Contents

1

Assimilation and Network Design

Thomas Kaminski (FastOpt) and Peter J. Rayner (LSCE)

1.1 Introduction

Information on the carbon cycle comes from a variety of sources. The methods described in this chapter provide a formalism for combining this information. Without such a formalism we are left making ad hoc choices about how to improve our understanding in the light of disagreements among various streams of information. The introduction of such methods into carbon cycle research, principally via the atmospheric studies of [ETFG93, ETF95] revolutionised the field and laid the groundwork for most of the subsequent investigations.

The methods in question are fundamentally statistical. They hence provide estimates of the confidence we should have in quantitative statements about the carbon cycle. These statements are usually couched as spreads of probability distributions or as confidence intervals. We refer to them generally as posterior uncertainties. These posterior uncertainties depend on the prior uncertainties of the various data streams that feed the estimation process, the method for combining these data streams (usually some kind of model) and on the particular state of the system. Of course an important aim of measurements is to reduce the posterior uncertainty.

The present chapter is concerned with quantitative network design, by which we understand the optimisation of a measurement strategy via minimisation of this posterior uncertainty for target quantities of particular interest. Examples of such target quantities are the long-term global mean terrestrial flux to the atmosphere over a period in the past or in the future. The computational tool that transforms the information provided by an observational network of the carbon cycle into an estimate of posterior uncertainty is a Carbon Cycle Data Assimilation System (CCDAS). Hence, network design is closely linked to assimilation both conceptually and computationally. Much of the work reviewed in this chapter lies in a small subset of possible network

design applications for the carbon cycle. In particular it uses a limited set of types of observations. This is not an inherent limitation of the approach but rather a limitation in modelling approaches that can combine many streams of measurements. This is changing now. Hence, much of the chapter looks forward to applications that combine different measurement approaches. It is useful, therefore, to describe the problem in general even if most cited examples are from simpler cases.

The first part of the chapter presents the formalism of carbon cycle data assimilation. It will describe the generation of posterior uncertainties for both simple and more complex cases. Next we review applications of that methodology to the simpler case of atmospheric transport inversion along with the presentation of important caveats. This is followed by a sketch of how network design might look in a more comprehensive CCDAS. Finally we give perspectives and recommendations.

1.2 Methodology

It is useful to look at a data assimilation system as a tool that combines various sources of information to form a consistent picture of the underlying system, which in our case is the global or regional carbon cycle. Among the pieces of available information are the observational data we would like to assimilate, estimates for various rate constants of the system and dynamical equations describing the system's evolution.

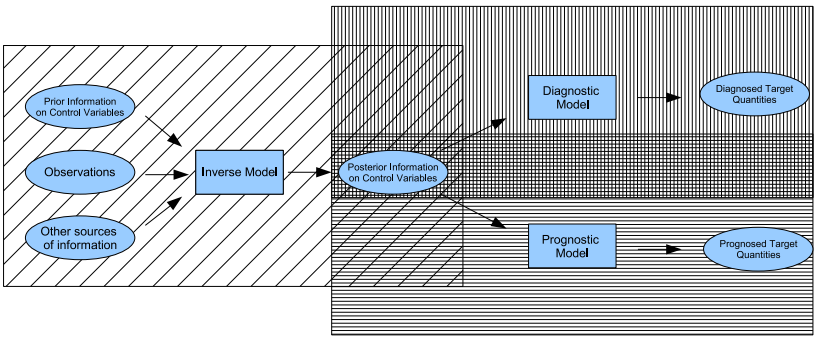

Fig. 1.1. Schematic overview on two-step procedure for inferring diagnosed and prognosed target quantities from data. Rectangular boxes denote processes, and oval boxes denote data. The diagonally hatched box includes the inversion (calibration) step, the vertically hatched box the diagnostic step, and the horizontally hatched box the prognostic step

One usually has one or more target quantities that one is interested in, for instance the terrestrial uptake of a continent such as Europe over a particular time interval. As we typically cannot observe the target quantity itself, we use a numerical model to link the target quantity to the observations. Often this is most conveniently achieved in the two-step procedure sketched in figure 1.1. First a set of control variables (a combination of initial- and boundary conditions and tuning parameters in the model equations) are nominated, and the underlying numerical model is run in inverse mode. This means prior information on the control variables, observational information and, if available, information from other sources are combined with the model to form posterior information on the control variables. If the set of control variables includes model parameters, this inversion process is also called calibration of the model. In figure 1.1 the inversion step corresponds to the diagonally hatched box.

In a second step the posterior information on the control variables is used to infer information on the target quantities. Figure 1.1 distinguishes between diagnosed and prognosed target quantities, i.e. between quantities computed in a diagnostic/prognostic model run. A diagnostic run occurs within the same domain (spatial or temporal) as the inversion step, whereas a prognostic run is (at least in part) outside this domain.

When nominating the control variables, it is not important that they are interesting in themselves. In the two-step approach they take the role of intermediate quantities on the way to the target quantities. The nomination of the control variables must rather attempt the minimisation of biases and other errors in the inversion step.

It is convenient to quantify the state of information on a specific physical quantity by a probability density function (PDF): the prior information is quantified by a PDF in the space of control variables, the observational information by a PDF in the space of observations, etc... [Tar87] describes the probabilistic framework in detail and provides examples. When the input to the inverse problem can be characterised by Gaussian PDFs, the model is linear, and the model error is Gaussian as well, [Tar87] shows that the posterior information is also quantified by a Gaussian PDF. The mean of that PDF is given by:

$$\mathbf{x} = \mathbf{x_0} + [\mathbf{M}^T\mathbf{C}(d)^{-1}\mathbf{M} + \mathbf{C}(x_0)^{-1}]^{-1}\mathbf{M}^T\mathbf{C}(d)^{-1}(\mathbf{d} - \mathbf{M}\mathbf{x_0}) \ , \qquad (1.1)$$

and the covariance of its uncertainty is given by:

$$\mathbf{C}(x)^{-1} = \mathbf{M}^T\mathbf{C}(d)^{-1}\mathbf{M} + \mathbf{C}(x_0)^{-1} \ , \qquad (1.2)$$

where $\mathbf{M}$ denotes (the Jacobian matrix of) the model (that links control variables to observations), x_0 and $\mathbf{C}(x_0)$ the mean and covariance of the prior information's PDF. On the observations' side, $\mathbf{d}$ and $\mathbf{C}(d)$ denote the mean and the covariance of uncertainty. In the inversion procedure the corresponding PDF has to reflect errors in both the observational process and our ability to correctly model the observations. We achieve this via

$$\mathbf{C}(d) = \mathbf{C}(d_{\text{obs}}) + \mathbf{C}(d_{\text{mod}}) \tag{1.3}$$

and by subtracting the mean model and observational errors from $\mathbf{Mx_0}$ and d, respectively.

