

Numerical Weather Prediction

Use of GPS Radio Occultation Data in Meteorological Services



Forecasting Research Technical Report No. 403

Christian Marquardt, John Eyre, Sean Healy, Adrian Jupp and David Offiler

email: nwp_publications@metoffice.com

©Crown Copyright

A decorative wavy line that starts on the left, dips down, rises up, and then dips down again towards the right.

USE OF GPS RADIO OCCULTATION DATA IN METEOROLOGICAL SERVICES

C. Marquardt, J. R. Eyre, S. B. Healy, A. Jupp, and D. Offiler

Met Office, Bracknell, London Road, RG12 2SZ, UK

ABSTRACT

GPS radio occultation data are expected to provide a wealth of new data on the atmosphere's temperature and humidity structure. In meteorological services, this information will primarily be used in data assimilation systems which provide an initial state of the atmosphere as starting point for their numerical weather forecast. The majority of data assimilation systems used today are based on variational methods. We therefore review the theoretical basis of variational data assimilation and discuss some implications for the use of radio occultation data. Since "optimal estimation" (or "1DVar") satellite retrievals share the same theoretical background, aspects of the theory can be illustrated by using examples from satellite retrievals. We use a variational retrieval of CHAMP radio occultation data for this purpose.

Key words: data assimilation; radio occultation.

1. INTRODUCTION

Numerical Weather Prediction (NWP) aims at forecasting the weather by running a numerical model of the atmosphere forward in time. Obviously, this is an initial value problem: the better the knowledge of the initial global state of the atmosphere, the better the forecast will be (note, however, that the non-linearity of the problem will still limit the predictability of the atmosphere). The process of constructing an appropriate initial state of the atmosphere from all available meteorological observations is called *data assimilation*. With an increasing number of new observation types becoming available in the future (especially from satellites), data assimilation is one of the major research areas at NWP centres. Newly available data sets – like radio occultation data – will be used as additional input to existing data assimilation systems, aiming at improving the quality of the initial global atmospheric state used for the forecast – and, thereby, of the weather forecast as well.

While satellites have been providing meteorological data operationally since the late 70's, it took until the 90's to show that the assimilation of satellite data does indeed have a consistent positive impact on numerical weather

forecasts in the Northern Hemisphere (Eyre et al., 1993). This proved to be especially difficult because a large number of conventional meteorological data is also available in the Northern Hemisphere. In fact, it was often found that the additional introduction of satellite data in these regions diminished the forecast quality rather than improving it (for a more detailed discussion, see Eyre et al., 1993, and references therein). Only when the variational data assimilation methods became technically feasible at NWP centres in the mid and late 90's, meteorological satellite data could prove its true potential for NWP.

Three dimensional variational data assimilation (3DVar) was introduced at the Met Office in early 1999 (Lorenc et al., 2000), and replaced the earlier "analysis correction" scheme. The impact of the introduction of the 3DVar on the quality of Met Office forecasts is illustrated in Fig. 1. The NWP index shown is a measure of the forecasting skill of the global NWP model over persistence. The time series is based on global model forecasts of selected parameters (mean sea-level pressure, 500 hPa height, and winds at 850 and 250 hPa) out to 5 days ahead for regions covering the whole globe verified by comparison with subsequent analyses. Each individual value is a combination of skill scores for each parameter and based on the preceding 12 months of data.

Also indicated in Fig. 1 are the major changes introduced to the operational suite, along with the improvements associated with them. The largest single improvement of the index occurred in early 1999, when 3DVar and data from the AMSU-A instrument of the ATOVS package were introduced simultaneously (English et al., 2000). The additional use of AMSU-A data over Siberia (mid 1999; also see English et al.), and later over sea ice (2001) also had positive impacts, as had other improvements in the assimilation of satellite data (e.g., the use of SSM/I in late 2000, AMSU-B and ATOVS on NOAA 16 in early 2001, and atmospheric motion vectors in late 2001). The interesting point is that the rate of change in the NWP index increased dramatically after the introduction of the 3DVar and ATOVS data.

We will see in section 3 that it was the introduction of the variational data assimilation scheme that made the improvements of the NWP index possible in first place. There's already a hint in Fig. 1: The second largest increment in the time series was related to the use of raw




Figure 1. Time series of the global NWP index of Met Office forecasts, based on 12 months of data. The dates of the introduction of significant changes to the operational NWP suite are indicated.

radiances from AMSU-A in late 1999 (Eyre et al., 2000). No additional data was assimilated – only the way the data was used changed. Apart from the generality of the variational framework (see the discussion in section 2), the major advantage the variational approach to data assimilation has over its predecessors is that satellite data can be assimilated in a raw, unprocessed form.

Similar experiences were made at other operational weather centres whenever variational data assimilation methods replaced the previously used ones (e.g., Andersson et al., 1994; Derber and Wu, 1998). This is the reason that variational methods have become common for data assimilation worldwide.

In section 2, we review the basic principles of variational data assimilation. We discuss the role of *a priori* data (or information) used in data assimilation and satellite retrievals in section 3. Section 4 is dedicated to a discussion of the importance of error estimates within variational data assimilation. We will also introduce the concept of information content of observational data, and use it to discuss the impact radio occultation soundings might have on NWP.

