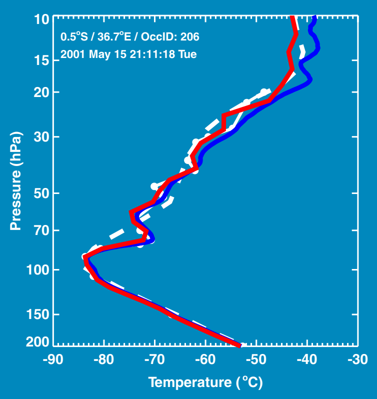

PROGRAMME DEVELOPMENT DEPARTMENT
TECHNICAL MEMORANDUM NO.12
FEBRUARY 2005

MEMORANDUM

TM.12

GRAS LEVEL 1B PRODUCT VALIDATION WITH 1D-VAR RETRIEVAL

CHRISTIAN MARQUARDT
SEAN HEALY
JUHA-PEKKA LUNTAMA
EAMONN MCKERNAN

EUMETSAT

Contents

1 Introduction

This Technical Memorandum presents the results of the “GRAS Level 1 Validation” study started by EUMETSAT with the Met Office in March 2003. The objectives of this study were to define a prototype algorithm for operational validation of the GRAS level 1 products, produce a prototype software package to test the algorithm, and demostrate the validation results. The study was completed in July 2004. The results of the study were used to produce the GRAS Calibration and Validation Plan for the operational validation of the GRAS level 1b data products (EUM, 2004).

One of the main interests in the “GRAS Level 1 validation study” was to assess whether it would be feasible to use 1D-Var retrieval to validate the bending angle profiles measured by the GRAS instrument. Because it was also considered important to verify the assessment with real RO data, a prototype 1D-Var software tool was developed as part of the study. This prototype was used to validate data from the CHAMP satellite mission. The validation study also indetified other possible validation methods (e.g. analysis of noise levels, direct intercomparisons) for GRAS products. The objective was to address a large number of methods without immediately screening out apparently difficult, expensive, or otherwise inconvenient techniques or data sources. The basis for later selection of the most promising methods is provided in the discussion in the memorandum. Finally, the study defined an overall approach for the GRAS level 1b product validation.

Section 2 this memorandum outlines the potential use of 1D-Var for the validation of Gras level 1b measurements. This section starts with a short introduction into the 1D-Var retrieval. However, the main purpose of this section is to describe the “diagnostics” provided by 1D-Var that are used in the product validation. This section also contains some discussion about the practical use of 1D-Var in data quality control and a concise description of the implementation of 1D-Var for RO bending angles. The main benefits of the 1D-Var retrieval in validation are:

- 1D-Var provides mathematically well defined “diagnostics” for identification of potential measurement or data processing errors;
- 1D-Var can be made fully automatic for operational use;
- Global NWP background allows validation of every GRAS measurement and rapid generation of validation statistics;
- 1D-Var allows verification of the GRAS product error characteristics;
- 1D-Var generates geophysical products that can be used in intercomparison with other measurement data.

The overall GRAS level 1b validation validation strategy outlined in Section 3 contains a stepwise approach for the validation activities during the pre-launch, commissioning, and operations phases of the mission. A key to successful validation is to establish a traceability chain for the data that is to be validated. In the case of GRAS, this means tracing the error characteristics from all elements of the GRAS measurement system to the level 1b products. Similar traceability must be established to the geophysical products if they are to be used in the validation (e.g. intercomparison with other remote sensing data). During the pre–launch phase of the EPS GRAS mission, emphasis in the validation activities is put on deriving a full error characterisation of calibrated and pre–processed GRAS level 1b data, and on testing the implemented software for consistency with the derived error characteristics. For the validation during the commissioning and operational phase of the GRAS instrument, the error characteristics of excess phase delays and bending angles should be assessed by statistical methods, and compared to the expected error characteristics. If consistency cannot be achieved, the theoretical error models may require revision.

Section 3 of the memorandum also discusses other validation methods that potentially can be used for the validation of both the level 1b products and the geophysical products retrieved from GRAS measurements. The techniques assessed here include direct intercomparison of independent RO measurements, noise level estimation, and intercomparisons of geophysical products from different sources. The key points in the assessment of these methods are the additional information provided by these methods on top of the 1D-Var based validation and the correct handling of the measurement error characteristics in the validation. It is also noted, that if it is necessary to compare temperature and humidity information derived from the RO measurements with NWP analyses and/or other measurements, a 1D-Var retrieval is the best way to derive the geophysical information (Marquardt et al., 2003a).

Finally, Appendix A provides an Overview, User Guide and Reference Manual for the bending angle 1D–Var software developed within the framework of the study. The purpose of the Appendix is to provide additional information about the 1D-Var retrieval described in Section 2. The software tool described in Appendix A was used during the study to validate one month of bending angle measurements from the CHAMP mission. The results from the validation are presented in Section 3.8. The bending angle 1D-Var software described in Appendix A is available from the authors on request.

2 The Use of 1D-Var for Validation

The one dimensional variational (1D-Var) method is widely used for estimating atmospheric information from satellite measurements. It has a clear theoretical basis, formally producing the most probable atmospheric state, given prior information on the state and the measurements. It also provides solution error estimates. The 1D–Var approach was also used to pre-process ('retrieve') satellite measurements prior to assimilation (e.g., Eyre1993a), but now with the advent of 3D/4D-Var assimilation systems it is more routinely used for monitoring and quality control (QC) purposes. Fundamentally, the 1D-Var approach relies on accurate error estimates for the a priori atmospheric state, forward model and the measurements in order to find the most probable solution. However, it also provides a number of diagnostics which indicate whether the error estimates assumed in the calculation are consistent with the actual errors. Given that the purpose of validation is to "confirm that the theoretical characterisation and error analysis actually represent the properties of the real data" (Rodgers, 2000) these diagnostics may be useful and 1D-Var could potentially be an important component in a validation strategy, similar to that outlined in Section 3.

The aim of this work is to demonstrate how 1D-Var retrievals can be used as part of the validation of GRAS level 1b measurements, whilst at the same time illustrating some practical difficulties that can arise through limited understanding of other error sources. We will outline an implementation of the method that uses profiles of ionospheric corrected bending angles, α, as a function of impact parameter a, but in principle the technique could also be applied to uncorrected (L1 and L2) bending angles, Doppler shifts and excess phase delays. The main difference with these approaches is that the "forward models" used to simulate the observed quantities from the NWP forecast information will be more complex, and require an ionospheric model. In section 2.1, the 1D-Var retrieval method will be outlined and the diagnostics useful for validation will be described. In section 2.2, we will outline some practical difficulties that arise when using 1D-Var for these purposes. In section 2.3 we will present an implementation for corrected bending angles, and a summary will be given in Section 2.4.

2.1 1D-Var retrievals

A major difficulty encountered in the validation of GRAS level 1b data is that other instruments, with the exception of other co-located RO observations, do not measure the same quantities. As a result it is necessary to forward model geophysical parameters (temperatures and humidities) to observation space and/or map the level 1b information to geophysical parameters, using a suitable inverse (retrieval) method. The 1D-Var retrieval technique essentially combines both of these approaches, because it is an inverse method that relies on forward modelling to the observation space. Consequently, many of the outputs routinely produced may be used for validation.

The 1D-Var retrieval method is based on a Bayesian statistical approach. The aim is to find the most probable atmospheric state vector, $\hat{\mathbf{x}}$, given the a priori (background) estimate, $\mathbf{x_b}$, with error covariance $\mathbf{B}$, and the observation vector, $\mathbf{y}^o$, with error covariance matrix $\mathbf{E}$. The method is simply a generalised, weighted least-squares approach, which attempts to find a solution that simultaneously fits both the a priori/background and measurement vectors to within their respective errors. If we assume the background and observation errors are unbiased and Gaussian, it can be shown that the most probable state, or best fit, is found by minimising the quadratic cost function, $J(\mathbf{x})$,

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x}-\mathbf{x_b})^T\mathbf{B}^{-1}(\mathbf{x}-\mathbf{x_b}) + \frac{1}{2}(\mathbf{y}^o - H(\mathbf{x}))^T(\mathbf{E}+\mathbf{F})^{-1}(\mathbf{y}^o - H(\mathbf{x})). \tag{1}$$

Here, H is the forward model and $H(\mathbf{x})$ is the simulated observation that would result for a given atmospheric state $\mathbf{x}$; the forward model error covariance matrix is $\mathbf{F}$. The forward model errors are usually estimated from simulations, and will be discussed in more detail in section 3. The method also provides an estimate of the solution error covariance matrix for $\hat{\mathbf{x}}$. In the linear limit this is given by,

$$\mathbf{P} \simeq (\mathbf{B}^{-1} + \mathbf{H}^T(\mathbf{E}+\mathbf{F})^{-1}\mathbf{H})^{-1} \tag{2}$$

where $\mathbf{H}$ is the gradient of the forward model, evaluated for the solution vector. The solution error covariance can be useful when comparing retrieval results with other measurements because it provides an estimate of the measurement errors mapped into geophysical parameter space.

Note that unlike 3D/4D-Var problems the relatively small vector dimensions of the 1D-Var problem mean that a Newtonian iteration can usually be implemented to minimise the cost function. If the problem is linear (the elements of $\mathbf{H}$ are independent of state vector) the solution will be found with one iteration.

2.1.1 "O-B" statistics

The differences between the observed values and those simulated with the background estimate, $\mathbf{x_b}$, are a result of the background errors mapped into observation space and the combined observation and forward model errors. Strictly, it is not necessary to use a 1D-Var retrieval to perform this test but we need the gradient of the forward model (i.e., a component of the 1D-Var) in order to estimate the background errors mapped into observation space.

In the linear limit, the consistency between the entire observation and model derived vectors can be analysed with the scalar value,

$$\chi^2 = (\mathbf{y}^o - H(\mathbf{x_b}))^T(\mathbf{HBH}^T + \mathbf{E} + \mathbf{F})^{-1}(\mathbf{y}^o - H(\mathbf{x_b})) \tag{3}$$

Ideally, χ^2 should have the theoretical chi-squared distribution with, m, the number of elements in the observation vector equalling the degrees of freedom. The expected, ensemble-averaged value of $\overline{\chi}^2 \simeq m$ with a standard deviation of $\sqrt{2m}$. If $\overline{\chi}^2$ is much smaller than m this indicates that the assumed errors are too large. Conversely, if $\overline{\chi}^2$ is much larger than m, this indicates that the assumed errors are too small. More generally, the chi-squared distribution can be used to determine the probability that the vector of differences, $\mathbf{y}^o - H(\mathbf{x_b})$, belongs to the distribution defined by the covariance matrices. Rodgers (2000) notes that the chi-squared distribution can be used to determine what fraction of members of the distribution have a probability density less than the observed vector difference. However, in practice it is unlikely that the background error covariance matrix is known to sufficient accuracy to interpret these results so strictly, as discussed in section 3.

