

Variational data assimilation for atmospheric CO₂

By DAVID F. BAKER^{1*}, SCOTT C. DONEY² and DAVID S. SCHIMEL¹, ¹National Center for Atmospheric Research, 1850 Table Mesa Drive, Boulder, CO 80307, USA; ²Woods Hole Oceanographic Institution, 266 Woods Hole Road, Woods Hole, MA 02543-1543, USA

(Manuscript received 13 January 2006; in final form 26 June 2006)

ABSTRACT

The sources and sinks of important climatic trace gases such as carbon dioxide (CO₂) are often deduced from spatial and temporal variations in atmospheric concentrations. Reducing uncertainties in our understanding of the contemporary carbon budget and its underlying dynamics, however, requires significantly denser observations globally than is practical with in situ measurements. Space-based measurements appear technically feasible but require innovations in data analysis approaches. We develop a variational data assimilation scheme to estimate surface CO₂ fluxes at fine time/space scales from such dense atmospheric data. Global flux estimates at a daily time step and model-grid spatial resolution (4° × 5° here) are rapidly achieved after only a few dozen minimization steps. We quantify the flux errors from existing, planned and hypothetical surface and space-borne observing systems. Simulations show that the planned NASA Orbital Carbon Observatory (OCO) satellite should provide significant additional information beyond that from existing and proposed in situ observations. Improvements in data assimilation techniques and in mechanistic process models are both needed to fully exploit the emerging global carbon observing system.

1. Introduction

Atmospheric concentration measurements provide critical information about sources and sinks of CO₂, CO, CH₄, N₂O and other important trace gases. Inverse methods are generally used to interpret the data, with atmospheric transport models providing the link between surface gas fluxes and their impact on atmospheric concentrations at later times. Over the past few decades, most CO₂ data have come from an expanding network of in situ sites, mostly surface flasks sampled at a weekly frequency (GLOBALVIEW-CO₂, 2004). These allow surface CO₂ fluxes to be estimated robustly on a monthly basis for only about a dozen globe-spanning emission regions (Bousquet et al., 2000; Rödenbeck et al., 2003; Baker et al., 2006). Data with higher spatial and temporal resolution are now becoming increasingly available from in situ continuous analysers, routine aircraft flights, and eddy flux towers. New types of measurements, such as column-averaged CO₂ from upward-looking Fourier transform spectrometers, are being introduced. Even more promising, concentration data from satellites may soon provide the greatest coverage boost. This is already the case for CO and CH₄ with the availability of MOPITT and SCIAMACHY products (Deeter et al., 2003; Frankenberg et al., 2005). For CO₂,

mid-troposphere column-averaged concentrations have been retrieved from AIRS (Crevoisier et al., 2004) and SCIAMACHY (Buchwitz et al., 2005), as well as upper-troposphere concentrations from TOVS records (Chedin et al., 2003). Though these have been of marginal use so far, both the Japanese and American space agencies plan to launch satellite missions in 2008 specifically designed to measure accurate column-integrated CO₂ concentrations, with sensitivity down to the surface.

High-density data should allow us to address a variety of science and policy questions that could not be answered before. Resolving surface fluxes to the regional biome level will help quantify the relative importance of the key driving processes; resolving them to the level of individual countries is needed for carbon management and for the verification of international emissions treaties (Dilling et al., 2003). To obtain this spatial resolution, daily to synoptic time variability must be resolved; this will also clarify the response of the different carbon cycle components to the short-term meteorological drivers (Geels et al., 2004). Resolving the fluxes diurnally could help separate land biospheric fluxes into their photosynthetic and respiratory components (Braswell et al., 2005).

The inverse methods most often used in the past become computationally impractical for processing this new high-density data. Here, we present a variational data assimilation approach that bypasses these computational limitations. Using the approach in a simulation framework, we test how well data from five CO₂ measurement networks constrain daily-mean CO₂ fluxes at model-grid spatial resolution (in our case, 4° × 5°).

*Corresponding author.
e-mail: dfb@ucar.edu
DOI: 10.1111/j.1600-0889.2006.00218.x

2. Method

In the batch least-squares methods used most commonly in the past (e.g. the ‘Bayesian synthesis’ approach; see Enting, 2002), a set of coupled linear equations relating the fluxes and the measured concentrations at different times is inverted in a single step. This approach has worked well for the limited data available through about the year 2000, but the increasing availability of new high-frequency data sources (including continuous in situ analysers and satellite-based column-integrated products), with the attendant possibility for estimating the fluxes at finer time and space resolution, is rendering this approach computationally infeasible. Three main problems arise as the problem size (the number of both estimated fluxes and measured concentrations) grows: (1) the matrix to be inverted becomes too large to fit into computer memory; (2) the direct solution of the linear equations becomes prohibitively expensive and (3) the greater number of atmospheric transport model integrations required to fill the transport matrix becomes too costly to run.