Both data assimilation and satellite retrieval aim at solving the same problem: Estimating the state of the atmosphere from one or many, possibly different, measurements. In the following, we will refer to this as the *data assimilation* (or retrieval) *problem*. Variational data assimilation, as formulated by Lorenc (1986), and satellite retrievals based on optimal estimation theory as introduced by Rodgers (1976, 1990) share the same mathematical background. Along with the theoretical discussion, we will use examples from a variational retrieval of CHAMP radio occultation data to illustrate some aspects of the theory.

2. VARIATIONAL DATA ASSIMILATION

In order to understand the basics of variational data assimilation, and also to discuss some implications for new satellite data like radio occultation soundings, we state the data assimilation problem in a formal, but very general way. Assume that we have a set of measurements \mathbf{y}_o which are related to the state of the atmosphere \mathbf{x} (i.e., a set of relevant atmospheric parameters). We then define the solution of the data assimilation problem as the *most likely* state $\tilde{\mathbf{x}}$ of the atmosphere in light of the given measurements. Formally, this can be expressed by means of conditional probabilities: Let $P(\mathbf{x}|\mathbf{y})$ denote the probability for an atmospheric state \mathbf{x} , given a set of measurements \mathbf{y} . The solution of the data assimilation problem for a specific set of measurements \mathbf{y}_o is then given by the specific state $\tilde{\mathbf{x}}$ which maximises the probability distribution function $P(\mathbf{x}|\mathbf{y}_o)$. This is sometimes referred to as the *maximum likelihood* solution of the data assimilation problem.

We also assume that we have some *a priori* (or “background”) information on the likely state of the atmosphere. We express this in terms of a prior probability density function $P(\mathbf{x})$. Bayes’ theorem allows us to write $P(\mathbf{x}|\mathbf{y}_o)$ as

$$P(\mathbf{x}|\mathbf{y}_o) = P(\mathbf{y}_o|\mathbf{x}) \cdot P(\mathbf{x})/P(\mathbf{y}_o). \quad (1)$$

Here, $P(\mathbf{y}_o|\mathbf{x})$ represents the probability of the measurement, given a specific atmospheric state. $P(\mathbf{y}_o)$ denotes the prior probability of the measurement, which is merely a normalisation factor¹. Taking the logarithm and multiplying by -1 , we find that $\tilde{\mathbf{x}}$ minimises the *cost function*

$$\begin{aligned} J(\mathbf{x}) &= -\log\{P(\mathbf{x}|\mathbf{y}_o)\} - c \\ &= -\log\{P(\mathbf{y}_o|\mathbf{x})\} - \log\{P(\mathbf{x})\}, \end{aligned} \quad (2)$$

where $c = \log(P(\mathbf{y}_o)) = \text{const.}$

In variational analysis, observations and meteorological variables are related by means of *forward operators*. A forward operator \mathbf{H} calculates how a measurement $\mathbf{y} = \mathbf{H}[\mathbf{x}]$ would look like for a given atmospheric state \mathbf{x} . Thus, forward operators contain the physics of the measurement. In case of radio occultation measurements, the forward operator used in an operational system might contain the calculation of the raw GPS signals, bending angle, or refractivity calculated from NWP fields.

In practice, it is often assumed that the probability functions introduced above are all Gaussian. By inserting the appropriate definitions of multivariate normal probability distribution functions into eq. (2), the cost function can be written as

$$\begin{aligned} J(\mathbf{x}) &= \frac{1}{2} (\mathbf{y}_o - \mathbf{H}[\mathbf{x}])^T (\mathbf{E} + \mathbf{F})^{-1} (\mathbf{y}_o - \mathbf{H}[\mathbf{x}]) \\ &\quad + \frac{1}{2} (\mathbf{x} - \mathbf{x}_b)^T \mathbf{B}^{-1} (\mathbf{x} - \mathbf{x}_b). \end{aligned} \quad (3)$$

¹ $P(\mathbf{y}_o) = \int P(\mathbf{y}_o|\mathbf{x}) P(\mathbf{x}) d\mathbf{x}$ is independent of \mathbf{x} . As we are interested in finding the location of maximum probability of $P(\mathbf{x}|\mathbf{y}_o)$ with respect to \mathbf{x} , we simply regard $P(\mathbf{y}_o)$ as an (unknown) constant ensuring that $\int P(\mathbf{x}|\mathbf{y}_o) d\mathbf{x} = 1$.

Here, \mathbf{E} , \mathbf{F} and \mathbf{B} denote *error covariance matrices*, describing the assumed uncertainties of the measurements, the forward model, and the background data, respectively. By specifying matrices we do not only take standard deviations of, e.g., the background's atmospheric parameters into account. Possible spatial correlations of the errors in one variable as well as cross correlations between different variables are also considered. Thus, the first term (obtained from $P(\mathbf{y}_o|\mathbf{x})$) measures the deviation between the actual measurement \mathbf{y}_o and the would-be measurement $\mathbf{H}[\mathbf{x}]$, weighted by the instrument's and forward model errors. The deviation between a current atmospheric state \mathbf{x} and the *a priori* \mathbf{x}_b , weighted by the expected background errors, is given by the second term (obtained from $P(\mathbf{x})$). Eq. 3 allows another common interpretation of the variational formalism: In the linear limit, we minimise the total error by minimising (3). Thus, the solution $\tilde{\mathbf{x}}$ obtained from a variational analysis is a least square fit to both the observations and the *a priori*.

