


TECHNICAL REPORTS

6



Estimating CO₂ sources and sinks
from atmospheric mixing ratio measurements
using a global inversion of atmospheric transport

by
Christian Rödenbeck

Technical Reports - Max-Planck-Institut für Biogeochemie 6, 2005

Max-Planck-Institut für Biogeochemie

P.O.Box 10 01 64

07701 Jena/Germany

phone: +49 3641 576354

fax: + 49 3641 577300

email: christian.roedenbeck@bgc-jena.mpg.de

<http://www.bgc-jena.mpg.de/~christian.roedenbeck/homepage.html>

Estimating CO₂ sources and sinks
from atmospheric mixing ratio measurements
using a global inversion of atmospheric transport

Christian Rödenbeck

Max Planck Institute for Biogeochemistry, Postfach 10 01 64, D-07701 Jena (Germany)

Abstract

This report describes algorithmic details used to estimate the spatio-temporal distribution of CO₂ sources/sinks from atmospheric concentration measurements. The algorithm extends that used in Rödenbeck et al. (2003) in the following ways:

Higher time resolution. The inversion uses individual mixing ratio measurements (rather than monthly mean concentrations): Either individual flask pair averages, or hourly values from continuous analyzers. Correspondingly, also the time resolution of fluxes was increased (to daily flux values). In order to combine weekly and hourly data, a data density weighting is applied.

Generalized process-oriented flux model. All the a-priori information is supplied in the form of a statistical linear ‘flux model’. Using this formulation, it is more easy to specify a-priori constraints based on direct process understanding in a flexible and transparent way. At the same time, the prominence of a-priori constraints can be reduced to a minimum. The ‘flux model’ indirectly determines a-priori correlations (both spatial and temporal), a spatio/temporal weighting, and an overall scaling.

Iterative solution. Because of the higher time resolution, the number of knowns and unknowns is increased to a level that the minimization cannot be done via the analytical matrix expressions. Therefore, an iterative algorithm is used to find the cost function minimum, even though the problem is still linear. As additional advantage, more flexibility is gained, because atmospheric transport is simulated ‘on-line’ (rather than being fixed as soon as transport basis functions have been pre-computed), such that choices like data selection can easily be changed. As disadvantage, each run takes considerable amounts of CPU time. A-posteriori (co)variances can be calculated for a number of selected scalar quantities.

Besides the general description of the mathematical algorithm, specific implementation details and results for the case of CO₂ are presented.

Contents

1	Inversion algorithm – Overview	1
1.1	General setting	1
1.2	Input/output quantities and model relations	1
1.3	Cost function	2
1.4	Analytic solution	2
1.5	A-posteriori flux uncertainties	3
1.6	Integrated and filtered fluxes	3
1.7	Scaling of the Bayesian a-priori constraint	4
2	Data	5
2.1	Measurement institutions, sites, and records	5
2.2	Data selection	5
2.3	Data weighting: concentration mismatch uncertainty	8
2.3.1	Measurement uncertainty	8
2.3.2	Model uncertainty	8
2.3.3	Data density weighting	9
2.3.4	Covariance matrix	11
3	Process-oriented flux model	12
3.1	General philosophy	12
3.2	Fixed a-priori term	13
3.3	Adjustable terms	14
3.3.1	Spatial coherency: $\{g_{m_s,i}^{\text{space}}(x,y)\}$	15
3.3.2	Temporal coherency, time scales: $\{g_{m_s,i}^{\text{time}}(x,y)\}$	15
3.3.3	Spatial/temporal weighting: $f_{\text{sh},i}(x,y,t)$	17
3.3.4	Scaling: α_i	18
3.3.5	Special flux components	20
3.3.6	The initialization flux component	20
3.4	A flux model for CO ₂	20
3.4.1	Standard case	20
3.4.2	Sensitivity tests	23
4	Results: CO₂ flux estimates	24
4.1	Standard results	24
4.2	Uncertainty	24
4.3	A-posteriori correlations	24
4.4	Remark: Time series filter	25
5	Iterative solution	26
5.1	Basic conjugate gradient algorithm	26
5.2	Numerical implementation of $\mathcal{A}$	28
5.3	Starting point	28
5.4	Diagnostics	28
5.4.1	Convergence diagnostics	28
5.4.2	Consistency diagnostics	30
5.5	Approximation of the inverse Hessian matrix	31
5.6	Ensuring conjugacy	31
5.7	Approximating the a-posteriori covariance matrix	32
5.8	Varying the a-priori scaling μ	32

A Time series filtering	34
A.1 Fourier decomposition	34
A.2 Specific filters	34
References	36
Color figures	37

1 Inversion algorithm – Overview

This section shortly reviews the inversion algorithm, and gives an overview about quantities and relations involved.

1.1 General setting

Time domain: The calculation is interannual, with a time period (from initial time t_i to end time t_e) of several years (≈ 5 to 20, depending on lengths of used data records). The time step of fluxes is 1 day.

Space domain: Global, spatial resolution of fluxes: $\approx 4^\circ$ latitude $\times 5^\circ$ longitude (land) or $\approx 8^\circ$ latitude $\times 10^\circ$ longitude (ocean).

1.2 Input/output quantities and model relations

The primary input is the *observed concentrations*. The total set of measured values (individual flask pair values or hourly values, respectively, see Sect. 2.2) is denoted by the vector $\mathbf{c}_{\text{obs}}$. The number of values is N .

Modelled concentrations are written as

$$\mathbf{c}_{\text{mod}} = \mathbf{A}\mathbf{f} + \mathbf{c}_{\text{ini}} \quad (1)$$

with the transport matrix $\mathbf{A}$ and the initial concentration¹ $\mathbf{c}_{\text{ini}}$. Fluxes $\mathbf{f}$ are written as a function of a vector of M free dimensionless parameters $\mathbf{p}$ in the form of a linear *flux model*

$$\mathbf{f} = \mathbf{f}_{\text{fix}} + \mathbf{F}\mathbf{p}. \quad (3)$$

This flux model is nothing else than defining the a-priori probability distribution of the fluxes in a different than the traditional way: The parameters $\mathbf{p}$ are assumed to have a-priori zero mean and unit variance and to be mutually uncorrelated,²

$$\langle \mathbf{p}_{\text{pri}} \rangle = 0, \quad \langle \mathbf{p}_{\text{pri}} \mathbf{p}_{\text{pri}}^T \rangle = \frac{1}{\mu} \mathbf{1} \quad (4)$$

(where $\mathbf{1}$ is the identity matrix, and the scaling factor μ should be thought of as 1 in the standard case, see Sect. 1.7 below). Then, the a-priori expectation value of the flux is represented by $\mathbf{f}_{\text{fix}}$,

$$\langle \mathbf{f}_{\text{pri}} \rangle = \mathbf{f}_{\text{fix}}, \quad (5)$$

and the a-priori covariance matrix of the fluxes follows from $\mathbf{F}$ as

$$\mathbf{Q}_{\mathbf{f},\text{pri}} = \frac{1}{\mu} \mathbf{F}\mathbf{F}^T. \quad (6)$$

¹The initial concentration is assumed fixed and corresponds to a well-mixed atmosphere, i.e., each element of the vector $\mathbf{c}_{\text{ini}}$ is equal to a value c_0 . For CO_2 , the value

$$c_0 = 368.0 \text{ ppm} + 1.5 \text{ ppm/yr} \cdot (t_i - t_{2000}) \quad (2)$$

has been chosen, where $(t_i - t_{2000})$ is the time difference (in years, possibly negative) between the start of the inversion period and the beginning of year 2000. The exact choice of the initial concentration does not matter; a consistent 3D concentration field will be achieved at the end of the spin-up period (first two years) by adjustment of the fluxes during the first months (compare Sect. 3.3). The assumption of a well-mixed initial atmosphere is convenient as this state is unchanged by transport and therefore simply leads to a constant offset c_0 to all modelled concentrations at any time and location.

²All stochastic quantities are assumed Gaussian. They are written as a sum of mean and error, e.g.,

$$\mathbf{f} = \langle \mathbf{f} \rangle + \Delta \mathbf{f}$$

Covariance matrices are written as

$$\mathbf{Q}_{\mathbf{f}} = \langle \Delta \mathbf{f} \Delta \mathbf{f}^T \rangle$$

Thus, the matrix $\mathbf{F}$ comprises all the a-priori information about flux uncertainties and correlations. (The a-priori covariance matrix $\mathbf{Q}_{\mathbf{f},\text{pri}}$ is only given for illustration, it does not appear explicitly anywhere in the algorithm.) The construction of $\mathbf{f}_{\text{fix}}$ and $\mathbf{F}$ is described in Sect. 3.

Combining transport model, Eqn. (1), and flux model, Eqn. (3), gives

$$\mathbf{c}_{\text{mod}} = \mathbf{c}_{\text{mod,fix}} + \mathbf{A}\mathbf{F}\mathbf{p} \quad (7)$$

where the abbreviation

$$\mathbf{c}_{\text{mod,fix}} = \mathbf{A}\mathbf{f}_{\text{fix}} + \mathbf{c}_{\text{ini}} \quad (8)$$

has been introduced. The concentration mismatch between observed and modelled values is defined as

$$\mathbf{m} = \mathbf{c}_{\text{obs}} - \mathbf{c}_{\text{mod}}. \quad (9)$$

1.3 Cost function

According to Eqn. (4), the a-priori (unconditional) probability distribution of the parameters is

$$\text{Prob}(\mathbf{p}) \propto \exp\left(-\frac{\mu}{2}\mathbf{p}^T\mathbf{p}\right). \quad (10)$$

If the true parameter values were given, the concentration mismatch $\mathbf{m}$ would be solely attributed to random errors in the measured and the modelled concentrations. These errors are assumed to have zero mean (i.e., no biases), and a given covariance matrix $\mathbf{Q}_{\mathbf{m}}$ (to be defined in Sect. 2.3). The (conditional) probability distribution of the mismatch therefore is

$$\text{Prob}(\mathbf{m}|\mathbf{p}) \propto \exp\left(-\frac{1}{2}\mathbf{m}^T\mathbf{Q}_{\mathbf{m}}^{-1}\mathbf{m}\right). \quad (11)$$

According to Bayes' theorem (Tarantola, 1987), the (conditional) probability distribution of the parameters, given the concentration mismatch, reads

$$\text{Prob}(\mathbf{p}|\mathbf{m}) = \frac{\text{Prob}(\mathbf{m}|\mathbf{p})\text{Prob}(\mathbf{p})}{\text{Prob}(\mathbf{m})}. \quad (12)$$

The inversion seeks the expectation value $\langle\mathbf{p}_{\text{post}}\rangle$ of this a-posteriori distribution. As $\langle\mathbf{p}_{\text{post}}\rangle$ coincides with the most probable parameter value, it is found by maximizing $\text{Prob}(\mathbf{p}|\mathbf{m})$ with respect to $\mathbf{p}$, which is equivalent to minimizing the cost function

$$J = -\ln\left(\text{Prob}(\mathbf{p}|\mathbf{m})\right) \quad (13)$$

$$= \frac{1}{2}\mathbf{m}^T\mathbf{Q}_{\mathbf{m}}^{-1}\mathbf{m} + \frac{\mu}{2}\mathbf{p}^T\mathbf{p} + C. \quad (14)$$

There is one term each for the data and the a-priori constraints. The additive constant C subsumes all parameter-independent terms, e.g., those arising from $\text{Prob}(\mathbf{m})$ and from the normalization of the distribution. The cost function minimum is calculated from

$$\left.\frac{\partial J}{\partial \mathbf{p}^T}\right|_{\mathbf{p}=\langle\mathbf{p}_{\text{post}}\rangle} = 0. \quad (15)$$

1.4 Analytic solution

The cost function J to be minimized with respect to $\mathbf{p}$ can be written in the form

$$J(\mathbf{p}) = \frac{1}{2}\mathbf{p}^T\mathcal{A}\mathbf{p} - \mathbf{b}^T\mathbf{p} + C' \quad (16)$$

with the Hessian matrix

$$\mathcal{A} = \mathbf{F}^T\mathbf{A}^T\mathbf{Q}_{\mathbf{m}}^{-1}\mathbf{A}\mathbf{F} + \mu\mathbf{1}, \quad (17)$$

the linear factor (the later ‘right-hand side’)

$$\mathbf{b} = \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} (\mathbf{c}_{\text{obs}} - \mathbf{c}_{\text{mod,fix}}), \quad (18)$$

and the constant

$$C' = \frac{1}{2} \mathbf{m}_{\text{fix}}^T \mathbf{Q}_m^{-1} \mathbf{m}_{\text{fix}} + C. \quad (19)$$

From

$$\frac{\partial J}{\partial \mathbf{p}^T} = \mathcal{A} \mathbf{p} - \mathbf{b} \quad (20)$$

and Eqn. (15) then follows

$$\mathbf{p}_{\text{post}} = \mathcal{A}^{-1} \mathbf{b}. \quad (21)$$

In the original quantities, the final result reads

$$\mathbf{f}_{\text{post}} = \mathbf{f}_{\text{fix}} + \mathbf{F} (\mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} + \mu \mathbf{1})^{-1} \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} (\mathbf{c}_{\text{obs}} - \mathbf{c}_{\text{mod,fix}}) \quad (22)$$

In most inversion calculations published so far, the transport matrix $\mathbf{A}$ has been pre-computed by the transport model (or its adjoint) and stored, and the matrix inverse has been computed by singular value decomposition (SVD) of the coordinate-transformed model matrix

$$\mathbf{M} = \mathbf{Q}_m^{-1/2} \mathbf{A} \mathbf{F}. \quad (23)$$

1.5 A-posteriori flux uncertainties

The uncertainties and correlations of the estimated parameter values are contained in the covariance matrix of the a-posteriori probability distribution $\text{Prob}(\mathbf{p}|\mathbf{m})$. It can be read off from Eqns. (13) and (16) to be $\mathcal{A}^{-1}$. The a-posteriori covariance matrix of the flux estimates therefore is

$$\mathbf{Q}_{\mathbf{f},\text{post}} = \mathbf{F} \mathcal{A}^{-1} \mathbf{F}^T. \quad (24)$$

1.6 Integrated and filtered fluxes

The full spatial and temporal resolution contained in the flux vector $\mathbf{f}$ (model pixels and daily time steps, respectively) are not normally the quantities of interest, and, more importantly, not well-constrained by the data. Rather, fluxes will always be integrated over certain spatial regions, and filtered in time. In this report, spatial aggregation is always based on the TransCom-3 set of regions shown in Fig. 11. Temporally, fluxes are either shown ‘as is’ (with seasonality and high-frequency variability), or deseasonalized and low-pass filtered for interannual variability (IAV). The standard IAV filter used in all time series figures deseasonalizes and has Gaussian spectral weights (‘Filt0.5gd’ as defined in Appendix A).

Integration/filtering are linear functionals (scalars) of the flux field which can be written as

$$f_{\mathbf{z}} = \mathbf{f}^T \mathbf{z}, \quad (25)$$

where the vector $\mathbf{z}$ can be thought of as a spatio-temporal pattern of weights. The corresponding a-priori and a-posteriori variances are given by

$$\sigma_{f_{\mathbf{z}},\text{pri}}^2 = \frac{1}{\mu} \mathbf{z}^T \mathbf{F} \mathbf{F}^T \mathbf{z}, \quad \sigma_{f_{\mathbf{z}},\text{post}}^2 = \mathbf{z}^T \mathbf{F} \mathcal{A}^{-1} \mathbf{F}^T \mathbf{z}. \quad (26)$$

The covariances between f_1 and f_2 corresponding to two different patterns $\mathbf{z}_1$ and $\mathbf{z}_2$ are

$$q_{f_1,f_2,\text{pri}} = \frac{1}{\mu} \mathbf{z}_1^T \mathbf{F} \mathbf{F}^T \mathbf{z}_2, \quad q_{f_1,f_2,\text{post}} = \mathbf{z}_1^T \mathbf{F} \mathcal{A}^{-1} \mathbf{F}^T \mathbf{z}_2. \quad (27)$$

1.7 Scaling of the Bayesian a-priori constraint

In the a-priori probability distribution of the parameters, a factor μ was introduced, but disregarded up to now ($\mu = 1$). As particularly obvious from Eqn. (14), this factor scales the impact of the a-priori constraints: In the limit $\mu \rightarrow 0$, the a-priori term would vanish, in accordance with $\mathbf{Q}_{f,pri}$ from Eqn. (6) going to infinity. More generally, however, the factor μ determines the *ratio* between a-priori and data constraints. This is seen by noting that the cost function J as a whole could be scaled by another arbitrary factor, without changing the a-posteriori fluxes (only the a-posteriori covariances would change then).

