Geoscientific Model Development Discussions

On the parallelization of atmospheric inversions of CO$_2$ surface fluxes within a variational framework

F. Chevallier

Laboratoire des Sciences du Climat et de l'Environnement, CEA-CNRS-UVSQ, L'Orme des Merisiers, UMR8212, Bat 701, 91191 Gif-sur-Yvette, France

Correspondence to: F. Chevallier (frederic.chevallier@lsce.ipsl.fr)

GMDD

6, 37–57, 2013

Parallel and variational CO$_2$ atmospheric inversion

F. Chevallier

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Abstract

The variational formulation of Bayes' theorem allows inferring CO$_2$ sources and sinks from atmospheric concentrations at much higher space-time resolution than the ensemble approach or the analytical one. However, it usually exhibits limited scalable parallelism. This limitation hinders global atmospheric inversions operated on decadal time scales and regional ones with kilometric spatial scales, because of the computational cost of the underlying transport model that has to be run at each iteration of the variational minimization. Here, we introduce a Physical Parallelisation (PP) of variational atmospheric inversions. In the PP, the inversion still manages a single physically and statistically consistent window, but the transport model is run in parallel overlapping sub-segments in order to massively reduce the computation wall clock time of the inversion. For global inversions, a simplification of transport modelling is described to connect the output of all segments. We demonstrate the performance of the approach on a global inversion for CO$_2$ with a 32-yr inversion window (1979–2010) with atmospheric measurements from 81 sites of the NOAA global cooperative air sampling network. In this case, we show that the duration of the inversion is reduced by a sevenfold factor (from months to days) while still processing the three decades consistently and with improved numerical stability.

1 Introduction

CO$_2$ mole fractions in the atmosphere are functions of the CO$_2$ surface fluxes, of the atmospheric transport fluxes and, to a smaller extent, of the photochemical CO$_2$ production in the atmosphere. Inferring (or inverting) the surface fluxes from mole fraction measurements is made challenging by the combination of these three inputs, but it is rigorously guided by the Bayesian paradigm. It requires the use of a model of the atmospheric transport. An additional complication lies in the diversity of the time-space scales that are involved. At one end, the exceptionally-long mean life time of CO$_2$ links

any mole fraction measurement to global CO_2 surface fluxes in the distant past. At the other end, point-wise surface measurements exhibit a strong sensitivity to the fluxes in the immediate vicinity of the detectors (Bocquet, 2005). This singularity is not much compensated (i.e. regularised) by the size of the prior errors, at least in the case of
5 terrestrial vegetation fluxes (Chevallier et al., 2006). The mixture of transport scales hinders the implementation of Bayes' theorem at reasonable computational cost. In practice, compromises are made between the technical complexity of the inversion code, the level of error induced by neglecting some of the flux variability (e.g. inferring continental monthly fluxes rather than high-resolution ones; Kaminski et al., 2001;
10 Thompson et al., 2011) and the length of the inversion window. Flux inversion has been traditionally performed based on a closed-form solution to Bayes' theorem (i.e. an algebraic formula; e.g. Gurney et al., 2002) but the computational cost of this analytical solution restricts it to small-size problems (typically less than 10^4 variables to infer or less than 10^4 measurements to exploit). In the context of high-performance parallel
15 computing, ensemble formulations of Bayes' theorem (i.e. using a statistical ensemble of simulations) have been increasingly popular (e.g. Peters et al., 2005; Zupanski et al., 2007; Feng et al., 2009) but the relatively small sizes that can be currently afforded for the statistical ensembles do not allow tackling problems of higher resolutions than the analytical approach. Imported from the data assimilation systems of Numeri-
20 cal Weather Prediction (NWP), the variational formulation of Bayes' theorem has been introduced in the mid-00s to lift this restriction (Chevallier et al., 2005; Rödenbeck et al., 2005), but it requires tedious initial developments (i.e. coding the adjoint of the atmospheric transport) and usually exhibits limited scalable parallelism (Hamrud, 2010).

This paper focusses on the technical performance of the variational approach.
25 Chevallier et al. (2010) performed a global 21-yr inversion at grid point (3.75° longitude and 2.5° latitude) and weekly scale with this approach, but they needed six weeks of sustained parallel computing, the parallelization being for the code of the transport model. The high cost of this 21-yr inversion implies prohibitive costs for extended inversion periods, higher resolution transport modelling or more sophisticated physical

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

⏮ | ⏭

◀ | ▶

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

parameterization of the transport model. Given the parallel structure of current super-computers, finding high levels of parallelism has become critical for the developers of the variational approach (see, e.g. Fisher et al., 2011; and Desroziers and Berre, 2012, in the field of NWP). Now, this approach is intrinsically sequential because it consists in iteratively minimizing a cost function, each iteration step relying on the results of the previous one. Higher resolutions of the transport models allow parallelizing on more processors, but the gain offered by the latter never fully compensates the cost of the former because of communication overheads. Moreover, transport models perform a lot of input operations (e.g. to read the transport mass fluxes), that are not well parallelizable. This problem is even bigger for their adjoint codes that read most of the information about the linearization point from the disk.