In addition to looking at the "O" and "B" vectors overall consistency, it is also useful to look at comparisons on an individual measurement value basis. This is because the degree of O-B agreement may vary with observation height. It also enables a gross error check prior to use in the 1D-Var retrieval. If the absolute difference exceeds a pre-specified limit, the value is removed from the measurement vector (see Section 3.8.2). Rodgers (2000) also notes that individual difference values should also have a chi-squared distribution, but once again the uncertainties in the background error estimates can cause practical difficulties.

Mean "O-B" vectors are used in bias correction schemes for passive sounder measurements. It is also interesting to note that ECMWF use the O-B method operationally for estimating and monitoring satellite measurement errors:

$$\mathbf{E} + \mathbf{F} = \overline{(\mathbf{y}^o - H(\mathbf{x_b}))(\mathbf{y}^o - H(\mathbf{x_b}))^T} - \mathbf{HBH}^T \tag{4}$$

This approach could be used to estimate the combined observation/forward error covariance matrix as part of a validation exercise.

In practical terms, the measurements should be interpolated to a fixed "impact height" vertical grid to derive the covariance of the differences. The impact height is defined as the impact parameter value minus the radius of curvature at the measurement location.

2.1.2 Cost at convergence

The cost function value at convergence, $J(\hat{\mathbf{x}})$, is also a useful scalar parameter which indicates whether the assumed errors are consistent with the actual errors in the a priori and observations. Similarly to the O-B test, the expected cost at convergence value is $2\overline{J(\hat{\mathbf{x}})} \simeq m$ with a standard deviation of $\sqrt{2m}$. Strictly, the cost at convergence test is only valid for linear retrievals, but simulation studies with non-linear retrievals (Healy and Eyre, 2000; Palmer et al., 2000) suggest it is applicable for RO problems. More generally, this approach is also used for investigating the consistency of error estimates used in 3D/4D-Var assimilation systems, although it is sometimes called the "Bennett-Talagrand ratio" in that context.

Note that if an observation vector has reasonable O-B statistics, but has a large cost at convergence, it indicates that the 1D-Var retrieval has converged to the wrong solution

(Rodgers, 2000). This consistency check can be useful for removing poor retrievals from validation statistics. Clearly, the number of profiles converging to the incorrect solution is of interest and should be monitored.

2.1.3 Validation of retrievals

This is discussed in more detail in Section 3. Briefly, it is often necessary to map level1b data to geophysical parameters for comparison against other measurement types and/or NWP information. Clearly, it is important to use the most accurate retrieval method available. We have found that the "traditional" RO temperature retrieval can be extremely sensitive to noise at altitudes greater than around 30 km. It also requires tropospheric a priori, but does not make any allowance for errors in this information. In contrast, the 1D-Var method provides the statistically optimal estimates of the geophysical variables in the state vector, given the measurements and a priori information. We can monitor the statistics of the 1D-Var increments, $\hat{\mathbf{x}} - \mathbf{x_b}$, because these should be unbiased and have a standard deviation given by the background error estimates. More generally, the 1D-Var retrievals can be compared with radiosondes and other measurements.

2.2 Practical experiences with 1D-Var

2.2.1 Uncertainty in other error estimates

It must be emphasised that we believe that 1D-Var retrievals with RO data are clearly superior to alternative inversion methods (which make no allowance for a priori and measurement errors!) even when the required error covariance matrices are not known precisely (Marquardt et al., 2003b). A particular advantage in using a 1D-Var retrieval method is that it provides diagnostics that are related to the quality of the retrieval. These indicate whether the error covariance matrices used in the retrieval are a good representation of the actual errors. If the diagnostics show that overall the errors estimates are incorrect, and we have good confidence in the background and forward model error estimates, then it seems reasonable to attribute the difficulties to poorly specified observation errors. However, it is rarely quite so straightforward and it is sometimes necessary to be pragmatic. For example, although broadly speaking the magnitude of NWP errors are understood, determining how the errors are vertically and horizontally correlated is a challenging problem and the generation of **B** matrices for retrieval and assimilation is an area of active research (e.g., Ingleby2001). Many operational centres (e.g., Met Office, NCEP) currently derive background error estimates with the "NMC method" (Parrish and Derber, 1992). Global estimates of errors in a six hour forecast are derived from differences in 24 and 48 hour NWP forecasts, valid at the same time. The theoretical reasoning behind such an approach is not obvious, but in practice it appears to give reasonable results. Nevertheless, the fact that background error estimates are being continually updated and refined should be noted.

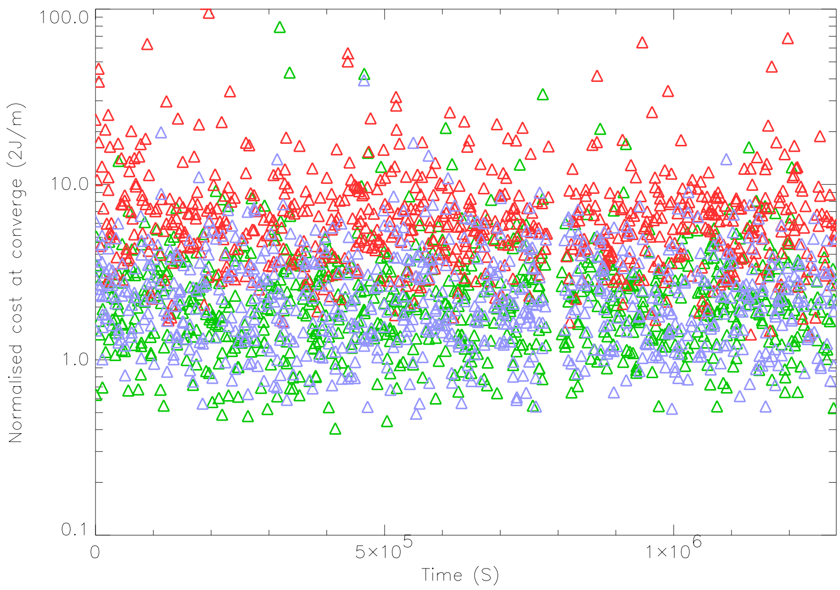

Fig. 1: The normalised cost-at-convergence as a function of time (seconds). Green NH, red Tropics and blue SH.

The forward model errors can also be difficult to characterise. They arise because:

1. the forward models used in variational retrievals contain approximations that simplify the physics and speed up the calculation. These approximations are usually investigated with simulation studies, by comparing the forward model performance with more sophisticated simulation tools. For example, Saunders et al. (1999) have tested the performance of the "RTTOV" fast radiative transfer code with a line-by-line model. Similarly, Healy and Eyre (2003) have compared an RO bending angle forward model against a 3D ray tracing code.

2. "Representation errors". These arise because the measurement may be sensitive to horizontal and vertical variations in the real atmosphere that cannot be represented by the NWP model. Representation errors are not measurement errors, because they arise as a result of limitations in the NWP, but they are effectively treated as measurement errors in the 1D-Var retrieval. Representation errors mean that some structures that can be measured correctly cannot be validated by comparing with NWP information.

2.2.2 1D-Var for QC: an example

A 1D-Var refractivity code is used at the Met Office to QC RO refractivity measurements made with the CHAMP satellite, prior to being passed to the 3D-Var system in a recent impact trial experiment. The observation error estimates used in the 1D-Var and 3D-Var were based on the values derived by Kursinski et al. (1997a), but included vertical correlations, modelled with a simple exponential form (Healy and Eyre, 2000).

Representation errors were neglected. The cost at convergence test was used to identify profile vectors that contained gross-errors and these profiles were subsequently assigned zero weight in the assimilation process. The residuals of the observations minus the values simulated with the 1D-Var solution were also used to set a "probability of gross error" (PGE) on an individual element by element basis within the measurement vector. If the absolute difference was greater than four times the expected observation error, the measured refractivity value was given zero weight in the 3D-Var assimilation.

Fig. 1 illustrates the normalised $2J(\hat{\mathbf{x}})/m$ cost at convergence as a function of time for 2213 retrievals, taken between May 21 and June 11, 2001. Ideally, the normalised value should be around unity, but this is not the case. This means the error estimates we are using are not optimal and should be re-assessed. However, we must be pragmatic when applying some aspects of the theory outlined above. Rigorously applying a PGE for a profile based on the chi-squared approach outlined by Rodgers (2000) to these data would result in the rejection of an unacceptably high number of profiles that contain useful information. In fact, we only rejected profiles outright if they were extreme outliers, defined as $2J(\hat{\mathbf{x}})/m \geq 20$. Note that we found that removing the vertical correlations from the observation errors reduced all the cost at convergence values by a factor of 3. The results show a latitudinal variation; the cost at convergence in the tropics (coloured red) have systematically higher values. Further analysis of these results indicated that refractivity representation errors are greater in the tropics than extra-tropics, (see Fig. 16 and Section 3.8.3 for more detailed discussion). In fact, representation errors in the tropics exceed assumed observation errors between 15-20 km, probably as a result of gravity waves which are not simulated by the NWP model.

This example illustrates a number of points. Firstly, the diagnostics indicate that we do not completely understand the errors, and therefore we must reconsider how they are derived. The results also suggest that the measurements are sensitive to scales that cannot be represented by the NWP model. Therefore, we cannot validate the measurement on these scales with NWP.