Various approaches have been tried to address these problems. Bousquet et al. (2000) broke the inversion span into several overlapping pieces, solving several smaller matrix inversions in place of a single large one. This idea may be taken further by performing the estimation sequentially, rather than in a single ‘batch’ inversion, using the Kalman filter (Baker, 2001; Bruhwiler et al., 2005). While this approach does speed up the computations, the need to account for multiple past flux emission times in the state vector forces the matrices operated upon to remain rather large. Ultimately, the need for the conventional Kalman filter to propagate and update a full-rank covariance matrix limits its usefulness for large problems. Recent work using ensemble Kalman filters (Peters et al., 2005) promises to surmount this problem. For the situation where the number of measurements is significantly smaller than the number of estimated fluxes, it is more efficient to generate the basis functions linking fluxes to concentrations by running the adjoint of the transport model backwards in time from the measurement site back to the emission region. Kaminski et al. (1999a,b) first used this approach to solve for fine-scale fluxes from flask data in a standard Bayesian inversion. Peylin et al. (2005) made a further improvement by noting that the size of the inversion is smaller when performed in the measurement space rather than in the flux space. But what can be done when both the number of measurements and estimated fluxes are large? In that case, the direct matrix inversion may be replaced by an iterative approach if the adjoint is available to allow the cost function gradient to be computed efficiently. Rödenbeck (2005) has used this approach to estimate CO₂ fluxes and Rayner et al. (2005) to estimate terrestrial model parameters in global atmospheric transport inversions.

Here, we develop these adjoint-based descent methods further in the context of variational data assimilation. Our variational data assimilation approach sidesteps all three problems noted above by casting the estimation of the fluxes as a minimization

problem and solving it with an iterative descent method. There is no matrix to store or invert, and the number of transport model runs required is proportional to the number of descent iterations rather than the number of estimated parameters. Iterative solution methods such as this achieve great computational efficiencies by not running the minimization to complete convergence, but by stopping at an approximate solution that is still acceptably close to the exact one. The key to using the minimization framework is the efficient computation of the cost function gradient, which is achieved in the variational approach through the derivation of the adjoint of the transport model (Bennett, 2002; Errico, 1997). The method is similar to the ‘4D-Var’ assimilation methods used in numerical weather prediction (NWP), though here it is used to solve at multiple time steps across a long span of data in a single retrospective analysis, rather than for just initial conditions in repeated near-real-time analyses across short spans (e.g. 6 h) as in NWP.

Suppose we have a transport model that computes atmospheric concentrations \mathbf{x}_i at times i across $i = 0, \dots, I$, forced by surface sources \mathbf{u}_k constant across emission spans $k = 1, \dots, K$. Given atmospheric measurements \mathbf{z}_j at times $j \in [0, I]$, we seek to estimate the surface sources \mathbf{u}_k plus the concentration field \mathbf{x}_0 at time $t = 0$, that together minimize a cost function of the form

$$2L = \sum_{i=0}^{I-1} \sum_j [\mathbf{h}_j(\mathbf{x}_i) - \mathbf{z}_j]^T \mathbf{R}_j^{-1} [\mathbf{h}_j(\mathbf{x}_i) - \mathbf{z}_j] \delta_{ij} + \sum_{k=1}^K (\mathbf{u}_k - \mathbf{u}_k^o)^T \mathbf{P}_{\mathbf{u}_k^o}^{-1} (\mathbf{u}_k - \mathbf{u}_k^o) + (\mathbf{x}_0 - \mathbf{x}_0^o)^T \mathbf{P}_{\mathbf{x}_0^o}^{-1} (\mathbf{x}_0 - \mathbf{x}_0^o) + \sum_{i=0}^{I-1} \lambda_{i+1}^T (\Phi_{i+1}^i \mathbf{x}_i + \mathbf{G}_i \mathbf{u}_k \delta_{ik} - \mathbf{x}_{i+1}), \quad (1)$$

where $\delta_{ij} = 1$ for measurements \mathbf{z}_j at all times j falling inside time step i and $\delta_{ij} = 0$ otherwise, $\delta_{ik} = 1$ when time i falls within emission span k and $\delta_{ik} = 0$ otherwise, \mathbf{u}_k^o and \mathbf{x}_0^o are a priori estimates of \mathbf{u}_k and \mathbf{x}_0 , and \mathbf{h}_j are measurement operators (possibly non-linear). \mathbf{R}_j , $\mathbf{P}_{\mathbf{u}_k^o}$ and $\mathbf{P}_{\mathbf{x}_0^o}$ are covariance matrices describing the errors in \mathbf{z}_j , \mathbf{u}_k^o and \mathbf{x}_0^o , respectively, which are assumed to be unbiased. The last term in (1) consists of a sum of dynamical mismatches adjoined to the cost function with Lagrange multipliers λ_{i+1} ; this constraint approach forces \mathbf{x}_i and \mathbf{u}_k to exactly satisfy a set of dynamical equations

$$\mathbf{x}_{i+1} = \Phi_{i+1}^i \mathbf{x}_i + \mathbf{G}_i \mathbf{u}_k \delta_{ik} \quad i = 0, \dots, I - 1, \quad (2)$$

where the transformation matrix Φ_{i+1}^i represents the linearized dynamics of our atmospheric transport model along the forward concentration trajectory (i.e. the tangent linear model), and matrix \mathbf{G}_i describes how fluxes \mathbf{u}_k perturb the concentrations in the lower levels of the model across time i .