In this paper, we introduce a parallel structure of the transport computations within variational atmospheric inversion systems in order to massively reduce the computation wall clock time, while still running the inversion on a single physically and statistically consistent window. This structure defines an ensemble of parallel computations that does not carry any statistical meaning, in contrast to the above-mentioned ensemble methods. We call it Physical parallelization (PP) here to distinguish it from the statistical methods. We illustrate the efficiency of the PP for a 32-yr inversion window (1979–2010) for CO_2 with atmospheric measurements from 81 sites of the NOAA global cooperative air sampling network. Section 2 introduces the PP. The inversion system and the measurements used to illustrate it are described in Sects. 3 and 4, respectively. The results are presented in Sect. 5. Section 6 concludes the paper.

2   Physical parallelization

An inversion system operates within a temporal window starting at a chosen time T (for instance, 1 January 1979 at 00:00 UTC in Sect. 4). Within this window, the transport model H simulates the mole fraction measurements made at any time t from the surface fluxes existing between times T and t and from the 3-D field of mole fraction at

the initial time T. In the case of regional inversions, H also uses the lateral boundary conditions of the mole fractions between times T and t.

Variational inversion systems minimize a Bayesian cost function and, to this end, exploit a linearized version of H, under the form of a tangent-linear code and of an adjoint code.

In the tangent-linear code of H, a perturbation in mole fraction caused by perturbations in surface fluxes, by perturbations in the lateral boundary conditions and by perturbations in initial mole fractions, is computed as:

$$\delta\boldsymbol{c}(t) = \sum_{t'=T}^{t} \mathbf{H}_{t't}^{\boldsymbol{\varphi}} \cdot \delta\boldsymbol{\varphi}(t') + \mathbf{H}_{t't}^{c} \cdot \delta\boldsymbol{c}(T) \tag{1}$$

with $\delta\boldsymbol{c}(t)$ the vector that contains the mole fraction increments at time t, $\delta\boldsymbol{\varphi}(t')$ the vector that contains the flux increments and the lateral boundary condition increments at time t', $\mathbf{H}_{t't}^{\boldsymbol{\varphi}}$ the linear operator (Jacobian matrix) that links $\delta\boldsymbol{\varphi}(t')$ and $\delta\boldsymbol{c}(t)$, and $\mathbf{H}_{t't}^{c}$ the linear operator that links $\delta\boldsymbol{c}(T)$ and $\delta\boldsymbol{c}(t)$. Note that in variational systems, the two Jacobian matrices are not explicitly computed: the multiplication by the input increment vector is obtained from the chain rule applied to the reference code, line by line.

The PP brings a small simplification to Eq. (1) by reducing the backward influence to a time τ which is later than T following:

$$\delta\boldsymbol{c}(t) = \sum_{t'=\tau}^{t} \mathbf{H}_{t't}^{\boldsymbol{\varphi}} \cdot \delta\boldsymbol{\varphi}(t') + \delta b(\tau) \tag{2}$$

with $\delta b(\tau)$ the global-mean mole fraction increment at time τ. The scalar $\delta b(\tau)$ corresponds to the perturbation of the global growth rate. In the case of regional inversions, $\delta b(\tau)$ is set to zero because the global growth rate is already accounted for by the lateral boundary conditions.

The advantage of Eq. (2) is that the transport model does not link successive temporal segments together any more, when computing mole fraction increments. The sum $\sum_{t'=\tau}^{t} \mathbf{H}_{t't}^{\boldsymbol{\varphi}} \cdot \delta\boldsymbol{\varphi}(t')$ can be computed in parallel for various segments of a long period: the bias $\delta b(\tau)$ is added a posteriori. For global inversions, the downside is that Eq. (2)

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

|◄ | ►|

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

introduces a bottleneck at time τ when all mole fraction increments are supposed to be uniformly mixed in the global atmosphere. For regional inversions, setting $\delta b(\tau)$ to zero implies neglecting the direct influence of the 3-D field of the compound within the regional domain at time τ. To minimize both effects, we overlap the various segments

5 so that they include a spin-up period without any observation. For global inversions, the spin-up time allows atmospheric mixing to operate. For regional ones, it allows the initial molecules to be advected outside the domain. The length of the spin-up period should therefore be chosen according to the corresponding time scales (the time scale of inter-hemispheric mixing for global inversions and the time scale of advection within

10 the domain for the regional ones). Once the spin-up time has been chosen, the PP forward algorithm can be described as follows:

1. Run the transport model over the full period without any simplification (i.e. serially). This run provides the linearization point for subsequent increment computations and the innovation vector (i.e. the vector of initial model-minus-observation
15 departures) at the best precision available.

2. Divide the full period into a series of overlapping segments.

3. Run the tangent-linear model for each segment in parallel without any perturbations to the mole fraction field at the segment initial time. Save the mole fraction increment $\sum_{t'=\tau}^{t} \mathbf{H}_{t't}^{\varphi} \cdot \delta\boldsymbol{\varphi}(t')$ at the time and location of each observation, except
20 for the spin up period when observations are ignored (since they are processed by the previous segment). The earliest segment (i.e. starting at T) is a particular case where Eq. (1) is applied directly.

4. In the case of global inversions, for each segment, compute the contribution to the global mass increment $\delta b(\tau)$ that corresponds to the initial state of the next
25 segment and add it to the mole fraction increments (Eq. 2).

The PP algorithm for the adjoint code simply stems from that of the tangent-linear code, reversing the order of the four operations and transposing them.