2.2.3 Gross errors in the observation vector and "automatic" QC

In some cases observations contain large errors that are extremely unlikely to occur if the Gaussian statistics associated with the $\mathbf{E}+\mathbf{F}$ matrix are correct. These are usually termed "gross errors". Gross errors can affect the performance of the 1D-Var code, often increasing the number of iterations required to minimise the cost function and degrading the quality of the geophysical retrieval. Ideally, gross errors should be "screened out" of the observation vector before the data is passed to the 1D-Var routine. A simple and effective screening approach is to reject data if the "O-B" difference exceeds a pre-determined factor multiplied by the combined observation and background errors mapped to observation space. For example, an observed bending angle value, y_o^i, might be rejected if the absolute difference from the simulated value fails the "10

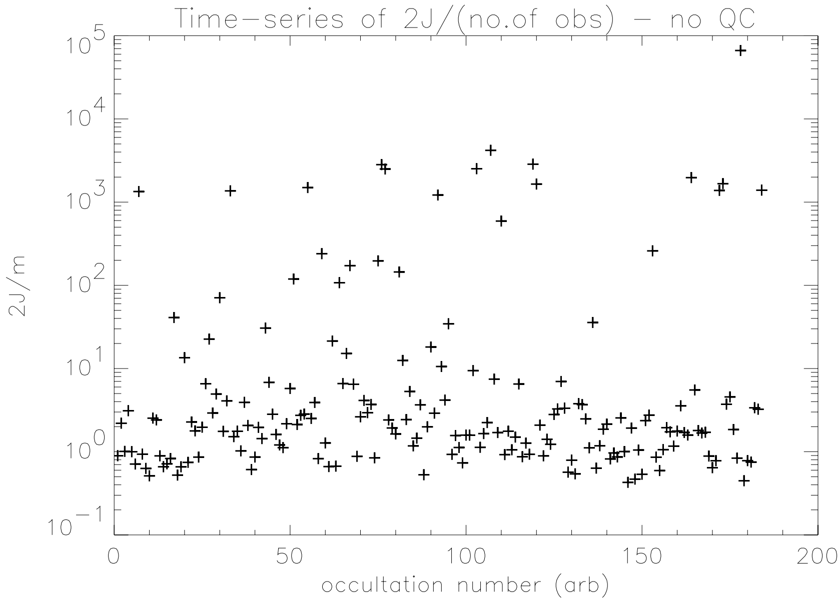

Fig. 2: The normalised cost-at-convergence for 1 day of CHAMP data (day 135, from 2001) determined with the bending angle 1D-Var code. The values greater than 1000 are a clear indication of gross-errors in the data.

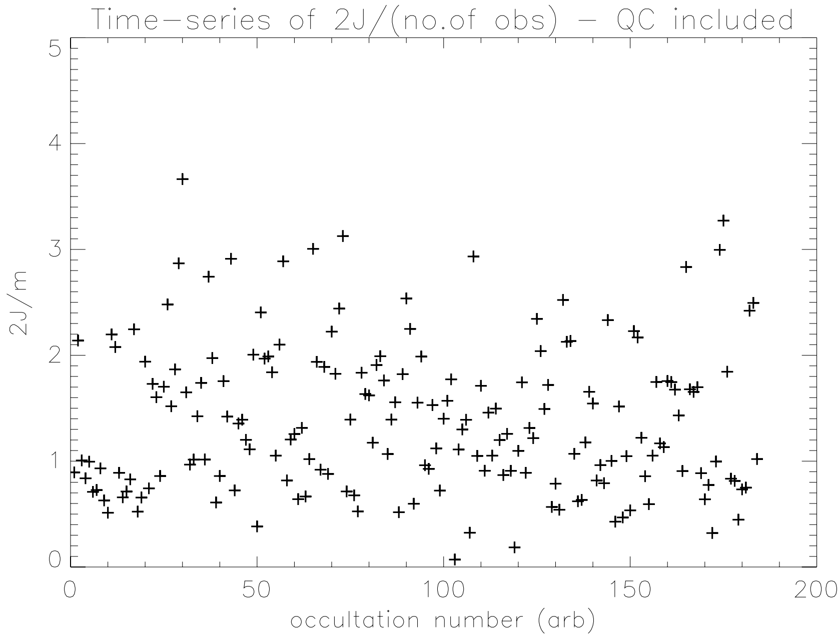

Fig. 3: As figure 4, but with QC on.

sigma" test whereby,

$$|y_o^i - y_b^i| \geq 10 \times ((\sigma_o^i)^2 + (\sigma_b^i)^2)^{1/2} \quad (5)$$

given y_b^i is the bending angle simulated from the background state and σ_o^i and σ_b^i are the observation (+ forward model) error and background error mapped to observation space, respectively. It should be emphasised that screening of the observations prior to passing the data to the 1D-Var is an important step and we advise that it should be implemented. However, note that in the context of a validation exercise it is essential to monitor the number of bending angles that fail the screening.

In addition to the screening, we can assign a probability of gross error (PGE) to each bending angle using a Bayesian approach. The method we have used is well established in NWP and is based on the work by Ingleby and Lorenc (1993) and Andersson and Jarvinen (1998). As with the screening step, the PGE is derived from the "O-B" differences. Broadly speaking, if the "O-B" is large when compared with the combined observation error and background error mapped to observation space, then the PGE is high. Ingleby and Lorenc (1993) and Andersson and Jarvinen (1998) assume that good observations (no gross error) have a Gaussian distribution and bad observations (containing gross errors) have a uniform distribution. The Bayesian estimate of the PGE value is found by evaluating the probability of finding the "O-B" difference *when there is a gross error*, divided by the *total* probability of finding the "O-B". Mathematically, the PGE for the ith observation can be written as,

$$PGE^i = 1 - \frac{1}{\gamma \exp(u) + 1} \tag{6}$$

given,

$$u = \frac{1}{2} \frac{(y_o^i - y_b^i)^2}{(\sigma_o^i)^2 + (\sigma_b^i)^2} \tag{7}$$

where the tunable parameter γ is given by eq. (11) in Andersson and Jarvinen (1998). It is related to the assumed characteristics of the *a priori* gross error probability density function. Note that in the limit $u \to 0$ then $PGE^i \to \gamma$. The 1D-Var code provides a PGE value for each bending angle in the profile. The profile of PGE values can be useful for identifying the individual bending angle values that produce a high cost-at-convergence. It may also be useful to monitor the number of bending angles with a PGE greater than a pre-determined threshold value. As with the screening step, the percentage of observed bending angles thought to contain a gross error will be a useful output of any validation exercise.

The PGE value can also be used to define a "QC weighting factor" given by,

$$W_{qc}^i = 1 - PGE^i = \frac{1}{\gamma \exp(u) + 1} \tag{8}$$

which is used to effectively[1] inflate the observation errors if the PGE estimate is high.

[1] In fact, the vector of observed minus simulated differences are multiplied by the corresponding elements of the $\mathbf{W_{qc}}$ vector.

This "automatic" QC within the 1D-Var routine is useful for ensuring reasonable retrieval results – particularly in the absence of a screening step – but it can lead to some difficulties in the context of a validation exercise. For example, Fig. 4 shows the normalised cost-at-convergence, $2J(\hat{\mathbf{x}})/m$, for 1 day of CHAMP data processed with the 1D-Var bending angle code. These calculations were performed without the QC weighting (or screening) of the observation values outlined above, and it is apparent that in a number of cases $2J(\hat{\mathbf{x}})/m > 1000$, clearly indicating that there is a problem with the retrieval. On examination, it was found that these observation vectors contained gross errors because of a processing error, which was straightforward to correct. In contrast, Fig. 5 shows the $2J(\hat{\mathbf{x}})/m$ values when the QC weighting has been implemented. The values are greater than unity, but they are reasonable because the QC has reduced the weight given to the observed values that are not consistent with the assumed errors. Although the automatic QC works well, it actually obscures the processing error. This may not be desirable during a validation exercise. Therefore, we would argue that initially the 1D-Var should be run without this form of QC if validation is the prime purpose of the 1D–Var (see Section 3.8.2).

2.3 1D-Var implementation for bending angle

The forward model and minimisation routine represent the two major components of a stand-alone 1D-Var package. Details of the implementation of the bending angle 1D-Var can be found in Appendix A.

2.3.1 The forward model

The forward model, H, used in the 1D-Var code will simulate level 1b ionospheric corrected bending angle values, α, as a function of impact parameter a. The model assumes spherical symmetry, because we are only using an NWP profile (rather than planar) information.

The forward model must evaluate the "bending angle integral" for each observed impact parameter,

$$\alpha(a) = -2a \int_a^{\infty} \frac{\frac{d\ln n}{dx}}{(x^2 - a^2)^{1/2}} dx \tag{9}$$

where n is the refractive index derived from the model, r is the radius value and $x = nr$ (x is the conventional notation used in the GPS literature for this product. It should not be confused with the 1D-Var state vector).

If we assume that the NWP model data is given on a set of fixed pressure levels the forward model must perform the following operations:

1. Calculate the geopotential height of the fixed pressure levels.

2. Calculate the geometric height and then radius values, r, by adding on the "radius of curvature" used in the evaluation of the bending angle and impact parameter values. (We should also consider the height difference between reference ellipsoid and geoid at the surface.)

3. Evaluate the refractivity values, N, on the pressure levels.

4. Evaluate $x = (1 + 10^{-6}N)r$ on the pressure levels.

5. Find the tangent point height using the impact parameter, where $x = a$.

6. Evaluate the bending angle integral for $x \geq a$ (i.e., at all heights above that corresponding to the impact parameter). We will assume that $d\ln n/dx$ varies exponentially with x between the model levels for evaluating the integral.

2.3.2 Minimisation routine

The Levenberg-Marquardt minimisation approach (Press et al., 1992) has proved successful for refractivity (Healy and Eyre, 2000) and bending angle (Palmer et al., 2000) 1D-Var retrievals. This method is easily developed in f90 and is implemented here. Matrix inversion and solution of matrix equations are be performed with Cholesky decomposition routines. For details, see the User Guide of the prototype 1D–Var package in Appendix A.

2.4 1D-Var Summary

We have illustrated the ways in which 1D-Var retrievals may be useful in the validation of level 1b information. These arise because the method provides a number of diagnostics that indicate whether the assumed observation errors are consistent with the actual errors. However, practical experience suggests that some probability of gross error estimates, based on chi-squared tests, need to be "tuned", otherwise too much data is rejected.

Above all, if it is considered necessary to compare temperature and humidity information derived from the RO measurements with NWP analyses and/or other measurements, we would argue that a 1D-Var approach is the best way to derive the geophysical information (Marquardt et al., 2003b).

Finally, we have also outlined a possible implementation of 1D-Var using ionospheric-corrected bending angles. A similar approach could be applied to uncorrected bending angles, Doppler shifts or excess phases, but the forward models will be more complex and an ionospheric model will be required.

3 A Strategy for the Validation of GRAS Level 1b Data

A measurement should always be accompanied with an estimate of the uncertainty in it. The uncertainty conveys the degree of confidence we have in the quoted value and enables potential users to assess whether it is good enough for their application. This is why Rodgers (2000) states:

> "The purpose of a validation in this general sense is to confirm that the theoretical characterisation and error analysis actually represent the properties of the real data".

Therefore, the fundamental aim of validation is to demonstrate that we understand the measurement errors, because we cannot claim to truly understand the measurement unless this is the case.