As discussed in more detail in Sect. 3.3.4, the overall magnitudes of both the model-data mismatch uncertainty $\mathbf{Q}_m$ and the a-priori flux covariance matrix $\mathbf{Q}_{f,pri}$ (as defined through $\mathbf{F}$) are not well determined in practice, even if there is some physically-based information on their internal structures. In this sense, μ may be considered just as a prototype for the many settings that have to be chosen when devising an inversion in a particular application, in particular as a prototype for settings in $\mathbf{F}$. This parameter μ has been highlighted here because (i) it is one of the most influential settings, (ii) it is directly linked to certain formal mathematical criteria of the Bayesian framework, discussed further in Sect. 3.3.4, and (iii) when solving the cost function minimization iteratively, there is a particularly easy way to vary μ for the purpose of sensitivity testing (Sect. 5.8).

2 Data

2.1 Measurement institutions, sites, and records

Air samples are collected and analyzed by several institutions, which are given in Table 2. Codes, names, and coordinates of sampling sites are listed in Table 1. A map is given in Fig. 1. In order to denominate data records in figure legends (and file names), the site codes (capital letters or digits) are extended by 2 prepended and 1 appended small letters. The first prepended letter denotes the kind of data (time resolution/selection, Table 3). The second prepended letter can indicate some additional data selection (if any, Table 4). At several sites, measurements are done by more than one institution. In order to distinguish the respective records, the appended letter indicates the data source, to the extent necessary for uniqueness (Table 2). The time coverage of all CO₂ concentration records used here is depicted in Fig. 4. This figure also indicates the selection of records in the different inversion runs presented later in Sect. 4. This record selection is done such that all included records span the entire target period of the respective run.

2.2 Data selection

Data as provided by the measurement institutions (Table 2) are first selected according to their accompanying ‘hard flags’ (validity flags or special values), if any. In some cases, concurrent records exist that differ in experimental aspects (e.g., sampling method); there, selection is done according to the recommendations by the data providers. An overview is given in Table 5. All other data sources not listed in Table 5 provided preselected data.

For flask sampling sites, the flask pair (or triplet) means are taken as individual data values. If individual flask values are provided, they are averaged into pair/triplet means if sampled within 1 hour from each other. For continuous (in-situ) measurement sites, hourly (or half-hourly, respectively, as provided) values are taken as individual data points.

Further data selection is done with the intention that, while as much valid data as possible should be retained, measurements should be avoided if they are very likely unrepresentative for larger source areas or particularly badly reproduced in the transport model (these two conditions are very similar, as the coarseness of the model implies that it cannot be able to capture any small-scale features). For flask sampling, this kind of selection is usually already done by the sampling institution, by selecting the sampling schedule. For continuous records from continental sites, mostly day-time values only are used (Table 3), because the transport model is expected to be particularly wrong during night when a stable boundary layer forms. Options for further automatic data selection are listed in Table 4. Normally, records with limited variability (selection ‘v’: disregard any pair of consecutive values if they represent a change of more than 1 ppm within 1 hr) are chosen for continuous sites.

In addition to this automatic selection, some manual selection was done (mainly removing ‘spikes’ in the concentration record; suspected to be due to some local emission or atmospheric transport pathway smaller in scale than the model resolution). A measurement is potentially disregarded if it considerably deviates from the mean seasonal cycle of that site. However, if the concentration modelled from the a-priori flux shows a similar feature in time, the excursion is considered transport-related and retained. As a further guidance, it is checked whether the considered measurement causes some unusual excursion in the estimated fluxes, and whether soft flags are set by the data provider. This procedure is clearly not free of arbitrariness, but the number of manually out-selected values is very small. The described selection is mainly done for flask sites; continuous records often contain very-high frequency variations of considerable amplitude, but the impact of a single value is low anyway due to the weighting described later.

At some sites, two or more alternative records exist, measured by different institutions or using different experimental techniques. As revealed in a quantitative assessment (Rödenbeck et al., 2005), differences between many of such records are small: Flux differences in response to systematic experimental errors are almost negligible compared to many other errors in the inversion. Therefore, it seems justified to merge such records, offering the advantage that data gaps can mutually be filled. Merging of records has been done here for a few sites where both flask records and continuous records exist (Fig. 4). Note that the data density weighting




Figure 1: Measurement sites.

Table 1: Sites measuring CO₂ concentrations that were used in the inversion (to be continued on next page). The locations are shown in Fig. 1.

Code	Name / Geographic location	Latitude (°)	Longitude (°)	Height (m a.s.l.)	Class (Tab. 6)
ALT	Alert, Canada	82.45	-62.52	210	S
AMS	Amsterdam Island, Indian Ocn.	-37.80	77.53	65	R
ASC	Ascension Island, Atlantic	-7.92	-14.42	54	R
ASK	Assekrem, Algeria	23.18	5.42	2728	R
AZR	Terceira Island, Azores	38.75	-27.08	30	R
BAL	Baltic Sea, Poland	55.50	16.67	7	C
BHD	Baring Head, New Zealand	-41.42	174.87	85	S
BME	St. David's Head, Bermuda	32.37	-64.65	30	R
BMW	Southampton, Bermuda	32.27	-64.88	30	R
BRW	Barrow, Alaska	71.32	-156.60	11	S
BSC	Black Sea, Constanta, Romania	44.17	28.68	3	C
CBA	Cold Bay, Alaska	55.20	-162.72	25	S
CFA	Cape Ferguson, Australia	-19.28	147.05	2	S
CGO	Cape Grim, Tasmania	-40.68	144.68	94	S
CHR	Christmas Island, Pacific	1.70	-157.17	3	R
CMO	Cape Meares, Oregon, USA	45.48	-123.97	30	S
CPT	Cape Point, South Africa	-34.35	18.48	230	S

Table 1: (continued) Sites measuring CO₂ concentrations.

Code	Name / Geographic location	Latitude (°)	Longitude (°)	Height (m a.s.l.)	Class (Tab. 6)
CRZ	Crozet, S Ocn.	-46.45	51.85	120	R
EIC	Easter Island, Pacific	-29.15	-109.43	50	R
ESP	Estevan Point, Canada	49.38	-126.53	39	S
GMI	Guam, Mariana Island, Pacific	13.43	144.78	2	R
HBA	Halley Bay, Antarctica	-75.67	-25.50	10	R
HUN	Hegyhatsal, Hungary	46.95	16.65	344	C
ICE	Heimaey, Iceland	63.25	-20.15	100	R
ITN	Grifton, North Carolina, USA	35.35	-77.38	505	C
IZO	Izaña, Tenerife	28.30	-16.48	2360	R
KEY	Key Biscayne, Florida	25.67	-80.20	3	S
KUM	Cape Kumukahi, Hawaii	19.52	-154.82	3	R
LEF	Park Falls, Wisconsin	45.93	-90.27	868	C
MAA	Mawson, Antarctica	-67.62	62.87	32	R
MBC	Mould Bay, Canada	76.25	-119.35	58	S
MHD	Mace Head, Ireland	53.33	-9.90	25	S
MID	Sand Island, Midway, Pacific	28.22	-177.37	4	R
MLO	Mauna Loa, Hawaii	19.53	-155.58	3397	RU
MNM	Minamitorishima, Japan	24.30	153.97	8	R
MQA	Macquarie Island, S Ocn.	-54.48	158.97	12	R
NWR	Niwot Ridge, USA	40.05	-105.58	3475	M
PSA	Palmer Station, Antarctica	-64.92	-64.00	10	R
RPB	Ragged Point, Barbados	13.17	-59.43	3	R
RYO	Ryori, Japan	39.03	141.83	230	S
SBL	Sable Island, Canada	43.93	-60.02	5	R
SEY	Mahe Island, Seychelles	-4.67	55.17	3	R
SHM	Shemya Island, Alaska	52.72	174.10	40	R
SIS	Shetland Islands, UK	60.17	-1.17	30	R
SMO	Tutuila, American Samoa, Pacific	-14.25	-170.57	42	R
SPO	South Pole	-89.98	-24.80	2810	R
STM	Station 'M', Atlantic	66.00	2.00	7	S
SYO	Syowa, Antarctica	-69.00	39.58	11	R
TAP	Tae-ahn Peninsula, Korea	36.73	126.13	20	S
TDF	Tierra del Fuego, Argentina	-54.87	-68.48	20	S
UTA	Wendover, Utah, USA	39.90	-113.72	1320	C
UUM	Ulaan Uul, Mongolia	44.45	111.10	914	C
WES	Westerland, Germany	54.93	8.32	12	S
WIS	Sede Boker, Israel	31.13	34.88	400	C
WLG	Mt. Waliguan, China	36.27	100.92	3810	M
YON	Yonagunijima, Japan	24.47	123.02	30	R
ZEP	Zeppelin, Spitsbergen	78.90	11.88	474	R

Table 2: Measurement institutions and references. The last column gives the appended letter in the record denomination.

CMDL	Climate Monitoring & Diagnostics Laboratory, USA (Conway et al., 1994)	c
CSIRO	Commonwealth Scientific and Industrial Research Organization, Australia (Francey et al., 2003)	a
JMA	Japan Meteorological Agency ^a (Watanabe et al., 2000)	j
LSCE	Laboratoire des Sciences du Climat et de l'Environnement, France www.lsce.cnrs-gif.fr	l
MSC	Meteorological Service of Canada (Worthy et al., 2003)	m
NIPR	National Institute of Polar Research and Tohoku University, Japan (Morimoto et al., 2003)	x
NIWA	National Institute of Water and Atmospheric Research, New Zealand ftp.niwa.co.nz/tropac/co2/bhd/	x
SAWS	South African Weather Service (Labuschagne et al., 2003)	x
UBA	Umwelt-Bundesamt, Germany ^a (Levin et al., 1995) www.umweltbundesamt.de/uba-info-daten/daten/mbm/	u

^aData source: World Data Centre for Greenhouse Gases (WDCGG),
gaw.kishou.go.jp/wdcgg.html

described in Sect. 2.3.3 below is essential here to compensate potentially varying temporal data density.

2.3 Data weighting: concentration mismatch uncertainty

According to Eqn. (14), each element of the data mismatch vector $\mathbf{m}$ is scaled inversely by a corresponding uncertainty interval. This acts as a weighting among the data values. Ideally, the uncertainty intervals should reflect the error statistics of both the measured and modelled concentrations. These are not known well, but in view on the weighting, at least the relative uncertainties should be reasonable. Therefore, a rather simple weighting was chosen, guided by assumptions on measurement and model errors, and taking into account a data density weighting.

2.3.1 Measurement uncertainty

For each individual flask pair mean or hourly mean, an uncertainty interval due to random measurement errors of

$$\sigma_{\text{obs}} = 0.3 \text{ ppm} \quad (28)$$

was assumed. It is based on the maximally allowed flask pair difference of 0.5 ppm (Conway et al., 1994) and intercomparison differences of 0.2 ppm (Masarie et al., 2001a), and assumed appropriate for all data.

2.3.2 Model uncertainty

Concerning the model uncertainty, it is expected that the ability of the transport model to realistically simulate the atmospheric concentration field, markedly varies according to the complexity of the local circulation as well as that of the surface fluxes surrounding the individual sites. In Rödenbeck et al. (2003), therefore, model uncertainty was related to temporal and spatial gradients at site location, as proxies. Due to some problems with this approach (it depends on the choice of a-priori fluxes; gradients become higher at finer model resolution

Table 3: Record denomination, first prepended letter: Kind of data.

f	Flask data
h	Hourly (or half-hourly) values
d	as ‘h’, Day-time selected (11:00–17:00 local time)
n	as ‘h’, Night-time selected (23:00–05:00 local time)
a	Aircraft (3D moving)

Table 4: Record denomination, second prepended letter: Additional data selection.

v	limited Variability (<1 ppm/hr)
x	no additional selection

although they can then be expected to be more realistic; counter-intuitive ratios between the uncertainties of certain sites –e.g., relatively high uncertainty for STM– occur), a simpler scheme is chosen here: All sites (Table 1) are classified according to their location with respect to continental vs. oceanic situations, as defined in Table 6. These site classes are then assigned different model uncertainties σ_{mod} (without any time dependence) for the standard inversion set-up. In addition, some sub-classes are defined, to be able to take into account special model problems at certain sites, if known (such as land/sea breeze systems). Apart from being simple (given the poor knowledge about the actual model errors), this scheme allows to test the impact of the concentration weighting in a sensitivity analysis. The standard choices of Table 6 are generally very similar to those set by the approach used in Rödenbeck et al. (2003).

2.3.3 Data density weighting

For a given site, data and model uncertainties as defined above are combined quadratically into a total mismatch uncertainty

$$\sigma_{\text{tot}} = \sqrt{\sigma_{\text{mod}}^2 + \sigma_{\text{obs}}^2} \quad (29)$$

with respect to the individual data values. However, the inversion algorithm is meant to make use of both flask and continuous concentration data. This needs some special consideration, to avoid that the much higher amount of data for continuous measurements would lead to a considerably stronger impact of these sites: Consider a certain site during some time interval, e.g., a particular week. Assume that there are N^* measurement values $c_{\text{obs},i}$. Then the contribution from these particular values to the cost function is

$$J^* = \frac{1}{2} \sum_{i=1}^{N^*} \frac{(c_{\text{obs},i} - c_{\text{mod},i})^2}{\sigma_i^2} \quad (30)$$

where the sum runs over all individual values in the considered week. If $\sigma_i = \sigma_{\text{tot}}$ for all these values, this contribution J^* would approximately increase in proportion to N^* (which is 1 for a typical flask site and up to $7 \times 24 = 168$ for an hourly site). Likewise, the contribution

$$\frac{\partial J^*}{\partial \mathbf{p}^T} = - \sum_{i=1}^{N^*} \frac{c_{\text{obs},i} - c_{\text{mod},i}}{\sigma_i^2} \cdot \frac{\partial c_{\text{mod},i}}{\partial \mathbf{p}^T} \quad (31)$$

to the cost function gradient would be approximately proportional to N^* (assuming that the parameter change $\partial \mathbf{p}$ changes all $c_{\text{mod},i}$ roughly in the same way, which in turn is the case if the flux model does not contain sub-weekly frequencies [cmp. Sect. 3] and if the area of influence of the $c_{\text{mod},i}$ [according to atmospheric transport

Table 5: Rules for automatic selection of CO₂ data, for cases where flagged raw data are provided by the measurement institution (see Sect. 2.2). For the meanings of the mentioned flags and other codes, refer to the respective file documentation by the data providers. Data sources not listed here provided preselected data.

Data source	Data values considered valid if ...
CMDL (flask)	Flag 1 is ‘.’ Filling method is ‘P’, ‘D’, or ‘N’ (at BAL, CHR, and KUM also ‘T’)
CMDL (in-situ)	Flag 1 is ‘.’ Flag 2 is not ‘V’, ‘U’
CSIRO (flask)	Flag 1 is ‘.’ Flag 2 is not ‘O’, ‘F’
JMA	Flag is 7

Table 6: Site classes with respect to the expected degree of realism of the transport model. The last column gives the assumed model uncertainties for the standard set-up of the CO₂ case.

Class	Explanation	Model uncertainty σ_{mod}
C	‘Continental’ – surface sites on continents; complex circulation and fluxes, strong vertical transport	3.0 ppm
M	‘Mountain’ – mountain sites (in continents); simpler circulation due to high elevation	1.5 ppm
R	‘Remote’ – remote (oceanic) sites, such as islands, island mountains, deserts, Antarctica; relatively simple circulation	1.0 ppm
S	‘Shore’ – at land/sea interface; influenced both by oceanic and continental tracer sources	1.5 ppm
RU	R with diurnal Upslope wind variations	1.0 ppm
SB	S with diurnal sea-Breeze system	1.5 ppm

pathways] is within the spatial correlation range). This dependence of J^* and $\partial J^*/\partial \mathbf{p}^T$ on N^* can be removed if an increased uncertainty of

$$\sigma_i = \sqrt{N^*} \cdot \sigma_{\text{tot}} \quad (32)$$

is used. To implement this in a time-homogeneous way, N_i^* could, for each value in the concentration time series, be taken as the number of values within the surrounding weekly interval. In the actual implementation, the number of values within a 3-week interval around each value, divided by 3, was used.

By this construction of the data density weighting, the implied uncertainty of concentration *averages* over weekly (or longer) intervals does not depend on the number of measurements any more, as long as there is at least one data value per week. Specifically, with the standard choices of the CO₂ case, the uncertainty of *yearly* averaged concentrations ranges between ≈ 0.15 ppm for a remote site and ≈ 0.45 ppm for a continental site.

The choice of 1 week as reference interval refers to the typical sampling frequency for flask sites. In addition, this corresponds to the highest frequency in the fluxes that is solved for (compare Sect. 3). There is yet another aspect: The factor $\sqrt{N^*}$ would also occur in the uncertainty of the average of N^* fully correlated values, relative to that of uncorrelated values. This fits to the expectation that correlation of transport model errors might be stronger within the time scale of synoptic weather phenomena, which is in the order of a week.