42

The wall clock time needed for this algorithm is about that of a segment, with a par-allelization structure on the number of segments.

At the highest level of the inversion system, the minimizer deals with the full inversion period at once with a single cost function (including correlated prior errors between segments if such correlations are assigned) and without any parallelization, therefore ensuring a consistent inversion from time T until the time of the latest observation available.

3 Inversion method

To illustrate the PP, we use the inversion scheme of Chevallier et al. (2005, 2010). This system mainly gathers a tracer transport model, its tangent-linear and adjoint codes, and a software to minimize the Bayesian cost function. It computes the Best Linear Unbiased Estimate (BLUE) of the CO_2 surface fluxes given the observations, the prior fluxes and their associated uncertainty. The BLUE fluxes are called inverted fluxes hereafter. Their uncertainty can also be estimated by the inversion system by a robust Monte Carlo method, but this facility is not exploited here.

In the configuration used in the result section, the transport model is the global general circulation model of the Laboratoire de Météorologie Dynamique (Hourdin et al., 2006) nudged to winds analysed by ECMWF, and the minimizer is the Lanc-zos version of the conjugate gradient algorithm, as described by Fisher (1998). The inverted fluxes are estimated on the global grid of the transport model (a reg-ular mesh of $3.75° \times 2.5°$) at a temporal resolution of 8 days, with day-time and night-time separated. The prior fluxes combine estimates of (i) annual anthropogenic emissions (EC-JRC/PBL, EDGAR version 4.2. http://edgar.jrc.ec.europa.eu/, 2011, http://cdiac.ornl.gov/ftp/ndp030/global.1751_2008.ems, accessed 6 July 2011, and http://cdiac.ornl.gov/trends/emis/meth_reg.html, accessed 8 January 2013), climato-logical monthly ocean fluxes (Takahashi et al., 2009), climatological monthly biomass burning emissions (taken as the 1997–2010 mean of the database of van der Werf

et al., 2010) and climatological 3-hourly biosphere-atmosphere fluxes taken as the 1989–2010 mean of a simulation of the ORganizing Carbon and Hydrology In Dynamic EcosystEms model (ORCHIDEE, Krinner et al., 2005), version 1.9.5.2. The mass of carbon emitted annually during specific fire events is compensated here by the same

5 annual flux of opposite sign representing the re-growth of burnt vegetation, which is distributed regularly throughout the year. These gridded prior fluxes exhibits 3-hourly variations but their inter-annual variations are caused by anthropogenic emissions only.

The uncertainty of the prior fluxes is described by a covariance matrix. Over land, we assume that the errors of the prior biosphere-atmosphere fluxes dominate the error

10 budget and the covariances are constrained by an analysis of in situ flux measurements (Chevallier et al., 2012): temporal correlations on daily mean Net Carbon Exchange (NEE) errors decay exponentially with a length of one month but night-time errors are assumed to be uncorrelated with daytime errors; spatial correlations decay exponentially with a length of 500 km; standard deviations are set to the climatological

15 daily-varying heterotrophic respiration flux simulated by ORCHIDEE with a ceiling of $4 \, \mathrm{g \, C \, m^{-2} \, day^{-1}}$. Over a full year, the total 1-sigma uncertainty for the prior land fluxes amounts to about $3.0 \, \mathrm{Gt \, C \, yr^{-1}}$. The error statistics for the open ocean correspond to a global air-sea flux uncertainty about $0.75 \, \mathrm{Gt \, C \, yr^{-1}}$ and are defined as follows: temporal correlations decay exponentially with a length of one month; unlike land, daytime

20 and night-time flux errors are fully correlated; spatial correlations follow an e-folding length of 1000 km; standard deviations are set to $0.15 \, \mathrm{g \, C \, m^{-2} \, day^{-1}}$. Land and ocean flux errors are not correlated.

Prior initial conditions (i.e. the 3-D field of CO_2 at the start of the inversion window) are taken from a previous simulation of the LMDZ model and are adjusted by the inver-

25 sion (with 0.3 % prior uncertainty).

4 Observations

In the present illustration, the assimilated observations are mole fractions of CO_2 in dry air, collected in flask air samples at various places in the world over land (from fixed sites) and over ocean (from commercial ships). We use 81 site records from the NOAA Earth System Research Laboratory archive (Conway et al., 2011) for the 1979–2010 period. The site location and the length of each record are shown in Fig. 1. The observed synoptic variability at each site is taken as a proxy of the observation uncertainty (which is driven by transport modelling errors) and is computed following Chevallier et al. (2010).

5 Results

In the present configuration of the PP inversion, we define computation segments of 15 months (from October of a given year until December of the following year) with a 3 month overlap from one segment to the next (from October until December). The 3-month overlap allows a large portion of the global mixing to operate (Bruhwiler et al., 2005; Peters et al., 2005) at least at the hemispheric scale. The first segment is a particular case: it gathers 12 months only and Eq. (1) is directly applied (i.e. without any simplification).