Variational data assimilation – or variational satellite retrievals – exploit the cost function formulation (3) directly by numerically minimising the cost function calculated from observations and *a priori*. For the global data assimilation problem, all available observations available within a window of, say, six hours centred around a synoptic time may be used to generate a global state of the atmosphere valid at that synoptic time (“3DVar”). For satellite data, which are not bound to traditional synoptic times, an even better approach is to take the temporal evolution of the atmosphere into account (“4DVar”). This requires the NWP model to become part of the minimisation cycle of the cost function, and is computationally significantly more expensive. Only few weather centres (like ECMWF) are presently able to run such a scheme operationally. The Met Office is currently running a 3DVar scheme, and is preparing the introduction of its 4DVar system. The use of satellite data, however, is identical in both 3D and 4DVar: their respective forward operators are used to calculate their contribution to the observations' part of the cost function (3). The setup of the data assimilation systems takes care of feeding the forward models with appropriate atmospheric variables, and of finding the maximum likelihood solution of the data assimilation problem.

In satellite retrievals, it is usually a single profile that is retrieved from a satellite sounding. Originally known as “optimal estimation” retrievals (Rodgers, 1976, 1990), variational retrievals have also been called “1DVar” retrievals (Eyre et al., 1993), obviously referring to the nomenclature used within the data assimilation community. For limb sounders scanning the atmosphere vertically within a 2 dimensional plane, a possible extension of a 1DVar scheme is the simultaneous retrieval of multiple profiles, all lying within the observation plane. Such a scheme (which could well be called a “2DVar”) has been proposed for the EOS-MLS instrument (Livesey and Read, 2000).

The assumption of normally distributed and unbiased errors is crucial in the derivation of the quadratic cost function (3). In order to ensure that these conditions are met by observations, NWP centres have extensive quality control procedures in place. These aim at identifying

and removing possible outliers from the observations. In addition, measurements are constantly monitored for biases. If necessary, bias corrections are applied before data enters data assimilation systems, primarily taking care of possible systematic errors in the forward operators (Eyre, 1992). Systematic errors of the observations (caused by, e.g., instrument drifts or degradations), however, can also be corrected. Thus, the physical origin of possible biases needs to be understood in order to be able to construct proper bias correction schemes. Understanding the cause of the tropospheric refractivity bias of radio occultation soundings which was found in both GPS/MET (Rocken et al., 1997) and CHAMP data (Marquardt et al., 2003) therefore is of particular importance for the assimilation of radio occultation data. Without being able to bias correct tropospheric radio occultation data, these probably won't be used in operational NWP systems, at least if the lower tropospheric bias should continue to be a problem.

The forward operators $\mathbf{H}(\mathbf{x})$ used within variational data assimilation can be nonlinear and may describe complicated physical relationships between meteorological variables and satellite measurements. Specifically, they allow the use of raw, unprocessed satellite measurements. This is an advantage compared to previously used optimal estimation or analysis correction schemes: those could only handle observations if the latter were converted to model variables. Thus, they required the use of atmospheric parameters retrieved from satellite soundings.

Forward simulations of raw satellite measurements can, of course, be quite complex and expensive in terms of CPU time. One possibility to assimilate radio occultation data is to simulate bending angles by using a ray tracer through the (model) atmosphere, matching both the receiving and the transmitting satellite. Zou et al. (2000) used this approach in a 3DVar, but the assimilation of 30 GPS/MET profiles increased the computation time for the analysis cycle by a factor of 10. This is prohibitive in an operational environment. The use of fast forward models like the one originally proposed by Eyre (1994), or the recently developed one by Healy and Eyre (2003), is clearly required. Another problem is that physical effects affecting the observations (like the ionosphere in case of radio occultation soundings) are currently not represented in NWP systems. Thus, the assimilation of refractivity, exhibiting a significant simpler and faster forward operator, might also be considered (Zou et al., for example, proposed a hybrid scheme which uses both bending angles and refractivity). This, of course, comes at the price of using preprocessed radio occultation data.

Finally, we emphasize the generality of the variational framework. As pointed out by, e.g., Rodgers (2000), the variational approach is not just another algorithm to retrieve atmospheric parameters from a satellite retrieval. It rather “encompasses all inverse methods by providing a way of characterising the class of possible solutions, considering all possible states, and assigning a probability to each”. In other words, the variational approach provides the most fundamental framework to satellite retrievals possible. In practice, this means that any reasonable well-tuned variational retrieval will be at least as “good” (when statistically compared with independent meteorological data over long periods) as any arbitrary non-variational