The requirement from the calculus of variations that $\partial L / \partial \mathbf{x}_i = \mathbf{0}$ at the minimum (for $i = 1, \dots, I$) provides the extra equations

required to solve for the introduced Lagrange multipliers:

$$\begin{aligned}
 \lambda_i &= (\Phi_{i+1}^i)^T \lambda_{i+1} + \sum_j \left(\frac{\partial \mathbf{h}_j}{\partial \mathbf{x}_i} \right)^T \mathbf{R}_j^{-1} [\mathbf{h}_j(\mathbf{x}_i) - \mathbf{z}_j] \delta_{ij} \\
 \lambda_I &= \sum_j \left(\frac{\partial \mathbf{h}_j}{\partial \mathbf{x}_I} \right)^T \mathbf{R}_j^{-1} [\mathbf{h}_j(\mathbf{x}_I) - \mathbf{z}_j] \delta_{Ij}.
 \end{aligned} \quad (3)$$

These are the Euler–Lagrange equations, which may be solved for λ_i by running the adjoint $(\Phi_{i+1}^i)^T$ of the forward tangent linear model backwards in time from λ_I , forced by weighted measurement mismatches $[\mathbf{h}_j(\mathbf{x}_i) - \mathbf{z}_j] \delta_{ij}$ at measurement times j . We have coded and tested a near-exact adjoint to our forward model code for use in computing these backward-in-time integrals.

With λ_i so determined, the gradient of L with respect to the control variables \mathbf{u}_k and \mathbf{x}_0 may be obtained as

$$\begin{aligned}
 \frac{\partial L}{\partial \mathbf{u}_k} &= \sum_i \mathbf{G}_i^T \lambda_{i+1} \delta_{ik} + \mathbf{P}_{\mathbf{u}_k}^{-1} (\mathbf{u}_k - \mathbf{u}_k^o) \equiv \nabla_{\mathbf{u}_k} \\
 k &= 1, \dots, K \\
 \frac{\partial L}{\partial \mathbf{x}_0} &= (\Phi_1^0)^T \lambda_1 + \mathbf{P}_{\mathbf{x}_0}^{-1} (\mathbf{x}_0 - \mathbf{x}_0^o) \\
 &+ \sum_j \left(\frac{\partial \mathbf{h}_j}{\partial \mathbf{x}_0} \right)^T \mathbf{R}_j^{-1} [\mathbf{h}_j(\mathbf{x}_0) - \mathbf{z}_j] \delta_{0j} \equiv \nabla_{\mathbf{x}_0}.
 \end{aligned} \quad (4)$$

At the stationary point, $\partial L / \partial \mathbf{u}_k = \mathbf{0}$ and $\partial L / \partial \mathbf{x}_0 = \mathbf{0}$. These equations, along with (3), define a coupled 2-point boundary value problem that must be solved to find the optimal control variables. For problems of small dimension, they may be solved directly using the sweep method (Bryson and Ho, 1975) or the representer method of Bennett (2002). For larger problems of the sort addressed here, this becomes computationally impractical, and approximate solutions must be sought instead. The gradient of the cost function with respect to our control parameters, defined by $\nabla = (\nabla_{\mathbf{u}_1}^T, \dots, \nabla_{\mathbf{u}_K}^T, \nabla_{\mathbf{x}_0}^T)^T$, may be used in any of a number of popular descent methods (e.g. the conjugate gradient method) to find values for \mathbf{u}_k and \mathbf{x}_0 that approximately minimize L . Here we use the variable metric method of Broyden, Fletcher, Goldfarb, and Shanno (BFGS) (Stoer and Bulirsch, 1980), preconditioned with the a priori covariance matrices $\mathbf{P}_{\mathbf{u}_k}^o$ and $\mathbf{P}_{\mathbf{x}_0}^o$. The BFGS method allows a low-order approximation to the a posteriori covariance matrix to be reconstructed efficiently from the updates to \mathbf{x} and τ saved at each descent step. Besides providing some insight into the information content of the a posteriori estimate, the low-order covariance matrix may be used to construct the covariance matrix for errors in the final concentration field, which may then be used as $\mathbf{P}_{\mathbf{x}_0}^o$ in an assimilation for the subsequent span. We have also obtained marginally faster convergence with the BFGS method than with the conjugate gradient method.

3. Results

To test our assimilation approach, we have set up simulation experiments in which the true fluxes are known and are used to simulate measurements from potential CO₂ monitoring networks, with random data errors added. These simulated ‘true’ measurements are fed into the variational data assimilation system and are used to correct an a priori estimate of the fluxes. The resulting a posteriori estimate is then compared to the (known) true fluxes to test the performance of the inversion method and to assess how well the observation networks constrain the true fluxes. Similar simulation studies have been performed before to assess the value of space-based measurements (Rayner and O’Brien, 2001; Houweling et al., 2004), but these have used the earlier direct inversion methods at lower resolution.