We also note that with a little linear algebra (1.2) can be reformulated to:

$$\mathbf{C}(x) = \mathbf{C}(x_0) - \mathbf{C}(x_0)\mathbf{M}^T(\mathbf{C}(d) + \mathbf{M}\mathbf{C}(x_0)\mathbf{M}^T)^{-1}\mathbf{M}^T\mathbf{C}(x_0) \tag{1.4}$$

One can easily verify that $\mathbf{x}$ (from (1.1)) also minimises the cost function (the exponent of the Gaussian posterior PDF)

$$\begin{aligned} J(\tilde{\mathbf{x}}) = \tfrac{1}{2}\,[&(\mathbf{M}\tilde{\mathbf{x}} - \mathbf{d})^T\mathbf{C}(d)^{-1}(\mathbf{M}\tilde{\mathbf{x}} - \mathbf{d}) \\ &+(\tilde{\mathbf{x}} - \mathbf{x_0})^T\mathbf{C}(x_0)^{-1}(\tilde{\mathbf{x}} - \mathbf{x_0})] \end{aligned} \tag{1.5}$$

and that the Hessian matrix $\mathbf{H}(\tilde{x})$ of J, i.e. the matrix composed of its second partial derivatives $\frac{\partial^2 J}{\partial x_i \partial x_j}$, is constant and given by

$$\mathbf{C}(x)^{-1} = \mathbf{H}(\tilde{x}) \ . \tag{1.6}$$

If the model is non linear or any of the PDFs of the inputs are non Gaussian, the Gaussian PDF with the minimum of (1.5) as mean and covariance given by (1.6) is an approximation of the posterior PDF.

For the second step, i.e. the estimation of a diagnostic or prognostic target quantity $\mathbf{y}$, its PDF can be approximated by a Gaussian with mean

$$\mathbf{y} = \mathbf{N}(\mathbf{x}) \tag{1.7}$$

and the covariance

$$\mathbf{C}(y) = \mathbf{D}(\mathbf{N})\mathbf{C}(x)\mathbf{D}(\mathbf{N})^T + \mathbf{C}(y_{mod}) \ , \tag{1.8}$$

where $\mathbf{N}$ is the model (in figure 1.1 denoted as diagnostic/prognostic model) that maps the control variables onto the target quantity, $\mathbf{D}(\mathbf{N})$ is its linearisation around the mean of the posterior PDF of the control variables, also denoted as the Jacobian matrix of $\mathbf{N}$, and $\mathbf{C}(y_{mod})$ is the uncertainty in the model result from errors in the model. Only if y coincides with one of the observations used in the inversion step, this uncertainty is already accounted for in $\mathbf{C}(x)$, and we omit the $\mathbf{C}(y_{mod})$ contribution. If $\mathbf{N}$ is linear and the posterior PDF of the control variables Gaussian, then the PDF of the target quantity is Gaussian as well, and completely described by (1.7) and (1.8).

1.3 Transport inversion

This section demonstrates network design for atmospheric transport inversion before addressing network design in an entire CCDAS. The setup for atmospheric transport inversion as introduced by [ETF95] uses a set of atmospheric

CO_2 observations provided by a regional or global sampling network to constrain surface fluxes. Adding this intermediate step before addressing network design in an entire CCDAS is useful for a number of reasons:

- Transport inversions allow us to introduce the main network design concepts for a single component of the carbon cycle.
- From the methodological point of view it is convenient that, at least for the typical time scales of interest (from diurnal to decadal), the atmospheric transport of CO_2 is modeled as linear.
- Transport inversions often use only a single data stream.
- There have been successful applications of network design concepts.
- Finally, transport inversion has been a typical and, hence, familiar starting point into CCDAS for many colleagues (and the authors).

In transport inversions the PDFs for priors and data are usually assumed Gaussian, i.e. the posterior PDF is given by (1.1) and (1.2), where $\mathbf{M}$ denotes the transport model, $\mathbf{x}$ and $\mathbf{d}$ the surface fluxes and data, respectively. [Ent02] provides details on transport inversion and an exhaustive list of examples and references.

As soon as the pioneering paper of [ETF95] introduced the calculation of posterior uncertainties into the atmospheric inversion problem, it became clear that they were disturbingly large. As a more recent example, the posterior estimates in Fig.2 of [GLD+02] show that even for Europe, a relatively well-sampled region, the posterior uncertainties are very large. Indeed there were large areas of the world where, according to that study, the atmospheric data added little information. In the optimistic case of the ocean this partly indicated the tight prior constraint. Here the atmosphere is really a consistency check. For many terrestrial regions the position is reversed. Even with large prior uncertainties, which imply enhanced sensitivity of the posterior flux to data, there is little reduction in uncertainty between the prior and posterior fluxes. Clearly more measurements are needed.

Such measurements are expensive and difficult. Either they require development and deployment of expensive instruments or the painstaking analysis of flasks returned to central measurement facilities. It would be helpful if augmentation of the network was guided so as to produce maximal return. We exploit the important property of (1.2) that, for linear systems and a given $\mathbf{C}(d)$, the posterior uncertainty depends only on prior uncertainty and the Jacobian. Thus we can predict the information content of a measurement without actually making it. We can use this property to study optimal deployment of measurements.

Like much else in the basic theory of atmospheric inversions, this idea was imported from solid geophysics. [HS92] (later published as [HS94]) optimised various quality measures of a seismographic network using station locations as parameters. This is a conventional minimisation problem but of a highly nonlinear function. They used the procedure of simulated annealing for the optimisation. The algorithm mimics the "shaking down" of ions into a crystal

on a sufficiently gradual phase transition. Put briefly, a function is calculated for a given vector of parameter values. The parameters are perturbed and the function recalculated. If the new function value is lower the new vector will be retained. There is also a finite probability that a higher value will be retained, avoiding the algorithm being stuck in a local minimum. The perturbations are proportional to a "temperature" and the probability of retension of higher values is based on a Boltzmann factor (a negative exponent of "temperature"). The "temperature" is reduced sufficiently gradually that the algorithm may find a global minimum. As the "temperature" is reduced, the probability of retaining suboptimal values is also reduced so that the iterations of the system converge.

[RET96] applied the same technique to atmospheric inversions using CO_2 and $\delta^{13}C$ measurements. The parameters were the longitude and latitude of the observing sites. They chose two objective functions [1], reflecting different selections of target quantities. The first target quantity was the total ocean uptake, and the corresponding objective function is given by $(\sum \mathbf{C}(\mathbf{S_{oc}}))^{0.5}$ where the subscript indicates the submatrix containing only ocean sources. The sum is over all elements of the submatrix. The second objective function uses actually a set of target quantities, namely the regional ocean uptakes. It expresses the uncertainty of regional ocean uptake summed over all ocean regions and is given by $(trace(\mathbf{C}(\mathbf{S_{oc}}))^{0.5})$. The difference between the two objective functions is the inclusion or exclusion of covariances.

The upper panel of figure 1.2 shows a plot of the first objective function along with the locations for adding up to three new stations to the network of [ETF95] as chosen by the optimisation algorithm. We see first that the simulated annealing algorithm performs well in finding the minimum in the function. Further, the positions of more than one station can also be predicted from the function map; for example the optimum locations for two stations are the two deepest minima on the map. The finding that the position of one new station does not distort the map sufficiently to influence the next choice greatly simplifies the problem if we are concerned with only small additions to the network. It allows us to optimise incrementally. The approach has been used by [PM02] and [LRW04].