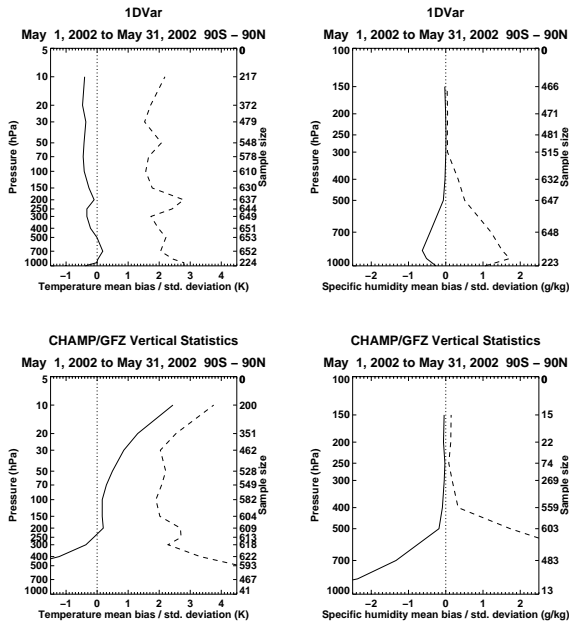


Figure 2. Statistical comparison of CHAMP temperature (left column) and specific humidity (right column) profiles obtained from a variational retrieval (top row) and from the classical retrieval as run at GFZ Potsdam (bottom row) against nearby radiosonde ascents during May 2002. Shown are mean error (thick) and standard deviation (dashed).

retrieval based on comparable information sources². We give an example using CHAMP data: Fig. 2 shows results from a validation of CHAMP data against nearby radio sonde ascents for May 2002. Retrievals were either obtained by a 1DVar (Marquardt and Healy, 2003), or by the classical retrieval operationally run at GFZ Potsdam (Wickert et al., 2001). The results discussed here for an individual month are representative for the complete CHAMP data set available so far. The statistics have been calculated from those cases where a radiosonde ascent and a CHAMP profile occurred within a 6 hour window, and were not separated by more than 300 km. The variational retrieval, which uses ECMWF analyses as *a priori*, provides stratospheric temperature retrievals with a mean bias of less than 0.5 K at all altitudes up to 10 hPa (about 30 km), with standard deviations of less than 2 K. The classical, non-optimal retrieval exhibits a warm bias above 50 hPa, exceeding 2 K at the 10 hPa level; standard deviations are also larger at all altitudes. Note that the classical “dry temperature” retrieval, in contrast to the 1DVar, does not provide any information on temperature as soon as water vapour abundance is significant in the troposphere (hence the tropospheric cold bias in the dry temperature statistics). Specific humidity profiles obtained from the variational retrieval have a smaller bias and exhibit significantly smaller standard deviations. For the non-optimal retrieval, both the temperature bias in the stratosphere and the poor performance of the water vapor retrieval were noted earlier (Marquardt et al., 2003), and

²This, of course, is only true if our original assumption that the true state of the atmosphere is indeed the most likely one. There are other ways of specifying how “the” solution can be characterized, e.g. in terms of maximising resolution (Backus and Gilbert, 1970).

were also present in the GPS/MET data set (Marquardt et al., 2001).

Note that we do not argue that the non-optimal retrieval cannot be improved: Marquardt et al. (2003) give an example of a non-optimal temperature retrieval tuned towards ECMWF analyses, exhibiting reduced biases and rms deviations for dry temperature. The authors emphasize that the improvement of the error statistics strongly depends on how and which external data or information is used within the retrieval. This highlights the importance of *a priori* data (which is discussed in the following section) within non-optimal radio occultation retrievals. While the non-optimal retrieval requires constant re-tuning, the variational retrieval yields much better validation statistics right from the start.

3. A PRIORI INFORMATION

It is important to keep in mind that satellite based instruments do not directly measure geophysical variables like temperature, humidity, or pressure. GPS radio occultations do not measure refractivity or even bending angles. Instead, measurements consist of amplitudes and phases of an electromagnetic signal. In case of nadir sounding instruments, radiances from a certain frequency range in the microwave or infrared portion of the electromagnetic spectrum are observed. This is the reason why a retrieval – the estimation of atmospheric parameters consistent with the measurements – is required.

Often, it is rather straightforward to construct the forward operator. Solving the inverse problem is usually more difficult, because: a) a physical ambiguity may exist, like the well-known water vapor ambiguity in GPS radio occultations (which is similarly common for infrared and microwave sounding instruments); b) measurements contain instrumental or measurement noise. Radio occultation soundings, for example, exhibit strong fluctuations in “measured” bending angles above 40 km (e.g., Hocke, 1997), commonly attributed to residual ionospheric effects. In both cases, the inverse problem is underdetermined. Some information, data or assumption independent of the measurement is required in order to solve the ambiguity, or to separate signal from noise. Within NWP, any such information is called *a priori*.