We show the impact of the simplification of the tangent-linear model by running a 32-yr inversion (1979–2010) both with the full system and with the simplified one. Each minimization is stopped after 30 iterations, as in Chevallier et al. (2010), when increments become stable. The full system required two months of sustained computing while the PP one took days only. The two inversions are compared by looking at the inverted fluxes at various time-space scales: from a 3-h time step to the year, from the model grid to the widely-used 22 TransCom3 regions of Gurney et al. (2002). The annual total fluxes aggregated at the scale of TransCom3 region Tropical Asia are shown in Fig. 2. The prior interannual variations (in black) are caused by fossil fuel emission

Title Page

Abstract	Introduction
Conclusions	References
Tables	Figures

◄ | ►|

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

increase only (see Sect. 3). The fit between the PP inversion (in red) and the one with the full system (in cyan) is surprisingly poor: the shape of the variations is similar but the absolute values differ by up to 600 MtC yr^{-1} (in 1990). Looking at the annual global scale (not shown), the PP inversion actually agrees with the mean CO_2 growth rate inferred by NOAA (ftp://ftp.cmdl.noaa.gov/ccg/co2/trends/co2_gr_gl.txt, accessed 6 September 2012) better than the reference inversion: the root mean square difference (RMS) of the misfit over the 32 yr is 0.19 ppm yr^{-1} for the former while it reaches 0.88 ppm yr^{-1} for the latter. To understand the better performance of the inversion that uses the simplified system, the reference inversion is repeated for a shorter period: 1979–1992. In this case, the RMS misfit of the reference inversion with the NOAA data goes down to 0.21 ppm yr^{-1}, while the PP misfit is slightly larger (0.22 ppm yr^{-1}) for the same period. This second reference inversion is also shown in Fig. 2, in standard blue: the curve is well fitted by the PP points, with differences less than 50 MtC yr^{-1}, even for the year 1990 when the longer reference inversion shows a large spike. From this test and from similar ones varying the length of the inversion window (not shown), we conclude that numerical instabilities spoil the reference system for very large inversion windows. The better behaviour of the PP inversion suggests that the small transport sensitivities for very long ranges cause these instabilities.

Figure 3 summarizes the differences between the PP inversion and the reference system (run with two inversion windows) at the scale of the year and of the 22 TransCom3 land regions. The RMS is very large (up to 570 MtC yr^{-1} in region South American Tropical) when comparing with the reference system run on the longer window, while it is less than 80 MtC yr^{-1} when comparing with the reference system run on the shorter window. The misfit with the shorter window system represents about 10 % of the inversion increments.

Figure 4 refines the comparison with the shorter-window reference system, by showing the map of the RMS at the weekly grid point scale. This time scale is chosen because it is the highest of the inversion system. The RMS is less than 0.08 g C m^{-2} day^{-1}

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

everywhere. The RMS represents less than 11 % of the flux inversion increments on any grid point, excluding those where the increments are negligible.

We have repeated the PP inversion with longer segments. In this test, the segments cover 18 months, from July of a given year until December of the following year, with a 6 month overlap from one segment to the next, from July until December. In this case, the 1979–1991 fit to the shorter-window reference system is improved by a few % points, with the largest relative improvement seen in the ocean basins of the Southern Hemisphere. In terms of annual budgets in the TransCom3 regions, the improvement is less than $20\,\mathrm{MtC\,yr^{-1}}$, with most values being a few $\mathrm{MtC\,yr^{-1}}$.

6 Discussion and conclusions

More than ten CO_2 monitoring sites have records that span more than 30 yr and tens of them span more than 20 yr. These archives offer the exceptional opportunity to yield consistent estimations of CO_2 surface fluxes over the last decades. However, the inversion problem involves relatively small spatial scales (in the transport properties and in the prior errors) that make such estimations technically difficult. Coarse resolution systems are usually chosen (e.g. Gurney and Eckels, 2011), which degrades the quality of the inversion (Kaminski et al., 2001). Regional inverse modelling for dense regional networks (e.g. Lauvaux et al., 2009; Broquet et al., 2011) faces a similar issue when it accounts for flux and transport variability on scales of the order of 10 km or less. In this paper we propose to parallelize some of the transport computations along overlapping temporal segments within variational inversion systems. This approach allows efficient "embarrassingly parallel" workload, i.e. tasks that do not need communicating, within each iteration step. In the case of global inversions, the transport information from one segment to the next one is carried out through a global bias term, while it is not needed for regional inversions. The initial simulation that provides the linearization point before the first iteration and the final simulation after the minimization are not parallelized.

CC BY

At the highest level, the inversion manages all segments at once with a single cost function.

For global inversions, the increment bottleneck that we introduce in each segment in the form of this bias term (on 1 October of each year in our illustration) reduces the accuracy of the tangent linear model with an effect that depends on the assimilated observations: a dense and precise observation network counterbalances the trend towards homogeneously-distributed inversion increments. In this context, the test with a single network (the NOAA global cooperative air sampling network), while tens of other sites are available (http://ds.data.jma.go.jp/gmd/wdcgg/), provides an upper bound on the method precision. The effect of the increment bottleneck is also mitigated by the overlap period from one segment to the next that plays the role of a spin-up time. In this study, we have tested our approach with a 3-month overlap. For a 14-yr inversion window, we have shown that the simplification degrades the inversion increments at various scales (weekly grid-scale and sub-yearly continental) by less than 20 %. For a 32-yr inversion window, the simplification makes the inversion numerically stable, which is not the case for the full system that suffers from round-off errors. Last, the 32-yr inversion does not take more time to run than for a 15-month inversion, i.e. days rather than months, while still processing the three decades consistently.