Validation is performed by comparison with independent quantities, which should have well characterised uncertainty. Ideally, the uncertainty in the quantities that are being validated against should have a clear "traceability". Traceability means that a measurement can be related to stated reference standards, through an unbroken chain of comparisons, each having a stated measurement uncertainty. Having established the traceablity of the independent measurement, a statistical analysis of the comparisons can be performed, establishing the traceability of the measurements being validated. The results of this validation exercise may lead in statements of the form:

> "The statistics of the differences are consistent with the assumed error levels, suggesting that the error model provides a reasonable representation of the actual errors."

or,

> "The statistics of the differences are inconsistent. The error modelling / propagation exercises must be revisited."

In practice, the validation of remotely sensed measurements is rarely as simple, for a number of reasons. Firstly, we do not validate the "raw measurements" made on board the satellite; they have usually been pre–processed to some degree. We define the term "observing system" to mean the instrument on board the satellite combined with any (pre–) processing steps. Thus, when validating pre–processed measurements (like excess phase delays; L1, L2 or ionospheric corrected bending angles, we are actually validating the full observing system, rather than just the instrument on board the satellite. Furthermore, the error propagation in the pre–processing steps can be complex and difficult to characterise in some cases. For example, errors associated with the assumption of local spherical symmetry in radio occultation (RO) measurements will depend on the atmospheric conditions. In such cases, numerical simulations might be the only way to estimate the errors.

A second problem often arises because new satellites produce information that is complementary to the existing observation network. This means that finding something to compare against directly is either impossible (e.g., finding an independent source for excess path delays in case of RO), or at least not a trivial problem (like finding an independent source of bending angles). Statistical methods may then have to be used to validate theoretical error characterisations. Another approach is to derive geophysical parameters, such as temperature and humidity. Validation statistics for retrieved geophysical parameters might be used to infer indirectly about possible problems in the level 1b data. However, in doing so we extend the scope and complexity of the observing system. We are then validating the instruments used to collect the raw measurements, combined with the pre–processing and the retrieval used to map to geophysical space. In order to be able to identify the most likely source of a possible problem, it is vital that the error propagation characteristics of pre–processing and retrieval are sufficiently well understood. To a large degree, this can be accomplished by numerical simulation studies, which are therefore a crucial part in any validation strategy of remotely sensed data.

The validation of retrieved geophysical quantities is usually achieved through direct comparison with in situ observations obtained from, e.g., radiosondes or aircraft measurements; with observations obtained from ground based (like lidar) or satellite borne remote sensing instruments; or even with profiles interpolated from meteorological fields obtained from Numerical Weather Prediction (NWP) fields. This seems to be a straightforward task, only requiring temporal and spatial coincidences.

However, producing accurate error estimates and establishing a clear traceability for each of the individual measurements is often difficult. Error characteristics of NWP data, for example, are complicated and depend on location and season. The intercomparison of satellite “measurements” is also not straightforward: Remote sensing instruments do not directly measure geophysical parameters, but require a retrieval. Since retrieval (or “inverse”) problems are often under–determined and ill–conditioned, their solution requires the use a priori data or assumptions. This contributes to the retrieval error. In regions where the information content of satellite measurements is low, retrieved geophysical parameters will be dominated by a priori rather than represent a true measurement. A priori used in different observing systems might be based on the same sources (like the same climatology, radiosonde statistics or NWP data). Thus, retrieved atmospheric parameters from different observing systems may not be truly independent.

The difficulty in establishing accurate and definitive error estimates for the independent data used in validation exercises leads to the pragmatic approach often adopted in satellite meteorology: we look at the consistency (or otherwise) of a number of inter–comparisons, making use of experience and expertise with the various independent data. Individually each intercomparison may have significant limitations, but these can be partially overcome by analysing the results from a number of comparisons. Many or all validation studies, therefore, try to use synergistic observations including as many relevant parameters as possible to compensate for the missing or unclear traceability of some of the involved measurements.

This approach works well as long as the retrieved atmospheric parameters agree well with the various auxiliary data sets, i.e. within their respective expected error estimates. Larger deviations, however, pose a problem: They may be related to problems in the raw measurements; or the calibration, pre–processing and retrieval steps might contribute more than expected to the errors of the full observing system; or some of the auxiliary data sets might be flawed. The only way to exclude the possibility that the cause for unexpected large deviations lies within the observing system itself is to go through each individual processing step and check that the error propagation is indeed as expected - in other words, to establish the traceability of the observing system.

GRAS measurements pose another problem: With not more than 250 to 500 daily observations which are globally distributed, the number of coincidences with observations taken at a single site will necessarily be small. This practically excludes single–site based measurements (or even a dedicated validation campaign) from the available options for the validation of GRAS level1b data. The number of co–located measurements that can be obtained in a short period of time (like a month) will simply be too small to provide meaningful validation statistics. Especially during the commisioning phase, the validation of GRAS level1b data has therefore to be based on other means.

In this paper, we propose a validation strategy based on the above considerations for level 1b data obtained from the GRAS instrument. In order to ensure traceability, our proposal contains a stepwise validation procedure, which is outlined in section 3.1. Details of specific tasks are given in the following sections: Benefits of conducting pre–launch numerical simulation studies are discussed in section 3.2, while statistical methods to estimate noise from actual calibrated (e.g., excess phase delays) and level 1b data are introduced in section 3.3. The direct comparison of co–located bending angle profiles (both ionospheric corrected and uncorrected) obtained from other radio occultation measurements is discussed in section 3.4. Forward modelling of bending angles from NWP fields or other measurements of meteorological parameters, and the use of of 1D–Var diagnostics are summarised in section 3.5; we note that a detailed description of the theory and its application in the validation context has already been provided in Section 2. The comparison of retrieved meteorological profiles against NWP fields and co–located measurements is discussed in section 3.6. This section also deals with the formally correct treatment of intercomparing remote sensing retrievals. Advantages and disadvantages of a possible field campaign aiming at the validation of GRAS level 1b data are discussed in section 3.7. Practical experiences and a demonstration of the usefulness of the proposed tools using CHAMP data form section 3.8. Conclusions in section 4 close the paper.

We finally note that in contrast to purely “physical” retrieval schemes such as the traditional dry temperature retrieval for radio occultation soundings, variational (or “optimal estimation") retrievals offer a plethora of diagnostics which are not available in other retrieval methods. These diagnostics can be used in a validation context and are, in fact, one of the core components of our proposed strategy. In addition, tools required in the variational retrieval context are also useful for validation. While their use does not require a variational retrieval to be actually performed, they are readily available if a variational retrieval is implemented, and have no counterpart in non–optimal retrievals. For the rest of this paper, we will therefore assume that the retrieval procedure is an implementation of a 1D–Var retrieval as outlined in Section 2. However, any other retrieval method which provides a similar set of retrieval diagnostics is equally well suited for the purpose of validation.

3.1 Validation strategy

Fig. 4 shows what we regard as the "GRAS observing system": The raw GPS measurements taken by the GRAS receiver will first be calibrated in order to correct for clock errors by single or double differencing, which includes the use of data from a global network of fiducial GPS ground stations. The calibration of the GPS data may include fixing of cycle slips or smoothing of data. Orbit information for METOP and the occulting GPS satellite will be utilised in order to calculate excess phase delays. These, along with Signal–to–Noise Ratio (SNR), or amplitude, data form a calibrated data set. In a further preprocessing step, precise orbit data and a priori information (like the assumption of spherical symmetry) are utilised to calculate excess doppler shift, L1 and L2 as well as ionospheric corrected bending angles. These form the level 1b data set. Calculating the various bending angles may consist of the basic phase–only retrieval of bending angles, or more complicated algorithms like the Canonical Transform or Full Spectral Inversion. Finally, level 1b data are used in a retrieval procedure to calculate geophysical parameters like refractivity, temperature, moisture and surface pressure.

It is already apparent from the above description that level 1b data has to be regarded as some processed form of the initial raw measurements provided by the GRAS instrument. Obviously, the error characteristics of the calibrated and preprocessed data will depend on the details of the calibration and preprocessing applied. As validation aims at checking if the actual error characteristics of the data are consistent with the theoretical estimates, or at least with prelaunch requirements, we require a comprehensive error model for the level 1b data. It will include an error covariance matrix for the calibrated measurements and each step of the preprocessing chain, e.g., the excess phase and amplitude (or SNR), doppler shift and bending angle step.

Provision of such an error model requires the computation of the error propagation through the calibration and preprocessing chain. The results of the error propagation can be used to test individual components of the preprocessing chain by letting them process simulated data with added random noise; the latter should exhibit the expected error characteristics. By this, we validate the preprocessing and retrieval software. These tasks can already be undertaken during the prelaunch phase of GRAS.

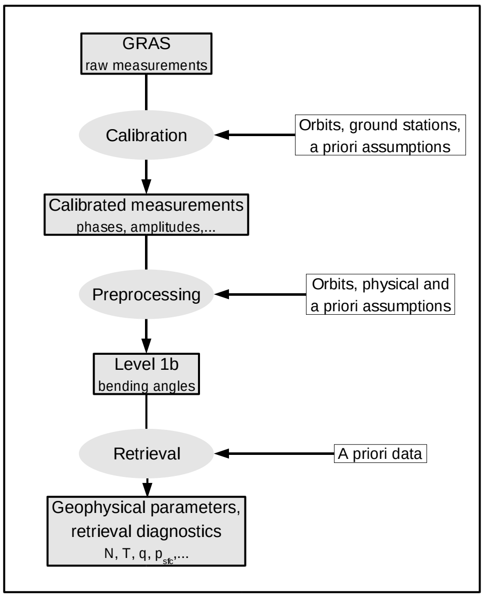

Fig. 4: The GRAS observing system.

Once a complete theoretical error characterisation of the GRAS observing system has been obtained and the implementation has been tested, validation of GRAS level 1b data shall proceed stepwise: For each step within the processing, it needs to be checked if the error estimates obtained from the respective error model are consistent with an independent estimate of the error characteristics of the actual data. If they are, the same procedure might be applied to the following step in the processing chain. If inconsistencies surface, these indicate a problem in the error model, the processing software for this specific step, the physical assumptions underlying both of them, or a problem of the validation data set. In either case, some revision is required.

Using this procedure, we effectively establish a traceability chain for GRAS data, beginning from the raw data to level 1b products. Applying the same strategy to the retrieval of geophysical parameters will establish a complete traceability chain for the GRAS observing system.

For the different processing steps within the GRAS observing system, different methods will obviously be required. To validate calibrated data (excess path delays and amplitudes), we will mainly have to rely on statistical methods estimating the noise component of the actual data, and compare it with the available error models. Well established statistical methods can be modified to work with correlated errors, but require some parameterised form of the underlying error characteristics and their correlations. For level1b data (excess doppler, uncorrected and corrected bending angles), the application of such statistical methods might also be useful. In addition, bending angles (both uncorrected and corrected) might be directly compared with co–located bending angles derived from other radio occultation experiments, if available. These comparisons will provide an independent check for the results of the statistical methods of noise estimation mentioned before.