An obvious example of *a priori* usage in a retrieval is the use of meteorological analysis or forecast data; another one is the use of a climatology. From an NWP perspective, there is no fundamental difference in using either of the two; apart from the fact that the analysis or forecast is very likely to be closer to “truth” than the climatology. More subtle examples of the use of *a priori* include

- any assumption on symmetry (e.g., using spherical symmetry for calculating bending angles from GPS measurements, or refractivity from bending angles)
- any assumption on smoothness (e.g., applying a filter to excess path delays prior to calculating excess doppler shifts)

- any assumption on a specific functional form of either measurements or atmospheric parameters (e.g., “fitting an exponential to bending angles / refractivity”, or “approximating tropospheric temperature by a quadratic (or other) polynomial”)
- any tuning of parameters in order to improve validation statistics (the data against which the validation has been undertaken becomes part of the retrieval, as it helped to determine the setting of the parameters)

As the above list indicates, any retrieval of geophysical parameters includes *a priori*. Thus, while often hidden, its use is unavoidable in satellite retrievals and data assimilation. This does not imply that the use of *a priori* will necessarily pose a problem: assuming refractivity varies exponentially with height over a few hundred meters might be a reasonable approximation of truth; assuming the same relationship over several ten km might introduce significant errors. Ideally, each instance of *a priori* use should be checked carefully for its effect on retrievals or analyses. Within the variational approach, the issue of *a priori* is, at least, dealt with openly, by making clear specifications on the assumed background information and its error statistics. Problems associated with the use of *a priori* are easily spotted.

When *a priori* is used in non-optimal retrievals, if noticed at all, it is often (and often silently) assumed that it has no degrading effect on the accuracy of retrievals. This assumption, however, is misleading. Following Eyre (1987), we write the linearised arbitrary satellite retrieval as

$$\tilde{\mathbf{x}} - \mathbf{x}_b = \mathbf{K}(\mathbf{y}_o - \mathbf{H}[\mathbf{x}_b]) , \quad (4)$$

where \mathbf{K} is the linearised “inversion” operator applied to the raw measurements of the instrument. We also linearize the forward equation to obtain

$$\mathbf{y}_o - \mathbf{H}[\mathbf{x}_b] = \mathbf{H}'(\mathbf{x}_t - \mathbf{x}_b) + \varepsilon , \quad (5)$$

where \mathbf{x}_t and $\mathbf{H}' = \nabla_{\mathbf{x}}\mathbf{H}[\mathbf{x}]$ denote the true state of the atmosphere and the linearised forward operator, respectively. ε denotes the measurement error. Combining (4) and (5) and rearranging gives an expression for the retrieval error (\mathbf{I} is the unit matrix):

$$\tilde{\mathbf{x}} - \mathbf{x}_t = (\mathbf{I} - \mathbf{K}\mathbf{H}')(\mathbf{x}_b - \mathbf{x}_t) + \mathbf{K}\varepsilon . \quad (6)$$

Thus, two terms contribute to the retrieval error: The measurement error, propagated by the retrieval procedure ($\mathbf{K}\varepsilon$); and an error related to the use of *a priori*. As $\mathbf{K}\mathbf{H}' \neq \mathbf{I}$ in general, any use of *a priori* will contribute to – and possibly increase – the retrieval error.

This has particular importance when using mean fields derived from satellite retrievals, e.g. for climatological purposes. Averaging eq. (6) over a large number of retrievals gives

$$\overline{\tilde{\mathbf{x}} - \mathbf{x}_t} = (\mathbf{I} - \mathbf{K}\mathbf{H}')(\overline{\mathbf{x}_b - \mathbf{x}_t}) + \mathbf{K}\overline{\varepsilon} . \quad (7)$$

Thus, climatologies based on satellite retrievals may suffer from systematic biases, even though they are based on perfectly unbiased measurements – simply because of the use of *a priori*.




Figure 3. Anomaly correlation coefficient during June 1999 for Southern Hemisphere 500 hPa geopotential heights when assimilating AMSU-A data as retrieved temperatures (solid) vs. radiances (dashed; from Eyre et al., 2000).

The obvious consequence for satellite retrievals is to use the best *a priori* which is available. In most cases, this will be a short range forecast from a NWP model rather than a climatology or some analytical assumptions.

For the assimilation of satellite data into NWP systems, the above finding suggests that any retrieved or preprocessed satellite data is contaminated with the *a priori* used in the retrieval, and has the potential to degrade the analysis’ quality. In order to make best use of satellite data within NWP, it seems advisable to use satellite data in its rawest, least processed form.

Using raw measurements rather than retrieved atmospheric quantities proved to be a key factor for the success of variational methods in data assimilation. Fig. 3 shows anomaly correlation coefficients (ACC’s) for the 500 hPa height field on the Southern Hemisphere during June 1999 for two different ways of exploiting AMSU-A observations within the Met Offices data assimilation system. The ACC is a commonly used measure for the forecast skill of numerical weather forecasts. It aims at quantifying the amount of similarity between forecast and analysed (“true”) synoptical patterns. As a rule of thumb, forecasters gain useful information from numerical forecast fields as long as the ACC is above 0.6 (for a more detailed discussion of ACC’s see, e.g., Wilks, 1995). By plotting the ACC of a series of forecasts against forecast period, some idea can be obtained on how far into the future a numerical weather prediction is useful.

In the early setup of the Met Office’s 3DVar, ATOVS temperature profiles were retrieved in a 1DVar; the resulting profiles were then assimilated in the 3DVar (English et al., 2000). According to Fig. 3, the assimilation of raw radiances instead increased the forecast period considered to be useful by more than 12 hours.