To set the data density weighting further into perspective, consider the hypothetical cost function contribu-

tion from the *average* concentration in the considered week:

$$J^{**} = \frac{1}{2} \frac{(c_{\text{obs}}^* - c_{\text{mod}}^*)^2}{\sigma_{\text{tot}}^2}, \quad c_{\text{obs}}^* = \frac{1}{N^*} \sum_{i=1}^{N^*} c_{\text{obs},i}, \quad c_{\text{mod}}^* = \frac{1}{N^*} \sum_{i=1}^{N^*} c_{\text{mod},i}. \quad (33)$$

Note that in this formula, the uncertainty σ_{tot}^2 of the individual values has still been used, rather than the correct one, $\sigma_{\text{tot}}^2/N^*$. The gradient of this hypothetical contribution,

$$\frac{\partial J^{**}}{\partial \mathbf{p}^T} = - \sum_{i=1}^{N^*} \frac{c_{\text{obs},i} - c_{\text{mod},i}}{N^* \cdot \sigma_{\text{tot}}^2} \cdot \frac{1}{N^*} \sum_{j=1}^{N^*} \frac{\partial c_{\text{mod},j}}{\partial \mathbf{p}^T}, \quad (34)$$

is essentially identical to Eqn. (31) with σ_i from Eqn. (32), as long as the above assumption holds that $\partial c_{\text{mod},i}/\partial \mathbf{p}^T$ is equal for all data points i in the considered week. Thus, to the extent that fluxes change only coherently within a week, and that all the $c_{\text{mod},i}$ share a common area of influence, data density weighting is essentially equivalent to filtering the concentration time series on a weekly time scale, but using time-constant uncertainties disregarding any data density fluctuations. In contrast to such a data filtering, however, data density weighting as used here preserves the link between possibly *different* areas of influence (due to changing transport paths) and correspondingly different signals among the $c_{\text{mod},i}$, which may help to better detect spatial structure in the fluxes³.

The artificial increase of data uncertainties also means that the effective number of data constraints is no longer the total number N of values, but is reduced to

$$N_{\text{eff}} \approx \sum_{i=1}^N \frac{1}{N_i^*} \quad (35)$$

This expression is based on the assumption that averages over N_i^* values represent only one degree of freedom.

2.3.4 Covariance matrix

As a final step, the covariance matrix $\mathbf{Q}_m$ of the concentration mismatch $\mathbf{m}$ needs to be constructed. The squared uncertainty intervals of the individual data values, σ_i^2 , represent its diagonal elements. Though there surely is statistical dependence among the errors of different measurements (through common influences of standard gases, apparatus functioning, etc.), and even more so between different modelled values (through conceptual deficiencies in parameterizations, as well as errors in the driving fields), these covariances are neglected. Thus, all off-diagonal elements of $\mathbf{Q}_m$ are set zero.

³On the other hand, filtering of (high-frequency) data on synoptic time scales, similar to Eqn. (33), was found by Law et al. (2003) to be beneficial in reducing the impact of model errors, when retrieving fluxes from synthetic data (known truth).

3 Process-oriented flux model

As mentioned in Sect. 1.2, the ‘flux model’

$$\mathbf{f} = \mathbf{f}_{\text{fix}} + \mathbf{F}\mathbf{p} \quad (3')$$

introduces all a-priori information into the inversion calculation. Like in Rödenbeck et al. (2003), the total surface-atmosphere flux $\mathbf{f}$ is thought of as the sum of N_{comp} different source/sink components,

$$\mathbf{f} = \sum_{i=1}^{N_{\text{comp}}} \left(\beta_i \mathbf{f}_{\text{fix},i} + \alpha_i \mathbf{F}_i \mathbf{p}_i \right). \quad (36)$$

For each such component, there is a fixed a-priori term $\mathbf{f}_{\text{fix},i}$, and a deviation term that linearly depends on a subset $\mathbf{p}_i$ of the adjustable parameters (vector $\mathbf{p}$) of the inversion. The construction of this flux model can be easier described by writing the fluxes as continuous (2+1)D fields $f(x, y, t)$ depending on longitude x , latitude y , and time t (in years, with integer values at Jan 1, 00:00). The flux vector $\mathbf{f}$ is just a temporally and spatially discretely sampled representation of this flux field⁴. The flux model can then be written as

$$f(x, y, t) = \sum_{i=1}^{N_{\text{comp}}} \left(\underbrace{\beta_i f_{\text{fix},i}(x, y, t)}_{\text{fixed}} + \underbrace{\alpha_i f_{\text{sh},i}(x, y, t) \sum_{m_t=1}^{N_{t,i}} \sum_{m_s=1}^{N_{s,i}} g_{m_t,i}^{\text{time}}(t) g_{m_s,i}^{\text{space}}(x, y) p_{m_t,m_s,i}}_{= \delta f_i(x, y, t) \text{ adjustable}} \right) \quad (37)$$

In the adjustable terms, the ‘shape’ $f_{\text{sh},i}(x, y, t)$ represents a weighting in space and time that defines the ‘domain of activity’ of the flux component, while the sets of functions $\{g_{m_t,i}^{\text{time}}(t)\}$ and $\{g_{m_s,i}^{\text{space}}(x, y)\}$ determine a temporal and spatial decomposition into statistically independent elements. The exact meanings of these various quantities, and the algorithms performed to assign values to them for the individual components, are described in Sects. 3.2 and 3.3. Prior to the detailed algorithm, however, Sect. 3.1 gives some general remarks about the construction of the flux model. An example of specific choices that could be used in a CO_2 inversion is finally described in Sect. 3.4.

3.1 General philosophy

As already mentioned in Sect. 1.2, the flux model determines the a-priori mean $\langle \mathbf{f}_{\text{pri}} \rangle$ and the a-priori covariance $\mathbf{Q}_{\mathbf{f},\text{pri}}$ of the fluxes (Eqns. (5) and (6)). Traditionally, $\mathbf{Q}_{\mathbf{f},\text{pri}}$ is specified directly. The definition via a simple flux model as described here offers the advantage to base these choices on process understanding in a more direct, flexible, and transparent way.

The individual flux components correspond, on the one hand, to physical processes, such as fossil fuel burning, carbon exchange of terrestrial vegetation (NEE), or ocean-atmosphere carbon fluxes. However, some of these processes are further split into their action on different time scales, such as their long-term mean, mean seasonal cycle, interannual or short-term variability. Each of these contributions is then treated as a separate component in the flux model. This is convenient as it allows to assign different shapes f_{sh} and different spatial decompositions g^{space} to these different time scales.

It needs to be stressed, however, that this ‘flux model’ is not a process model. The ultimate reason to include a-priori information into the Bayesian inversion at all, is to remedy the fact that the problem is otherwise highly under-determined. Corresponding to this aim, the flux model is only meant to prescribe certain statistical properties of the fluxes, to the degree that seems necessary to stabilize the estimation mathematically. As a general guidance, it is tried to prescribe as little a-priori flux constraints as possible. In particular, the fixed a-priori flux terms $f_{\text{fix},i}(x, y, t)$ do not have those kinds of variability that are meant to be assessed by the inversion,

⁴Imagine $f(x, y, t)$ to be given in $\text{Tmol/m}^2/\text{s}$, and the elements of $\mathbf{f}$ to be the integrals of $f(x, y, t)$ over the respective grid cells and time steps, given in Tmol/cell/day (or Tmol/cell/month , resp.). With this convention, integrals over $f(x, y, t)$ then correspond to sums over elements of $\mathbf{f}$.

in order that all the variability found in the result be clearly traceable back to the atmospheric information. This is a slightly different view than the classical Bayesian interpretation of $f_{\text{fix},i}(x, y, t)$ being an a-priori best-guess flux estimate and the adjustable term just being its error. Rather, as the ‘best-guess’ deliberately lacks the variability of interest, the adjustable term (its ‘error’) plays the role of the actual flux component, while $f_{\text{fix},i}(x, y, t)$ just pre-subtracts some part of it that probably cannot be retrieved by the inversion or that is not of interest in the particular study.

Therefore, it is the adjustable terms that receive most attention here. Whenever possible, it is tried to use process understanding in defining them, e.g., by statistical analysis of suitable input data sets as described below, or by reasoning about spatial and temporal scales of events that may lead to coherent flux anomalies (for CO₂, e.g., drought conditions on land, or upwelling events in the ocean). However, if the available information is rather ambiguous, the most simple choice (e.g., spatially or temporally unstructured shapes, globally valid correlation lengths, etc.) are preferred (“Ockham’s razor”).

As a further construction principle, it is tried to also have flux components with implied correlations of very long spatial or temporal range (long term fluxes, mean seasonal cycles, global fluxes, etc.) besides those for short-term/small-scale variability, whenever it seems justified from the process perspective. This is because these large scale components are constrained by a (much) greater amount of data than the small scales; therefore, it can be expected that splitting them off will tend to stabilize the calculation in a similar way as a best-guess prior flux would do. Yet, these large-scale estimates will reflect information from the atmospheric data, not a-priori information. This treatment was inspired by the Geostatistical approach, applied to the CO₂ inversion problem by Michalak et al. (2004). In the Geostatistical approach, a long-term value is estimated in addition to the variability term, effectively changing the correlation structure in such a way that the inversion is stabilized even without an explicit a-priori term. Extending this idea, the different time scales of one physical process can be implemented as separate components in the flux model here. This is particularly relevant for the seasonal cycle: The frequencies $\nu_k = 1/\text{yr}, 2/\text{yr}, 3/\text{yr}, \dots$ are known a-priori to have much higher amplitudes than all other time scales, immediately justifying their special role. In addition, there are processes for which distinctly different behaviour (in terms of variability amplitudes or spatial correlations) can be expected at low (e.g., interannual) or high (e.g., synoptic) frequencies, or in the long-term mean. The time scale separation allows to exploit this information.

In the following, algorithms for the construction of a flux model are described in a rather general way, leaving many alternative choices at various occasions. Yet more possibilities arise since different input data sets may be used. For a particular application, such as the case of CO₂ considered here, it is generally not possible to decide which choices will be ‘best’. In view of “Ockham’s razor” mentioned above, it probably does not make sense to actually exploit all the available flexibility to create a sophisticated flux model. Rather, the view is advocated that several different statistical ‘flux models’ (i.e., different values of parameters like correlation lengths, scaling factors, etc.), based on rational choices, should be tried and compared. The resulting range of flux estimates is then an indication of the robustness of the result.

3.2 Fixed a-priori term

For some components, there is no parameter-independent term, i.e., the implicit a-priori flux is zero, $\beta_i f_{\text{fix},i}(x, y, t) = 0$. Otherwise, the flux field $f_{\text{fix},i}(x, y, t)$ is derived from a suitable input data set $f_{\text{input},j}(x, y, t)$. A scaling factor β_i may be needed if this input data are not directly fluxes of the considered species, but those of another species linked through some conversion factor, or other proxy data. According to the general construction philosophy outlined above, the chosen input $f_{\text{input},j}(x, y, t)$ is generally filtered, to use part of the information only. There are the following options, with decreasing degrees of filtering:

- long-term mean $f_{\text{input},j}^{\text{mean}}(x, y)$ [see Eqn. (113)],
- mean seasonal cycle $f_{\text{input},j}^{\text{seas}}(x, y, t)$ [using e.g. a filter such as ‘Filt7.0Ts’, see Appendix A],
- low-pass filtered time series $f_{\text{input},j}^{(\text{filter})}(x, y, t)$, or
- time series $f_{\text{input},j}(x, y, t)$ as is.

In addition, many available input data sets are given at monthly time steps. Then, some low-pass filtering is necessary just to obtain smoothly interpolated flux values at the daily time steps. In these cases, generally ‘Filt7.0T’ (i.e., essentially monthly variations) was chosen.

3.3 Adjustable terms

In the flux model Eqn. (37), the adjustable term

$$\delta f_i(x, y, t) = \alpha_i f_{\text{sh},i}(x, y, t) \sum_{m_t=1}^{N_{t,i}} \sum_{m_s=1}^{N_{s,i}} g_{m_t,i}^{\text{time}}(t) g_{m_s,i}^{\text{space}}(x, y) p_{m_t,m_s,i} \quad (38)$$

of flux component i is defined by a (scaled) spatio-temporal ‘shape’ function $\alpha_i f_{\text{sh},i}(x, y, t)$, by a decomposition into a set of temporal elements $\{g_{m_t,i}^{\text{time}}(t)\}$, and by a decomposition $\{g_{m_s,i}^{\text{space}}(x, y)\}$ into spatial elements. The (relative) strength of all combinations of elements is then determined by the individual adjustable parameters $p_{m_t,m_s,i}$. A-priori, the parameters (Gaussian stochastic variables) are assumed to be statistically independent of each other (and of the parameters of all other components), and to have zero mean and unit variance,

$$\langle p_{m_t,m_s,i} \rangle = 0, \quad \langle p_{m_t,m_s,i}^2 \rangle = 1 \quad (39)$$

(Eqn. (4)). Consequently, the adjustable term has zero mean itself, and the a-priori covariance of the flux component at location (x, y) at time t with that at location (x', y') at time t' is

$$\begin{aligned} \langle \delta f_i(x, y, t) \delta f_i(x', y', t') \rangle &= \alpha_i^2 f_{\text{sh},i}(x, y, t) f_{\text{sh},i}(x', y', t') \\ &\times \sum_{m_t=1}^{N_{t,i}} g_{m_t,i}^{\text{time}}(t) g_{m_t,i}^{\text{time}}(t') \sum_{m_s=1}^{N_{s,i}} g_{m_s,i}^{\text{space}}(x, y) g_{m_s,i}^{\text{space}}(x', y') \end{aligned} \quad (40)$$

If the normalization condition

$$\sum_{m_s=1}^{N_{s,i}} \left(g_{m_s,i}^{\text{space}}(x, y) \right)^2 = 1 \quad (41)$$

is fulfilled at every point (x, y) , and

$$\sum_{m_t=1}^{N_{t,i}} \left(g_{m_t,i}^{\text{time}}(t) \right)^2 = 1 \quad (42)$$

at any time t , then the a-priori uncertainty of the flux at any point/time is simply given by the shape,

$$\sqrt{\langle (\delta f_i(x, y, t))^2 \rangle} = \alpha_i f_{\text{sh},i}(x, y, t). \quad (43)$$

(As detailed below, it can also be advantageous to replace this pointwise normalization by a normalization with respect to larger time and/or space scales.) The correlation coefficients between simultaneous fluxes at different places are determined by the set of spatial elements,

$$\frac{\langle \delta f_i(x, y, t) \delta f_i(x', y', t) \rangle}{\sqrt{\langle (\delta f_i(x, y, t))^2 \rangle \langle (\delta f_i(x', y', t))^2 \rangle}} = \sum_{m_s=1}^{N_{s,i}} g_{m_s,i}^{\text{space}}(x, y) g_{m_s,i}^{\text{space}}(x', y'), \quad (44)$$

with an analogous relation for the correlation coefficients between fluxes at different times, but same location:

$$\frac{\langle \delta f_i(x, y, t) \delta f_i(x, y, t') \rangle}{\sqrt{\langle (\delta f_i(x, y, t))^2 \rangle \langle (\delta f_i(x, y, t'))^2 \rangle}} = \sum_{m_t=1}^{N_{t,i}} g_{m_t,i}^{\text{time}}(t) g_{m_t,i}^{\text{time}}(t'). \quad (45)$$

3.3.1 Spatial coherency: $\{g_{m_s,i}^{\text{space}}(x, y)\}$

As seen from Eqn. (44), the spatial decomposition $\{g_{m_s,i}^{\text{space}}(x, y)\}$ determines the spatial coherency of the flux. If the set of space elements would be a set of non-overlapping regions

$$g_{m_s,i,\text{uncorr}}^{\text{space}}(x, y) = \begin{cases} 1 & (\text{inside region}(m_s, i)), \\ 0 & (\text{elsewhere}), \end{cases} \quad (46)$$

then fluxes in different such regions would be a-priori uncorrelated (and fully correlated within each region). These spatial elements might be larger regions like those of TransCom-3 (Fig. 11) which leads to a traditional setup. Here, the elements mostly are pixels of a regular grid (e.g., identical to the grid of the transport model on land, or aggregates of 2×2 cells over ocean, see Fig. 6). Spatial correlations between regions arise as soon as the space elements $g_{m_s,i}^{\text{space}}(x, y)$ overlap each other. Exponentially decaying overlap of the form

$$g_{m_s,i}^{\text{space}}(x, y) = a_{m_s,i} \iint \exp\left(-\frac{\varrho(x, y; x', y')}{r_i^{(x)}}\right) \exp\left(-\frac{\varrho(x, y; x', y')}{r_i^{(y)}}\right) g_{m_s,i,\text{uncorr}}^{\text{space}}(x', y') dx' dy' \quad (47)$$

was chosen, leading approximately to exponentially decaying spatial correlations. Here, ϱ gives the distance between two points at the Earth surface, and $r_i^{(x)}$ and $r_i^{(y)}$ determine the length scales of the correlations in longitudinal and latitudinal direction, respectively. The integration effectively extents over the individual (small) regions considered independent previously. Note that Eqn. (47) represents a spatial convolution of the regions (here: pixels). It is therefore nothing else than a spatial smoothing of the previously uncorrelated a-priori flux field.