For the land regions, the ‘true’ fluxes are taken from the Lund–Potsdam–Jena (LPJ) land biosphere model (Sitch et al., 2003) and the ‘prior’ ones from the Carnegie–Ames–Stanford Approach (CASA) land biosphere model (Randerson et al., 1997). The ‘true’ air–sea fluxes are from the NCAR ocean carbon model (Doney et al., 2004) and the ‘prior’ from the Takahashi et al. (1999) data synthesis. All these fluxes vary seasonally but not synoptically or diurnally. We run our transport model (PCTM; Kawa et al., 2004) at a 1 h time step and a $4^\circ \times 5^\circ$ resolution, solving for fluxes at a daily (24 h) resolution across a 10 d span. We have used a linear version of PCTM in our simulations, with our adjoint constructed directly from this, rather than through the intermediary of a tangent linear model.

Both the measurement error covariance matrices \mathbf{R}_j and the a priori flux covariance matrices $\mathbf{P}_{\mathbf{u}_k}^o$ have been chosen to be consistent with the applied random measurement errors and the difference between the true and prior fluxes, respectively. In other words, the simulation experiments are not ‘mis-tuned’ due to inaccuracies in \mathbf{R}_j and $\mathbf{P}_{\mathbf{u}_k}^o$; the impact of such errors will be left to a subsequent paper. $\mathbf{P}_{\mathbf{u}_k}^o$ is calculated directly from the absolute value of the truth minus prior flux field differences.

3.1. Convergence and error analysis with highly-dense data coverage

To test the implementation and the convergence of the method, we examine a network with hourly measurements in the bottom model level of every $4^\circ \times 5^\circ$ grid box—a density of in situ observations unlikely to be achieved globally in the near future. Four cases are considered, found by adding (or not adding) a 1 ppm gaussian error to the measurements, and by applying (or not applying) an a priori constraint with a 1.0 PgC yr⁻¹ uncertainty on each daily gridded flux value. The measurement and prior errors in \mathbf{R} and $\mathbf{P}_{\mathbf{u}_k}^o$ are considered to be correlated in time with a 1 day time constant, but no spatial correlations are considered. We do not assess the impact of transport model error here; this important error source will be examined in a subsequent study.

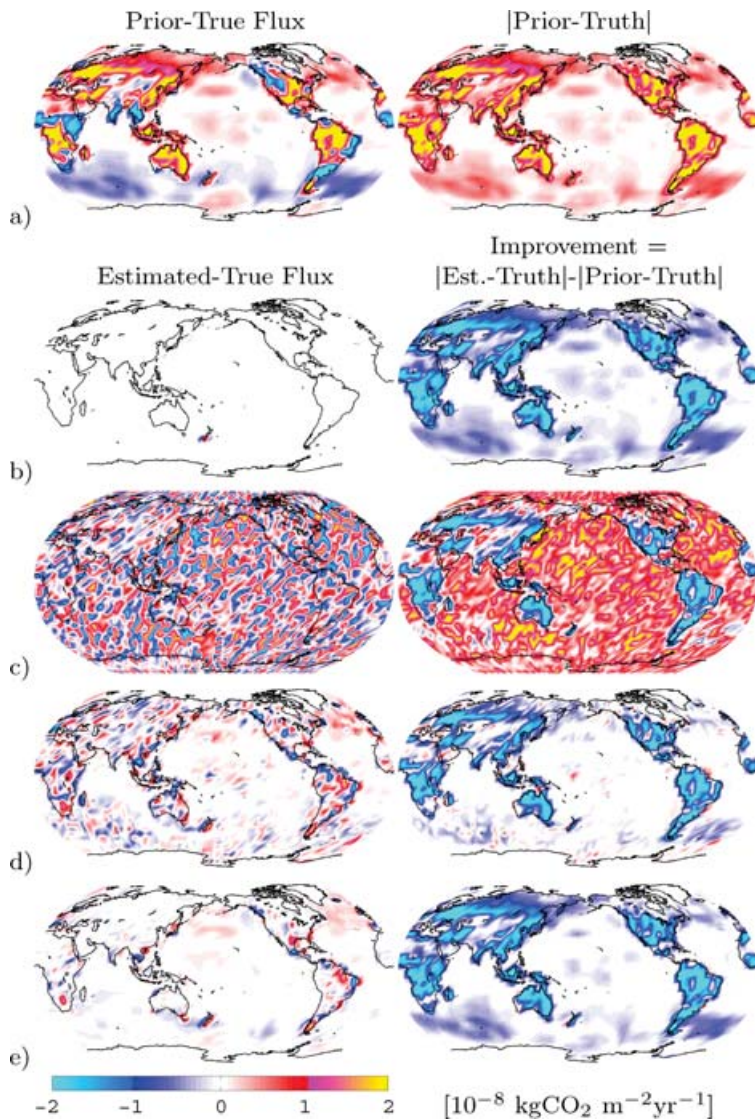


Fig. 1. (Estimate–Truth) flux differences and the ‘improvement’ in the estimated fluxes, defined by $|\text{Estimate} - \text{Truth}| - |\text{Prior} - \text{Truth}|$, for Day 1 of a 10 d assimilation using highly dense surface data (hourly at $4^\circ \times 5^\circ$) after 60 descent iterations [$\text{kgCO}_2 \text{ m}^{-2} \text{ yr}^{-1}$]. A posteriori flux estimates for four cases are considered: (b) no data errors added and no prior constraint applied; (c) data errors added but no prior constraint applied; (d) both data errors and an a priori constraint applied and (e) the a priori constraint applied but no data errors added. The pre-assimilation difference between the a priori and true fluxes is given in (a) for comparison, along with its absolute value.