Note that even an optimal 1DVar retrieval (as in the Met Office early setup) is not able to cure the *a priori* problem. The reason is that error correlations between adjacent retrievals were not taken into account properly. While this can be done in principle (Joiner and Dee, 2000), the process is tedious; it is certainly much simpler to use raw satellite data and avoid the difficulties in setting up appropriate error correlations for the retrieved quantities.

For radio occultation soundings, this suggests to use bending angles (Eyre, 1994) or even excess path delays (and amplitudes) as variables to be assimilated in NWP. On the other hand, from an operational point of view it is advantageous to use simple forward operators, i.e. using retrieved refractivity profiles. Research at data assimilation centres will have to identify the optimal balance between the simplicity of forward operators and the degrading effects of *a priori* used in the preprocessing / retrieval of the data.

4. ERROR ESTIMATES AND INFORMATION CONTENT

In the linear case, the solution of the data assimilation problem can be written as (e.g., Rodgers, 1976; Lorenc, 1986)

$$\tilde{\mathbf{x}} = (\mathbf{H}'^T(\mathbf{E} + \mathbf{F})^{-1}\mathbf{H}' + \mathbf{B}^{-1})^{-1} \cdot (\mathbf{H}'^T(\mathbf{E} + \mathbf{F})^{-1}\mathbf{y}_o + \mathbf{B}^{-1}\mathbf{x}_b) \quad (8)$$

Thus, the analysis provided by a variational data analysis (or 1DVar) can be understood as a weighted sum of observations and *a priori*, where the weighting factors are determined by the error covariances: It is the balance of the error statistics that determines where and how much the observations will affect the analysis. The theory also provides an estimate of the error covariance \mathbf{P} of the analysis, given by

$$\mathbf{P}^{-1} = \mathbf{B}^{-1} + \mathbf{H}'^T(\mathbf{E} + \mathbf{F})^{-1}\mathbf{H}' \quad (9)$$

Note that the accuracy (i.e., the inverse error) of the analysis is better than any of the accuracies of the background or the measurement alone.

The above interpretation of eq. (8) raises the question on how much a variational retrieval actually depends on the *a priori*. After all, two different *a priori*'s will give two different retrievals. This seems to be contradictory to our previous statement that a variational data assimilation (or retrieval) system will deliver a solution close to “the true state” of the atmosphere.

However, the resulting analyses (or retrievals) will agree within their errors, as given by (9). As the error of the variational solution is smaller than that of the *a priori*'s involved, the two solutions will agree better with each other than the two *a priori*'s (at least if the different *a priori*'s errors have been specified sufficiently realistically and are of comparable accuracy). An example using CHAMP data is given in Fig. 4. Apart from the already introduced ECMWF-based retrieval, a 1DVar based on Healy and Eyre (2000), adapted to the Met Office's NWP model, was used to retrieve a temperature profile from CHAMP refractivity data. The *a priori* data used in the retrieval are also shown. Note that the two variational retrievals use different vertical levels, were independently developed, and use different *a priori* data and error estimates. As the theory predicts, the two retrieved temperature profiles agree better with each other than the two *a*




Figure 4. Two 1DVar retrievals of a CHAMP sounding from May 30, 2001, 15:03 UT, located at 48.0°N, 86.6°W, using *a priori* data sets from the Met Office (red) and ECMWF (black), respectively. The background temperature profiles used in the retrievals are also shown (Met Office blue, ECMWF green).

priori's do, especially in the tropopause and lower stratosphere region. They also agree well within their error estimates (not shown). Thus, the additional information provided by the CHAMP refractivity data has improved the different *a priori* data sets towards a very similar solution.

Higher up, however, the different retrievals follow their respective backgrounds rather closely, even in cases when these differ substantially from each other. While the CHAMP measurements have obviously added some information to our knowledge in the tropopause and lower stratosphere region, less additional information can be obtained in the mid or upper stratosphere. This introduces the concept of *information content* of remote sensing measurements.

“Information” is relative to what we already know (i.e., our *a priori* knowledge of the atmosphere; see Eyre, 1990). One way of measuring the information content of a remote sensing measurement compared to the *a priori* is to relate the uncertainty of the retrieved meteorological quantity to the error assumed for the *a priori*. Fig. 5 shows the retrieval error as percentage of ECMWF background errors for the CHAMP profile already known from Fig. 4. Similar figures based on a theoretical study were presented by Healy and Eyre (2000). The reduction of the temperature error is quite significant in the upper troposphere and lower stratosphere, but decreases at mid stratospheric altitudes above 20 hPa because refractivity errors increase above. For humidity, the largest improvement can be expected between the 700 and the 500 hPa level. Above, humidity's contribution to refractivity is weak; below, refractivity errors are large because of multipath effects, while background errors are small. Therefore, we expect radio occultation data to have the largest impact in data assimilation systems around the tropopause and in the lower stratosphere for temperature, and in the mid troposphere for humidity (but see below).