The second striking thing about the upper panel of figure 1.2 is the preference for land sites, despite the target quantity being ocean uptake. This arises from the propagation of information in the presence of global constraints. In the formulation of [ETF95] the global growth-rate is specified and given a fairly small uncertainty. The land and ocean uptakes must sum to this figure. Thus an improvement in knowledge of one of these gives a corresponding improvement in knowledge of the other. Such an improvement is best achieved

[1] In network design, we are dealing with a nested optimisation problem. The inverse problem of solving the cost function of (1.5) is nested into the problem of optimising the network quality. This quality is expressed by a second function, which we will call objective function in order to avoid confusion with (1.5)

by constraining the flux with the highest posterior uncertainty in the current set-up. That is the Amazonian region and so this is chosen. The next favourite region is also over land. This coupling of regions was the major conclusion of the study.

The final thing to note about figure 1.2 is the radical difference between the two objective functions. Unsurprisingly, when asked to optimise the summed regional uncertainty (second objective function) rather than uncertainty of the total ocean uptake (first objective function), the algorithm chooses to improve estimates for the least constrained ocean regions. This is quite a different network to the first example. It is clear that quantitative network design requires precisely formulated questions. Requirements like "the best knowledge of biosphere function in a region" will produce very different networks depending on how they are mathematically formulated. It is also obvious that the network depends on the prior uncertainties in use. If, for example, some external information greatly reduced the prior uncertainty on the Amazon region, the posterior uncertainty would (according to (1.2)) also decrease and possibly eliminate the region from the favoured list. The influence of the prior cannot be circumvented by using an objective function such as the uncertainty reduction (ratio of posterior to prior uncertainty). Here one has the opposite problem so that a very weak prior will yield a more dramatic reduction with the addition of further stations.

After these caveats one is left with robust general findings about the behaviour of optimal networks. For example [RET96] did show that the existence of global constraints couples information from disparate regions. Thus improvements in knowledge of Amazonian fluxes improves knowledge of both total land and ocean fluxes. A corollary is that the "hot-spot" strategy of choosing the least known region for improved observations is globally efficient.

[RET96] also touched on the impact of different descriptions of data uncertainty on network design. Rather than the uniform data uncertainty used to generate figure 1.2 they used an uncertainty proportional to the atmospheric signature of the terrestrial biospheric flux, meant to approximate the model error contribution $\mathbf{C}(d_{\mathrm{mod}})$ to the over-all uncertainty in (1.3). The different choice did not overturn the main findings but the choice of scaling for the new error term was arbitrary. [GFPS00] greatly expanded this aspect of the study and produced networks that compromised between the strongest possible observations of a region (often gained by placing a station in the centre) and large data uncertainty.

[GFPS00] made an interesting contrast with the earlier study of [RET96]. First they rejected the use of prior information at all. As such they required larger numbers of stations (around 150) to saturate the information needed for their inversion. They also studied a wider range of sampling strategies than [RET96], especially aerial profiles and upper tropospheric transects. Again it is the generalities of the study rather than specific information on placement that will endure. Particularly (and controversially) they noted that the "fence post" strategy of surrounding a target region with measurement sites

was less effective than sampling of the concentration gradients among neighbouring regions by placing stations near their centres. They also noted that aerial profiles were more useful than airborne transects since measurements in the upper troposphere were only weakly connected to surface sources by atmospheric transport.

These early studies were made in the context of inversion configurations common at the time. In particular they solved for a relatively small number of regions and used data taken as monthly or annual averages. Both of these have since changed and require some revising of earlier results. The use of large regions in these early inversions gave measurements an unnaturally large influence since a measurement could constrain parts of the large region that might never be connected to it by atmospheric transport. [KRHE01] noted the dangers of this for biasing inversions and current practice is to increase the number of regions, preferably to the full spatial resolution of the transport model. This must have implications for network design. A hint of this was given by [PMB$^+$03] who noted that the range of preferred sites in their network optimisation expanded as they increased the number of regions in their inversion. They performed an incremental inversion in which sites were added one at a time and the optimal site list for n sites was used for the placement of site $n + 1$. In general, the approach suggested by the optimisation is to place a site in each region. Obviously this is impossible in the context of inversions at the resolution of the transport model. And even that resolution can only provide a discrete approximation of the actual two-dimensional flux field. Much of the underdetermined nature of the inverse problem is, thus, hidden by suppressing most of its degrees of freedom [KH01]. It is unclear yet how this affects the optimal network.

The other major change in inversion formulation since the initial network design studies is the trend towards using data at higher time resolutions. In the context of formal inversions this was pioneered by [LRSE02, LRSE03]. They showed that the differential sampling afforded by synoptic variations in advection could act as a more precise regional constraint than the highly diffusive monthly mean responses. Further, provided the biases due to the use of large regions could be avoided, the estimates were usefully accurate, that is errors were consistent with their uncertainties. This work still used the "large region" approach although with considerably higher spatial resolution in the source space than was traditional. [PRB$^+$05] has since demonstrated inversions of such high-frequency data in combination with the transport model resolution of sources.

In a network design context, [LRSE03] investigated a network of high-frequency monitoring stations intended to elucidate the regional patterns of carbon flux within Australia. They use 12 regions over Australia, still much coarser resolution than the transport model but unusually high for the large-region approach. They used incremental optimisation to design the network. They noted again the tendency to place a station in each region.

Another application field for network design is the evaluation of spaceborne sensors providing integrals of carbon dioxide over vertical columns. This type of application does not require an optimisation algorithm. It was pioneered by [RO01], who estimated the sensor precision required by a transport inversion to achieve a specified minimim posterior uncertainty for surface flux field. A more recent example of this approach is provided by [HBA+04], who evaluated the performance of a set of existing and planned satellite instruments.

What has not yet been tried is the network design problem with inversions at the transport model resolution and high-frequency data. The problem appears computationally challenging as the transport model Jacobian $\mathbf{M}$ in (1.1) and (1.2) is large. Here the alternative formulation (1.4) helps. In the case where we wish to optimise the constraint on an integrated flux we can use (1.8), where $\mathbf{N}$ is a linear operator. Further we note that in the network optimisation calculations only $\mathbf{M}$ and $\mathbf{C}(d)$ change as the network changes. Substituting (1.4) into (1.8) we note that most terms like $\mathbf{D}(\mathbf{N})\mathbf{C}(x_0)$ can be precomputed. If we are considering additions to a fixed network we can also replace $\mathbf{C}(x_0)$ by the posterior covariance from that network so most of the term $\mathbf{M}\mathbf{C}(x_0)\mathbf{M}^T$ is also precomputed. We are then left inverting matrices of the dimension corresponding to the additional observations, which may be feasible for diurnal or even hourly data.

An accurate specification of the uncertainty contribution from model error in (1.3) is notoriously difficult. Quantifying the impact of differences in modelled transport is, however, less difficult and provides a popular surrogate for model error. An ideal environment for this task is provided by the TRANSCOM project (see http://www.purdue.edu/transcom) of the International Geosphere-Biosphere Programme (IGBP), which compares forward simulations and inverse simulations of a set of transport models. [GLD+02] describe the TRANSCOM inversion experiment, in which sixteen different transport models are used with all other inputs to the inversion held constant. They reported two uncertainty expressions as covariance matrices

$$\mathbf{C}^B_{j,k} = \frac{1}{16}\sum_{i=1}^{16}(\mathbf{S}^i_j - \overline{\mathbf{S}}_j)(\mathbf{S}^i_k - \overline{\mathbf{S}}_k) \tag{1.9}$$

and

$$\mathbf{C}^W_{j,k} = \overline{\mathbf{C}_{j,k}} \tag{1.10}$$

where i counts the sixteen models, j and k their 22 source regions. The overbar denotes the ensemble average across the models and the superscripts B and W stand for between-uncertainty and within-uncertainty, respectively. (1.9) expresses the uncertainty in actual estimates from the inversion while (1.10) expresses the average posterior uncertainty.