In order to fulfill the normalization condition Eq. (41), the individual spatial elements $g_{m_s,i}^{\text{space}}(x, y)$ are scaled through $a_{m_s,i}$. Since there are more grid cells than spatial elements, the values of the $a_{m_s,i}$ are overdetermined, however, such that the normalization can only be satisfied approximately. Therefore, a few iterations are performed in each of which the left hand side of Eq. (41) resulting from the current set of elements is calculated, and each element $g_{m_s,i}^{\text{space}}(x, y)$ is divided by the mean left hand side (average weighted by the respective element itself).

3.3.2 Temporal coherency, time scales: $\{g_{m_t,i}^{\text{time}}(x, y)\}$

Analogously to the spatial decomposition, the temporal decomposition corresponds to uncorrelated time intervals if

$$g_{m_t,i,\text{uncorr}}^{\text{time}}(t) = \begin{cases} 1 & (\text{during time interval}(m_t, i)), \\ 0 & (\text{otherwise}). \end{cases} \quad (48)$$

If the time intervals are the time steps, then the resulting a-priori flux at each location (x, y) is essentially a ‘white noise’ process. Time correlations can be implemented by overlapping elements, such as triangular pulses of width τ_i centered around each time step (here: each day). As in the spatial case, this way of introducing time correlations corresponds to a smoothing of the ‘white noise’ time series, by convolution with a (triangular) pulse response function.

However, if time scale separation is desired as outlined above, it can be more convenient to define time correlations via a Fourier expansion, where $g_{m_t,i}^{\text{time}}(t)$ are Fourier components (sine and cosine waves, plus a constant). Both the temporal coherence and the selection of time scales is then defined by choosing appropriate weights $w_i(\nu_k)$ for the amplitudes of these Fourier components. To make this quantitative, note that the uncorrelated (‘white noise’) time series corresponding to Eqn. (48) can be written as the Fourier series

$$\delta f_i(t) = \alpha_i f_{\text{sh},i}(t) \left[\sqrt{\frac{1}{N}} p_i^{\text{ave}} + \sum_{k=1}^{\infty} \sqrt{\frac{2}{N}} \left(\cos(2\pi\nu_k t) p_{k,i}^{\cos} + \sin(2\pi\nu_k t) p_{k,i}^{\sin} \right) \right] \quad (49)$$

where N is the number of time intervals, and all the parameters p_i^{ave} , $p_{k,i}^{\cos}$, and $p_{k,i}^{\sin}$ are independent and have zero mean and unit variance. This follows from Eqns. (113), (108), and (109) applied to white noise (where

the flux from each interval has been approximated by an independent Dirac δ pulse, and the shape $f_{\text{sh},i}(t)$ is assumed constant). To construct a time-varying component, the adjustable term is written as a filtered Fourier series,

$$\delta f_i(x, y, t) = \alpha_i \phi_i(x, y) \sum_{k=1}^{\infty} a_i w_i(\nu_k) u_i(\nu_k) \left(\cos(2\pi \nu_k t) p_{k,i}^{\cos}(x, y) + \sin(2\pi \nu_k t) p_{k,i}^{\sin}(x, y) \right). \quad (50)$$

This has the general structure (38) of the adjustable terms, by the identifications $f_{\text{sh},i}(x, y, t) = \phi_i(x, y)$, and $g_{m_t,i}^{\text{time}}(t) = a_i w_i(\nu_k) u_i(\nu_k) \cos(2\pi \nu_k t)$ or $a_i w_i(\nu_k) u_i(\nu_k) \sin(2\pi \nu_k t)$ with a suitable correspondence between m_t and k . The parameter terms

$$p_{k,i}^{\cos|\sin}(x, y) = \sum_{m_s=1}^{N_{s,i}} g_{m_s,i}^{\text{space}}(x, y) p_{m_t, m_s, i} \quad (51)$$

have zero mean and unit variance at each spatial point, due to the spatial normalization (41).

Considering an individual point (x, y) , the a-priori frequency spectrum of the component is determined by the factors $a_i w_i(\nu_k) u_i(\nu_k)$. In this expression, $w_i(\nu_k)$ denotes the filter to be applied in order to select certain time scales and to impose temporal correlations: Depending on which set of weights [such as Eqns. (118) or (119)] is chosen, the component will be sensitive to the low-frequency non-seasonal variability or to the mean seasonal cycle, where the parameter ν^{low} determines the temporal correlation length. Remember that the chosen low-pass filter has the same effect as convolution of the time series with a (triangular) impulse-response function.

In addition to the filter, one might also want to provide a-priori information about the general form of the spectrum, which can be done by the factor $u_i(\nu_k)$. Depending on the physical process, spectra like

$$u_i(\nu) = \begin{cases} \nu^{\varepsilon_i} & (\nu \geq 1) \\ 1 & (\text{else}) \end{cases} \quad (52)$$

could be prescribed, with typical exponents $\varepsilon = -1/2$ or $\varepsilon = -1$. However, if the sampling frequency of the data is high enough compared with the frequencies of interest, an a-priori white spectrum, $u_i(\nu_k) = 1$, may be most appropriate. If the resulting (unnormalized) weight $w_i(\nu_k) u_i(\nu_k)$ is less than 0.02 (i.e., if the frequency ν_k is damped to less than 2% of the maximum), then it is omitted completely, in order to avoid almost ‘dead’ parameters.

Finally, the constant a_i is set to

$$\frac{1}{a_i^2} = \sum_{k=1}^{\infty} \left(w_i(\nu_k) u_i(\nu_k) \right)^2 \quad (53)$$

in order to fulfill the normalization (42) (due to the statistical independence of the parameters attached to the individual Fourier components)⁵. It can, however, also be convenient to modify the normalization as

$$\frac{1}{a_i^2} = \sum_{k=1}^{\infty} \left(w_i^{(\text{norm})}(\nu_k) w_i(\nu_k) u_i(\nu_k) \right)^2 \quad (54)$$

with some additional low-pass filter $w_i^{(\text{norm})}(\nu)$. Of course, the normalization (42) is no longer fulfilled then. The effect is that the shape no longer determines the standard deviation of the *instantaneous* flux as in Eqn. (43), but that of the flux *filtered* according to $w_i^{(\text{norm})}(\nu)$, while the instantaneous flux has a higher standard deviation. For example, this can be used during sensitivity testing to ensure identical standard deviations for interannual variations when comparing different flux models with different spectra for the high frequencies. Another application is discussed in Sect. 3.3.3.

⁵For a ‘white noise’ component ($w_i u_i = 1$), Eqn. (53) gives $1/a_i^2 = N/2$, because there are $N/2$ Fourier components in a time series of N steps. This confirms consistency between Eqns. (49) and (50).

In the flux component according to Eqn. (50) described so far, any constant term is missing. This is because long-term fluxes will be treated as separate flux components. Obviously, such a long-term component has only one time element ($N_{t,i} = 1$) which is simply constant, $g_{1,i}^{\text{time}}(t) = a_i$. However, the value of a_i needs some attention. If this long-term component i is meant to be the missing constant term of some variability component i' of the form Eqn. (50), then $a_i = a_{i'}/\sqrt{2}$, where $a_{i'}$ is the normalization constant from Eqn. (53) of that other component, and the factor $1/\sqrt{2}$ follows from Eqn. (49). In practice, the other component i' does not necessarily need to exist, but it is important to note that a_i constructed in this way depends on the length $T = t_e - t_i$ of the time series in the same way as the implied standard deviation of the long-term flux when using a conventional specification of time variability as a sequence of emission pulses: namely, proportionally to $1/\sqrt{T}$. Such a dependence on T is essential to ensure that the behaviour of the inversion is independent on the length (remember that the number of data constraints also increases with T). It is further important to note that both a_i and the implied long-term standard deviation of a conventional term depend on the temporal correlations (if the instantaneous standard deviation is kept constant). This might be desired, but if not, then a_i can just be chosen to correspond to some imaginary variability component i' the power spectrum of which is chosen independently from any existing variability components, e.g., using $w_i^{(\text{norm})}(\nu)$ from Eqn. (54). This freedom is one of the advantages of time scale separation over the conventional way imposing fixed relations between long-term behaviour and anomalies.

An intuitive objection against Fourier modes might be that they are non-local in nature. However, as shown, a (filtered) Fourier series has the same statistical properties as a conventional (time-correlated) a-priori term defined in the time domain. The only exception is that the Fourier decomposition imposes periodicity (i.e., start and end values are equal), which can be avoided in the time domain by special treatment of the borders. To decouple start and end values in the Fourier case, and to avoid undue influence of the spin-up period onto the long-term mean or mean seasonal cycle, special precaution is taken by adding an initialization flux component described in Sect. 3.3.6. On the other hand, Fourier modes have the advantage that the number of adjustable parameters only depends on the highest included frequency, which can be considerably fewer than in the time domain where each time step requires its own parameters. Despite this lower number of parameters, the resulting time series will be smooth.

3.3.3 Spatial/temporal weighting: $f_{\text{sh},i}(x, y, t)$

According to Eqn. (43), the shape $\alpha_i f_{\text{sh},i}(x, y, t)$ determines the local/instantaneous a-priori standard deviation of the adjustable term, implying a weighting: Signals in the data will preferably be projected into locations/times with large a-priori standard deviation. Therefore, the shape is chosen to reflect the spatio-temporal distribution of the action of the process thought to underly the flux component. Depending on the intention, this can be done in different ways, e.g., by statistical analysis of a suitable –possibly filtered– input data set, or from proxy information.

If the interest focuses on the *time variations* of fluxes at the individual locations (or within certain areas), the appropriate shape is constant in time, $\alpha_i f_{\text{sh},i}(x, y, t) = \alpha_i \phi_i(x, y)$, where the spatial pattern $\alpha_i \phi_i(x, y)$ gives the amplitude (time-wise standard deviation with respect to the autocorrelation implied by the chosen $\{g_{m_t,i}^{\text{time}}(t)\}$) at each location. A natural way to assign such a spatial pattern, then, is to use the time-wise standard deviation of a suitable input data set $f_{\text{input}}(x, y, t)$, the statistical properties of which are assumed to represent the statistical properties of the considered flux process. The amplitude is conveniently calculated as the square root of its total spectral power,

$$\phi_{\text{input}}^{(\text{filter})}(x, y) = \sqrt{\sum_{k=1}^{\infty} \left(w^{(\text{filter})}(\nu_k) f_{\text{input}}^{\cos,k}(x, y) \right)^2 + \sum_{k=1}^{\infty} \left(w^{(\text{filter})}(\nu_k) f_{\text{input}}^{\sin,k}(x, y) \right)^2}, \quad (55)$$

where $f_{\text{input}}^{\cos,k}(x, y)$ and $f_{\text{input}}^{\sin,k}(x, y)$ are the Fourier coefficients of $f_{\text{input}}(x, y, t)$ from Eqns. (108) and (109). To set this amplitude into perspective: For a purely sinusoidal input (and no filter damping, $w = 1$), it would be related to the peak-to-peak amplitude by

$$\phi_{\text{input}}(x, y) = \frac{1}{2} \left(\max_t \{f_{\text{input}}(x, y, t)\} - \min_t \{f_{\text{input}}(x, y, t)\} \right). \quad (56)$$

The filter weights $w^{(\text{filter})}(\nu)$ included here can serve one of two purposes:

- If time scales were separated (i.e., if the considered adjustable term is meant to represent only a certain part of the temporal variability of the underlying process), then the input data set $f_{\text{input}}(x, y, t)$ needs to be filtered with the same weights as used in the adjustable term, in order to calculate the time-wise standard deviations only from these very frequencies.
- If the input data set $f_{\text{input}}(x, y, t)$ does only provide lower frequencies (e.g., when only monthly values are available), then it needs to be filtered as explained in Sect. 3.2. In that case, the normalization of the adjustable term has to be done using Eqn. (54), with $w_i^{(\text{norm})}(\nu) = w^{(\text{filter})}(\nu)$.

The amplitudes could further be scaled by the scalar α_i , if there is extra information indicating that the input data set has too little or too large variability (or for sensitivity testing). Otherwise, $\alpha_i = 1$ is appropriate.

It should be noted that this treatment sets the *local* amplitude (time-wise standard deviation) to that found in the input data set. However, this does not automatically ensure the same correspondence between amplitudes at larger spatial scales (e.g., for fluxes averaged over larger areas, or for the global flux). This is only the case to the extent that the spatial correlation structure chosen for the flux component via $\{g_{m_s, i}^{\text{space}}(x, y)\}$ is representative for the spatial correlation structure of the input data set $f_{\text{input}}(x, y, t)$ (the underlying process, respectively). Theoretically, the spatial correlation structure of $f_{\text{input}}(x, y, t)$ could be determined numerically, but this would be quite expensive, and it is questionable how well it would represent the true correlation structure of the process. As an easier fix, the scaling α_i can be adjusted, as discussed in Sect. 3.3.4.

Though the described approach seems attractive, suitable input data sets are often not available. Another possibility to define the shape $\alpha_i f_{\text{sh}, i}(x, y, t)$ is to employ proxy information. For example, for a CO₂ component related to terrestrial vegetation, an appropriate shape will be zero over deserts and oceans, but high in rain-forests – as the long-term mean of NPP has this property, it could serve as a (crude) proxy (Rödenbeck et al., 2003). If no solid information is available, a flat shape (constant in space) may be most appropriate. In both cases, the scaling factor α_i is still free, and needs to be defined by the methods of Sect. 3.3.4.

3.3.4 Scaling: α_i

The final element to be defined is the overall scaling factor α_i for the flux component. On the one hand, this scaling has the largest impact on the results, as will be shown in Sect. 4.2 below. On the other hand, unfortunately, it is the least known of all flux model elements, for a number of reasons:

- In some cases (shape obtained from input standard deviations, Sects. 3.3.3 above, or as a relative error, Sects. 3.3.5 below), α_i has a conceptual meaning which defines its value, but in many other cases (e.g., shape proportional to some proxy), additional information is required.
- The flux estimates do not actually depend on the a-priori standard deviations themselves, but on the *ratios* between the standard deviations of different flux components (land vs. ocean, Tropics vs. high latitudes, etc.), as well as between a-priori fluxes and data (covariance matrix $\mathbf{Q}_m$). Thus, unrealistic scaling in only one flux component, or in the model-data mismatch uncertainty, spoil the set-up as a whole.
- As mentioned before, due to the normalizations (41) and (42) the shape determines the a-priori standard deviations specifically at the local/instantaneous scales, while the a-priori standard deviations on all other scales then follow according to the chosen correlation structure. However, correlation structure itself is not too well-known in most practical cases. Normalization could also be made referring to other scales (as done, e.g., when using Eqn. (54)), but then the local/instantaneous normalization is lost instead. It is not always obvious which scale should be chosen. However, picking the spatial/temporal scale of normalization is essentially equivalent to changing the scaling factor α_i .
- The pure existence of a Bayesian a-priori term in the cost function is damping all variability. This is because, by construction of the Bayesian formalism, a zero adjustable term is declared as most likely a-priori, while here the adjustable term carries the actual variability of interest. Suppose that the flux

model $\mathbf{F}$ would contain enough degrees of freedom to model the true flux $\mathbf{f}_{\text{true}}$, that the true parameters $\mathbf{p}_{\text{true}}$ to obtain this true flux would be known, that the transport model $\mathbf{A}$ would be perfect, and that there is no measurement error, such that $\mathbf{c}_{\text{obs}} = \mathbf{c}_{\text{true}} = \mathbf{A}\mathbf{F}\mathbf{p}_{\text{true}}$. Then the retrieved parameters would be (cmp. Sect. 1.4)

$$\mathbf{p}_{\text{post}} = (\mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} + \mu \mathbf{1})^{-1} \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{p}_{\text{true}} \quad (57)$$

(where $\mathbf{f}_{\text{fix}} = 0$ has been assumed for simplicity). This would only coincide with the true parameters in the limit $\mu \rightarrow 0$, which however cannot be taken as the inverse $(\mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F})^{-1}$ does not exist (after all, this is the very reason to consider using a-priori constraints).⁶ However, one may try to reduce this damping by choosing μ smaller than unity, or, equivalently, increase (some of) the α_i (such an artificial increase of a-priori uncertainties over what seems to result from physical arguments has been termed ‘reluctant Bayesian’).

In order to solve these problems, there are two groups of remedies, one of them based on additional external information, the other on internal consistency arguments:

Scaling according to given scalars. If a-priori information about some (generally larger-scale) feature of the flux is available, such as the scalars mentioned in Sect. 1.6, α_i can be used to set the corresponding a-priori standard deviation. Examples are:

- Standard deviation of the *long-term global flux*,

$$\sigma_i^{\text{lt, glob}} = \sqrt{\frac{T}{T_\sigma} \left\langle \left(\frac{1}{T} \int_{t_i}^{t_e} \iint_{\text{Globe}} \delta f_i(x, y, t) \, dx dy dt \right)^2 \right\rangle} \quad (58)$$

(T_σ is the time interval the standard deviation refers to). Scaling with respect to $\sigma_i^{\text{lt, glob}}$ makes sense for long-term flux components.