1D–Var diagnostics and forward modelled Numerical Weather Prediction (NWP) data will then be used to validate (ionospheric corrected) bending angles. As a by–product, retrieved geophysical parameters become available. Because the 1D–Var framework allows for the calculation of the retrieval error based on the errors of observations and a priori, the complete traceability chain for the GRAS operating system will have been established.

Finally, geophysical parameters might be validated against data interpolated from NWP fields as well as against against co–located meteorological measurements, both from in-situ as well as from remote sensing instruments. The latter will especially be useful to validate aspects of the level 1b data which are not accessible by the forward modelling approach, e.g. due to errors of representativeness.

3.2 Numerical simulation studies

The aim of numerical simulation studies, which should be carried out before the launch of the METOP spacecraft, is twofold: First, theoretical estimates of the error characteristics of the various processing steps of the GRAS observing system will be provided. These will include error covariance matrices, which will later be validated by comparison against independent information. Further numerical simulations may be carried out to assess the contribution of physical assumptions made in the preprocessing and retrieval software (like spherical symmetry, the role of the ionosphere under strongly disturbed conditions, etc.) to the error budget. Second, the implemented software performing these processing steps will be tested (by Monte–Carlo type simulations) for being consistent with the theoretical error estimates. This will ensure that the software used for the calibration, preprocessing and retrieval of GRAS data has been implemented correctly.

3.2.1 Error propagation

The statistical error characteristics of the output of an individual processing step is composed of two parts: noise already present in the input data, propagated through the processing algorithm (which can be estimated by linearising the algorithm); and errors contained in external data or related to physical assumptions.

As an example, consider the preprocessing required to derive ionospheric corrected bending angles from excess phase delays. In a phase based retrieval of bending angles, excess L1 and L2 phase delays are first smoothed and then numerically differentiated, to provide L1 and L2 Doppler shifts. L1 and L2 bending angles are derived from the Doppler shifts, using orbit data for both satellites involved in the measurement. Finally, an ionospheric correction is applied. The error in the Doppler shifts $\Delta\mathbf{d}$, to first order, can formally be written as

$$\Delta \mathbf{d} = \mathbf{K}_d \Delta \mathbf{p} + \varepsilon_d \tag{10}$$

where $\mathbf{K}_d$ a linearised matrix representation of the code that derives Doppler values from phase delays $\mathbf{p}$. Note that the matrix $\mathbf{K}_d$ includes smoothing and numerical differentiation; both will introduce error correlations. $\Delta\mathbf{p}$ is a vector containing the errors on the phase delays and ε_d denotes any additional errors introduced by the algorithm, e.g. by the finite difference numerical differentiation. Thus, errors in the Doppler shifts are a result of phase errors mapped to Doppler space, plus a new error term. If the errors in the excess path delays are characterised by an error covariance matrix $\mathbf{O}_p$, and the newly introduced error by a covariance matrix $\mathbf{O}_{\varepsilon_d}$, the covariance matrix $\mathbf{O}_d$ of the doppler shift errors is given by

$$\mathbf{O}_d = \mathbf{K}_d \mathbf{O}_p \mathbf{K}_d^T + \mathbf{O}_{\varepsilon_d} \ , \tag{11}$$

provided that the various errors ($\Delta\mathbf{p}$ and ε_d) are statistically independent from each other. Additional complications will arise if errors are correlated.

Similar expressions can be written down for the L1 and L2 bending angles, i.e.

$$\Delta\alpha = \mathbf{K}_\alpha \Delta \mathbf{d} + \varepsilon_\alpha \ , \tag{12}$$

and for corrected bending angles:

$$\Delta\alpha_c = \mathbf{K}_{\alpha_c} \Delta\alpha + \varepsilon_{\alpha_c} \ . \tag{13}$$

ε_α will be caused by, e.g., assuming spherical symmetry. ε_{α_c} will arise because of limitations in the ionospheric correction. While some of these error terms might be negligible in practice (e.g., ε_d), others may only be investigated by numerical simulations. Fig. 5 shows estimates for ionospheric corrected bending angle error and its covariance structure for a randomly selected CHAMP profile. The calculation, similar Syndergaard (1999), assumes uncorrelated Gaussian noise ε_p with a standard deviation of 1.5 mm for L1 and L2 excess phase delays. Errors due to the numerical differentiation (ε_d) and of the spherical symmetry assumption (ε_α) have been neglected; the standard deviation of ε_{α_c} has been set to to a value of 1.5×10^{-7}rad. This represents a conservative estimate of the error in the ionospheric correction for altitudes of up to 80 km for standard ionospheric conditions (Vorob'ev and Krasil'nikova, 1994). Note the sync–like vertical structure in the error covariances; their vertical extent is determined by the width of a filter applied to the excess phase delays during the processing. Error correlation lengths decrease at lower altitudes, because the downward propagation of the tangential point of the radio occultation is slowed down in the troposphere. The increase of the bending angle errors in the troposphere is related to the defocusing which is incorporated in the error estimate. We note that similar qualitative results were obtained by Syndergaard (1999); the results shown here differ quantitatively, however, because more realistic assumptions on the error characteristics of the excess path delays

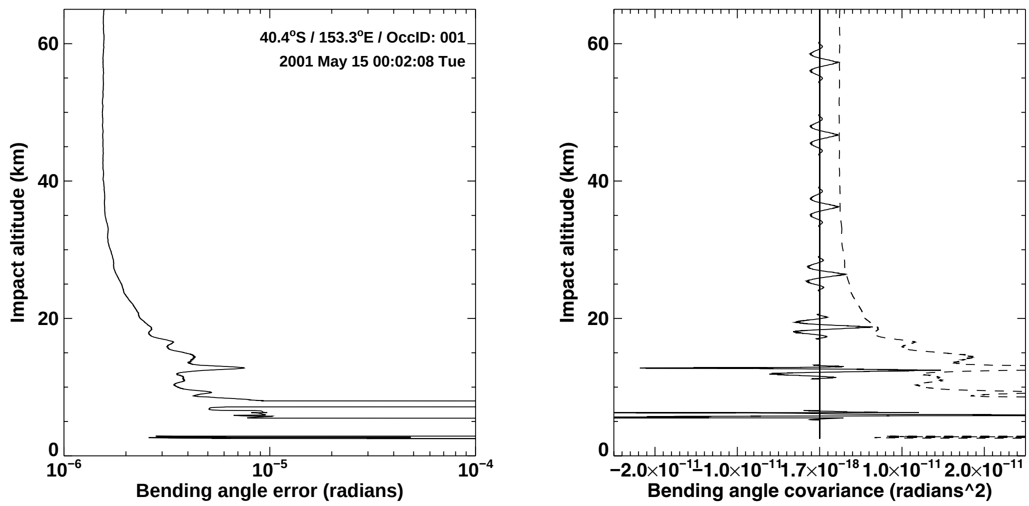

Fig. 5: Estimated standard deviation (smoothed, left) and vertical error correlations (right) of ionospheric corrected bending angles for a CHAMP occultation. The dashed line in the right figure denotes the squared standard deviation for comparison. Calculations are based on a phase only retrieval of bending angles, assuming a Gaussian uncorrelated excess phase delay noise of 1.5 mm for L1 and L2 phase delay data.

were made. Our estimate of 1.6×10^{-6} rad for the random error of upper stratospheric bending angles is also consistent with other error estimates in the literature.

Contributions of horizontal gradients (which are thought to cause the main contribution to bending angle errors in the lower stratosphere and upper troposphere; see, e.g., Kursinski et al., 1997b) on the error budgets are missing in the above study. They may be investigated by exploiting 3D ray-tracing codes, which are able to simulate excess L1 and L2 observable phase delays and signal amplitudes through relatively smooth atmospheres and ionospheres with sub–mm accuracy (Healy and Eyre, 2003). More importantly, they also provide the "true" Doppler shifts and bending angles that are consistent with phase delays. For example, the latter are simply given by the angle between the ray vectors at the start and end of the ray-path, without having to assume spherical symmetry. Alternatively, analytic expressions for estimated errors due to horizontal gradients given by Healy (2001) might be applied.

Similar procedures have to be followed when investigating the error propagation through advanced preprocessing algorithms like the Canonical Transform (Gorbunov, 2002b,c,a) or Full Spectral Inversion (Jensen et al., 2003). This will not only require error models for excess phase, but also for the amplitude (or SNR) data provided by the GRAS receiver. In addition, the error propagation will probably be more complicated than in the above example. The use of automated source code differentiation tools like TAPENADE[2] or TAF[3] (or it's freely usable predecessor TAMC) might ease the tedious process of coding the linearised versions of these algorithms. It must be kept in mind,

[2] See `http://www-sop.inria.fr/tropics/tapenade.html`
[3] See `http://www.fastopt.de`

though, that such tools are not free of errors, and usually require manual tuning and testing of the generated code. Restrictions on the use of code derived using these tools may apply for commercial applications.

For the retrieval of geophysical parameters, error characteristics of the solution obtained from a 1D–Var retrieval are well known from the literature (e.g., Rodgers, 1976, 1990, 2000), given that error covariances are available for both the preprocessed observations and the a priori data used. They are discussed in Section 2. If a non–variational retrieval is implemented instead, error estimates have to be obtained by procedures similar to the ones outlined above.

3.2.2 Monte–Carlo simulations

Once theoretical estimates of the expected error characteristics are available, the software implemented for the various calibration and preprocessing steps can be tested using simulated raw, calibrated and preprocessed measurements. Testing will usually be carried out by adding random noise to the respective simulated data set; the added noise will be consistent with the expected error characteristics of the respective measurements. In a Monte–Carlo type setup, the response of each software component can then be evaluated, and the statistics obtained from a sufficiently large ensemble of simulations may be compared with the theoretical error estimates for that processing step. Further simulations may be carried out to assess the sensitivity of statistical diagnostics (see section 3.3) to assumptions on error characteristics, and on how relevant deviations of the raw data's error characteristics are for the practical application of these tools.

Simulated data sets should be generated prior to launch, e.g. (for simulated level 1b data) by ray–tracing through 3 dimensional high resolution NWP fields. Apart from idealised atmospheric conditions which are consistent with physical assumptions made in the preprocessing and retrieval software (like spherical homogeneity), these simulations should also cover a wide range of atmospheric conditions, e.g. conditions of tropospheric multi-path or super-refraction as well as large horizontal gradients. Stratospheric measurements should be simulated not only for average, but also for strongly disturbed conditions like major stratospheric warmings. This will ensure that a possible degrading influence of a priori obtained from some climatology on the performance of the observing system under unusual conditions can be detected early.