[PMB+03] used the above-sketched incremental optimisation approach to extend the [GLD+02] network iteratively. A site was added to the current

network when it performed best at reducing the average posterior uncertainty over all source regions for any of the models or the within uncertainty (1.10).

[Ray04] also based a network design study on the [GLD+02] setup (with one additional model, as described by [GLD+03]). He defined the sum of the between (1.9) and within (1.10) uncertainties as total uncertainty and used that as an additional objective function. The inclusion of the variance of actual estimates from (1.9) immediately circumscribes the study since one must have real data to calculate these. One is, hence, left choosing networks from a discrete list of stations. [Ray04] used the list of 110 compiled by [GLD+02] from which they chose the 76 used in that study. The use of discrete lists also demands a change in algorithm since simulated annealing requires continuous fields. [Ray04] used the techniques of genetic algorithms. Briefly these maintain populations of potential solutions which are allowed to "breed", "mutate" and "compete". The iterations of the scheme can be compared to generations of a population and we hope that the overall fitness of the population will improve. See `http://csep1.phy.ornl.gov/CSEP/MO/MO.html` and references therein for more details.

[Ray04] optimised networks of 76 (like TRANSCOM) and 110 stations and with the within-uncertainty or total-uncertainty objective functions. Figure 1.3 shows the cases for 76 stations. One might expect that the networks for total-uncertainty would be smaller than for within-uncertainty as the optimisation rejected stations with large model-model differences. In fact the opposite is the case, the networks generally dispersed in the presence of model-model difference. It appears that we should consider model-model difference like a small-scale heterogeneity. For example two models may place the maximum response to a flux at each of two neighbouring stations. A network designed to reduce the impact of model-model difference will average across this heterogeneity. If we optimise only for within-uncertainty we risk large between-uncertainty and ultimately large total-uncertainty. The need for real data limits the use of this approach, however. [Ray04] noted that stations with highly variable (among models) responses to input fluxes would be penalised in an optimisation of total uncertainty. This generality seems likely to survive beyond the details of the study.

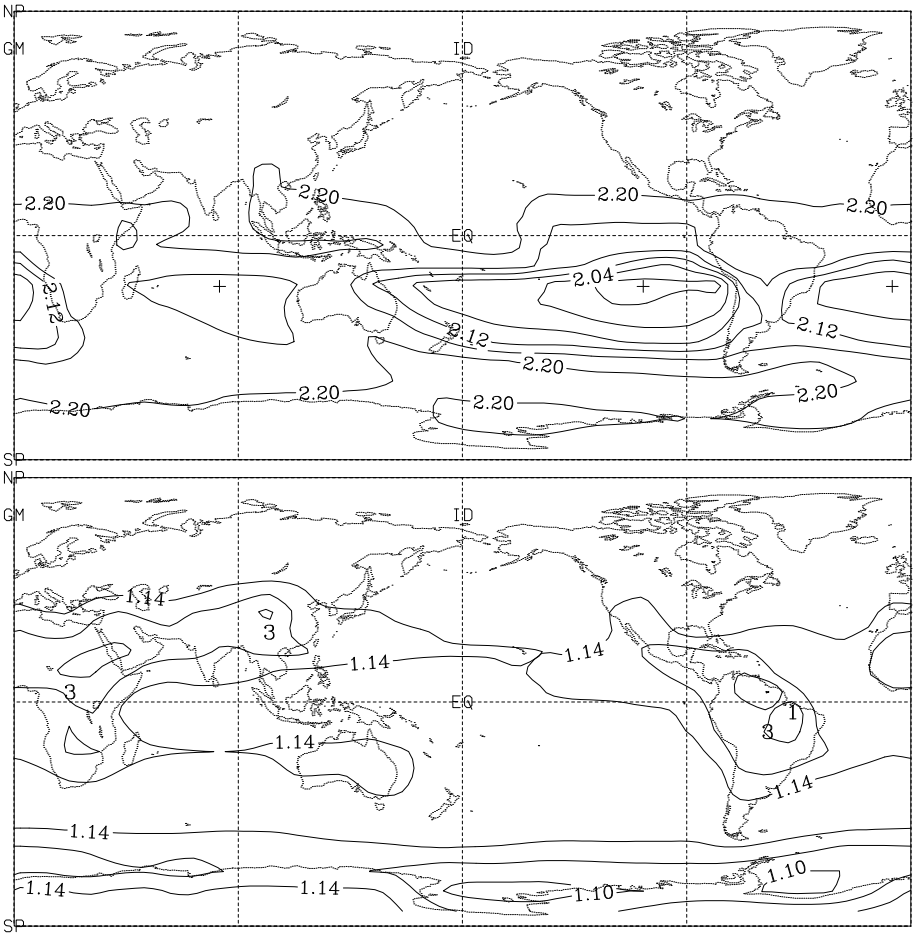

Fig. 1.2. Each panel shows a separate objective function of the location of one additional station to the current network. The station measures the same species with the same uncertainties as the existing station at Cape Grim, Tasmania. The upper panel's objective function is the standard deviation (in GTC) of global ocean uptake, and the lower panel's objective function is the square root of the sum of the variance of individual oceanic source components (in GTC), which provides a separate constraint for each ocean region. The upper panel has a contour interval of 0.02 and the lower panel of 0.04. Contouring is curtailed at 80^o to avoid problems arising from multi-valued responses at the poles. The upper panel indicates by crosses the optimal locations of three additional stations, and the lower panel uses the numbers 1 and 3 to indicate optimal locations of 1 or 3 more stations. Both figures taken from [RET96].