Qualitatively, the effect of the transport truncation is balanced by the limited robustness of transport models for long ranges. At the global scale, this limitation was clearly seen in the validation of the interhemispheric exchange times simulated by a pool of models, even though the precision was shown to have improved compared to older models (Patra et al., 2011). We argue that performing long inversions is important despite transport weaknesses because long inversions ensure both the statistical consistency and the physical consistency of the inferred long-term CO_2 budgets. Indeed, a sequential filter (e.g. Peters et al., 2007) can either conserve mass throughout the inversion cycles or account for the uncertainty of the past cycles, but cannot do both. Furthermore, long inversions benefit from the long temporal structures of the prior errors (Chevallier et al., 2006), thereby propagating observation information well beyond

measurement time in a way that is complementary to transport. This second property is particularly interesting for regional inversions when high-resolution flux estimates are sought while the network is sparse at the scale of interest.

For regional (i.e. domain-limited) inversions, the information about the global growth rate is carried by the atmospheric lateral boundary conditions around the inversion domain. In this case, the bias term disappears from Eq. (2): for a given segment, our method neglects the direct influence of the initial conditions of this segment. The spin-up period (i.e. the between-segment overlap period) damps this effect that becomes negligible when the spin-up period is long enough. Adequate spin-up period mainly depends on the domain latitude range and on the speed of the zonal flow in the domain.

To further improve the inversion speed, additional levels of parallelism can be found in two ways. First, parallel computing can be introduced within each segment (for instance with an OpenMP structure), in addition to between-segment parallelism (for instance with an MPI driver). This hybrid parallel strategy would fit high performance computing systems that are usually made of clusters of shared memory nodes. Second, advantage can be taken from ensembles of perturbed inversions because they accelerate convergence when performed together (Desroziers and Berre, 2012) while yielding rigorous inversion uncertainty computation (Chevallier et al., 2007). The parallelization of such ensembles is straight-forward because the ensemble members do not communicate.

Among Bayesian systems, the variational ones are the fittest to deal with large inversion problems. The opportunities for parallelism that we see in the case of global and regional CO_2 atmospheric inversions reinforce their attractiveness in the context of high performance cluster computing. Our approach also applies when the atmospheric inversion infers parameters of flux process models, like in variational Carbon Cycle Data Assimilation Systems (Rayner et al., 2005), rather than just the fluxes. We have focused here on CO_2 because our method exploits its passive transport in the global atmosphere. It can also be applied for reactive species like methane, nitrogen and chlorofluorocarbons within regional inversions. For global inversions of these species, the

distant past cannot be summarized by a constant bias: in this case, Eq. (2) has to be adapted by introducing a model of the temporal evolution of the global offset.

References

Bocquet, M.: Grid resolution dependence in the reconstruction of an atmospheric tracer source, Nonlinear Proc. Geoph., 12, 219–233, 2005.

Chevallier, F., Fisher, M., Peylin, P., Serrar, S., Bousquet, P., Bréon, F.-M., Chédin, A., and Ciais, P.: Inferring CO_2 sources and sinks from satellite observations: method and application to TOVS data, J. Geophys. Res., 110, D24309, doi:10.1029/2005JD006390, 2005.

Chevallier, F., Viovy, N., Reichstein, M., and Ciais, P.: On the assignment of prior errors in Bayesian inversions of CO_2 surface fluxes, Geophys. Res. Lett., 33, L13802, doi:10.1029/2006GL026496, 2006.

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Chevallier, F., Bréon, F.-M., and Rayner, P. J.: The contribution of the Orbiting Carbon Observatory to the estimation of CO_2 sources and sinks: theoretical study in a variational data assimilation framework, J. Geophys. Res., 112, D09307, doi:10.1029/2006JD007375, 2007.

Chevallier, F., Wang, T., Ciais, P., Maignan, F., Bocquet, M., Arain, A., Cescatti, A., Chen, J.-Q., Dolman, H., Law, B. E., Margolis, H. A., Montagni, L., and Moors, E. J.: What eddy-covariance flux measurements tell us about prior errors in CO_2-flux inversion schemes, Global Biogeochem. Cy., 26, GB1021, doi:10.1029/2010GB003974, 2012.

Ciais, P., Conway, T. J., Aalto, T., Anderson, B. E., Bousquet, P., Brunke, E. G., Ciattaglia, L., Esaki, Y., Fröhlich, M., Gomez, A. J., Gomez-Pelaez, A. J., Haszpra, L., Krummel, P., Langenfelds, R., Leuenberger, M., Machida, T., Maignan, F., Matsueda, H., Morguí, J. A., Mukai, H., Nakazawa, T., Peylin, P., Ramonet, M., Rivier, L., Sawa, Y., Schmidt, M., Steele, P., Vay, S. A., Vermeulen, A. T., Wofsy, S., and Worthy, D.: CO_2 surface fluxes at grid point scale estimated from a global 21-year reanalysis of atmospheric measurements, J. Geophys. Res., 115, D21307, doi:10.1029/2010JD013887, 2010.