3.3 Statistical methods for noise estimation

Because no independent measurements of excess phase delays (and possibly also for excess doppler and bending angles) will be available, the only way to compare theoretical error estimates with the error characteristics of actual data is to use statistical methods. One method of separating noise from signal is to fit a smooth function to the noisy data, and to calculate a noise estimate from the residuals.

A well established method for the objective estimation of parameters for smoothing splines is known as Generalised Cross Validation (GCV; Craven and Wahba, 1979). Using GCV estimates of uncorrelated random noise on simulated excess phase data provides surprisingly accurate estimates of the noise (see below). This suggests that GCV (or a generalised form of objectively fitting smoothing splines to noisy data, taking correlated errors into account if necessary), may be a promising technique for validating theoretical error estimates of calibrated and level 1b GRAS data. Alternative algorithms, for example those based on Generalised Least Squares fitting procedures, might also be considered.

In the following, we discuss the basic idea behind smoothing splines and GCV, and review a generalisation of the method for correlated errors.

3.3.1 Smoothing splines for noise estimates with uncorrelated errors

When trying to fit a smooth curve through a given set of discrete noisy "measurements" $y_i = y(t_i)$ for times t_i, a noise model of the following form

$$y_i = g(t_i) + \varepsilon(t_i) \ , \quad i = 1 \ldots n \tag{14}$$

is usually assumed, where g_i represents the discrete counterpart of a "smooth" function. $\varepsilon(t)$ denotes a white (i.e., uncorrelated) noise process with standard deviation σ. A popular estimate for the underlying smooth function g is the function $g_{n,\lambda}$ which minimises the cost function

$$J = \frac{1}{n} \sum_{i=1}^{n} (f(t_i) - y_i)^2 + \lambda \int_{t_1}^{t_n} (f^{(m)}(t))^2 \, dt \ . \tag{15}$$

The minimisation is carried out over the space of sufficiently differentiable functions f. It is well known (e.g. Reinsch, 1967, 1971; Wahba, 1975) that $g_{n,\lambda}$ is a polynomial smoothing spline of degree $2m-1$. Note that the solution $g_{n,\lambda}$ is a linear function of the observations, i.e. with $\mathbf{y} = (y_1, \ldots, y_n)^T$ and $\mathbf{g}_{n,\lambda} = (g_{n,\lambda}(t_1), \ldots, g_{n,\lambda}(t_n))^T$, the influence matrix $\mathbf{A}(\lambda)$ with

$$\mathbf{g}_{n,\lambda} = \mathbf{A}(\lambda)\mathbf{y} \tag{16}$$

can be calculated.

The parameter λ controls the tradeoff between the roughness of the smoothing spline (as measured by the second term in (15)) and the goodness–of–fit to the observations, as measured by the first term in (15), and needs to be chosen somehow. Reinsch (1967) suggests, roughly, that if σ is known, λ should be chosen so that the goodness–of–fit satisfies

$$\frac{1}{n} \sum_{i=1}^{n} (f(t_i) - y_i)^2 = \sigma^2 \ .$$

Wahba (1975), however, indicated that Reinsch's suggestion leads to systematic over-smoothing, and showed that λ should actually be chosen so that the goodness–of–fit is slightly smaller than σ^2. The value of this "true mean square error" can be estimated if σ is known.

For the case of unknown σ, Craven and Wahba (1979) applied the idea of "cross validation" to the fitting problem (hence the name of the method): If $g_{n,\lambda}^{[k]}$ denotes the smoothing spline for all but the k^{th} observation, the ability of $g_{n,\lambda}^{[k]}$ to predict the missing value y_k is taken as a measure for the goodness of λ. Craven and Wahba then showed that for $n \to \infty$, a λ minimising the generalised cross validation function

$$V(\lambda) = \frac{\frac{1}{n}||(\mathbf{I} - \mathbf{A}(\lambda))\mathbf{y}||^2}{\left[\frac{1}{n}Tr(\mathbf{I} - \mathbf{A}(\lambda))\right]^2}\,, \tag{17}$$

will also minimise the estimated true mean square error. Here, $Tr(\cdot)$ denotes the trace of a matrix, and $||\cdot||$ the usual L^2–Norm over the n–dimensional Euclidean space.

Thus, by minimising (17) and calculating a smoothing spline with λ as smoothing parameter, an objective method is available to calculate a smoothing spline fitting the data in an optimal (cross validation) sense without prior knowledge of the noise in the measurement set. We note that GCV in the form presented here assumes both uncorrelated and stationary noise.

3.3.2 Application to simulated excess phase delay noise

To illustrate the possible benefits of using objectively determined smoothing splines within the GRAS level 1b validation, we have conducted a Monte-Carlo study. 1000 realisations of uncorrelated Gaussian noise with a standard deviation of $\sigma = 1$ mm have been added to a simulated time series of L1 excess phases, and were analysed using the GCV algorithm. Fluctuations in excess phases are dominated by measurement noise during the first few seconds of an individual occultation; in the lower stratosphere and upper troposphere, the level of fluctuations might be increased due to the atmosphere's signal in the measurements. GCV, assuming a stationary noise, can therefore be expected to overestimate the noise when being applied to tropospheric data. Fig. 6 shows the mean noise estimate obtained by GCV, as function of the length of the data window subject to the GCV analysis. Note that after about 40 s, when the occultation ray enters the lower atmosphere, noise estimates do indeed increase. Also note the excellent agreement between the mean noise estimate provided by the GCV and the true noise. We have found similar good agreements over a wide range of σ's.

For comparison, Fig. 6 also shows noise estimates obtained from calculating the standard deviation of residuals obtained from a sliding polynomial fit to the simulated data, for different lengths of the filter window. This is a standard way of smoothing raw excess phase data in a preprocessing chain for radio occultation data, and filter widths

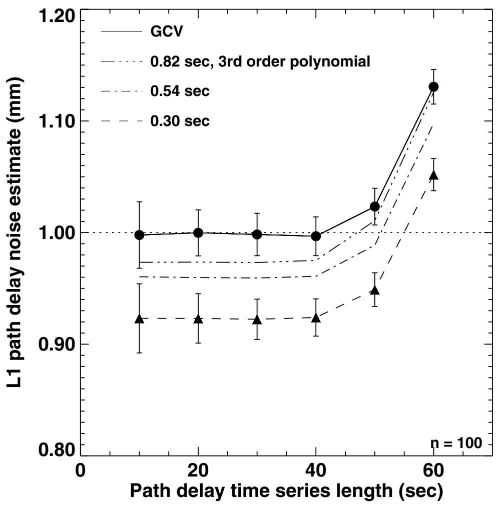

Fig. 6: Mean phase delay noise estimated by GCV as function of data window length subject to the GCV analysis (measured from the beginning of the occultation), for an ensemble of 100 realisations of uncorrelated Gaussian noise with $\sigma = 1$mm. The additional lines show similar noise estimates, which are based on various 3^{d} order polynomial fits to the simulated phase delay data. Times in the legend refer to the window width of the polynomial filter.

correspond to parameters used in, e.g., GFZ's CHAMP retrieval (Wickert, pers. comm.; also see Tsuda and Hocke, 2002, for another set of filter widths published in the literature); in all cases, the true noise is underestimated.

Real excess phase delay data, however, already has undergone some preprocessing, and may therefore exhibit significant error correlations. The GCV method may fail when this is the case (see, e.g., Wang, 1998, and refererences therein). For monitoring purposes, residuals obtained by a standard GCV therefore should undergo an autocorrelation analysis. Numerical simulations like the one described above, but using more complex error models for the excess phase delay noise, might give some indication what level of autocorrelation is acceptable for the applicability of GCV, or if more elaborate statistical methods have to be applied. Other standard time series analysis methods, like estimators of power spectra and the like, might also prove useful in identifying certain error characteristics of GRAS data during the commissioning phase, and should be provided accordingly.

3.3.3 Noise estimates for correlated errors

If error correlations of excess phase delays should turn out to be a problem, a generalised form of the GCV approach might be applied. Wang (1998) writes the cost function for a smoothing spline in the case of correlated errors as

$$J = \frac{1}{n}(\mathbf{y}-\mathbf{f})^T\mathbf{W}(\mathbf{y}-\mathbf{f}) + \lambda \int_{t_1}^{t_n} (f^{(m)}(t))^2\, dt \ , \tag{18}$$

where $\mathbf{W}$ is related to the error covariance matrix of the observations. $\mathbf{y}$ and $\mathbf{f}$ denote arrays containing the values of y_i and $f(t_i)$, respectively. The corresponding GCV function is then given as

$$V(\lambda) = \frac{\frac{1}{n}||\mathbf{W}(\mathbf{I}-\mathbf{A}(\lambda))\mathbf{y}||^2}{\left[\frac{1}{n}Tr(\mathbf{W}(\mathbf{I}-\mathbf{A}(\lambda)))\right]^2} \ . \tag{19}$$

Assuming that $\mathbf{W}$ can be represented by a fixed amount of parameters τ, Wang proposes to construct the smoothing spline based on a parameter set (λ, τ) which minimises (19). Alternatively, the original spline cost function might be utilised directly by minimising (18) with respect to the spline coefficients and the parameter set (λ, τ) simultaneously[4]. This is usually referred to as "Generalised Least Squares" approach.

The objective estimation of a smoothing spline for, e.g., excess path delay data, might provide a technique to validate the structure of theoretical error characterisations. How well this works in practice is currently unclear, but deserves further research. We also note the similarity of (15) and (18), respectively, to the regularisation approach introduced by Twomey and Tichonow, which is well known in the the retrieval literature.

If the error covariance matrix is thought to be known sufficiently well, a simpler approach might be to calculate a symmetric square root of the assumed error covariance matrix (e.g., see Appendix A in Rodgers, 2000), and then transform the observation vectors according to

$$\hat{\mathbf{y}} = \mathbf{O}^{-1/2}\mathbf{y} \ . \tag{20}$$

If the assumed errors are correct, the transformed observation vector will exhibit uncorrelated errors with standard deviation of unity. A GCV estimator for uncorrelated noise can be used to check this, if the correlation structure of the noise has indeed been removed successfully.

We would like to stress that all methods described require some knowledge on the error covariance of the observations. Thus, the numerical simulation of the expected error characterstics of the raw GRAS measurements and simulations assessing the usefulness and limits of the proposed statistical methods are important task to be performed during the prelaunch phase of GRAS.