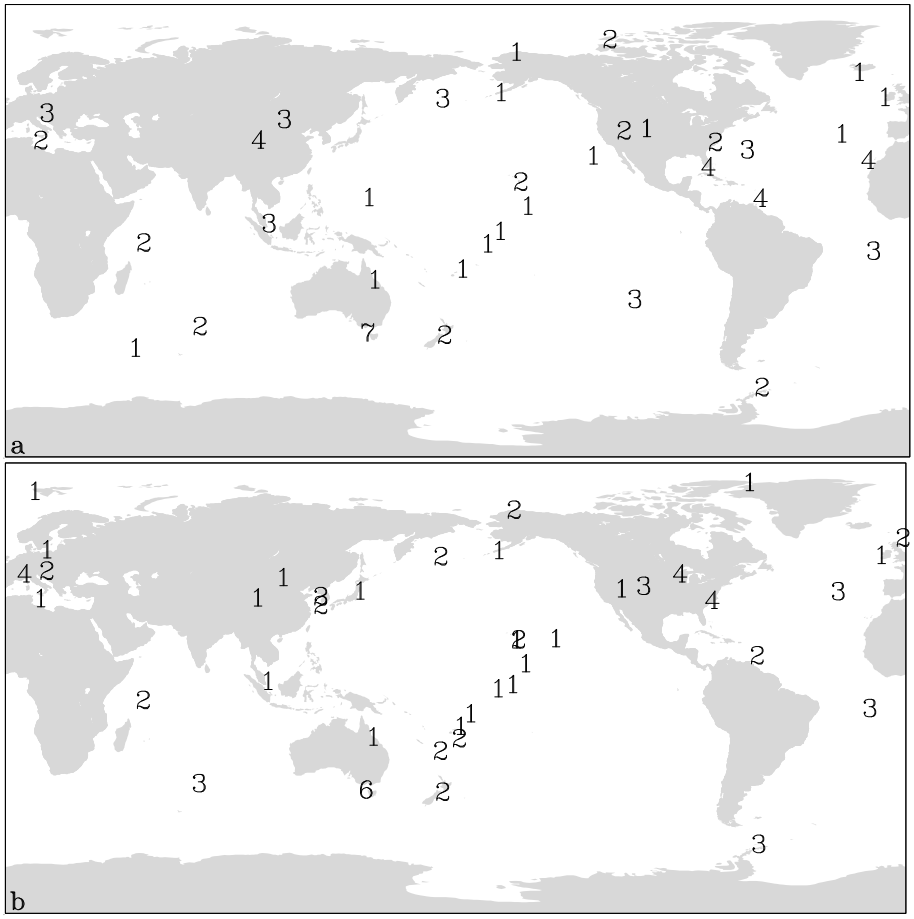

Fig. 1.3. Optimised networks for within-model **(a)** and total **(b)** uncertainty for a 76-site network. Numbers indicate the number of times the site is included. Figure taken from [Ray04].

1.4 CCDAS

The previous section has demonstrated network design concepts for the special case of atmospheric transport inversions. Transport inversions deliver posterior information on the exchange fluxes at the Earth's surface over the atmospheric sampling period. Information on the processes behind the fluxes can only enter the inversion via the priors, for instance by using the output of a process-based model as a prior flux field. The transport inversion, however, has no means of conserving the process dynamics. Thus the posterior flux field will generally be inconsistent with the output of process-based models. The second significant restriction of transport inversions consists in the fact that possible target quantities (for use in (1.7) and (1.8)) can only be diagnostic. Extending the modelling system by a process model that can prognose the process dynamics avoids these two restrictions. The corresponding inverse model is usually termed a carbon cycle data assimilation system (CCDAS). Another benefit of including a process model is that it enables the access to additional observations (see examples below). The CCDAS control variables can be initial or boundary conditions as well as parameters in the formulation of the processes. The capability of estimating process parameters provides a means to directly improve our process understanding. The present section generalises the design concepts that the previous section introduced and illustrated for transport inversions to the design of networks that provide observations to a CCDAS.

[KKRH02] present an early CCDAS version built around the Simple Diagnostic Biosphere Model (SDBM, [KH95]) coupled to the atmospheric transport model TM2 ([Hei95]). CCDAS was later upgraded [KGS$^+$03, Sch03, RSK$^+$05] by replacing SDBM with the Biosphere Energy Transfer HYdrology Scheme (BETHY, [Kno97, Kno00]). In contrast to the diagnostic SDBM, which is driven by observed vegetation index data, BETHY is a prognostic model that uses a set of meteorological driving data to integrate the model state forward in time.

Adding the terrestrial component to the transport model renders the composite model non-linear. Hence, we must now use equations (1.5) and (1.6) to infer posterior information on the control variables of a CCDAS. While the CCDAS built around SDBM has 24 control variables, the one around BETHY has 57 to about 1000, depending on the setup. For a control space of that size, (1.5) is most efficiently minimised via an iterative gradient algorithm. At the minimum of the cost function, CCDAS evaluates the cost function's Hessian to approximate the posterior uncertainties via (1.6). Typical CCDAS target quantities are regional and global means of fluxes such as NEP, averaged over diagnostic or prognostic integration periods. The target quantities' posterior uncertainties are derived via (1.8). The Jacobian matrix N is the derivative of the target quantities with respect to the control variables. Efficient derivative code providing the gradient (adjoint model), the Hessian and the Jacobian are generated automatically from BETHY's source code via the automatic

differentiation tool TAF ([GK98]). This automation has proven useful, as it allows quick updates of CCDAS after modifications of BETHY.

The above-cited CCDAS applications run globally and assimilate atmospheric carbon dioxide flask samples provided by a global network and (in a less formal pre-step) remotely sensed vegetation index data. There are a number of assimilation systems that differ from CCDAS in the assimilation approach or in the type of data that are assimilated. For instance, the system of [Pak04] uses the same approach but assimilates eddy flux measurements of heat and carbon into a one point version of the CSIRO Atmosphere Biosphere Land Exchange (CABLE) model. [VBS01] also minimised (1.5) by a gradient method, using the adjoint of their k-model to provide the gradient of J with respect 16 parameters. They assimilated atmospheric temperature and carbon dioxide but did not calculate posterior uncertainties.

Lacking the adjoint, gradient information is usually approximated by finite differences. This means each component of the control vector is perturbed in turn, and the perturbation's impact on the cost function is evaluated in a separate model run. This impact may be dominated by higher order effects, if the perturbation is too large. If it is too small, the impact is dominated by numerical noise. The approximate gradient usually slows down the convergence of the gradient algorithm. The computational cost of the approximation typically limits the complexity of the model, its spatial resolution, or the number of control variables that can be treated. [WLCC01] used finite difference approximations for the gradient and Hessian (actually via Jacobians) with respect to up to 7 model parameters for assimilation of eddy flux measurements into a one point version of the CSIRO Biospheric Model (CBM). [SPVC03] applied the same strategy to estimate 5 parameters of a one point version of their model ORCHIDEE from eddy flux measurements. [WSL$^+$05] use an ensemble Kalman filter embedded in a finite difference gradient algorithm to estimate 14 control variables of a box model from eddy flux and carbon stock measurements.

The CCDAS approach finds the most likely value for the control parameters by maximising the posterior PDF. This is possible for a large variety of distributions. However the identification of the inverse Hessian with the covariance, key to the network design application, requires the assumption of a Gaussian posterior. [RSB$^+$02] avoid such assumptions by directly sampling a three parameter posterior PDF which quantifies the fit of their coupled biosphere-transport model to flask samples of atmospheric carbon and its isotopic composition. The model domain are the northern high-latitudes, and in complexity the model is similar to the above-mentioned SDBM. [KK05] apply a guided Monte Carlo sampling method ([MRR$^+$53]) of the 14 (C4) and 23 (C3) posterior parameter PDFs, for one point setups of BETHY. Their posterior PDFs are generally close to Gaussian.

[Bar02] used a genetic algorithm to estimate 22 parameters of the conceptual VAST model from measurements of NPP and carbon pool sizes over Australia but did not calculate posterior uncertainties.

The Optimisation InterComparison project (OptIC, see http://www.globalcarbonproject.org/ACTIVITIES/OptIC.htm) evaluates different optimisation techniques for estimating four parameters of a simple terrestrial biosphere model.

Fig. 1.4. Evaluating a network control vector in terms of the posterior uncertainty of the target quantity. Oval boxes denote data, and rectangular boxes denote processing

In a CCDAS context, the network design problem can be tackled very much in the same manner as for a pure atmospheric transport inversion. The basis is again the specification of a target quantity (see figure 1.4). The CCDAS can then quantify the quality of a given network via (1.6) and (1.8) by the target quantity's posterior uncertainty. In contrast to the case of transport inversion, this target quantity can also be a prognostic one, i.e. a network can be optimised in order to reduce the uncertainty in a particular aspect of the model prediction.