Comparisons of the information content of different remote sensing instruments were undertaken by, e.g., Eyre (1990), in order to assess the benefits of combining in-




Figure 5. Estimated retrieval error of temperature (top) and specific humidity (bottom) relative to the ECMWF background for the CHAMP sounding from Fig. 4.

struments with different physical properties. A study on the combined effect of radio occultations and an upcoming nadir looking infrared instrument (the Infrared Atmospheric Sounding Interferometer IASI) was recently undertaken by Collard and Healy (2003). They concluded that radio occultation and IASI data will not only provide similarly accurate, but also complementary, data to NWP.

Studies on information content depend on the error estimates given for observations and background; the ones presented here and in Collard and Healy (2003) are based on error estimates by Kursinski et al. (1997). Recent experiences with the GPS instrument onboard CHAMP (Marquardt et al., 2003) suggest that tracking errors in the receivers onboard software are likely to be responsible for systematic biases and large errors in the bending angle and refractivity profiles in the mid and lower troposphere. Additional error sources like horizontal gradients (e.g., Healy, 2001) and superrefraction (Kursinski et al., 2000) also increase radio occultation errors in these altitude regions. Thus, the above conclusions with regard to the value of radio occultation data in the troposphere might be overly optimistic.

It cannot be overstressed that accurate error estimates for radio occultations taking instrument issues as well as atmospheric conditions properly into account are crucial for the assimilation of radio occultation data. Too optimistic error estimates, while looking promising on first sight, will eventually cause quality control mechanisms to reject an unnecessary large part of radio occultation data, and in effect reduce the possible importance of radio occultation data in NWP. Overly pessimistic error estimates, on the other hand, will unnecessarily limit the

impact radio occultations could have in data assimilation.

5. SUMMARY

Variational data assimilation has become the de-facto standard in meteorological services, mainly because the large positive impact on weather forecasts the introduction of variational schemes has proven to have. Especially, it allowed meteorological satellite data to develop its full potential within NWP.

Discussing the role of *a priori* in satellite retrievals, we found that its use is usually unavoidable, but will contribute to the retrieval error. Variational data assimilation schemes allow the use of raw satellite data, and thereby minimise the potentially degrading influence of *a priori* data used within the retrieval of atmospheric variables from satellite data. The generality of the variational approach also allows its use in the retrieval of individual atmospheric profiles from satellite data. Variational retrievals, also known as “optimal estimation” retrievals, usually deliver better error statistics than their non-optimal counterparts. The reason is that the difficult issue of dealing with *a priori* (which is done openly in the variational approach) is strictly based on the relative accuracy of observations and *a priori*, as it is within data assimilation. A central role within variational data assimilation, therefore, is held by error estimates for both observations and *a priori*, as the balance of the two determines if and where a certain data set will affect the analysis.

New types of remote sensing data, like radio occultations, will be used as additional input to the already existing data assimilation schemes. Apart from the forward operators describing the physics of the measurement, error estimates for the observations are required. For radio occultation data, the above suggests that bending angles or even the excess path delay and amplitude data measured by GPS receivers might be the preferred “raw” observations to be assimilated. Having in mind that the forward simulation of these, when done with a full ray tracer, is expensive in terms of computing time, significantly faster forward models will be used. A simpler alternative could be the use of refractivity, even though this will come at the price of dealing with the additional *a priori* required to calculate refractivity from bending angle.

Finally, the variational framework provides a number of useful diagnostics tools; we have discussed the information content of radio occultation data as an example. A comparison of the merits of different remote sensing instruments indicates that these do not compete against each other, but rather provide complementary information to data assimilation systems. It is the combination of the various strengths of different remote sensing systems that has made data assimilation so successful.

ACKNOWLEDGMENTS

We are grateful to GFZ Potsdam for providing the CHAMP data on which the illustrating examples are

based.