- Standard deviation of the *filtered global flux* at a reference time t_r ,

$$\sigma_i^{(\text{filter}), \text{glob}} = \sqrt{\left\langle \left(\int_{t_i}^{t_e} z^{(\text{filter})}(t|t_r) \iint_{\text{Globe}} \delta f_i(x, y, t) \, dx dy dt \right)^2 \right\rangle}, \quad (59)$$

where $z^{(\text{filter})}(t|t_r)$ is the pulse response function (with respect to time t_r) of the chosen filter, assumed to be normalized as $\int z(t|t_r) dt = 1$. Scaling with respect to $\sigma_i^{(\text{filter}), \text{glob}}$ may be appropriate for flux components representing variability.

For interannual variability, an appropriate filter is the IAV filter of Eqn. (115) with $\nu^{\text{sigma}} = 0.5/\text{yr}$ (‘Filt0.5gd’), which is also used then to smooth the time series plots. The choice of the reference time t_r is generally arbitrary, as long as the shape $f_{\text{sh}, i}(x, y, t)$ is time-constant and t_r is sufficiently far from t_i and t_e to avoid edge effects. Here, $t_r = t_e - 1.5 \text{ yr}$ (i.e., in the middle of the one-but-last year) was taken. The standard deviation of the so-filtered flux at t_r will be denoted by σ_i^{IAV} .

In the numerical implementation, the pulse response $z^{\text{IAV}}(t|t_r)$ is calculated by Fourier transformation of the Dirac $\delta(t|t_r)$, weighting by $w^{\text{gd}}(\nu; \nu^{\text{sigma}})$ from Eqn. (115), and back-transformation.

- Most generally, the standard deviation of any *linear functional* of the flux,

$$\sigma_i^{(z)} = \sqrt{\left\langle \left(\int_{t_i}^{t_e} \iint_{\text{Globe}} z(x, y, t) \delta f_i(x, y, t) \, dx dy dt \right)^2 \right\rangle}, \quad (60)$$

e.g., filtered flux from a certain region.

⁶In a simple conceptual example one could sample a sine wave of a given amplitude at N time instants and retrieve the amplitude with a Bayesian zero prior. The estimated amplitude will always be smaller than the true one, unless in the limit of infinite a-priori amplitude. The underestimation becomes less with increasing N .

Again, if α_i is chosen such that a certain one of these standard deviations has a given value, then the a-priori standard deviations of the *local/instantaneous* flux depend on the chosen spatial and temporal correlation structure (i.e., decompositions $\{g_{m_t,i}^{\text{time}}(t)\}$ and $\{g_{m_s,i}^{\text{space}}(x,y)\}$). This has to be kept in mind when varying correlation lengths during sensitivity testing: A-priori standard deviations will be constant only with respect to the particular space and time scales determined by the above choice of scalar, but will vary on any larger or smaller scales (cmp. discussion after Eqn. (54)).

Mathematical criteria. On a more formal level, there are theoretical criteria whether the assumed uncertainty structure and the data information are compatible. Such criteria can, in turn, also be used to choose parameters in the flux model. The following possibilities have been suggested in the CO₂ inversion literature:

- The doubled cost function $2J$ at the minimum should be statistically χ^2 -distributed with N degrees of freedom (Tarantola, 1987). If the actually obtained value is unlikely with respect to this distribution, the assumed uncertainties of model-data mismatch and/or of a-priori fluxes may be inadequate. Then, all scaling factors α_i simultaneously (or, equivalently, the overall scaling μ) could be adjusted to enforce $2J_{\text{post}}/N \approx 1$ (Press et al., 1992; Rayner et al., 1999).
- Krakauer et al. (2004) suggested Generalized Cross-Validation (e.g., Hansen, 1998) to determine individual scaling factors α_i (or other ‘flux model’ parameters).
- Michalak et al. (2005) proposed a Maximum Likelihood approach to determine a-priori variances (equivalent to scaling factors α_i in the ‘flux model’). [In a previous study (Michalak et al., 2004), the application of methods from Geostatistics has been proposed to also estimate parameters like correlation lengths etc.]

3.3.5 Special flux components

For some flux components, the adjustable term may play the role of a relative error to the fixed term. Then $f_{\text{sh},i}(x,y,t) = f_{\text{fix},i}(x,y,t)$, and α_i is given by the relative error. In the simplest case, only one global and constant scaling parameter is estimated ($N_{t,i} = 1$, $N_{s,i} = 1$, $g_{1,i}^{\text{time}}(t) = 1$, $g_{1,i}^{\text{space}}(x,y) = 1$). If the relative error is assumed to vary with space and/or time, then again the decompositions of Sects. 3.3.1 and/or 3.3.2 apply. A potential example is fossil fuel CO₂ emissions.

3.3.6 The initialization flux component

The first few months of the target period are special because the fluxes estimated there also serve to obtain the ‘right’ implicit 3D initial concentration field in the model atmosphere (remember that the initial concentration is just set uniformly to some crude value given by Eqn. (2)). Therefore, a special flux component i_0 was added, that only acts during the first half a year, with $N_{t,i_0} = 6$ independent emission pulses each of which is linearly declining during the first 1, 2, ..., or 6 months after initial time t_i :

$$g_{m_t,i_0}^{\text{time}}(t) = \begin{cases} 1 - (t - t_i)/(m_t/12)\text{yr} & (t < t_i + (m_t/12)\text{yr}) \\ 0 & (\text{else}). \end{cases} \quad (61)$$

Spatially, the flux component is divided into emissions from a regular grid of $\approx 8^\circ$ latitude $\times 10^\circ$ longitude pixels (as the ocean grid Fig. 6, but global coverage), spatially correlated at short scale (0.05 Earth radii). The shape has no structure, $f_{\text{sh},i_0}(x,y,t) = \text{const}$, and the scaling α_{i_0} is determined by setting an uncertainty $\sigma_{i_0}^{\text{lt, glob}}$ for the total emissions from this component and using Eqn. (58).

3.4 A flux model for CO₂

3.4.1 Standard case

The statistical ‘flux model’ described here is similar to the inversion set-up of Rödenbeck et al. (2003). It mainly differs in that less a-priori information has been used, by separating land and ocean fluxes into 3 components

each (representing long-term, mean seasonal, and interannual variations), while avoiding the fixed terms except for a time-mean pattern.

The following list gives the specific choices for the individual components in a standard set-up. At a number of places, details had to be specified in an ad hoc fashion for lack of solid a-priori information. Many of these choices, indicated by (★), will be varied in the sensitivity cases introduced below.

NEE (long-term):

- The fixed term $f_{\text{fix,nee,lt}}(x, y, t)$ equals the 1980–1992 mean NEE flux from the CCMLP-LPJ biosphere model (Sitch et al., 2000; McGuire et al., 2001). Despite the general intention to avoid fixed terms, a time-mean pattern is taken here because of the limited trust in the inversion’s ability to get long-term fluxes right (★).
- The adjustable term has a spatial weighting (shape $f_{\text{sh,nee,lt}}(x, y, t)$, see Fig. 5) proportional to the mean NPP of CCMLP-LPJ. This NPP weighting is taken as a crude proxy for vegetation activity.
- According to the discussion in Sect. 3.3.2 above, the single (constant) temporal element $g_{1,\text{nee,lt}}^{\text{time}}$ is chosen to refer to a Fourier series, which is here assumed to contain all frequencies slower than 1/yr (unweighted). Since the Fourier spectrum actually used in the interannual NEE term contains much more frequencies (see below), this choice leads to enhanced uncertainties for the long-term flux (due to the normalization of total spectral power).
- Spatial correlations arise from exponentially decaying elements $g_{m_s,\text{nee,lt}}^{\text{space}}(x, y)$ with e-folding length of 0.075 Earth radii (480 km) in both latitude and longitude directions (Fig. 7) (★). There are spatial elements for each land pixel of the transport model ($\approx 4^\circ$ latitude $\times 5^\circ$ longitude, Fig. 6).
- Scaling $\alpha_{\text{nee,lt}} = 1.2$ (★) as in the interannual NEE component (see below).

NEE (seasonal): Although temporal variability is already covered by the ‘interannual’ component below, the ‘seasonal’ component provides additional freedom for the large-scale mean seasonal cycle (★). One rationale is that geographic areas of large/small seasonality need not at all coincide with areas of large/small interannual variability (i.e., the power spectrum –with respect to the ratio of seasonal and nonseasonal frequencies– is space dependent, which cannot be accommodated within one flux component of the form (38)). Further, as seasonality is assumed to be coherent over larger spatial ranges (especially on land along latitude circles), spatial correlations can be chosen longer than those for IAV. Thus, this component may improve constraints in regions far from atmospheric measurements.

- This component has no fixed term, $f_{\text{fix,nee,seas}}(x, y, t) = 0$, because a fixed NEE term already exists in the non-seasonal component.
- The shape $f_{\text{sh,nee,seas}}(x, y, t)$ of the adjustable term is constant in time, its spatial structure (Fig. 5) equals the amplitude of the mean seasonal cycle of NEE in the CCMLP-LPJ simulations, low-pass filtered with $\nu^{\text{low}} = 6/\text{yr}$, Eqn. (55).
- The time correlation of the adjustable term are defined in terms of its Fourier spectrum (containing seasonal frequencies only), weighted by a low-pass filter with $\nu^{\text{low}} = 6/\text{yr}$ (Fig. 9). Additionally, high frequencies are dampened proportional to ν^{-1} , which roughly corresponds to the power spectrum of NEE in CCMLP-LPJ. The normalization of total spectral power involves an additional low-pass filter with $\nu^{\text{low}} = 6/\text{yr}$, which enhances uncertainties for the seasonal component.
- The spatial coherence is similar to that of the long-term (and interannual) NEE components, except that correlation lengths in longitude direction are 3 times as long.
- No additional scaling, $\alpha_{\text{nee,seas}} = 1$.

NEE (interannual): Despite the name ‘interannual’, this is the component to accommodate the temporal variability.

- Again, no fixed term.

- The shape $f_{\text{sh,nee,iav}}(x, y, t)$ is NPP weighted, as in the long-term NEE component.
- The time correlation corresponds to the Fourier spectrum of a low-pass filter with $\nu^{\text{low}} = 12/\text{yr}$ (i.e., it contains frequencies up to half-monthly) (★). As the main rationale, this choice does not contain any frequencies higher than the weekly or two-weekly sampling density of most flask sites. It represents a compromise between the need to avoid aliasing of the sampling schedule into the higher frequencies, and the desire to minimize aggregation errors from not having enough degrees of freedom on short time scales. Apart from the low-pass filter, the spectrum is ‘white’, i.e., no a-priori damping is prescribed. Moreover, it also contains the seasonal frequencies (again), in order to have a full spectrum even where the seasonal component is very small (e.g. in rain forests) and at the shorter spatial correlation.
- Spatial coherence like long-term NEE.
- The scaling is $\alpha_{\text{nee,iav}} = 1.2$ (★). It is chosen rather ad-hoc, leading to a standard deviation of IAV filtered global flux of $\approx 3.6 \text{ PgC/yr}$.

Ocean exchange (long-term):

- The fixed term $f_{\text{fix,oc,lt}}(x, y, t)$ (★) is constant in time, with large-scale values from the ocean inversion by Gloor et al. (2003). Small-scale spatial patterns are proportional to the $\Delta p\text{CO}_2$ based flux estimates by Takahashi et al. (2002).
- In the adjustable term, the shape $f_{\text{sh,oc,lt}}(x, y, t)$ is chosen constant and flat (Fig. 5), scaled to $0.3 \mu\text{mol/m}^2/\text{s}$ (★).
- Normalization of the temporal element as in the long-term NEE component (see above).
- The spatial correlations are defined analogously to NEE, though with an e -folding length of 0.1 Earth radii (630 km) (★). The longer correlation range allows a twice as coarse grid (aggregated to $\approx 8^\circ$ latitude $\times 10^\circ$ longitude, Fig. 6), in order to reduce the number of adjustable parameters.

Ocean exchange (seasonal): An extra seasonal component has been specified also for the ocean exchange (★), although benefits are less clear because ocean seasonality is much less pronounced and much less coherent spatially, than on land.

- There is no fixed term, $f_{\text{fix,oc,seas}}(x, y, t) = 0$.
- The shape $f_{\text{sh,oc,seas}}(x, y, t)$ equals the seasonal cycle amplitude of the $\Delta p\text{CO}_2$ based ocean flux by Takahashi et al. (2002) [data version including the correction for 10 m winds].
- Time structure identical to that of seasonal NEE.
- Spatial structure identical to that of long-term ocean exchange, but twice as long correlations in longitude.

Ocean exchange (interannual):

- Again, no fixed term, $f_{\text{fix,oc,iav}}(x, y, t) = 0$.
- The shape $f_{\text{sh,nee,iav}}(x, y, t)$ is identical to that of the long-term ocean component.
- Time correlation as interannual NEE component (for the same rationale, which is not process-related).
- The spatial coherence like long-term ocean component.

Fossil fuel emissions: Fossil fuel emissions are generally considered much better known than the NEE or Ocean components. Moreover, they have a very peaked spatial structure that surely cannot be resolved by the atmospheric data. Therefore, emissions estimates are taken from fuel use statistics.

- Pre-subtracted fossil fuel emissions $f_{\text{fix,ff}}(x, y, t)$: Yearly totals and geographical distribution according to global statistics from the EDGAR data base (Olivier et al., 2001). After year 2001, extrapolation of the 2001 pattern using a global factor of $\approx 1.4\%/yr$ (linear increase of 0.1035 PgC/yr).

- No adjustable term.

Initialization component: As described in Sect. 3.3.6, with $\sigma_{i_0}^{\text{lt, glob}} = 4 \text{ PgC/yr}$ (with respect to a time scale of $T_\sigma = 1 \text{ yr}$).

3.4.2 Sensitivity tests

Individual choices of the set-up were changed to test their impact:

‘all2’: All a-priori standard deviations multiplied by 2 ($\mu = 0.25$).

‘all0.5’: All a-priori standard deviations divided by 2 ($\mu = 4$).

‘ld2’: A-priori standard deviations of land flux components multiplied by 2.

‘ld0.5’: A-priori standard deviations of land flux components divided by 2.

‘oc2’: A-priori standard deviations of ocean flux components multiplied by 2.

‘oc0.5’: A-priori standard deviations of ocean flux components divided by 2.

‘long’: Spatial correlation lengths of all components longer (Fig. 10):

- 4 times in longitude for long-term and interannual
- 2 times in latitude for long-term and interannual
- 3 times for seasonal NEE
- 2 times for seasonal Ocean exchange

(amplitude normalized to ensure σ_i^{IAV} [Sect. 3.3.4] identical to standard, by $\mu = 5$).

‘slow’: Stronger damping of high frequencies (Fig. 9): $\nu^{\text{low}} = 6/\text{yr}$ in both interannual components (normalized to ensure σ_i^{IAV} identical to standard, by $\mu = 2$).

‘fast’: Allow more high frequencies (Fig. 9): $\nu^{\text{low}} = 26/\text{yr}$ in both interannual components (normalized to ensure σ_i^{IAV} identical to standard, by $\mu = 0.46$).

‘noFix’: No fixed terms (except fossil fuel).

‘noSeas’: No additional seasonal components.

4 Results: CO₂ flux estimates

4.1 Standard results

The algorithm described here is illustrated by the resulting surface-atmosphere CO₂ flux estimates. Fig. 12 shows the interannual variations of regionally integrated fluxes (deseasonalized and filtered as described in Sect. 4.4). The full temporal variability is seen from Fig. 13. All these runs use the standard data treatment of Sect. 2, as well as the standard flux model of Sect. 3.4.1. They differ in their sets of observation sites (see Fig. 4). For the larger sets, differences on the displayed spatial resolution into 22 continental-scale regions are relatively small, suggesting that there is already some beneficial redundancy, and the possibility that more spatial detail could be resolved.

4.2 Uncertainty

There is two kinds of measures for the uncertainty in the flux estimates, both shown in Fig. 14 (IAV) and Fig. 15 (full variability):

A-posteriori standard deviations. The width of the gray shading around the standard case gives the standard deviation of the interannual flux anomaly (i.e., of the variability flux components only) at reference time t_r (the standard deviation is time dependent in principle, but variations are small). This standard deviation reflects the propagation of model-data mismatch errors and a-priori flux errors onto the result, according to their respective assumed magnitudes, and given the transport and flux models. As these magnitudes are rather arbitrary (cmp. in particular Sect. 3.3.4), the a-posteriori standard deviations *only* have a meaning *relative* to the corresponding a-priori standard deviations (given in Fig. 14 as error bar around the black line). The larger the reduction in standard deviation from a-priori to a-posteriori, the better constrained by data the region is.