Next, a set of candidate networks needs to be defined. The methodological framework of section 1.2 requires the network to provide its observations in the form of a mean value and a covariance matrix of uncertainties. Now, one of the major tasks of network design is to explore and evaluate not yet existing, fictive networks. This means for many candidate networks only a fraction of the observations is real and the remainder fictive. The most convenient way of accounting for fictive observations is to ignore them in (1.5) and then assume in (1.6) that the observed value equals the value that the model simulates from the posterior mean. The inclusion of a fictive observation then only affects the posterior uncertainty and not the posterior mean value. This means, for each candidate network, we need to extend the real observations' covariance of uncertainty (1.3) by the components corresponding to the fictive observations.

The strategy for solving the network design problem will usually be selected according to the set of candidate networks. If it comprises only a few elements, the posterior uncertainty of the target quantity can be evaluated for each of them and the networks be ranked accordingly. As an example imagine the extension of an existing network for budgeting the Amazon rain forest by a flux tower. It is likely that logistic and political constraints limit the set of candidate networks to just a few.

Larger sets of candidate networks have to be searched by an algorithm that is more efficient than testing candidate by candidate. Such search (or optimisation) algorithms typically operate on a subset of R^n. Hence, the set

of candidate networks is to be parametrised, i.e. it is represented as the range of a function of a vector of network [2] control variables. For instance, when the network design problem consists in adding a fixed number of atmospheric sampling sites to a given network, these control variables can be the vector of the station's coordinates on the globe. The control vector could also have components with a discrete domain, e.g. when the set of candidate networks contains elements that differ by a data stream that can be switched on or off.

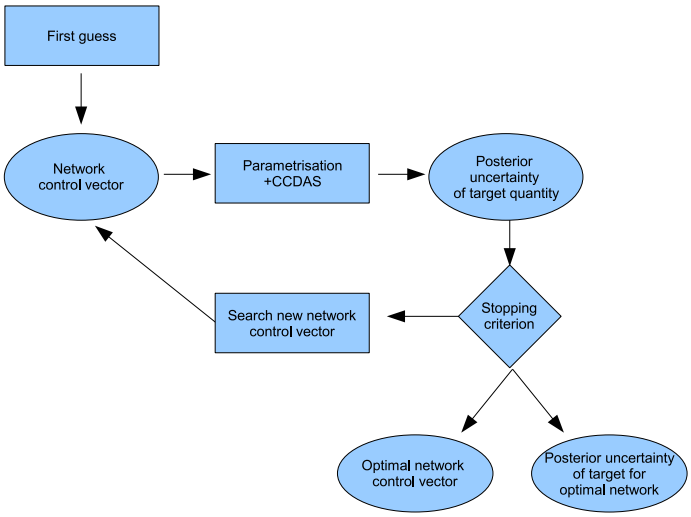

Fig. 1.5. Flow diagram for an iterative optimisation algorithm. Oval boxes denote data, and rectangular boxes denote processing. figure 1.4 details box "parametrisation + CCDAS".

The task of the optimisation algorithm is to search the domain of the control vector for an element such that the corresponding network minimises the target quantity's posterior uncertainty. Section 1.3 provides examples for applications of optimisation algorithms to network design. [RET96] and [Ray04] demonstrated use of simulated annealing and a genetic algorithm. Another important class of optimisation algorithms are the powerful gradient algorithms [GMW81]. The principal functioning of these optimisation algorithms

[2] To avoid confusion we prepend the term 'network' and thus distinguish the control variables determining the network from those determining the behaviour of the model inside CCDAS (denoted by $\mathbf{x}$ in (1.5)).

is sketched in figure 1.5. Starting by a first guess of the control vector, the algorithms work iteratively through the following loop:

1. Perform at least one evaluation of the target quantity's posterior uncertainty as a function of the control vector (as depicted in figure 1.4). This evaluation involves the composition of the network from the network control vector and then a CCDAS run for that network. Gradient algorithms require, in addition, at least one evaluation of the gradient of the posterior uncertainty with respect to the control vector.
2. Check for convergence and exit in case of convergence. A suitable convergence criterion might be a threshold for the target quantity's posterior uncertainty. Gradient algorithms typically use a threshold for the gradient, which approaches zero as the posterior uncertainty reaches a minimum.
3. Propose new network control vector.

The outputs of the optimisation algorithm are the optimal network control vector, and thus the optimal network, plus the corresponding posterior uncertainty of the target quantity.

A first demonstration of network design in a CCDAS context is provided by the above-sketched study of [KKRH02]. They did, however, not enter into an optimisation loop but tested just two candidate networks. Both networks differ by a flux measurement in the model's broadleaf evergreen (BE) biome. The study did not select a particular target quantity, but looked instead at the posterior uncertainties of the control variables. Figure 1.6 shows the posterior values and uncertainties of the soil model parameter Q10 for both networks.

The above approach can be easily extended to network design problems where more than one target quantity are to be considered. The network can then be optimised for a function, e.g. a weighted sum, of the posterior uncertainties of the individual uncertainties. Recall the above-mentioned example of [RET96] who chose the regional ocean fluxes a target quantities and the sum of squares as weighting function. This means they optimised their atmospheric network to minimise the sum of squares of the regional ocean fluxes posterior uncertainties.

1.5 Perspective and Recommendations

The previous section has made clear that a CCDAS constitutes an ideal tool for network design, as it is able to rigorously quantify the performance of an observational network in terms of the posterior uncertainty in specified target quantities. These target quantities may be regional or global net fluxes integrated over present or future periods, as required by policy makers as a basis for their decisions. Alternative target quantities more relevant for model developers may be fluxes resulting from particular processes such as photosynthesis or heterotrophic respiration integrated over the distribution of plant functional types.