Conway, T. J., Lang, P. M., and Masarie, K. A.: Atmospheric Carbon Dioxide Dry Air Mole Fractions from the NOAA ESRL Carbon Cycle Cooperative Global Air Sampling Network, 1968–2010, Version: 2011-06-21, available at: ftp://ftp.cmdl.noaa.gov/ccg/co2/flask/event/ (last access: 5 November 2011), 2011.

Desroziers, G. and Berre, L.: Accelerating and parallelizing minimizations in ensemble and deterministic variational assimilations, Q. J. Roy. Meteor. Soc., 138, 1599–1610, doi:10.1002/qj.1886, 2012.

Feng, L., Palmer, P. I., Bösch, H., and Dance, S.: Estimating surface CO_2 fluxes from spaceborne CO_2 dry air mole fraction observations using an ensemble Kalman Filter, Atmos. Chem. Phys., 9, 2619–2633, doi:10.5194/acp-9-2619-2009, 2009.

Fisher, M.: Minimization algorithms for variational data assimilation, in: Seminar on Recent Developments in Numerical Methods for Atmospheric Modelling, ECMWF, Reading, 7–11 September 1998, available at: http://www.ecmwf.int/publications/library/ecpublications/_pdf/seminar/1998/seminar1998_fisher.pdf (last access: 8 January 2013), 1998.

Fisher, M., Tremolet, Y., Auvinen, H., Tan, D., and Poli, P.: Weak-Constraint and Long Window 4DVAR, ECMWF Technical Memorandum No. 655, available at: http://www.ecmwf.int/publications/library/ecpublications/_pdf/tm/601-700/tm655.pdf (last access: 8 January 2013), 2011.

Gurney, K. R. and Eckels, W. J.: Regional trends in terrestrial carbon exchange and their seasonal signatures, Tellus B, 63, 328–339, doi:10.1111/j.1600-0889.2011.00534.x, 2011.

Gurney, K. R., Law, R. M., Denning, A. S., Rayner, P. J., Baker, D., Bousquet, P., Bruhwiler, L., Chen, Y. H., Ciais, P., Fan, S., Fung, I. Y., Gloor, M., Heimann, M., Higuchi, K., John, J., Maki, T., Maksyutov, S., Masarie, K., Peylin, P., Prather, M., Pak, B. C., Randerson, J., Sarmiento, J., Taguchi, S., Takahashi, T., and Yuen, C. W.: Towards robust regional estimates of CO_2 sources and sinks using atmospheric transport models, Nature, 415, 626–630, 2002.

Hamrud, M.: Report from IFS Scalability Project, ECMWF Technical Memorandum No. 616, available at: http://www.ecmwf.int/publications/library/ecpublications/_pdf/tm/601-700/tm616.pdf (last access: 8 January 2013), 2011.

Hourdin, F., Musat, I., Bony, S., Braconnot, P., Codron, F., Dufresne, J.-L., Fairhead, L., Filiberti, M.-A., Friedlingstein, P., Grandpeix, J.-Y., Krinner, G., and Levan, P.: The LMDZ4 general circulation model: climate performance and sensitivity to parametrized physics with emphasis on tropical convection, Clim. Dynam., 27, 787–813, doi:10.1007/s00382-006-0158-0, 2006.

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

Krinner, G., Viovy, N., de Noblet-Ducoudré, N., Ogée, J., Polcher, J., Friedlingstein, P., Ciais, P., Sitch, S., and Prentice, I. C.: A dynamic global vegetation model for studies of the coupled atmosphere-biosphere system, Global Biogeochem. Cy., 19, GB1015, doi:10.1029/2003GB002199, 2005.

Patra, P. K., Houweling, S., Krol, M., Bousquet, P., Belikov, D., Bergmann, D., Bian, H., Cameron-Smith, P., Chipperfield, M. P., Corbin, K., Fortems-Cheiney, A., Fraser, A., Gloor, E., Hess, P., Ito, A., Kawa, S. R., Law, R. M., Loh, Z., Maksyutov, S., Meng, L., Palmer, P. I., Prinn, R. G., Rigby, M., Saito, R., and Wilson, C.: TransCom model simulations of CH_4 and related species: linking transport, surface flux and chemical loss with CH_4 variability in the troposphere and lower stratosphere, Atmos. Chem. Phys., 11, 12813–12837, doi:10.5194/acp-11-12813-2011, 2011.

Peters, W., Jacobson, A. R., Sweeney, C., Andrews, A. E., Conway, T. J., Masarie, K., Miller, J. B., Bruhwiler, L. M. P., Petron, G., Hirsch, A. I., Worthy, D. E. J., van der Werf, G. R., Randerson, J. T., Wennberg, P. O., Krol, M. C., and Tans, P. P.: An atmospheric perspective on North American carbon dioxide exchange: carbon tracker, P. Natl. Acad. Sci. USA, 48, 18925–18930, 2007.

Rayner, P. J., Scholze, M., Knorr, W., Kaminski, T., Giering, R., and Widmann, H.: Two decades of terrestrial carbon fluxes from a carbon cycle data assimilation system (CCDAS), Global Biogeochem. Cy., 19, GB2026, doi:10.1029/2004GB002254, 2005.

Rödenbeck, C.: Estimating CO_2 sources and sinks from atmospheric mixing ratio mea-
5   surements using a global inverse of atmospheric transport, Technical Reports, Max-Planck-Institut für Biogeochemie 6, avialable at: http://www.bgc-jena.mpg.de/mpg/websiteBiogeochemie/Publikationen/Technical_Reports/tech_report6.pdf (last access: 10 February 2011), 2005.