Figure 1(a) gives an example of the initial Prior–Truth flux difference ($\mathbf{u}_k^o - \mathbf{u}_k^{\text{true}}$, at $k = 1$) that the assimilation seeks to correct; in a perfect assimilation, this difference would be driven to zero. First, to test convergence, no random data errors are added and no a priori constraint applied; after 60 descent steps, almost all the errors in the a priori fluxes are corrected (Fig. 1b). Figure 2 shows the convergence in the error; for this 10 d run, 60 iterations are sufficient to drive the convergence error to under 1% of the initial errors.

We then add 1 ppm random errors with the same time correlations assumed in R to each measurement, again without applying an a priori constraint to the fluxes. The resulting flux estimate (Fig. 1c) suffers from fine-scale estimation errors. We may more easily see whether the assimilation has improved the a priori estimate or not by examining maps of the flux statistic $|\text{Estimate} - \text{Truth}| - |\text{Prior} - \text{Truth}|$ (second column of Fig. 1). ‘Improvement’

(negative values) occurs when the assimilation drives the final estimate more closely towards the truth; positive values arise when the estimate is driven farther away (either in the wrong direction or by overshooting in the right direction). The assimilation improves the estimate where the initial errors are large (cf. Fig. 1a) but not much elsewhere (Fig. 1c).

For the case in Fig. 1c, the assimilation is trying to over-fit the data; in the absence of an a priori constraint on the fluxes, an unrealistic ‘noise’ is added to the fluxes in an attempt to optimally fit the ‘noise’ in the data. Figure 2 shows that the assimilation initially reduces the flux errors overall (as the worst errors over the continents are reduced), but that the overall error grows again after descent step 8 when the fine-scale flux errors overwhelm the large-scale improvements. The magnitude of this random error in the fluxes ($\sim 10^{-8} \text{ kgCO}_2 \text{ m}^{-2} \text{ s}^{-1}$) is a measure of the a posteriori error in the fluxes; if the RMS

flux error were to be computed across a statistically significant sample of inversions with independent sets of random errors, the result would be the same as that given by the full-rank covariance matrix given by more traditional inversion methods. While we do not compute this full error estimate here, it is clear from Fig. 1 that the measurements are only sufficient to improve the worst of the initial flux errors (those greater than about 10^{-8} kgCO₂ m⁻² s⁻¹, over the land regions) at least at this resolution.

When an a priori flux constraint is applied along with the data errors, the random flux errors due to the over-fitting are largely damped out, but at the cost of a systematic error in the direction of the a priori flux (Fig. 1d). This error can be seen explicitly in Fig. 1e for the case in which the a priori flux constraint is applied but no random flux errors added. The lower level of random error in Fig. 1d is a measure of the a posteriori information from both the measurements and the a priori flux constraint. Both this random error and the systematic error towards the prior must be considered in computing the total errors for the method.

An additional error is incurred if the descent towards the minimum is terminated too early (Fig. 2). Limited computational resources may prevent full convergence for larger problems. This convergence error must be considered along with the other random and systematic errors.




Fig. 2. Flux error ($|\text{Estimate}-\text{Truth}|$) as a function of descent iteration. Solid lines show the results for the dense surface measurement case (red: no data error added, no prior constraint; blue: data error added, no prior constraint; green: no data error added, with prior constraint; and black: data error added, with prior constraint). Cases in which both data errors and the prior constraint were applied are given for: the current in situ network (black dash-dotted line), the expanded version of the current in situ network (black dotted line), the OCO satellite case (black dashed line), as well as the dense ($4^\circ \times 5^\circ$) surface measurement case (solid black line).

3.2. Error analysis for five potential networks

To explore the impact of using sparser, more-realistic sampling, we evaluate five measurement networks with data errors added and an a priori constraint applied (the case in Fig. 1d). Network #1 is the current one from GLOBALVIEW-CO₂ (2004), with 35 continuous sites (with data applied at the 1 h time step here) and 101 weekly/biweekly flask sites (applied daily here, with an appropriate deweighting through the data errors). Network #2 is an expanded version of the current network composed of 186 continuous sites, mostly on tall towers and buoys (applied hourly). Network #3 consists of measurements similar to those from the planned sun-synchronous Orbital Carbon Observatory (OCO): column-averaged CO₂ concentrations (with 1 ppm uncertainty) from 14.6 day lit orbital passes per day, each separated by 24.7° in longitude (Crisp et al., 2004), with a 16 d repeat cycle for the ground track. Network #5 is the same one used in Section 3.1 and Fig. 1, with measurements in the surface layer of the model for each $4^\circ \times 5^\circ$ grid box, once per hour. Network #4 is similar to #5, but uses OCO column-averages instead of surface measurements. All measurements are assumed to be unbiased.