Sensitivity testing. The ensemble of thin lines illustrates the spread in fluxes when ‘flux model’ settings are varied according to the list in Sect. 3.4.2. This ensemble has been devised by changing individual elements (scaling of flux components, correlation lengths, etc.) in both possible directions around the standard case. Settings that, according to other criteria, seem less trustworthy than the standard, have been marked by broken lines. All solid-line cases seem as reasonable as the standard case. Selected examples with rather large deviations from the standard have been highlighted by color.

4.3 A-posteriori correlations

In contrast to the a-posteriori standard deviations, a-posteriori correlation coefficients between scalars (e.g. regional integrals) of the flux field (Fig. 16) have a direct interpretation. The most favorite situation are small correlation coefficients, indicating that the pair of scalars is constrained independently from one another by the data. A negative value means anti-correlation: For example, the sum of two anti-correlated regions is, in terms of relative uncertainty, better constrained than the regions individually. A positive value indicates that both regions depend on the same signals in the data, and that information distinguishing them is potentially missing.

As a new piece of information with respect to Rödenbeck et al. (2003), Fig. 16 gives separate sets of correlation coefficients for *interannual anomalies* (upper-left block) and *long-term fluxes* (lower-right block) among the TransCom-3 regions, as well as the correlation coefficients between anomaly and long-term mean (identical lower-left or upper-right blocks). These latter are seen to be small, suggesting that temporal and long-term (spatial) signals are largely decoupled. Thus, chances are that the larger errors in the long-term values do not spoil the temporal variability estimated. The split between anomaly and average has become possible through the time scale separation in the flux model, such that correlations between the individual flux components could easily be calculated.

4.4 Remark: Time series filter

Apart from the actual sensitivity of the results to error sources and choices made, the figures also depend on the degree of smoothing and filtering. For time series, Fig. 17 compares three different choices:

- The widely used Box-car smoother (running yearly sums) has the advantage of an intuitive interpretation as an annual total, and of its built-in deseasonalizing effect. However, it fails to remove much of the high-frequency variability not actually desired in IAV plots. Even more importantly, due to its bad spectral properties (involving negative weights for many Fourier components, including all between the first and second seasonal harmonic), it flips the phase of some anomalies, such as those in 1997 and 1998 in Europe.
- The Triangular filter ‘Filt1.0T’ (spectral weights $w^{\text{Td}}(\nu|\nu^{\text{low}})$ from Eqn. (118) with $\nu^{\text{low}} = 1/\text{yr}$, corresponding in the time domain to an approximately triangular pulse response function with HWHH of 1 yr) also deseasonalizes by itself, while having only non-negative weights in frequency space. High-frequency variability is efficiently dampened. However, it also considerably dampens interesting IAV frequencies, even in the frequency range below $0.5/\text{yr}$.
- The standard filter ‘Filt0.5gd’ used in all time series figures has Gaussian spectral weights, Eqn. (115), with $\nu^{\text{sigma}} = 0.5/\text{yr}$. It does not damp as heavily as the triangular filter. On IAV time scales, it is seen to behave quite similar to the intuitive box-bar filter, while avoiding its disadvantages. As a drawback, seasonal frequencies have to be removed explicitly, though all but the main seasonal harmonic ($\nu = 1/\text{yr}$) are essentially damped away already (Fig. 18). A further disadvantage is that the pulse response function extends over the full time domain, which increases computation demands on convolution.

5 Iterative solution

This is a technical section on the numerical minimization of

$$J(\mathbf{p}) = \frac{1}{2}\mathbf{p}^T \mathcal{A} \mathbf{p} - \mathbf{b}^T \mathbf{p} + C' \quad (16')$$

with respect to $\mathbf{p}$, which is equivalent to finding the solution of

$$\mathbf{p} = \mathcal{A}^{-1} \mathbf{b}. \quad (21')$$

Due to the high temporal resolution, the number of data values and parameters is very large, such that the involved matrices cannot be handled any more by present-day computers. An iterative solution of the cost function minimization remedies this situation because it only involves vectors (and products of matrices and vectors), but no explicit storage of the matrix $\mathcal{A}$ itself in memory. As the gradient of the cost function is available via the adjoint of the transport model, the conjugate gradient algorithm was chosen (Press et al., 1992). Standard implementations of this algorithm are generally meant to minimize non-linear functions $J(\mathbf{p})$; in every iteration, they require several function calls to minimize the non-linear function along the current search direction. As, in the present case, each function call involves a rather expensive transport model run, the actual linearity of the function was exploited to perform only a single function call per iteration, as described in the following. However, the large number of parameters leads to round-off problems. These had to be counteracted by additional correction steps. Further, this section describes diagnostical quantities used to measure progress and consistency of the iteration, ways to calculate a-posteriori covariance matrices (approximation of $\mathcal{A}^{-1}$), and to vary the parameter μ in $\mathcal{A}$ for sensitivity testing.

5.1 Basic conjugate gradient algorithm

The conjugate gradient algorithm minimizes the cost function $J(\mathbf{p})$ successively along certain directions $\mathbf{h}_i$ (Press et al., 1992). Any iteration step i ($i = 1, \dots$) arrives at the parameter value $\mathbf{p}_i$ by minimization along $\mathbf{h}_{i-1}$,

$$\mathbf{p}_i = \mathbf{p}_{i-1} + x_i \mathbf{h}_{i-1} \quad (62)$$

with $\mathbf{p}_{i-1}$ and $\mathbf{h}_{i-1}$ from the previous step (and the step length x_i yet to be found). The new direction for the next step is

$$\mathbf{h}_i = \mathbf{g}_i + \gamma_i \mathbf{h}_{i-1} \quad (63)$$

where $\mathbf{g}_i = \mathbf{g}(\mathbf{p}_i)$ is the negative gradient

$$\begin{aligned} \mathbf{g}(\mathbf{p}) &:= -\frac{\partial J(\mathbf{p})}{\partial \mathbf{p}^T} \\ &= -\mathcal{A} \mathbf{p} + \mathbf{b} \end{aligned} \quad (64)$$

at point $\mathbf{p}_i$. The essence of the method is the factor

$$\gamma_i = \frac{\mathbf{g}_i^T \mathbf{g}_i}{\mathbf{g}_{i-1}^T \mathbf{g}_{i-1}}. \quad (65)$$

It ensures that the directions of the individual iteration steps are pairwise conjugate,

$$\mathbf{h}_i^T \mathcal{A} \mathbf{h}_j = 0 \quad \forall i \neq j. \quad (66)$$

This conjugacy has an illustrative interpretation: In the absence of the ‘prior term’ (the $\mathbf{p}^T \mathbf{p}$ in J , or the $\mathbf{1}$ in $\mathcal{A}$, Eqn. (17), which could be neglected if a-priori uncertainties are large), conjugacy of the parameter directions $\mathbf{h}_i$ implies orthogonality of the corresponding directions $\mathbf{Q}_m^{-1/2} \mathbf{A} \mathbf{F} \mathbf{h}_i$ in (normalized) concentration space.

This means that each conjugate step is maximally efficient because it explores completely new concentration information. In addition to the conjugacy of the $\mathbf{h}_i$, the gradients $\mathbf{g}_i$ are orthogonal to each other,

$$\mathbf{g}_i^T \mathbf{g}_j = 0 \quad \forall i \neq j. \quad (67)$$

The algorithm is started at some point $\mathbf{p}_0$ into the direction of the negative gradient there,

$$\mathbf{h}_0 = \mathbf{g}_0. \quad (68)$$

For non-linear problems, the step length x_i would have to be searched for by line-minimization along $\mathbf{h}_{i-1}$. In the linear problem with the purely quadratic cost function, the step length x_i can be found analytically from the condition

$$\frac{\partial}{\partial x_i} J(\mathbf{p}_{i-1} + x_i \mathbf{h}_{i-1}) = 0 \quad (69)$$

which is equivalent to

$$\mathbf{g}_i^T \mathbf{h}_{i-1} = 0. \quad (70)$$

The step length is thus given by

$$x_i = \frac{1}{k_i} \mathbf{h}_{i-1}^T \mathbf{g}_{i-1} \quad (71)$$

with k_i being the curvature of the cost function along $\mathbf{h}_{i-1}$,

$$k_i = \frac{\partial^2}{\partial x_i^2} J(\mathbf{p}_{i-1} + x_i \mathbf{h}_{i-1}) \quad (72)$$

$$= \mathbf{h}_{i-1}^T \mathcal{A} \mathbf{h}_{i-1} \quad (73)$$

The curvature does not depend on $\mathbf{p}$, but only on the direction $\mathbf{h}$. The negative gradient at the new point is

$$\mathbf{g}_i = \mathbf{g}_{i-1} - x_i \mathcal{A} \mathbf{h}_{i-1} \quad (74)$$

The iteration steps can be calculated without having to explicitly generate and store the large matrix $\mathcal{A}$, because all what is needed in Eqns. (73) and (74) is the products

$$\mathbf{l}_i = \mathcal{A} \mathbf{h}_{i-1}. \quad (75)$$

At start-up of the iteration, one also needs the initial gradient

$$\mathbf{g}_0 = -\mathcal{A} \mathbf{p}_0 + \mathbf{b}. \quad (76)$$

For later use, the following matrix notation is introduced. After iteration step i has been completed, all directions explored so far are summarized into the matrix

$$\mathbf{H}_i = (\mathbf{h}_0, \dots, \mathbf{h}_{i-1}), \quad (77)$$

and the corresponding products into

$$\mathbf{L}_i = (\mathbf{l}_1, \dots, \mathbf{l}_i) = \mathcal{A} \mathbf{H}_i. \quad (78)$$

In this matrix notation, the conjugacy simply means that the $(i \times i)$ -matrix

$$\mathbf{K}_i := \mathbf{H}_i^T \mathbf{L}_i \quad (79)$$

is diagonal. Further, the state arrived at by iteration i can be written as

$$\mathbf{p}_i = \mathbf{p}_0 + \mathbf{H}_i \mathbf{x}_i. \quad (80)$$

5.2 Numerical implementation of $\mathcal{A}$

With Eq. (17), the product in Eq. (75) reads

$$\mathbf{l}_i = \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1} + \mathbf{h}_{i-1}. \quad (81)$$

It is obtained by the following steps:

1. $\mathbf{F} \mathbf{h}_{i-1}$ by matrix multiplication (which is not as involved as a full multiplication because $\mathbf{F}$ is defined as a product of block matrices, see Sect. 3);
2. $\mathbf{A} \mathbf{F} \mathbf{h}_{i-1}$ by a forward run of the atmospheric tracer transport model;
3. $\mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1}$ by matrix multiplication (which is fast because $\mathbf{Q}_m$ is assumed diagonal);
4. $\mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1}$ by an adjoint run of the transport model;
5. $\mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1}$ again by matrix multiplication ('adjoint flux model').

The initial gradient

$$\mathbf{g}_0 = \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} (\mathbf{c}_{\text{obs}} - \mathbf{A}[\mathbf{f}_{\text{fix}} + \mathbf{F} \mathbf{p}_0] - \mathbf{c}_{\text{ini}}) - \mathbf{p}_0 \quad (82)$$

is calculated in a similar way.

5.3 Starting point

In most cases, the iteration has been started at $\mathbf{p}_0 = 0$, i.e., from the a-priori fluxes. Then, $\mathbf{h}_0 = \mathbf{g}_0 = \mathbf{b}$, and the first iterate $\mathbf{p}_1$ is proportional to $\mathbf{b}$ which corresponds to approximating the Hessian by its a-priori part, $\mathcal{A} \approx 1$. Many of the method extensions described below take easiest form if $\mathbf{p}_0 = 0$.

There is some potential to save iterations (computing time) by starting at some pre-optimized state, such as an estimate obtained with a transport model on lower spatial resolution. During sensitivity testing, one may also start at the result of some previous run with similar settings (flux model, data, etc.). However, contrary to optimistic expectations, CPU time savings are only moderate. This is because the majority of iterations are typically spent on relatively small changes at finer spatial or temporal scales.

5.4 Diagnostics

5.4.1 Convergence diagnostics

The values $J_i = J(\mathbf{p}_i)$ of the *cost function* itself are not actually needed in the iteration, but can be used as a diagnostic for its progress. By iteration step i , the cost function is diminished by

$$J_i = J_{i-1} - \frac{1}{2} x_i^2 k_i. \quad (83)$$

The absolute value of the cost function is only defined up to an additive constant; at the starting point, it is set to

$$J_0 = \frac{1}{2} (\mathbf{c}_{\text{obs}} - \mathbf{A}[\mathbf{f}_{\text{fix}} + \mathbf{F} \mathbf{p}_0] - \mathbf{c}_{\text{ini}})^T \mathbf{Q}_m^{-1} (\mathbf{c}_{\text{obs}} - \mathbf{A}[\mathbf{f}_{\text{fix}} + \mathbf{F} \mathbf{p}_0] - \mathbf{c}_{\text{ini}}) + \frac{1}{2} \mathbf{p}_0^T \mathbf{p}_0 \quad (84)$$

(see Eqn. (19)), calculated as a by-product of Eqn. (82). The iteration could be terminated if

$$2|J_i - J_{i-1}| \leq \theta (|J_i| + |J_{i-1}|) \quad (85)$$

with a relative tolerance of, e.g., $\theta = 10^{-5}$. This is not an optimal criterion, as it involves some arbitrariness (due to the undefined additive constant in J , and the dependence of Eqn. (84) on $\mathbf{p}_0$), and as it is not directly linked to precision.

A more direct way to measure the iteration's progress, would be to look at the evolution of selected functionals of the flux (cmp. Sect. 1.6). The selection of such functionals should reflect the quantities on interest.




Figure 2: Diagnostical quantities in the course of the iteration. Left: Convergence diagnostics. Right: Consistency diagnostics.

The set chosen here comprises the interannual flux anomaly in all TransCom-3 regions in a certain reference year (specifically, the IAV-filtered flux [standard filter ‘Filt0.5gd’] taken at time $t_e - 1.5$ yr, i.e., in the middle of the one-but-last year), plus the time-mean flux in all TransCom-3 regions (44 functionals in total). The iteration may be terminated if any of the relative flux functionals from Eqn. (25)

$$\frac{\mathbf{f}_i^T \mathbf{z}_j}{\sqrt{\frac{1}{\mu} \mathbf{z}_j^T \mathbf{F} \mathbf{F}^T \mathbf{z}_j}} \quad (86)$$

(normalized by the a-priori standard deviation from Eqn. (26)) attains a stable value. The disadvantage of this criterion is that it is specific to the chosen functionals, though it seems that convergence is quite similar among the relevant choices $\mathbf{z}_j$.

Another, again global, measure of convergence is the length of the negative gradient $\mathbf{g}_i$, because, according to Eqn. (64), it is also the residual of Eqn. (21) to be solved. The iteration may be terminated as soon as the relative *residual index*

$$RI_i = \frac{\mathbf{g}_i^T \mathbf{g}_i}{\mathbf{g}_0^T \mathbf{g}_0} \quad (87)$$

becomes smaller than a given tolerance. Again, this does not provide a fully objective criterion, because it depends on the starting point $\mathbf{p}_0$. If $\mathbf{p}_0 = 0$ (and consequently $\mathbf{g}_0 = \mathbf{b}$), a useful tolerance is 10^{-5} , which has actually been used in the runs presented here. However, the value of the tolerance was chosen based on the previously described convergence indices.




Figure 3: Development of selected scalar functionals of the flux, normalized as in Eqn. (86), in the course of the iteration. Left: Functionals are IAV anomalies in the 22 TransCom-3 regions at some reference time. Right: Functionals are long-term fluxes in the 22 TransCom-3 regions.

Generally, more iterations than necessary are spent. For security, termination is also invoked if the gradient becomes exactly zero, or if more than a maximum number of iterations (e.g., 150) have been performed.

An additional diagnostic of progress is the *null-space index*, defined as

$$NSI_i = \frac{\mathbf{h}_{i-1}^T \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1}}{\mathbf{h}_{i-1}^T \mathbf{h}_{i-1}}. \quad (88)$$

It quantifies the relative contribution of the data part of the cost function in the step along $\mathbf{h}_{i-1}$. It will be zero if the direction $\mathbf{h}_{i-1}$ lies in the null-space of the matrix $\mathbf{M} = \mathbf{Q}_m^{-1/2} \mathbf{A} \mathbf{F}$. A small NSI_i implies that any change along $\mathbf{h}_{i-1}$ in the parameter space only corresponds to a small change $\mathbf{M} \mathbf{h}_{i-1}$ in data space, i.e., this search direction is not well ‘seen’ by the concentration data (given the atmospheric transport and the flux model), but only constrained by the a-priori term. Therefore, if NSI_i stops to decrease, the solution will not be improved substantially any more by further iterations.

