R. G.: Feasibility of determining surface emissions of trace gases using an inverse method in a three-dimensional chemical transport model, J. Geophys. Res., 98, 5183–5197, 1993.

Heimann, M. and Koerner, S.: The global atmospheric tracer model TM3, model description and users manual, release 3.8a, Tech. Rep. No. 5, Max Planck Institute for Biogeochemistry (MPIBGC), 2003.

Kalman, R.: A new approach to linear filtering and prediction problems, J. Basic Eng. (ASME), 82D, 35–45, 1960.

Kaminski, T., Rayner, P., Heimann, M., and Enting, I.: On aggregation errors in atmospheric transport inversions, J. Geophys. Res.-A, 106, D5 4703–4715 2001.

Kaminski, T., Heimann, M., and Giering, R.: A coarse grid three-dimensional global inverse model of the atmospheric transport – 2. Inversion of the transport of CO2 in the 1980s, J. Geophys. Res.-A, 104, 18 555–18 581, 1999.

Law, R. M., Chen, Y. H., and Gurney, K. R.: TransCom 3 CO2 inversion intercomparison: 2. Sensitivity of annual mean results to data choices, Tellus Series B-Chemical and Physical Meteorology, 55, 580–595, 2003.

Louis, J. F.: A parametric model of vertical eddy fluxes in the atmosphere, Boundary Layer Meteorol., 17, 178–202, 1979.

Peters, W.: An ensemble data assimilation system to estimate CO2 surface fluxes from atmospheric trace gas observations, J. Geophys. Res., doi:2005JD006157R, in press, 2005.

Randerson, J., Thompson, M., Conway, T., Fung, I., and Field, C.: The contribution of terrestrial sources and sinks to trends in the seasonal cycle of atmospheric carbon dioxide, Global Biogeochem. Cycles, 11, 535–560, 1997.

Rayner, P. J. and Law, R. M.: The interannual variability of the global carbon cycle, Tellus Series B-Chemical and Physical Meteorology, 51, 210–212, 1999.

Rodenbeck, C., Houweling, S., Gloor, M., and Heimann, M.: CO2 flux history 1982-2001 inferred from atmospheric data using a global inversion of atmospheric transport, Atmos. Chem. Phys., 3, 1919–1964, 2003,
SRef-ID: 1680-7324/acp/2003-3-1919.

Russell, G. and Lerner, J.: A new finite-differencing scheme for the tracer transport equation, J. Appl. Meteorol., 20, 1483–1498, 1981.

Takahashi, T., Wanninkhof, R., Feely, R. A., Weiss, R. F., Chipman, D., Bates, N., Olafsson, J., Sabine, C., and Sutherland, S. C.: Net sea-air CO2 flux over the global oceans: An improved estimate based on the sea-air pCO2 difference, in: 2nd CO2 in Oceans Symposium, Natl. Inst. for Environ. Stud., Tsukuba, Japan, 1999.

Tiedke, M.: A Comprehensive Mass Flux Scheme for Cumulus Parameterization in Large-Scale Models, Monthly Weather Review, 117, 1779–1800, 1989.