Figure 3 shows the difference between the true and estimated fluxes and the improvement over the prior for these five sampling strategies after 60 iterations of the BFGS method. At the resolution solved for here (daily fluxes at $4^\circ \times 5^\circ$), the current network (#1) improves the initial estimate in only a few locations (Fig. 3a). The expanded surface network (#2, Fig. 3b) and the OCO data (#3, Fig. 3c) give larger improvements, with those for OCO covering a larger area. However, the improvement for both networks is still only a small fraction of the initial error. Despite being much denser than the current network, both these networks are still too sparse to fully constrain fluxes at these fine scales. Moving to dense spatial coverage allows about half of the initial errors to be corrected (compare the dense-column and OCO networks, #4 and #3, in Figs 3d and c), while moving the measurements closer to the surface is required to obtain the rest of the correction (compare the dense-surface and dense-column networks, #5 and #4, in Figs. 3e and d).

4. Conclusions

We have tested a new method for estimating surface trace gas fluxes at fine time/space scales using dense atmospheric concentration data. Using a variational data assimilation system, we achieve near-exact convergence after 60 iterations of our descent method in our no-noise, perfect model case. In the presence of random data errors, an a priori constraint on the fluxes is effective in damping random errors in the estimate, at the cost of adding a systematic error in the direction of the prior. Our need of this constraint highlights the importance of accurate ‘bottom-up’ models of the basic processes for use as the prior estimate; dense data alone are not enough, given the likely future data errors. Despite the great efficiency of the method, computational constraints on

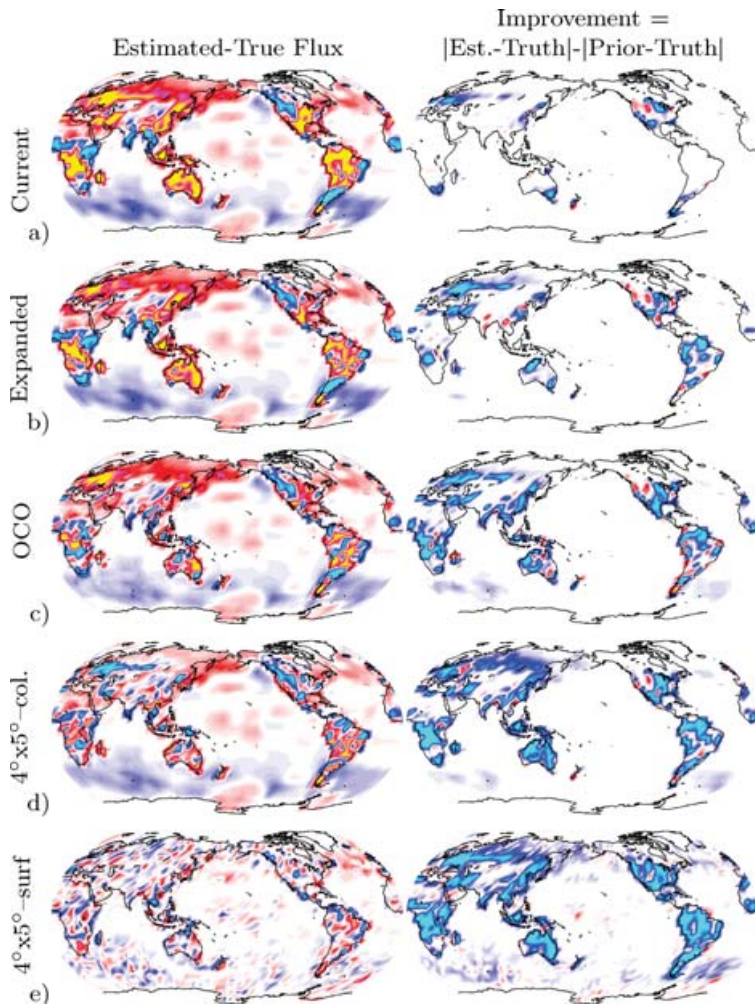


Fig. 3. (Estimate–Truth) flux differences and the ‘improvement’ in the estimated fluxes, defined by $|\text{Estimate} - \text{Truth}| - |\text{Prior} - \text{Truth}|$, for Day 1 of a 10 d assimilation after 60 descent iterations [$\text{kgCO}_2 \text{ m}^{-2} \text{ yr}^{-1}$]. Random data errors and an a priori constraint on the fluxes have been applied (as for the case in Fig. 1d). Results for five measurement networks are shown: (a) the current in situ network from GLOBALVIEW-CO (2004), including 35 continuous sites and 101 weekly sites; (b) an expanded version of the current network with 186 continuous sites; (c) column-averaged observations similar to those from the upcoming OCO satellite; (d) hourly column-averaged data at $4^\circ \times 5^\circ$ resolution across the globe and (e) hourly surface data at $4^\circ \times 5^\circ$ resolution (i.e. the case presented in Section 3.1 and Fig. 1d).

the degree of convergence will result in convergence errors that will become more significant at finer time/space resolutions.