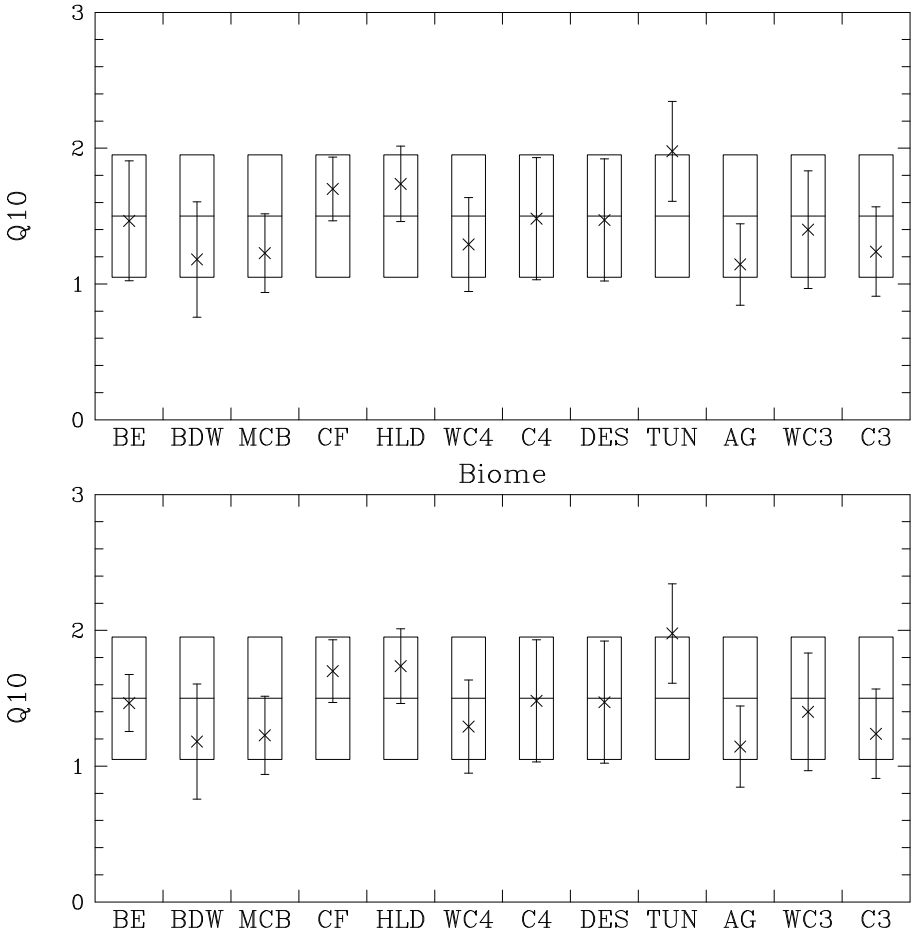

Fig. 1.6. Impact of a fictive flux measurement with uncertainty of 10 $gC/m^2/year$ in the broadleaf evergreen (BE) biome onto the posterior uncertainty in the soil respiration parameter Q10 as quantified by a CCDAS based on the Simple Diagnostic Biosphere Model ([KH95]) and using atmospheric carbon dioxide samples. Upper panel: Without fictive measurement. Lower panel: With fictive measurement. Figure taken from [KKRH02].

The previous section has, however, also made clear that there are hardly any CCDAS applications to network design. The primary reason is that as yet only few of these systems exist, and none of them can yet handle all available streams of observations relevant to the carbon cycle. Our first recommendation is, hence, to upgrade the existing CCDASs (or build new ones) so as to handle the available data streams.

Restricting the focus to the atmospheric component of the carbon cycle, the situation is fortunately much better. There are a number of atmospheric

transport inversion systems in operation, and some of them have been applied in design studies for atmospheric networks, including air- and space-borne instrumentation. Most of our insights are, hence, based on the experience from these studies, and the following remarks generalise these insights to CCDAS-based network design.

One of these generalised findings is that the optimal network usually depends strongly on how one chooses the target quantity, that formalises the quality of the network: For instance, it is likely that reducing uncertainties in the 21st century's terrestrial net flux either over Europe or globally requires significantly different networks. When no decision can be made between competing target functions, it may be useful to optimise for a weighting function of their posterior uncertainties, or use a multi-criteria optimisation.

Also, the optimal network depends on the model underlying the CCDAS. This has an interesting implication. Assume a network is to be constructed to provide observations for assimilation into a particular CCDAS. We then recommend to use that very CCDAS for the network design, because optimality of the network is only guaranteed with respect to the CCDAS it has been designed with. Of course continual improvements in modelling and delays in building network infrastructure mean that the optimal network will always be slightly out of date. This is another reason why general properties of the optimal network are more reliable than specific choices.

The reason for the dependency of the optimal network on the model is that each model can only approximate the truth, i.e. there is model error. Another generalisation from transport inversion systems is that incorporating model error into CCDAS will impact the optimal network. A further recommendation is, hence, first to quantify this uncertainty most accurately and second to reduce it by improving the models underlying the CCDASs. Both is simplified by keeping the CCDAS flexible with respect to the process formulations in the model. Future studies must also pay more careful regard to the statistical coherence of the modelling systems we use.

References

[Bar02] D. J. Barrett. Steady state turnover time of carbon in the Australian terrestrial biosphere. *Glob. Biogeochem. Cycles*, 16(4):doi:10.1029/2002GB001860, 2002.

[Ent02] I. G. Enting. *Inverse Problems in Atmospheric Constituent Transport.* Cambridge University Press, 2002.

[ETF95] I. G. Enting, C. M. Trudinger, and R. J. Francey. A synthesis inversion of the concentration and δ^{13}C of atmospheric CO_2. *Tellus*, 47B:35–52, 1995.

[ETFG93] I. G. Enting, C. M. Trudinger, R. J. Francey, and H. Granek. Synthesis inversion of atmospheric CO_2 using the GISS tracer transport model. Tech. Paper No. 29, CSIRO Div. Atmos. Res., 1993.

[GFPS00] M. Gloor, S. M. Fan, S. Pacala, and J. Sarmiento. Optimal sampling of the atmosphere for purpose of inverse modeling: A model study. *Global Biogeochem. Cycles*, 14:407–428, 2000.

[GK98] R. Giering and T. Kaminski. Recipes for Adjoint Code Construction. *ACM Trans. Math. Software*, 24(4):437–474, 1998.

[GLD+02] K. R. Gurney, R. M. Law, A. S. Denning, P. J. Rayner, D. Baker, P. Bousquet, L. Bruhwiler, Y.-H. Chen, P. Ciais, S. Fan, I. Y. Fung, M. Gloor, M. Heimann, K. Higuchi, J. John, T. Maki, S. Maksyutov, K. Masarie, P. Peylin, M. Prather, B. C. Pak, J. Randerson, J. Sarmiento, S. Taguchi, T. Takahashi, and C.-W. Yuen. Towards robust regional estimates of CO_2 sources and sinks using atmospheric transport models. *Nature*, 415:626–630, 2002.

[GLD+03] K. R. Gurney, R. M. Law, A. S. Denning, P. J. Rayner, D. Baker, P. Bousquet, L. Bruhwiler, Y.-H. Chen, P. Ciais, S. Fan, I. Y. Fung, M. Gloor, M. Heimann, K. Higuchi, J. John, E. Kowalczyk, T. Maki, S. Maksyutov, P. Peylin, M. Prather, B. C. Pak, J. Sarmiento, S. Taguchi, T. Takahashi, and C.-W. Yuen. TransCom 3 CO_2 inversion intercomparison: 1. Annual mean control results and sensitivity to transport and prior flux information. *Tellus*, 55B(2):555–579, 2003. doi:10.1034/j.1600-0560.2003.00049.x.

[GMW81] Philip E. Gill, Walter Murray, and Margret H. Wright. *Practical Optimization.* Academic Press, New York, 1981.

[HBA+04] S. Houweling, F.-M. Breon, I. Aben, C. Rödenbeck, M. Gloor, M. Heimann, and P. Ciais. Inverse modeling of CO_2 sources and sinks using satellite data: a synthetic inter-comparison of measurement techniques and their performance as a function of space and time. *Atmos. Chem. Phys.*, 4:523–538, 2004.

[Hei95] M. Heimann. The global atmospheric tracer model TM2. Technical report no. 10, Deutsches Klimarechenzentrum, Hamburg, Germany (ISSN 0940-9327), 1995.

[HS92] M. Hardt and F. Scherbaum. Optimizing the station distributions of seismic networks for aftershock recordings by simulated annealing. In *EOS Supplement, Transactions of the Fall Meeting, 1992*, page 351. A.G.U., 1992.