5.4.2 Consistency diagnostics

The degree to which directions are conjugate numerically, is measured by the *conjugacy index*

$$CI_i = \max_{j < i} \frac{\mathbf{h}_i^T \mathcal{A} \mathbf{h}_j}{\sqrt{\mathbf{h}_i^T \mathbf{h}_i} \sqrt{\mathbf{h}_j^T \mathbf{h}_j}}. \quad (89)$$

As described later, precautions are necessary to ensure that conjugacy is sufficiently valid.

The algorithm assumes a symmetric matrix $\mathcal{A}$, which is the case if the adjoint model $\mathbf{A}^T$ was a perfect transpose of the forward model $\mathbf{A}$. The *symmetry index*

$$SYI_i = \left| \frac{(\mathbf{Q}_m^{-1/2} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1})^T (\mathbf{Q}_m^{-1/2} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1})}{\mathbf{h}_{i-1}^T \mathbf{F}^T \mathbf{A}^T \mathbf{Q}_m^{-1} \mathbf{A} \mathbf{F} \mathbf{h}_{i-1}} - 1 \right| \quad (90)$$

measures any deviation from this symmetry, because the denominator depends on both the forward and adjoint models, while the –analytically identical– numerator only depends on the forward model. The symmetry index is calculated as a by-product of $\mathbf{l}_i$.

5.5 Approximation of the inverse Hessian matrix

For the calculation of a-posteriori uncertainties, Eqn. (24), as well as for the numerical improvements of the iteration described below, (approximate) solutions $\mathbf{p}^{[j]}$ to

$$\mathcal{A}\mathbf{p}^{[j]} = \mathbf{b}^{[j]} \quad (91)$$

are needed, where the $\mathbf{b}^{[j]}$ represent arbitrary given right-hand sides (not necessarily related to the original right-hand side $\mathbf{b}$ of Eqn. (21)). After iteration i has been completed (i.e., all the directions $\mathbf{H}_i$ have been explored), all available information about $\mathcal{A}$ is contained in $\mathbf{L}_i = \mathcal{A}\mathbf{H}_i$ (and the resulting $\mathbf{K}_i = \mathbf{H}_i^T \mathbf{L}_i$). Analogously to Eqn. (80), the ansatz

$$\mathbf{p}_i^{[j]} = \mathbf{p}_0^{[j]} + \mathbf{H}_i \mathbf{x}^{[j]} \quad (92)$$

is used, where $\mathbf{p}_0^{[j]}$ is a ‘first guess’ about $\mathbf{p}^{[j]}$ (not necessarily related to $\mathbf{p}_0$), and $\mathbf{H}_i \mathbf{x}^{[j]}$ is a correction of $\mathbf{p}_0^{[j]}$ along all the directions $\mathbf{H}_i$ available up to now. If this correction is optimal, then the residual between this approximation $\mathbf{p}_i^{[j]}$ and the true solution $\mathbf{p}^{[j]}$ lies in the space of those directions not yet explored by the iteration; therefore, it has to be conjugate to $\mathbf{H}_i$,

$$\mathbf{H}_i^T \mathcal{A}(\mathbf{p}_i^{[j]} - \mathbf{p}^{[j]}) = 0. \quad (93)$$

This leads to

$$\mathbf{p}_i^{[j]} = \mathbf{p}_0^{[j]} + \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{b}^{[j]} - \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{L}_i^T \mathbf{p}_0^{[j]}. \quad (94)$$

Since $\mathbf{K}_i$ is diagonal due to conjugacy, the inverse $\mathbf{K}_i^{-1}$ is easy to compute.

Below, we also encounter the case that $\mathbf{b}^{[j]}$ is given indirectly through a quantity

$$\mathbf{g}_0^{[j]} = -\mathcal{A}\mathbf{p}_0^{[j]} + \mathbf{b}^{[j]}. \quad (95)$$

Then, Eqn. (94) becomes

$$\mathbf{p}_i^{[j]} = \mathbf{p}_0^{[j]} + \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{g}_0^{[j]}. \quad (96)$$

5.6 Ensuring conjugacy

Though in theory the conjugacy Eqn. (66) follows from the construction of the algorithm, there is a numerical complication, arising from the large condition number of $\mathcal{A}$. The lengths of the vectors $\mathbf{h}_i$ (as well as those of $\mathbf{l}_i$, $\mathbf{g}_i$, and the curvature k_i) change by several orders of magnitude in the course of the iteration. This is then also the ratio between the individual components of these vectors. In the straightforward implementation of the algorithm described above, this leads to roundoff errors, also by summing up smaller and smaller contributions to $\mathbf{p}$ and $\mathbf{g}$. This causes the conjugacy and orthogonality conditions to be increasingly violated. Although the cost function minimum will, nevertheless, still be found, it will take the algorithm a considerably greater number of iterations.

As a remedy, one can enforce conjugacy, as follows. Assume that iteration i has been completed, such that $\mathbf{H}_i$, $\mathbf{L}_i$, and $\mathbf{K}_i$ are available. The current state $\mathbf{p}_i$, the negative gradient $\mathbf{g}_i$, and the cost function value J_i are then recomputed by

$$\mathbf{p}_i = \mathbf{p}_0 + \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{g}_0, \quad (97)$$

$$\mathbf{g}_i = \mathbf{g}_0 - \mathbf{L}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{g}_0, \quad (98)$$

and

$$J_i = J_0 - \frac{1}{2} \mathbf{g}_0^T \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{g}_0 \quad (99)$$

which follows from Eqn. (94) with the initial point $\mathbf{p}_0$ of the iteration as ‘first guess’, together with Eqn. (76). The new direction $\mathbf{h}_i$ is then calculated by

$$\mathbf{h}_i = \left[\mathbf{1} - \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{L}_i^T \right] \mathbf{g}_i \quad (100)$$

which is equivalent to Eqn. (63). The matrix in brackets is the projection onto the space conjugate to all previous directions $\mathbf{H}_i$.

To improve numerical accuracy, the sums implied in the matrix products of Eqns. (97), (98), and (99) are done with decreasing indices (smallest terms first). Further, noting that $(\mathbf{p}_0, \mathbf{g}_0)$ can be replaced by any $(\mathbf{p}_j, \mathbf{g}_j)$ along the course of the iteration, the formulas are applied twice. Moreover, the projection of Eqn. (100) is done in the form

$$\mathbf{h}_i = \prod_{j=0}^{i-1} [\mathbf{1} - \mathbf{h}_j \mathbf{k}_j^{-1} \mathbf{l}_j^T] \mathbf{g}_i \quad (101)$$

with increasing indices j , and repeated, too.

5.7 Approximating the a-posteriori covariance matrix

As the a-posteriori (co)variances from Eqn. (27) involve $\mathcal{A}^{-1}$, they can be obtained from a solution of Eqn. (91), with the r.h.s.

$$\mathbf{b}^{[j]} = \mathbf{F}^T \mathbf{z}_j \implies \mathbf{p}^{[j]} = \mathcal{A}^{-1} \mathbf{F}^T \mathbf{z}_j. \quad (102)$$

Here, the $\mathbf{z}_j$ denote a chosen set of spatio-temporal integration patterns for which (co)variances are to be calculated, such as the scalars mentioned in Sect. 1.6. In the above equation, the adjoint flux model $\mathbf{F}^T$ is applied to each of the patterns $\mathbf{z}_j$.

If the ‘first guess’ $\mathbf{p}_0^{[j]} = 0$ would be used, Eqn. (94) would give the approximation

$$\mathbf{z}_1^T \mathbf{F} \mathcal{A}^{-1} \mathbf{F}^T \mathbf{z}_2 \approx \mathbf{z}_1^T \mathbf{F} \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{F}^T \mathbf{z}_2$$

where $\mathbf{z}_1$ and $\mathbf{z}_2$ stand for two arbitrary ones of the chosen patterns $\mathbf{z}_j$. This approximation will underestimate the (co)variances, because it only contains contributions from the directions contained in $\mathbf{H}_i$, while all other directions appear to be infinitely well constrained. With the ‘first guess’ $\mathbf{p}_0^{[j]} = \frac{1}{\mu} \mathbf{b}^{[j]} = \frac{1}{\mu} \mathbf{F}^T \mathbf{z}_j$, one obtains

$$\mathbf{z}_1^T \mathbf{F} \mathcal{A}^{-1} \mathbf{F}^T \mathbf{z}_2 \approx \frac{1}{\mu} \mathbf{z}_1^T \mathbf{F} \mathbf{F}^T \mathbf{z}_2 - \frac{1}{\mu} \mathbf{z}_1^T \mathbf{F} \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{L}_i^T \mathbf{F}^T \mathbf{z}_2 + \mathbf{z}_1^T \mathbf{F} \mathbf{H}_i \mathbf{K}_i^{-1} \mathbf{H}_i^T \mathbf{F}^T \mathbf{z}_2. \quad (103)$$

This first guess is more relevant because $\mathbf{p}_0^{[j]}$ then corresponds to the approximation $\mathcal{A} \approx \mu \mathbf{1}$, which, by comparison of Eqns. (6) and (24), means $\mathbf{Q}_{f,\text{post}} \approx \mathbf{Q}_{f,\text{pri}}$. Eqn. (103) now corrects this first guess along any directions spanned by $\mathbf{H}_i$, where uncertainty has been reduced by the inversion. Along all other directions of the parameter space, the approximate a-posteriori covariance matrix will still coincide with the a-priori covariance matrix, which is indeed appropriate for all directions not constrained by the data (prior-constrained, small null-space index) and therefore not gone by the iteration. Still, if $\mathbf{H}_i$ misses any relevant directions, Eqn. (103) will overestimate uncertainties. As has been checked with a few-parameter case where uncertainties can be calculated exactly by the traditional SVD, this overestimation is by far not negligible. This means that, unfortunately, the set of directions $\mathbf{h}_i$ sufficient to solve for $\mathbf{p}_{\text{post}}$ does not contain all the directions relevant to solve for $\mathbf{p}^{[j]}$.

However, the approximate (co)variances can be improved to any accuracy by continuing the iteration: That way, information about $\mathcal{A}$ along any missing directions $\mathbf{h}_{i'}$ is obtained. This can be done using the algorithm of Sect. 5.6: Since the recomputed state $\mathbf{p}_i$ as well as the next direction $\mathbf{h}_i$ just depend on the present $\mathbf{H}_i$, $\mathbf{L}_i$, and on the pair of initial state and gradient $(\mathbf{p}_0, \mathbf{g}_0)$, this pair can be exchanged by $(0, \mathbf{b}^{[j]})$ to switch to a new r.h.s. $\mathbf{b}^{[j]}$ (and zero starting point, cmp. Eqn. (76)). Specifically, after completion of the main iteration solving for $\mathbf{p}_{\text{post}}$, a new r.h.s. $\mathbf{b}^{[1]} = \mathbf{F}^T \mathbf{z}_1$ corresponding to the first of the chosen integrations/filters is switched to, and iterated into convergence. This is continued in turn for all $\mathbf{z}_j$. Of course, the number of these patterns has to be small in order to be feasible. However, the number of additional iterations needed decreases from pattern to pattern, because of the information (products $\mathcal{A} \mathbf{h}_i$ along directions $\mathbf{h}_i$) already present.

5.8 Varying the a-priori scaling μ

As discussed in Sect. 1.7, the scaling parameter μ of the Bayesian a-priori term is interesting to be varied in the course of sensitivity testing. There is a way to do this without repeating the expensive full iteration.

Suppose that the iteration has been conducted with value μ of the parameter, and the matrices $\mathbf{H}_i$ and $\mathbf{L}_i$ are available. Changing the parameter into $\tilde{\mu}$ changes the Hessian into

$$\tilde{\mathcal{A}} = \mathcal{A} + (\tilde{\mu} - \mu)\mathbf{1}. \quad (104)$$

If the set of directions is kept, $\tilde{\mathbf{H}}_i = \mathbf{H}_i$, a new matrix replacing $\mathbf{L}_i$ is

$$\begin{aligned} \tilde{\mathbf{L}}_i &= \tilde{\mathcal{A}}\tilde{\mathbf{H}}_i \\ &= \mathbf{L}_i + (\tilde{\mu} - \mu)\mathbf{H}_i. \end{aligned} \quad (105)$$

Of course, the (old) directions $\mathbf{H}_i$ will not be conjugate with respect to the new Hessian $\tilde{\mathcal{A}}$, i.e., the matrix $\tilde{\mathbf{K}}_i = \tilde{\mathbf{H}}_i^T \tilde{\mathbf{L}}_i$ is no longer diagonal. Nevertheless, $\tilde{\mathbf{p}}_i$ can still be computed by Eqn. (97) using the tilde'd quantities (note that $\tilde{\mathbf{g}}_0 = \mathbf{g}_0 - (\tilde{\mu} - \mu)\mathbf{p}_0$). The only difference is that the evaluation of terms containing $\tilde{\mathbf{K}}_i^{-1}$ now requires a library call to some linear systems solver. This is no problem, however, because the dimension of $\tilde{\mathbf{K}}_i$ is only the number of iterations, which is much smaller than the number of parameters. Likewise, the formula of Sect. 5.7 to compute (co)variances can still be used.

The question remains whether $\mathbf{H}_i$ from the converged μ -iteration actually also contains all the search directions necessary for the case of $\tilde{\mu}$. Luckily, as long as a-priori damping is *increased* (i.e., $\tilde{\mu} > \mu$), this requirement is fulfilled, because a larger μ requires less directions. This has also been confirmed by direct comparison⁷.

In order to simultaneously calculate solutions for a variety of $\tilde{\mu}$ values, the actual iteration can even be done with $\mu = 0$ (even though the analytic solution would diverge then). The larger $\tilde{\mu}$, the earlier the corresponding solution will converge (such that the iteration can be stopped according to the smallest $\tilde{\mu}$ of interest). This is also related to the regularizing property of the conjugate gradients method (Hansen, 1998).

⁷Unfortunately, this easy change of scaling only works for the overall scaling parameter μ , not however for the scaling parameters α_i of the individual flux components, or any other detail of the flux model. This is because any change in the ratios of different a-priori constraints would require further state space directions to be sampled.

A Time series filtering

A.1 Fourier decomposition

The construction of both fixed and variability terms in the flux model, as well as the presentation of results, may involve time series filtering. For a given (possibly space-dependent) flux time series $f(x, y, t)$, filtering is done here in terms of weighted Fourier expansions,

$$\begin{aligned} f^{(\text{filter})}(x, y, t) &= w_0^{(\text{filter})} f^{\text{ave}}(x, y) \\ &+ \sum_{k=1}^{\infty} w^{(\text{filter})}(\nu_k) \left(f^{\cos, k}(x, y) \cos(2\pi\nu_k t) + f^{\sin, k}(x, y) \sin(2\pi\nu_k t) \right) \end{aligned} \quad (106)$$

where $w^{(\text{filter})}(\nu)$ is the chosen filter function. The frequencies are

$$\nu_k = k/T \quad (107)$$

(in 1/yr) where $T = t_e - t_i$ is the time period of the inversion (in years). In Eqn. (106), the sum over k does not actually go to ∞ but stops at a highest frequency set by the time resolution of fluxes (here 1 day). The Fourier coefficients are obtained from the original time series $f(x, y, t)$ by

$$f^{\cos, k}(x, y) = \frac{2}{t_2 - t_1} \int_{t_1}^{t_2} \cos(2\pi\nu_k t) f(x, y, t) dt, \quad (108)$$

$$f^{\sin, k}(x, y) = \frac{2}{t_2 - t_1} \int_{t_1}^{t_2} \sin(2\pi\nu_k t) f(x, y, t) dt, \quad (109)$$

if $f(x, y, t)$ is given over the time interval t_1 through t_2 . Generally, t_1 and t_2 coincide with the start and end times of the inversion period, t_i and t_e . (If filtering is used to get mean fluxes or mean seasonal cycles from external input data, and if the inversion period $T = t_e - t_i$ is short, one might want to improve statistics by using a longer part of the input data set $f(x, y, t)$, if available. Also, it might be desirable to fix t_1 and t_2 in order to make the Fourier coefficients independent of the particular inversion period, if inversions of different lengths are to be compared.) However, these formulas only work if

$$2k(t_2 - t_1) \quad (110)$$

is an integer multiple of T (for the seasonal frequencies $\nu_k = 1/\text{yr}, 2/\text{yr}, 3/\text{yr}, \dots$, this condition is sure to be fulfilled). We preferred to use the definitions Eqns. (108) and (109) over the usual FFT method, because they are more simple to be handled given the variety of time intervals and/or sampling distances occurring in available input data sets. They can be calculated efficiently by employing the addition theorems

$$\cos(2\pi\nu_k t) = \cos(2\pi\nu_{k-1} t) \cos(2\pi\nu_1 t) - \sin(2\pi\nu_{k-1} t) \sin(2\pi\nu_1 t) \quad (111)$$

$$\sin(2\pi\nu_k t) = \sin(2\pi\nu_{k-1} t) \cos(2\pi\nu_1 t) + \cos(2\pi\nu_{k-1} t) \sin(2\pi\nu_1 t) \quad (112)$$

in the calculation of all higher-order triangular functions (Press et al., 1992). The long-term average is

$$f^{\text{ave}}(x, y) = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} f(x, y, t) dt. \quad (113)$$

A.2 Specific filters

All filters used in the present calculations are defined here in terms of their spectral weights. They are shown in Fig. 18, together with their corresponding pulse response functions.