[4]Wang (1998) actually provides some evidence that the latter might be advantageous in practice.

3.3.4 Noise estimates for bending angles

Statistical estimates of the noise associated in bending angles may be estimated by applying the same methods as described above. A bending angle profile obtained from forward modelling a representative atmospheric state or NWP information may be used as an alternative to the smoothing spline, and statistics may be calculated from the residuals between retrieved and forward modelled bending angle over the first few seconds / upper impact altitude range of an individual sounding. This approach is similar to the noise estimation of corrected bending angles as proposed by Sokolovskiy and Hunt (1996). We note that, based on the error propagation results presented in Fig. 5, bending angle errors will exhibit correlations as soon as some form of filtering is applied. Thus, methods allowing for error correlations are probably required.

3.4 Direct comparisons of bending angles

Atmospheric bending angles like those obtained from GRAS are only observed by radio occultation instruments. The only way to validate bending angle profiles with independent data directly, therefore, is to compare GRAS bending angle profiles with other radio occultation measurements. This could be done by comparing co–located bending angles obtained from either GRAS itself (if sufficiently many close–by occultations occur within a short period of time), or from a different GPS radio occultation receiver on another satellite (like CHAMP, SAC-C, COSMIC or EQUARS). Comparisons between GRAS bending angles and co–located airborne or mountain based radio occultations could also be made if such measurements are available or provided within a dedicated validation campaign. We will first discuss some general aspects of a direct comparison, and then the different options in turn.

3.4.1 General aspects

Some insight into what can be gained by direct comparisons of bending angle profiles from identical or reasonable similar instruments can be obtained as follows: We assume that a retrieved (vector of) bending angle(s) $\hat{\mathbf{x}}$ differs from its true value $\mathbf{x}$ by

$$\hat{\mathbf{x}} = \mathbf{x} + \mathbf{b} \pm \varepsilon \; . \tag{21}$$

Here, $\mathbf{b}$ describes a systematic bias of the observation system, while ε denotes an unbiased random (or noise) error (in the remote sensing literature, these are sometimes referred to as “accuracy” and “precision”, respectively). Alternatively, $\mathbf{b}$ may be viewed as describing errors on a time scale which is long compared to the period over which measurements are taken, while ε denotes a random error component varying from measurement to measurement. In Rodger’s (2000) notation, ε refers to the “retrieval error” associated to the actual instrumental noise, propagated by the processing, while $\mathbf{b}$ summarises the smoothing error and any bias possibly introduced by the retrieval.

If two "measurements" $\hat{\mathbf{x}}_1$ and $\hat{\mathbf{x}}_2$ of (nearly) the same air mass are taken by the same instrument, their difference

$$\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1 = (\mathbf{b}_2 - \mathbf{b}_1) \pm \varepsilon_1 \pm \varepsilon_2 \tag{22}$$

will follow the covariance statistics

$$\mathbf{O}_{\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1} = \mathbf{O}_{\mathbf{b}_2 - \mathbf{b}_1} + \mathbf{O}_{\varepsilon_1} + \mathbf{O}_{\varepsilon_2} \ . \tag{23}$$

In the most simple case in which the the two measurements exhibit identical error characteristics, i.e. $\mathbf{b}_1 = \mathbf{b}_2 = \mathbf{b}$ and $\mathbf{O}_{\varepsilon_1} = \mathbf{O}_{\varepsilon_2} = \mathbf{O}_{\varepsilon}$, this becomes particularly simple, i.e.

$$\mathbf{O}_{\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1} = 2\mathbf{O}_{\varepsilon} \ . \tag{24}$$

In the case of uncorrelated errors, this may further be reduced to

$$\sigma_\varepsilon = \sqrt{\frac{\langle \hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1 \rangle^2}{2}} \tag{25}$$

where σ_ε denotes a vector of standard deviations of the noise component of the measurements. $< \cdots >$ denotes the mean over a sufficiently large ensemble of such intercomparisons. Note that the **b**–term cancels out: the intercomparison of retrievals from two nearly identical observing systems does not provide information on possible systematic errors, but only on the random part of the measurement noise. This has implications for the interpretation of the obtained random error statistics:

- If two raw measurements have been obtained by the same instrument, and have been processed by the same processing software, the comparison will provide an estimate of the random noise error of the instrument, as propagated through the software components of the observing system. This provides a possibility to validate the previously discussed statistical methods for noise estimation with independent data.

- If the raw measurements stem from different (though still very similar) instruments, but have undergone the same processing, the estimated random error characteristics will reflect the combined random error characteristics of the different instruments, each one independently propagated through the preprocessing. In this case, $\mathbf{O}_{\varepsilon_1}$ and $\mathbf{O}_{\varepsilon_1}$ will be different, but $\mathbf{b}_2 - \mathbf{b}_1$ *might* still be neglible, only reflecting instrumental biases (this has to be checked, though). Otherwise, the full eq. (23) needs to be considered.

- If the raw measurements stem from different instruments, and have been processed by different processing systems, the estimated error characteristics will also reflect differences between the processing components of the two observing systems. As any use of a priori contributes to the errors (Eyre, 1987), differences in a priori use in the two processing systems might cause an additional increase in the errors. Again, the full eq. (23) needs to be exploited. Because some physical assumptions or a priori data used in the systems will be similar, common systematic errors might still be present and will remain undetected.

Thus, the direct comparison of bending angle profiles will provide useful information on the random error component of the measurement, but is not suited for estimating the "true" error of bending angle profiles.

An additional complication arises because the profiles to be compared will not be exactly co–located, but exhibit a certain spread in both time and location. The estimated random error will therefore contain a component related to the temporal and spatial variability of the atmosphere, which, strictly speaking, has to be subtracted. Statistics describing the atmospheric variability on appropriate spatial and temporal scales may be obtained from NWP fields. Since atmospheric variability strongly depends on season and location, care must be taken that the ensemble of profiles to be compared stems from a period and region with sufficiently homogeneous atmospheric variability. It is not useful, for example, to calculate a single statistic from both tropical and extra-tropical tropospheric bending angle comparisons, or from extra-tropical stratospheric comparisons obtained on both hemispheres during, e.g., NH winter. In both cases, the atmospheric variability differs significantly between parts of the data set, and can therefore not be subtracted properly.

The same methodology may also be applied to retrieved vertical profiles of refractivity, temperature and humidity, obtained from different radio occultation experiments. Similar limitations apply.

3.4.2 Space borne radio occultation measurements

Comparing retrieved profiles of atmospheric variables and constituents taken at nearly the same location by the same instrument, but separated in time by about 90 minutes, was extensively used in the validation of Upper Atmosphere Research Satellite (UARS) data (for temperature profiles see, e.g., Dudhia and Livesey, 1996; Fishbein et al., 1996; Gille et al., 1996). This was possible because the orbital configuration of the UARS satellite caused the limb viewing tracks of some of its instruments to intersect themselves at 80°N and 32°S (or 80°S and 32°N, depending of the phase of the spacecraft yaw cycle). Limitations of this approach were clearly recognised.

For the validation of GRAS level 1b data, bending angle profiles obtained from GPS receivers onboard radio occultation satellites can be intercompared. Possible candidates are profiles obtained from CHAMP and SAC-C (if they still provide data when METOP is launched), the COSMIC constellation (scheduled to be launched in 2005), or the EQUARS mission (scheduled for early 2006). If data from COSMIC is available, a sufficiently large amount of matches can be obtained within a short period of time.

This will also allow the stratification of matches according to season and region, in order to avoid the co–location / variability problem. The aim of these comparisons is to get an independent estimate of the instrumental noise, as propagated through the preprocessing software. Thus, raw data should be processed by the same processing system, in order to avoid complicating the interpretation of the obtained statistics.

3.4.3 Airborne radio occultation measurements

Zuffada et al. (1999) were the first to illustrate that a "space-based equivalent" bending angle profile between the surface and the receiver's altitude can be derived from a GPS receiver within the Earth's atmosphere. A straightforward application is to mount a GPS receiver on a research airplane and take tropospheric bending angle observations. Because the time and approximate location of individual radio occultations only depend on the orbit parameters of the spacecrafts involved, they can be predicted for a few days in advance. Accordingly, flights can be scheduled and directed to match GRAS occultations, and possibly even aligned with their viewing directions. Thus, a sufficient number of matching observations can probably be obtained within a comparatively short period (one month) over a sufficiently large region like Europe.

In order to obtain the best possible characterisation of the atmospheric states at the time of the measurements, matching GRAS and airborne occultations should be accompanied by additional correlative measurements, e.g. from radiosondes and lidars, if available. Because of the irregular distribution of occultations, however, the number of such correlative measurements will be limited. The most cost effective solution to this problem, at least when it comes to radiosondes, is to ask operational weather stations to launch additional radiosondes when a matching GRAS occultation occurs close to their site. Met Services involved in such a validation campaign might provide improved mesoscale data assimilation products, using the research data along with the operationally available data.

Matching bending angle measurements accompanied with a complete characterisation of the atmosphere along the line of sight(s) of the occultations would in particular be useful for checking the GRAS receiver's operation under (confirmed) atmospheric multipath and/or ducting conditions. It could also be checked if both the airborne and the space borne receiver's switching from phase locked to open loop tracking occurs according to the atmospheric conditions, and if the preprocessing / retrieval system is able to interpret the raw data accordingly.

On the other hand, it has not yet been demonstrated in practice that airborne radio occultations do indeed work; the errors of airborne radio occultations are currently not well understood. The major obstacle probably is the exact determination of the airplane's position, especially under turbulent atmospheric conditions along the flight path. The errors of airborne radio occultations could well be larger than the expected errors on the GRAS measurements. This would limit their usefulness for validation. Another disadvantage is that measurements are potentially very costly. However, it is also possible that the airborne technique itself could be more thoroughly investigated before the METOP launch in 2005. For example, Prof. Toshitaka Tsuda at the University of Kyoto is planning a first airborne occultation experiment for 2003 or 2004; Dr. Jennifer Haase at the University of Purdue is seeking funding for a similar experiment.

3.4.4 Mountain based radio occultation measurements

Mountain based radio occultation measurements, also inspired by the work of Zuffada et al. (1999), have successfully been carried out by scientists at the University of Kyoto and the Japanese Met Agency. It has also been demonstrated that such measurements can indeed be taken and processed. The calculated partial bending angles can be converted towards the "space-based equivalent" bending angle, as described in Zuffada et al. (1999).