[HS94] M. Hardt and F. Scherbaum. The design of optimum networks for aftershock recordings. *Geophys. J. Int.*, 117:716–726, 1994.

[KGS+03] Thomas Kaminski, Ralf Giering, Marko Scholze, Peter Rayner, and Wolfgang Knorr. An example of an automatic differentiation-based modelling system. In V. Kumar, L. Gavrilova, C. J. K. Tan, and P. L'Ecuyer, editors, *Computational Science – ICCSA 2003, International Conference Montreal, Canada, May 2003, Proceedings, Part II*, volume 2668 of *Lecture Notes in Computer Science*, pages 95–104, Berlin, 2003. Springer.

[KH95] W. Knorr and M. Heimann. Impact of drought stress and other factors on seasonal land biosphere CO_2 exchange studied through an atmospheric tracer transport model. *Tellus, Ser. B*, 47(4):471–489, 1995.

[KH01] T. Kaminski and M. Heimann. Inverse modeling of atmospheric carbon dioxide fluxes. *Science*, 294:5541, October 2001.

[KK05] W. Knorr and J. Kattge. Inversion of terrestrial biosphere model parameter values against eddy covariance measurements using Monte Carlo sampling. *Global Change Biology*, 11:1333–1351, 2005.

[KKRH02] Thomas Kaminski, Wolfgang Knorr, Peter Rayner, and Martin Heimann. Assimilating atmospheric data into a terrestrial biosphere model: A case study of the seasonal cycle. *Glob. Biogeochem. Cyc.*, 16:1066, 2002. doi:10.1029/2001GB001463.

[Kno97] W. Knorr. *Satellitengestützte Fernerkundung und Modellierung des globalen CO_2-Austauschs der Landvegetation: Eine Synthese*. PhD thesis, Max-Planck-Institut für Meteorologie, Hamburg, Germany, 1997.

[Kno00] W. Knorr. Annual and interannual CO_2 exchanges of the terrestrial biosphere: process-based simulations and uncertainties. *Glob. Ecol. and Biogeogr.*, 9:225–252, 2000.

[KRHE01] Thomas Kaminski, Peter J. Rayner, Martin Heimann, and Ian G. Enting. On aggregation errors in atmospheric transport inversions. *J. Geophys. Res.*, 106:4703–4715, 2001.

[LRSE02] R. M. Law, P. J. Rayner, L. P. Steele, and I. G. Enting. Using high temporal frequency data for CO_2 inversions. *Global Biogeochem. Cycles*, 16:1053, 2002. doi:10.1029/2001GB001593.

[LRSE03] R. M. Law, P. J. Rayner, L. P. Steele, and I. G. Enting. Data and modelling requirements for CO_2 inversions using high frequency data. *Tellus*, 55B(2):512–521, 2003. doi:10.1034/j.1600-0560.2003.0029.x.

[LRW04] R. M. Law, P. J. Rayner, and Y. P. Wang. Inversion of diurnally-varying synthetic CO_2: network optimisation for an Australian test case. *Global Biogeochem. Cycles*, 18(1):GB1044, 2004. doi:10.1029/2003GB002136.

[MRR+53] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations for fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[Pak04] B. Pak. Parameter optimization using the adjoint of a biosphere model. Abstract A11F-02. *Eos Trans. AGU*, 85(47), December 2004.

[PM02] P. K. Patra and S. Maksyutov. Incremental approach to the optimal network design for CO_2 surface source inversion. *Geophys. Res. Lett.*, 29(10):1459, 2002. doi:10.1029/2001GL013943.

[PMB+03] P. K. Patra, S. Maksyutov, D. Baker, P. Bousquet, L. Bruhwiler, Y-H. Chen, P. Ciais, A. S. Denning, S. Fan, I. Y. Fung, M. Gloor, K. R. Gurney, M. Heimann, K. Higuchi, J. John, R. M. Law, T. Maki, P. Peylin, M. Prather, B. Pak, P. J. Rayner, J. L. Sarmiento, S. Taguchi, T. Takahashi, and C-W. Yuen. Sensitivity of optimal extension of CO_2 observation networks to model transport. *Tellus B*, 55(2):498–511, 2003.

[PRB+05] P. Peylin, P. J. Rayner, P. Bousquet, C. Carouge, F. Hourdin, P. Ciais, P. Heinrich, and AeroCarb Contributors. Daily CO_2 flux estimate over Europe from continuous atmospheric measurements: Part 1 inverse methodology. *Atmos. Chem. Phys.*, 5:3173–3186, 2005.

[Ray04] P. J. Rayner. Optimizing CO_2 observing networks in the presence of model error: results from transcom 3. *Atmos. Chem. Phys.*, 4:413–421, 2004.

[RET96] P. J. Rayner, I. G. Enting, and C. M. Trudinger. Optimizing the CO_2 observing network for constraining sources and sinks. *Tellus*, 48B:433–444, 1996.

[RO01] P. J. Rayner and D. M. O'Brien. The utility of remotely sensed CO_2 concentration data in surface source inversions. *Geophys. Res. Let.*, 28:175–178, 2001.

[RSB+02] J. T. Randerson, C. J. Still, J. J. Balle, I. Y. Fung, S. C. Doney, P. P. Tans, T. J. Conway, J. W. C. White, B. Vaughn, N. Suits, and A. S. Denning. Carbon isotope discrimination of arctic and boreal biomes inferred from remote atmospheric measurements and a biosphere- atmosphere model. *Glob. Biogeochem. Cycles*, 16(3):doi:10.1029/2001GB001435, 2002.

[RSK+05] P. Rayner, M. Scholze, W. Knorr, T. Kaminski, R. Giering, and H. Widmann. Two decades of terrestrial Carbon fluxes from a Carbon Cycle Data Assimilation System (CCDAS). *Global Biogeochemical Cycles*, 19:doi:10.1029/2004GB002254, 2005.

[Sch03] M. Scholze. *Model studies on the response of the terrestrial carbon cycle on climate change and variability*. Examensarbeit, Max-Planck-Institut für Meteorologie, Hamburg, Germany, 2003.

[SPVC03] D. Santaren, P. Peylin, N. Viovy, and P. Ciais. Parameter Estimation in Biogeochemical Surface Model using Nonlinear Inversion: Optimization with Measurements of a Pine Forest. *Geophysical Research Abstracts*, 5:12020, 2003.

[Tar87] A. Tarantola. *Inverse problem theroy - Methods for data fitting and model parameter estimation*. Elsevier Science, New York, USA, 1987.

[VBS01] T. Vukićević, B. H. Braswell, and D. Schimel. A diagnostic study of temperature controls on global terrestrial carbon exchange. *Tellus*, 53B(2):150–170, 2001.

[WLCC01] Y. P. Wang, R. Leuning, H. Cleugh, and P. A. Coppin. Parameter estimation in surface exchange models using non-linear inversion: How many parameters can we estimate and which measurements are most useful? *Glob. Change Biol.*, 7:495–510, 2001.

[WSL+05] M. Williams, P. A. Schwarz, B. E. Law, J. Irvine, and M. R. Kurpius. An improved analysis of forest carbon dynamics using data assimilation. *Glob. Change Biol.*, 11:89–105, 2005.