Gaussian ('g'):

$$w^g(\nu; \nu^{\text{sigma}}) = \exp\left(-\frac{\nu^2}{2(\nu^{\text{sigma}})^2}\right) \quad (114)$$

Gaussian, deseasonalized ('gd'):

$$w^{\text{gd}}(\nu; \nu^{\text{sigma}}) = \begin{cases} 0, & (\nu = 1/\text{yr}, 2/\text{yr}, 3/\text{yr}, \dots) \\ w^g(\nu; \nu^{\text{sigma}}), & (\text{else}) \end{cases} \quad (115)$$

This is the standard IAV filter in time series figures (then with $\nu^{\text{sigma}} = 0.5/\text{yr}$).

Triangular ('t'): The following low-pass filter (with parameter ν^{low}) corresponds in the time domain to a triangular response function that linearly rises/falls away from the central peak and reaches zero at $\pm 1/\nu^{\text{low}}$.

$$w^t(\nu; \nu^{\text{low}}) = \frac{2}{(2\pi)^2} \left(\frac{\nu^{\text{low}}}{\nu}\right)^2 \left(1 - \cos\left(2\pi \frac{\nu}{\nu^{\text{low}}}\right)\right). \quad (116)$$

It removes essentially all frequencies ν^{low} and higher (the cut-off frequency ν^{low} itself and all its multiples are removed completely, other frequencies higher than ν^{low} pass with at most 5% amplitude).

Truncated triangular ('T'): The standard low-pass filter used in the adjustable terms of the flux model is the 't' one, where additionally all frequencies above ν^{low} are removed completely.

$$w^T(\nu; \nu^{\text{low}}) = \begin{cases} w^t(\nu; \nu^{\text{low}}), & (\nu < \nu^{\text{low}}) \\ 0, & (\text{else}) \end{cases} \quad (117)$$

Truncated triangular, deseasonalized ('Td'):

$$w^{\text{Td}}(\nu; \nu^{\text{low}}) = \begin{cases} 0, & (\nu = 1/\text{yr}, 2/\text{yr}, 3/\text{yr}, \dots) \\ w^T(\nu; \nu^{\text{low}}), & (\text{else}) \end{cases} \quad (118)$$

Truncated triangular, mean seasonal cycle ('Ts'):

$$w^{\text{Ts}}(\nu; \nu^{\text{low}}) = \begin{cases} w^T(\nu; \nu^{\text{low}}), & (\nu = 1/\text{yr}, 2/\text{yr}, 3/\text{yr}, \dots) \\ 0, & (\text{else}) \end{cases} \quad (119)$$

In numerical filtering, spectral components are omitted (filter weights are set zero) if $w(\nu) < 0.01$.

The nomenclature used to refer to these filters in this report, consists of 'Filt' followed by the value of the characteristic frequency (ν^{sigma} or ν^{low} , respectively) and the short-hand label given above. Examples are 'Filt0.5gd' or 'Filt1.0T'.

References

- Bousquet, P., Peylin, P., Ciais, P., Le Quéré, C., Friedlingstein, P., and Tans, P.: Regional changes in carbon dioxide fluxes of land and oceans since 1980, *Science*, **290**, 1342–1346 (2000).
- Conway, T., Tans, P., Waterman, L., Thoning, K., Kitzis, D., Masarie, K., and Zhang, N.: Evidence for interannual variability of the carbon cycle from the national oceanic and atmospheric administration climate monitoring and diagnostics laboratory global air sampling network, *J. Geophys. Res.*, **99**, 22 831–22 855 (1994).
- Francey R.J., Steele L.P., Spencer D.A., Langenfelds R.L., Law R.M., Krummel P.B., Fraser P.J., Etheridge D.M., Derek N., Coram S.A., Cooper L.N., Allison C.E., Porter L., and Baly S.: The CSIRO (Australia) measurement of greenhouse gases in the global atmosphere, report of the 11th WMO/IAEA Meeting of Experts on Carbon Dioxide Concentration and Related Tracer Measurement Techniques, Tokyo, Japan, September 2001, S.Toru and S. Kazuto (editors), World Meteorological Organization Global Atmosphere Watch, 97-111, 2003.
- Gloor, M., Gruber, N., Sarmiento, J., Sabine, C., Feely, R., and Rödenbeck, C.: A first estimate of present and preindustrial air-sea CO₂ flux patterns based on ocean interior carbon measurements and models, *Geophys. Res. Lett.*, **30**, 10.1029/2002GL015 594 (2003).
- Gurney, K., Law, R. M., Denning, A. S., et al.: Towards robust regional estimates of CO₂ sources and sinks using atmospheric transport models, *Nature*, **415**, 626–630 (2002).
- Hansen P.C.: *Rank-Deficient and Discrete Ill-Posed Problems. Numerical Aspects of Linear Inversion*, Soc. for Industrial and Applied Mathematics, Philadelphia, 1998.
- Krakauer N.Y., Schneider T., Randerson J.T., Olsen S.C.: Using generalized cross-validation to select parameters in inversions for regional carbon fluxes, *Geophys. Res. Lett.* **31**, 10.1029/2004GL020323 (2004).
- Labuschagne C., Brunke E.-G., and Scheel H.E.: Global Trace Gas Trends as Observed at the Cape Point Global Atmosphere Watch (GAW) Station. Presented at: Global Change and Regional Sustainability in South Africa. Kirstenbosch, Cape Town, 2003.
- Law R.M., Rayner P.J., Steele L.P., Enting I.G.: Data and modelling requirements for CO₂ inversions using high-frequency data, *Tellus* **55B**, 512 (2003).
- Levin, I., R. Graul, and N.B.A. Trivett: Long term observations of atmospheric CO₂ and carbon isotopes at continental sites in Germany, *Tellus* **47B**, 23-34 (1995).
- Masarie, K., Langenfelds, R., Allison, C., Conway, T., Dlugokencky, E., Francey, R., Novelli, P., Steele, L., Tans, P., Vaughn, B., and White, J.: NOAA/CSIRO flask air intercomparison experiment: A strategy for directly assessing consistency among atmospheric measurements made by independent laboratories, *J. Geophys. Res.*, **106**, 445–464 (2001).
- McGuire A., Sitch S., Clein J. S., Dargaville R., Esser G., Foley J., Heimann M., Joos F., Kaplan J., Kicklighter D.W., Meier R.A., Melillo J.M., Moore III B., Prentice I.C., Ramankutty N., Reichenau T., Schloss A., Tian H., Williams L.J., Wittenberg U.: Carbon balance of the terrestrial biosphere in the 20th century: Analysis of CO₂, climate and land use effects with four process-based ecosystem models, *Global Biogeochemical Cycles*, **15**, 183–206 (2001).
- Michalak A., Bruhwiler L., and Tans P.: A Geostatistical approach to surface flux estimation of atmospheric trace gases, *J. Geophys. Res.* **109**, 10.1029 (2004).
- Michalak A., Hirsch A., Bruhwiler L., Gurney K.R., Peters W., Tans P.P.: Maximum likelihood estimation of covariance parameters for Bayesian atmospheric trace gas surface flux inversions, *J. Geophys. Res.*, submitted (2005).

- Morimoto S., Nakazawa T., Aoki S., Hashida G., and Yamanouchi T.: Concentration variations of atmospheric CO₂ observed at Syowa Station, Antarctica from 1984 to 2000, *Tellus 55B*, 170-177 (2003).
- Olivier, J. G. J. and Berdowski, J. J. M.: Global emissions sources and sinks. In: Berdowski, J., Guicherit, R. and Heij, B. J. (eds.) "The Climate System", A. A. Balkema Publishers/Swets & Zeitlinger Publishers, Lisse, The Netherlands, ISBN 90 5809 255 0, 33–78 (2001).
- Peylin P., Bousquet P., Le Quéré C., Sitch S., Friedlingstein P., McKinley G., Gruber N., Rayner P., Ciais P.: Multiple constraints on regional CO₂ flux variations over land and oceans, *Glob. Biogeochem. Cycles* **19**, 10.1029 (2005).
- Press, W., Teukolsky, S., Vetterling, W., and Flannery, B.: *Numerical Recipes*, Cambridge Univ. Press, 1992.
- Rayner, P., Enting, I., Francey, R., and Langenfelds, R.: Reconstructing the recent carbon cycle from atmospheric CO₂, $\delta^{13}\text{CO}_2$ and O₂/N₂ observations, *Tellus B*, *51*, 213–232 (1999).
- Rödenbeck 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).
- Rödenbeck C., Conway T.J., Langenfelds R.: The effect of systematic measurement errors on atmospheric CO₂ inversions: A quantitative assessment *Atmos. Chem. Phys.*, submitted, (2005).
- Sitch, S., Prentice, I., Smith, B., Cramer, W., Kaplan, J., Lucht, W., Sykes, M.; Thonicke, K., and Venevsky, S.: LPJ – a coupled model of vegetation dynamics and the terrestrial carbon cycle, in "The role of vegetation dynamics in the control of atmospheric CO₂ content" (S. Sitch, PhD Thesis), Lund University, Lund (Sweden), 2000.
- Takahashi, T., Sutherland, S. C., Sweeney, C., Poisson, A., Metzl, N., Tilbrook, B., Bates, N., Wanninkhof, R., Feely, R. A., Sabine, C., Olafsson, J., and Nojiri, Y.: Global sea-air CO₂ flux based on climatological surface ocean pCO₂, and seasonal biological and temperature effects, *Deep Sea Res. II*, *49*, 1601–1623 (2002).
- Tarantola, A.: *Inverse Problem Theory, Methods for Data Fitting and Model Parameter Estimation*, Elsevier, New York, 1987.
- Watanabe F., O. Uchino, Y. Joo, M. Aono, K. Higashijima, Y. Hirano, K. Tsuboi, and K. Suda: Interannual variation of growth rate of atmospheric carbon dioxide concentration observed at the JMA's three monitoring stations: Large increase in concentration of atmospheric carbon dioxide in 1998. *J. Meteor. Soc. Japan* *78*, 673-682 (2000).
- Worthy D.E.J. (ed): Canadian Baseline Program, Meteorological Service of Canada, Downsview, Ontario, 2003.

Acknowledgment

I would like to thank T. Conway and P. Tans (CMDL), P. Steele, R. Langenfelds, and R. Francey (CSIRO), I. Levin (IUP-H), Y. Tsutsumi (JMA), M. Ramonet (LSCE), D. Worthy (MSC), T. Nakazawa and S. Morimoto (NIPR), A. Gomez, G. Brailsford, and D. Lowe (NIWA), E. Brunke and C. Labuschagne (SAWS), K. Uhse (UBA), and all their coworkers for contributing their CO₂ measurements and for their help in using the data. Fruitful discussions and a lot of help by C. Le Quéré, C. Gerbig, M. Gloor, S. Houweling, and M. Heimann are gratefully acknowledged. I thank S. Körner for preparing TM3's meteorological driver data. I am obliged to the computing centers Deutsches Klimarechenzentrum Hamburg and Gesellschaft für wissenschaftliche Datenverarbeitung Göttingen for their kind support. H.-O. Beismann and U. Körner (DKRZ) were very helpful in improving the performance of the computations.





Figure 4: CO₂ concentration records (to be continued on next page). The left part shows the time coverage. The coloured stars on the right side group the records into sets used in different inversion runs.

Figure 4: (continued) CO₂ concentration records.




Figure 5: Local/instantaneous a-priori standard deviations [$\mu\text{mol}/\text{m}^2/\text{s}$] for the adjustable terms of the flux components in the standard CO₂ flux model. They are defined through the ‘shapes’ $f_{\text{sh},\text{nee},\text{lt}}(x, y, t)$, etc. The stripe structure in some maps only results from small numerical deviations in the normalization due to the aggregated pixels (cmp. Sect. 3.3.1). A-priori standard deviations only have spatial structure, but are constant in time.




Figure 6: Regular grid of regions used in the standard CO₂ flux model. Colours repeat every 5 or 11 regions; in some lines, neighbouring regions happen to have equal colours.




Figure 7: Top: Two representative examples of individual spatial elements $g_{m_s,i}^{\text{space}}(x,y)$ for long-term and interannual NEE or Ocean flux components, respectively. Bottom: Correlation coefficients with respect to the center locations of these elements, as implied by the respective sets $\{g_{m_s,i}^{\text{space}}(x,y)\}$ of spatial elements.




Figure 8: As Fig. 7, but for the seasonal NEE or Ocean flux components.




Figure 9: A-priori Fourier spectra used for the seasonal and interannual components in the standard set-up. The interannual spectrum has been varied in two sensitivity cases ('fast' and 'slow', Sect. 3.4.2).




Figure 10: Correlation coefficients with respect to two example locations (as in Fig. 7) for the ‘long’ sensitivity set-up (Sect. 3.4.2).




Figure 11: Map of the TransCom-3 regions (Gurney et al., 2002) over which the estimated fluxes are integrated to obtain time series. Note that neighbouring land and ocean regions actually overlap each other beyond the coast line, with the integration of fluxes being done separately for the individual flux components, according to whether they refer to land or ocean processes. The apparent coast lines shown here correspond to a land/water map at the standard resolution of the transport model.




Figure 12: Estimated fluxes based on different sets of observation sites. Fluxes are IAV-filtered ('Filt0.5gc'). Part I. Fluxes integrated over the TransCom-3 land regions.




Figure 12: Part II. Fluxes integrated over the TransCom-3 ocean regions.




Figure 13: As Fig. 12, but full temporal variability, shown for a 4 years' period. Part I. Fluxes integrated over the TransCom-3 land regions.




Figure 13: Part II. Fluxes integrated over the TransCom-3 ocean regions.




Figure 14: Several measures of uncertainty of the interannual variations estimated by the inversion. The red line surrounded by the gray shading gives the standard result (interannual components only, deseasonalized and filtered), together with its a-posteriori $\pm 1\sigma$ interval (calculated at the reference time 2001.5). The a-priori fluxes and their corresponding $\pm 1\sigma$ intervals are given by the black line with error bar. The ensemble of thin (gray) lines shows the spread of results due to change of individual set-up details in the sensitivity testing. Three particular sensitivity cases are marked by thin color lines. For efficiency, these runs have been done using the coarse grid transport model.

Part I. Fluxes integrated over the TransCom-3 land regions.




Figure 14: Part II. Fluxes integrated over the TransCom-3 ocean regions.




Figure 15: As Fig. 14, but full temporal variability. $\pm 1\sigma$ intervals are not available here. Part I. Fluxes integrated over the TransCom-3 land regions.




Figure 15: Part II. Fluxes integrated over the TransCom-3 ocean regions.




Figure 16: Matrices of a-priori and a-posteriori error correlation coefficients (standard set-up, coarse grid transport). Row and column indices refer to the fluxes integrated over individual TransCom-3 regions, further split into the interannual flux components (filtered with Filt1.0T, reference time 2001.5, indices 1–22) and the long-term components (indices 23–44):

1	23	North American Bor.
2	24	North American Temp.
3	25	South American Trop.
4	26	South American Temp.
5	27	Northern Africa
6	28	Southern Africa
7	29	Eurasian Boreal
8	30	Eurasian Temperate
9	31	Tropical Asia
10	32	Australia
11	33	Europe
12	34	North Pacific Temp.
13	35	West Pacific Tropics
14	36	East Pacific Tropics
15	37	South Pacific Temp.
16	38	Northern Ocean
17	39	North Atlantic Temp.
18	40	Atlantic Tropics
19	41	South Atlantic Temp.
20	42	Southern Ocean
21	43	Indian Trop. Ocean
22	44	S.Indian Temp. Ocean




Figure 17: Comparison of different time series filters: Standard filter used here ('Filt0.5gd'), box-car filter (running annual sums), and self-deseasonalizing triangular filter ('Filt1.0T'). Filters are defined in Appendix A.2 and illustrated in Fig. 18. Fluxes shown for 2 selected TransCom-3 regions.




Figure 18: Example filters used in the 'flux model' or for smoothing. Filter properties are given as spectral weights (left) or corresponding pulse response functions (right).

Upper row: Theoretical relations from Appendix A.

Lower row: Discrete numerical representations used in Fig. 17, and their numerical Fourier back-transformations obtained by filtering a Dirac pulse $\delta(t - t_r)$ centered at the reference time $t_r = 2001.5$ [for boxcar, filtering is numerically done in the time domain]. High-frequency values missing in the numerical spectra are zero.

Note that the Gaussian filter (red) is deseasonalizing in the lower row (Filt0.5gd), but not in the upper (Filt0.5g).