Takahashi, T., Sutherland, S. C., Wanninkhof, R., Sweeney, C., Feely, R. A., Hales, B.,
10   Friederich, G., Chavez, F., Watson, A., Bakker, D. C. E., Schuster, U., Metzl, N., Yoshikawa-Inoue, H., Ishii, M., Midorikawa, T., Sabine, C., Hoppema, M., Olafsson, J., Arnarson, T. S., Tilbrook, B., Johannessen, T., Olsen, A., Bellerby, R., De Baar, H. J. W., Nojiri, Y., Wong, C. S., Delille, B., and Bates, N. R.: Climatological mean and decadal changes in surface ocean pCO_2, and net sea-air CO_2 flux over the global oceans, Deep-Sea Res. Pt. II, 56, 554–577,
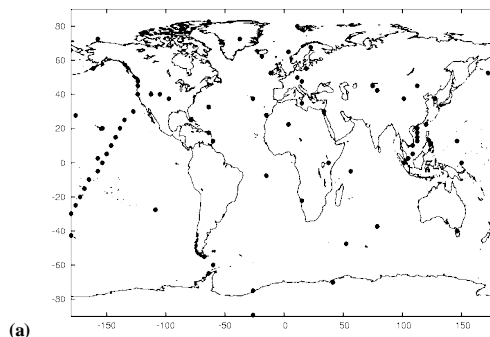15   doi:10.1016/j.dsr2.2008.12.009, 2009.

Thompson, R. L., Gerbig, C., and Rödenbeck, C.: A Bayesian inversion estimate of N_2O emissions for western and central Europe and the assessment of aggregation errors, Atmos. Chem. Phys., 11, 3443–3458, doi:10.5194/acp-11-3443-2011, 2011.

van der Werf, G. R., Randerson, J. T., Giglio, L., Collatz, G. J., Mu, M., Kasibhatla, P. S., Mor-
20   ton, D. C., DeFries, R. S., Jin, Y., and van Leeuwen, T. T.: Global fire emissions and the contribution of deforestation, savanna, forest, agricultural, and peat fires (1997–2009), Atmos. Chem. Phys., 10, 11707–11735, doi:10.5194/acp-10-11707-2010, 2010.
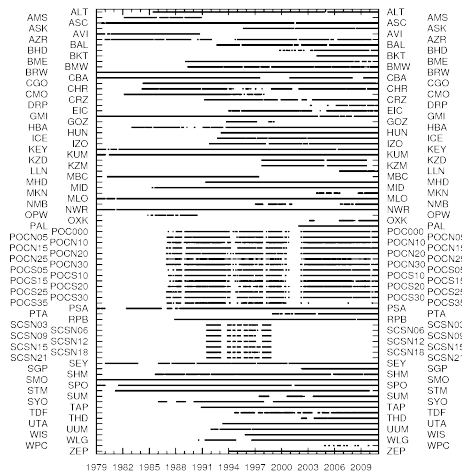
Zupanski, D., Denning, A. S., Uliasz, M., Zupanski, M., Schuh, A. E., Rayner, P. J., Peters, W., and Corbin, K. D.: Carbon flux bias estimation employing Maximum Likelihood Ensemble
25   Filter (MLEF), J. Geophys. Res., 112, D17107, doi:10.1029/2006JD008371, 2007.