The broad spatial coverage provided by column-averaged CO_2 measurements from sun-synchronous satellites should enable us to estimate surface gas fluxes at daily to synoptic time scales and model-grid spatial scales. This is a substantial improvement over current in situ measurements. Continuous near-surface in situ measurements may ultimately provide a more effective constraint on the surface fluxes, given their diurnal coverage and closer proximity to the fluxes. Achieving such spatial coverage is a challenge, requiring low-cost, calibrated analysers distributed globally across both land and ocean. Until near complete surface coverage is achieved, the broad column-averaged spatial coverage of the satellites and the spatially sparse, but continuous, near-surface coverage of the in situ data will be complementary, each filling in gaps in the other network.

5. Acknowledgments

This work was made possible through support from the Office of Global Programs (OGP) of the National Atmospheric and

Oceanographic Administration (NOAA) (Grant NA16GP2935 at NCAR, NA16GP2008 at WHOI). NCAR is sponsored by the National Science Foundation.

References

- Baker, D. F. 2001. Sources and Sinks of Atmospheric CO_2 Estimated from Batch Least Squares Inversions of CO_2 Concentration Measurements, PhD dissertation, 414 pp., Program in Atmospheric and Oceanic Sciences, Princeton University, Princeton, New Jersey, January.
- Baker, D. F., Law, R. M., Gurney, K. R., Rayner, P., Peylin, P. and co-authors. 2006. TransCom3 inversion intercomparison: Impact of transport model errors on the interannual variability of regional CO_2 fluxes, 1988–2003. *Global Biogeochem. Cycles*, **20**, GB1002, doi:10.1029/2004GB002439.
- Bennett, A. F. 2002. *Inverse Modeling of the Ocean and Atmosphere*, Cambridge University Press, Cambridge, 234 pp.
- Bousquet, P., Peylin, P., Ciais, P., Le Quééré, C., Friedlingstein, P. and co-authors. 2000. Regional changes in carbon dioxide fluxes of land and oceans since 1980. *Science*, **290**, 1342–1346.