REFERENCES

- Andersson, E., Pailleux, J., Thépaut, J.-N., Eyre, J., McNally, A., Kelly, G., and Courtier, P., Use of cloud-cleared radiances in three/four-dimensional variational data assimilation, *Quart. J. Roy. Meteorol. Soc.*, **120**, 627–653, 1994.
- Backus, G. and Gilbert, J., Uniqueness in the inversion of inaccurate gross earth data, *Phil. Trans. R. Soc. Lond.*, **266**, 123, 1970.
- Collard, A. and Healy, S., The combined impact of future space-based atmospheric sounding instruments on numerical weather prediction, accepted for publication in *Quart. J. Roy. Meteorol. Soc.*, 2003.
- Derber, J. and Wu, W., The use of TOVS cloud-cleared radiances in the NCEP SSI analysis, *Mon. Wea. Rev.*, **126**, 2287–2299, 1998.
- English, S., Renshaw, R., Dibben, P., Smith, A., Rayer, P., Poulsen, C., Saunders, F., and Eyre, J., A comparison of the impact of TOVS and ATOVS satellite sounding data on the accuracy of numerical weather forecasts, *Quart. J. Roy. Meteorol. Soc.*, **126**, 2911–2931, 2000.
- Eyre, J., On systematic errors in satellite sounding products and their climatological mean values, *Quart. J. Roy. Meteorol. Soc.*, **113**, 279–292, 1987.
- Eyre, J., The information content of data from satellite sounding systems: A simulation study, *Quart. J. Roy. Meteorol. Soc.*, **116**, 401–434, 1990.
- Eyre, J., A bias correction scheme for simulated TOVS brightness temperatures, ECMWF Technical Memorandum No. 186, 1992.
- Eyre, J., Assimilation of radio occultation measurements into a numerical weather prediction system, ECMWF Technical Memorandum No. 199, 1994.
- Eyre, J., Kelly, G., McNally, A., Andersson, E., and Persson, A., Assimilation of TOVS radiance information through one-dimensional variational analysis, *Quart. J. Roy. Meteorol. Soc.*, **119**, 1427–1463, 1993.
- Eyre, J., English, S., P. Butterworth, Renshaw, R., Ridley, J., and Ringer, M., Recent progress in the use of satellite data in NWP, NWP Technical Report No. 296, available from the National Meteorological Library, Bracknell, UK, 2000.
- Healy, S., Radio occultation bending angle and impact parameter errors caused by horizontal refractive index gradients in the troposphere: A simulation study, *J. Geophys. Res.*, **106**, 11 875–11 889, 2001.
- Healy, S. and Eyre, J., Retrieving temperature, water vapor and surface pressure information from refractive-index profiles derived by radio occultation: A simulation study, *Quart. J. Roy. Meteorol. Soc.*, **126**, 1661–1683, 2000.
- Healy, S. and Eyre, J., A forward model for the assimilation of GPS radio occultation bending angle measurements, submitted to *Quart. J. Roy. Meteorol. Soc.*, 2003.
- Hocke, K., Inversion of GPS meteorology data, *Ann. Geophys.*, **15**, 443–450, 1997.
- Joiner, J. and Dee, D., An error analysis of radiance and suboptimal retrieval assimilation, *Quart. J. Roy. Meteorol. Soc.*, **126**, 1495–1514, 2000.
- Kursinski, E., Hajj, G., Schofield, J., Linfield, R., and Hardy, K., Observing earth's atmosphere with radio occultation measurements using the Global Positioning System, *J. Geophys. Res.*, **102**, 23 429–23 465, 1997.
- Kursinski, E., Hajj, G., Leroy, S., and Herman, B., The GPS radio occultation technique, *Terr. Atmos. Ocean. Sci.*, **11**, 53–114, 2000.
- Livesey, N. and Read, W., Direct retrieval of line-of-sight atmospheric structure from limb sounding observations, *Geophys. Res. Lett.*, **27**, 891–894, 2000.
- Lorenc, A., Analysis methods for numerical weather prediction, *Quart. J. Roy. Meteorol. Soc.*, **112**, 1177–1194, 1986.
- Lorenc, A., Ballard, S., Bell, R., Ingleby, N., Andrews, P., Barker, D., Bray, J., Clayton, A., Dalby, T., Li, D., Payne, T., and Saunders, F., The Met. Office global three-dimensional variational data assimilation scheme, *Quart. J. Roy. Meteorol. Soc.*, **126**, 2991–3012, 2000.
- Marquardt, C. and Healy, S., Variational retrieval of CHAMP radio occultation data, in prep., 2003.
- Marquardt, C., Labitzke, K., Reigber, C., Schmidt, T., and Wickert, J., An assessment of the quality of GPS/MET radio limb soundings during February 1997, *Phys. Chem. Earth*, **26**, 125–130, 2001.
- Marquardt, C., Schöllhammer, K., Beyerle, G., Schmidt, T., Wickert, J., and Reigber, C., Validation and data quality of CHAMP radio occultation data, in *First CHAMP mission results for gravity, magnetic and atmospheric studies*, 384–396, Springer, 2003.
- Rocken, C., Anthes, R., Exner, M., Hunt, D., Sokolovsky, S., Ware, R., Gorbunov, M., Schreiner, W., Feng, D., Herman, B., Kuo, Y.-H., and Zou, X., Analysis and validation of GPS/MET data in the neutral atmosphere, *J. Geophys. Res.*, **102**, 29 849–29 866, 1997.
- Rodgers, C., Retrieval of atmospheric temperature and composition from remote sounding measurements of thermal radiation, *Rev. Geophys. Space Phys.*, **14**, 609–624, 1976.
- Rodgers, C., Characterization and error analysis of profiles retrieved from remote sounding measurements, *J. Geophys. Res.*, **95**, 5587–5595, 1990.
- Rodgers, C., *Inverse methods for atmospheric sounding: Theory and practice*, World Scientific Publishing, Singapore, New Jersey, London, Hong Kong, 2000.
- Wickert, J., Reigber, C., Beyerle, G., König, R., Marquardt, C., Schmidt, T., Grunwaldt, L., Galas, R., Meehan, T., Melbourne, W., and Hocke, K., Atmosphere sounding by GPS radio occultation: First results from CHAMP, *Geophys. Res. Lett.*, **28**, 3263–3266, 2001.
- Wilks, D., *Statistical methods in the atmospheric sciences*, Academic Press, San Diego, 1995.
- Zou, X., Wang, B., Liu, H., Anthes, R., Matsumura, T., and Zhu, Y.-J., Use of GPS/MET refraction angles in three-dimensional variational analysis, *Quart. J. Roy. Meteorol. Soc.*, **126**, 3013–3040, 2000.