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Fig. 1. Location of the NOAA ESRL flask measurements used in this study **(a)** and measurement availability at each site **(b)**. In **(b)**, the sites are designated by their identifier in the NOAA database.
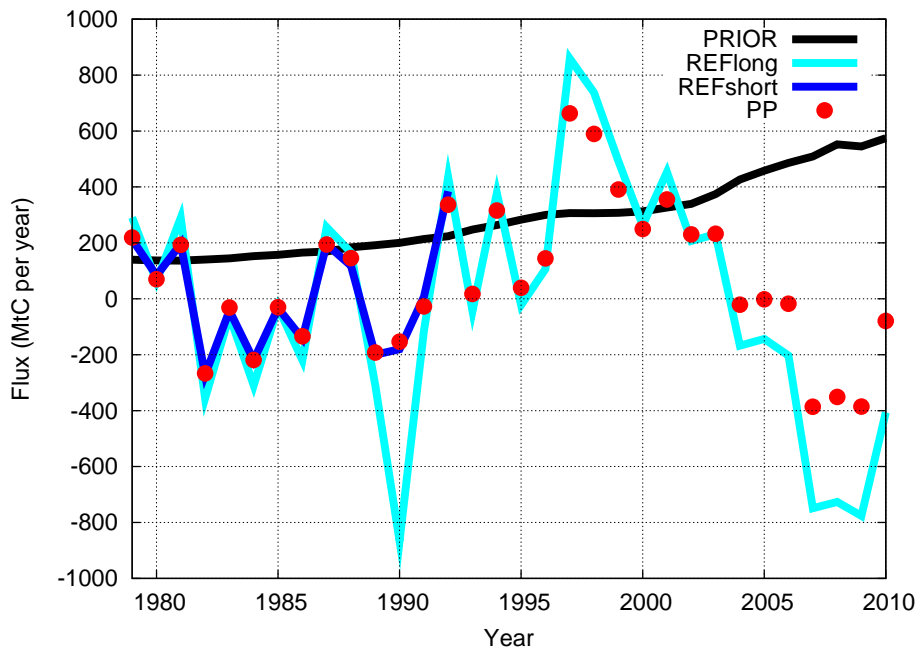
Fig. 2. Time series of the annual total CO$_2$ fluxes (including fossil fuel) in region Tropical Asia. In the sign convention, positive fluxes correspond to a net carbon source into the atmosphere. The black line (marked PRIOR) corresponds to the prior fluxes. The two coloured lines show the inverted fluxes for the reference 1979–2010 inversion (in cyan, marked REFlong) and for the reference 1979–1992 inversion (in regular blue, marked REFshort). The 1979–2010 PP inversion appears with red disks (PP).
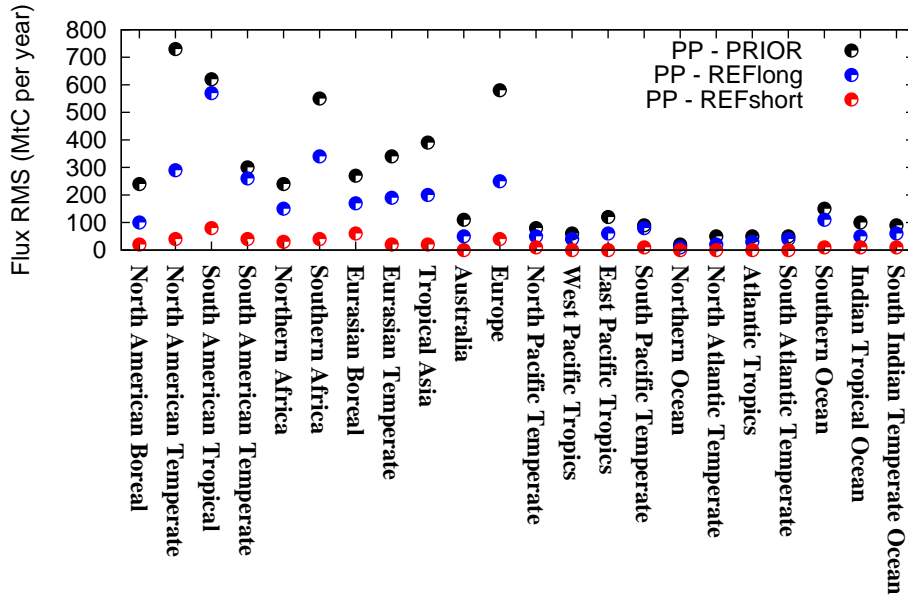
Fig. 3. RMS differences for annual CO$_2$ fluxes between the PP inversion and its prior for the period 1979–2010 (black, PP-PRIOR), between the 1979–2010 PP inversion and the 1979–2010 reference inversion for the same period (blue, PP-REFlong) and between the 1979–2010 PP inversion and the 1979–1992 reference inversion for 1979–1991 (red, PP-REFshort). The world is split into the 22 TransCom3 regions.
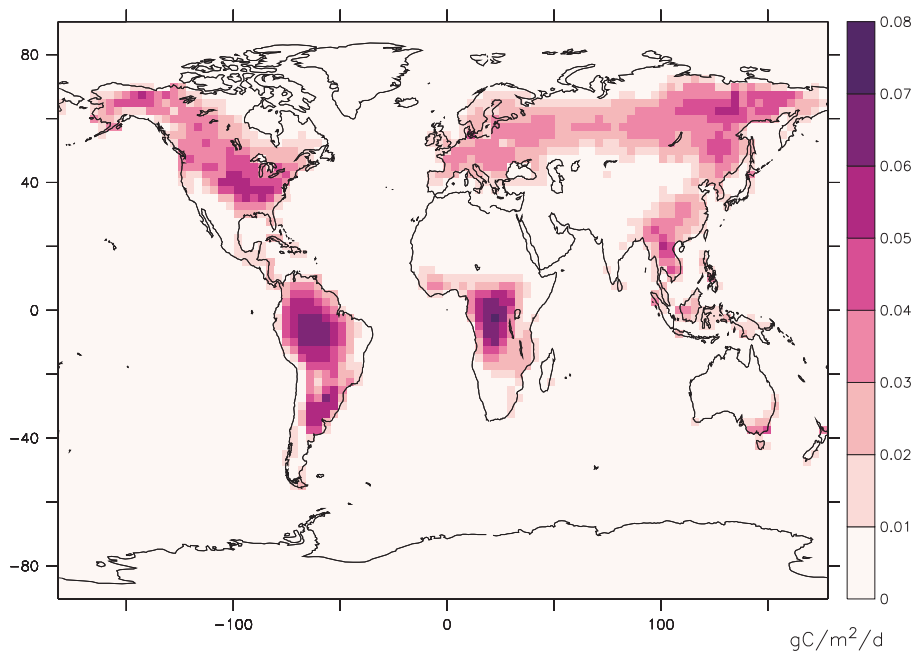
gC/m^2/d

Fig. 4. RMS differences between the 1979–2010 PP inversion and the reference 1979–1992 inversion when computing the inverted weekly fluxes. The statistics are computed for the 1979–1991 period.