- Braswell, B. H., Sacks, W. J., Linder, E. and Schimel, D. S. 2005. Estimating diurnal to annual ecosystem parameters by synthesis of a carbon flux model with eddy covariance net ecosystem exchange observations. *Global Change Biol.* **11**, 335–355, doi:10.1111/j.1365-2486.2005.00897.x.
- Bruhwyler, L. M. P., Michalak, A. M., Peters, W., Baker, D. F. and Tans, P. 2005. An improved Kalman Smoother for atmospheric inversions. *Atmos. Chem. Phys.* **5**, 2691–2702, SRef-ID: 1680-7324/acp/2005-5-2691.
- Bryson, A. E. and Ho, Y.-C. 1975. *Applied Optimal Control*. Hemisphere Publ. Co., New York, 481 pp.
- Buchwitz, M., de Beek, R., Burrows, J. P., Bovensmann, H., Warneke, T. and co-authors. 2005. Atmospheric methane and carbon dioxide from SCIAMACHY satellite data: initial comparison with chemistry and transport models. *Atmos. Chem. Phys.* **5**, 941–962.
- Chedin, A., Serrar, S., Scott, N. A., Crevoisier, C. and Armante, R. 2003. First global measurement of midtropospheric CO₂ from NOAA polar satellites: tropical zone. *J. Geophys. Res., Atmos.* **108**(D18), 4581 doi:10.1029/2003JD003439.
- Crevoisier, C., Heilliette, S., Chedin, A., Serrar, S., Armante, R. and Scott, N. A. 2004. Midtropospheric CO₂ concentration retrieval from AIRS observations in the tropics. *Geophys. Res. Lett.* **31**(17), L17106, doi:10.1029/2004GL020141.
- Crisp, D., Atlas, R. M., Breon, F.-M., Brown, L. R., Burrows, J. P. and co-authors. 2004. The Orbiting Carbon Observatory (OCO) Mission. *Adv. Space Res.* **34**, 700–709.
- Deeter, M. N., Emmons, L. K., Francis, G. L., Edwards, D. P., Gille, J. C. and co-authors. 2003. Operational carbon monoxide retrieval algorithm and selected results for the MOPITT instrument. *J. Geophys. Res.* **108**(D14), 4399, doi:10.1029/2002JD003186.
- Dilling, L., Doney, S. C., Edmonds, J., Gurney, K. R., Harriss, R. and co-authors. 2003. The role of carbon cycle observations and knowledge in carbon management. *Ann. Rev. Environ. Resour.* **28**, doi:10.1146/annurev.energy.28.011503.163443, 521–558.
- Doney, S. C., Lindsay, K., Caldeira, K., Campin, J.-M., Drange, H. and co-authors. 2004. Evaluating global ocean carbon models: the importance of realistic physics. *Global Biogeochem. Cycles.* **18**, GB3017, doi:10.1029/2003GB002150.
- Enting, I. 2002. *Inverse Problems in Atmospheric Constituent Transport*, Cambridge Univ. Press, Cambridge, UK, 392 pp.
- Errico, R. M. 1997. What is an adjoint model? *Bull. Am. Meteorol. Soc.* **78**, 2577–2591.
- Frankenberg, C., Meirink, J. F., van Weele, M., Platt, U. and Wagner, T. 2005. Assessing methane emissions from global space-borne observations. *Science.* **308**(5724), 1010–1014.
- Geels, C., Doney, S. C., Dargaville, R., Brandt, J. and Christensen, J. H. 2004. Investigating the sources of synoptic variability in atmospheric CO₂ measurements over the Northern Hemisphere continents: a regional model study. *Tellus.* **56B**, 35–50.
- GLOBALVIEW-CO₂ 2004. Cooperative Atmospheric Data Integration Project—Carbon Dioxide. CD-ROM, NOAA CMDL, Boulder, Colorado [Also available on Internet via anonymous FTP to ftp.cmdl.noaa.gov, Path: ccg/co2/GLOBALVIEW].
- Houweling, S., Breon, F.-M., Aben, I., Rödenbeck, C., Heimann, M. and co-authors. 2004. Inverse modelling of CO₂ sources and sinks using satellite data: a synthetic inter-comparison of measurement techniques and their performance as a function of space and time. *Atmos. Chem. Phys.* **4**, 523–538.
- Kaminski, T., Heimann, M. and Giering, R. 1999a. A coarse grid three-dimensional global inverse model of the atmospheric transport, 1: adjoint model and Jacobian matrix. *J. Geophys. Res.* **104**, 18 535–18 553.
- Kaminski, T., Heimann, M. and Giering, R. 1999b. 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.* **104**, 18 555–18 581.
- Kawa, S. R., Erickson, D. J., Pawson, S. and Zhu, Z. 2004. Global CO₂ transport simulations using meteorological data from the NASA data assimilation system. *J. Geophys. Res.-Atmos.* **109**(D18), D18312, doi:10.1029/2004JD004554.
- Peters W., Miller, J. B., Whitaker, J., Denning, A. S., Hirsch, A. and co-authors. 2005. An ensemble data assimilation system to estimate CO₂ surface fluxes from atmospheric trace gas observations. *J. Geophys. Res.* **110** (D24), Art. No. D24304.
- Peylin, P., Rayner, P. J., Bousquet, P., Carouge, C., Hourdin, F. and co-authors. 2005. Daily CO₂ flux estimates over Europe from continuous atmospheric measurements. 1: inverse methodology. *Atmos. Chem. Phys. Discuss.* **5**, 1647–1678, SRef-ID: 1680-7375/acpd/2005-5-1647.
- Randerson, J. T., Thompson, M. V., Conway, T. J., Fung, I. Y. and Field, C. B. 1997. The contribution of terrestrial sources and sinks to trends in the seasonal cycle of atmospheric carbon dioxide. *Global Biogeochem. Cycles.* **11**, 535–560.
- Rayner, P. J. and O'Brien, D. 2001. The utility of remotely sensed CO₂ concentration data in surface source inversions. *Geophys. Res. Lett.* **28**, 175–178.
- Rayner, P. J., Scholze, M., Knorr, W., Kaminski, T., Giering, R. and co-authors. 2005. Two decades of terrestrial carbon fluxes from a carbon cycle data assimilation system (CCDAS). *Global Biogeochem. Cycles.* **19**, GB2026, doi:10.1029/2004GB002254.
- Rödenbeck, C., Houweling, S., Gloor, M. and Heimann, M. 2003. CO₂ flux history 1982–2001 inferred from atmospheric data using a global inversion of atmospheric transport. *Atmos. Chem. Phys.* **3**, 1919–1964.
- Rödenbeck, C. 2005. Estimating CO₂ sources and sinks from atmospheric mixing ratio measurements using a global inversion of atmospheric transport. *Max-Planck-Institut für Biogeochemie*: Technical Paper 6.
- Sitch, S., Smith, B., Prentice, I. C., Arneth, A., Bondeau, A. and co-authors 2003. Evaluation of ecosystem dynamics, plant geography and terrestrial carbon cycling in the LPJ dynamic global vegetation model. *Global Change Biol.* **9**(2), 161–185.
- Stoer, J. and Bulirsch, R. 1980. *Introduction to Numerical Analysis*, Springer-Verlag, New York, 609 pp.
- Takahashi, T., Wanninkhof, R. H., Feely, R. A., Weiss, R. F., Chipman, D. W. and co-authors. 1999. Net sea-air CO₂ flux over the global oceans: an improved estimate based on the sea-air pCO₂ difference, Proceedings of the 2nd International Symposium: CO₂ in the Oceans, the 12th Global Environmental Tsukuba, 18–22 January 1999, Tsukuba Center of Institutes, (ed. Y. Nojiri). Natl. Inst. for Environ. Studies, Environ. Agency of Japan.