Measurements co–located to GRAS occultations might give bending angle estimates with similar benefits as those obtained from airborne radio occultations. Because the location and probably also viewing direction of the mountain based GPS receiver are fixed, the launch of dedicated radiosondes in the viewing direction of the mountain based receiver and/or the space based occultation can be more easily accomplished than in the airborne case. The fixed position of the receiver on a mountain top certainly solves the problem of precise positioning inherent in the airborne radio, although no work has yet been done on the error characteristics of such measurements.

However, the fixed location of the GPS receiver also restricts the number of possible matches, or requires the comparison of profiles being further apart in both distance and time. The latter would render the interpretation of the statistical results more difficult because of the increased role of atmospheric variability. Based on our experiences with radiosonde matches for CHAMP, we would not expect more than one to four matches per month between radio occultations and a single location with a mountain based GPS receiver for a single radio occultation satellite like METOP. As airborne radio occultations, mountain based measurements are limited to altitudes below the receivers height, i.e. to the lower tropophere.

Based on the rare occurance of matches for an individual site, we believe that a single (or even a few) mountain based GPS receiver(s) may not contribute significantly to the direct validation of GRAS bending angles in the troposphere.

3.5 Forward modelling and 1D–Var diagnostics

If no or only few direct comparisons of bending angle profiles are available, "measured" bending angle profiles can be compared with NWP fields or conventional meteorological measurements by forward modelling these measurements towards bending angles. The appropriate tools, i.e. a forward model for the calculation of bending angles, and its linearised version for the forward propagation of the background errors into bending angle space will be readily available if a 1D–Var based on bending angles has been implemented. Apart from comparing the observed bending angle measurements with those calculated from the background directly, or by calculating the appropriate χ^2 statistics or probabilities of gross error (PGE), the variational retrieval framework offers additional diagnostic tools. Examples are the value of the cost function at convergence, or χ^2 statistics regarding deviations between retrieval and background or retrieval and observations. The details of these diagnostics are discussed in Section 2, and will not be repeated here. We would like to stress (as in Section 2), that the theoretical behaviour of these diagnostics can only be achieved in practice if correct error estimates for both the background and the observations are available. Even if this is not the case, however, as it might happen in the beginning of a new instrument's lifetime, the variational diagnostics may still be applied for validation and quality control purposes. Then, heuristically derived thresholds may be used instead of the strict theoretical ones. We demonstrate the usefullness of this approach in section 3.8.

Because of the limited options for directly validating GRAS level1b products with independent measurements of, e.g., bending angles, the validation using 1D–Var diagnostics, together with establishing reliable error models by the methods described in the previous sections, is the core component of our validation strategy. This implies that our proposed strategy heavily relies on NWP data used as background in the actual 1D–Var calculations. Therefore, our approach may raise concerns about the quality and reliability of NWP "model" data compared to "real" observations (like those obtained from other remote sensing measurements) which are used in more traditional validation approaches.

We note, however, that today's NWP products (like analyses or short range forecasts) are usually generated using variational data assimilation methods. The mathematical theory is identical to the one used in the usual 1D–Var (or "statistically optimal"; see, e.g., Rodgers, 1976, 1990, 2000) retrievals. Thus, variational data assimilation products can be understood as a global retrieval of the atmospheric state using *all* available conventional and remote sensing data *simultaneously*. Compared to, say, a vertical profile obtained from a remote sensing measurement, a NWP analysis at the same location and time using the same raw remote sensing measurement is based on significantly more data. NWP products also use (at least a simplified representation of) the physical laws governing the atmospheric flow. From that point of view, it is not immediately clear why a single retrieval from a given remote sensing instrument should give any *additional* information compared to the use of an NWP profile, assuming that the same data has been assimilated. Furthermore, the theory of variational data assimilation predicts that, at least in data rich regions, the actual error in the analysis and short range forecast is smaller than that of the individual observations. This has recently been confirmed in practice for the leading data assimilation systems like those implemented at ECMWF or the Met Office (Simmons and Hollingsworth, 2002). Meteorological analyses are also constantly monitored and validated against, e.g., radiosondes; traceability is given to a large degree. Over data void regions, of course, NWP analysis (or short range forecasts) provide the only available information on the state of the atmosphere. In some sense, therefore, it can be argued that NWP data provides the most accurate information available at any arbitrary point on the globe.

On the other hand, it must also be noted that NWP analyses and short range forecasts exhibit complicated error characteristics, which are not always well understood. Using meteorological analysis data for the validation of a new observation system can, therefore, be misleading, if systematic problems of the analyses or the seasonal characteristics of the uncertainties in various parts of the atmosphere are not properly addressed. We present an example in the following section.

We would also like to repeat (as in Section 2) that errors of representativeness in the NWP background may limit the applicability of variational diagnostics to the validation of GRAS level 1b data. An example are tropical small vertical scale waves including gravity waves which might be resolved by GRAS soundings, but which are not represented in NWP products. In this case, the comparison of (sufficiently well resolved) retrievals of geophysical parameters with co-located in-situ or remote sensing data may be the only way to validate small scale structures in the observed data. We provide an example in section 3.8.

Thus, even if the use of 1D–Var diagnostics provides the core component of our validation strategy, the intercomparison of geophysical parameters retrieved from GRAS level1b data with other data may be important, and therefore has its place in the strategy. The data sets chosen, however, should be carefully selected to make sure it is indeed providing the relevant information. For example, in the gravity wave scenario discussed above, co–located high resolution radiosonde data will give useful additional information. Comparisons with co–located retrievals from nadir sounding instruments like AMSU will not, as their vertical resolution is too poor.

3.6 Validation of retrieved atmospheric parameters

It is common practice in the literature for all types of remote sensing instrumentation to compare retrieved geophysical quantities with other co–located “measurements” of the respective quantities. These may include NWP products, ground based and air borne in-situ and remote sensing data, as well as retrievals obtained from satellite instruments. NWP products are also commonly used as transfer reference, by comparing the analyses with other measurements. In the following, we discuss a number of options relevant for GRAS products, along with possible limitations and interpretation caveats. We will specifically address the particularities involved in comparing retrievals of one remote sensing instrument with retrievals from another.

3.6.1 Comparisons with NWP products

The intercomparison of retrieved atmospheric parameters with data interpolated from NWP fields is the easiest way (in technical terms) of validating remote sensing products.

Because NWP data can be interpolated to any location and time, their use (superficially) also avoids the problem of imperfect co–locations, mentioned in section 3.4. As discussed in the previous section, NWP analyses and short range forcasts do provide a thouroughly quality checked data set to compare with.

As also discussed above, NWP analyses and short range forecasts exhibit complicated error characteristics. Positioning and timing errors for synoptic events, for example, re–introduce the issue of the co–location problem in a slightly different way, at least over data poor regions of the Earth. Systematic errors in NWP data, which may depend on season, and the seasonal variability in the random error characteristics need to be taken into account, especially if a validation exercise only covers a short period.

An example for the complexities of NWP based validation can be found in the interpretation of early GPS/MET comparisons with meteorological analyses in the stratosphere. Kursinski et al. (1996), by comparing a small number of profiles against ECMWF analyses on the Southern Hemisphere, found larger (mean) deviations than normal in the tropopause region over the Southern East Pacific. In a careful analysis, the authors attributed these to systematic deficiencies in the ECMWF's analyses tropopause representation in this particular region and during this period of time, related to the small amount of observational data in this part of the world.

The argument was then extended by Rocken et al. (1997), who found a better agreement (both in terms of bias and standard deviations) between GPS/MET and NCEP stratospheric analyses over data rich (Europe and the US) than over data poor region (South Pacific) for GPS/MET data during October 1995. *"Because there is no reason to believe that GPS/MET data quality is different for the two regions"*, the authors concluded that the larger deviations are due to deficiencies in the analysis, and that radio occultation data would therefore be able to improve global NWP analyses. The argument was repeated by Steiner et al. (1999), based on a comparison of 2 days' worth of data during the same period, but with ECMWF analyses.

While we do not question the results of Kursinski et al. (1996), we do believe that later interpretations like those from Rocken et al. and Steiner et al. are somewhat oversimplified. To begin with, Marquardt et al. (2001) demonstrated that the distribution of both biases and standard deviations between GPS/MET and NWP analysis data do not solely depend on region, but also on season. Fig. 7 shows the meridional distribution of latitudinally binned root–mean–square (RMS) temperature deviations between GPS/MET and ECMWF analyses for the prime times in October 1995 and February 1997, respectively[5]. The RMS deviations during October 1995 certainly support the interpretation of Rocken et al. and Steiner et al., as they clearly show a better agreement between GPS/MET and ECMWF analyses in the lower and mid stratosphere on the data rich NH compared to the data poor SH. During February 1997, however, the situation is more or less reversed: In the stratosphere, RMS deviations smaller than, say, 1.5 K,

[5]The increase in both tropical mean and RMS deviations towards the lower boundary of the plotted data is related to the increase in tropospheric water vapour, which has not been corrected for in the plots.

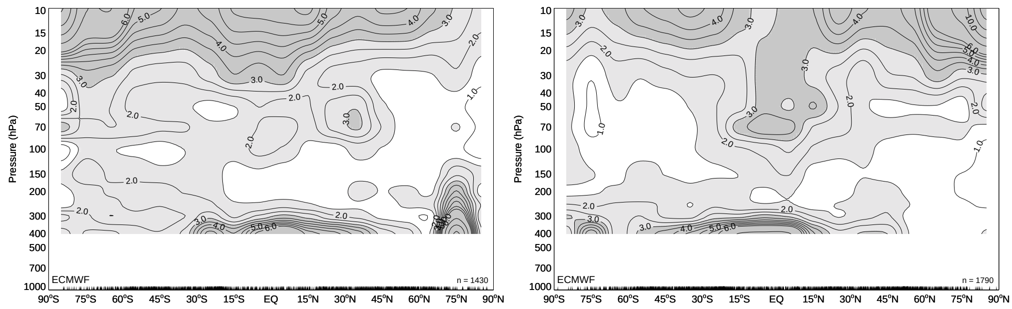

Fig. 7: Root–Mean–Square (RMS) deviations between GPS/MET dry and ECMWF analysed temperatures during October 1995 (left) and February 1997 (right). Light (dark) shading indicates RMS deviations above 1.5 (3) K. The contour line interval is 0.5 K (after Marquardt et al., 2001).

actually cover a larger area / altitude range on the SH compared to the NH. Based on the February distribution of RMS deviations, and repeating the argument of Rocken et al. and Steiner et al., one would have to conclude that the data poor SH stratosphere is at least as well, if not better represented in the ECMWF analysis than the data rich NH stratosphere.

A possible explanation for the RMS deviations in the SH stratosphere during SH summer (February 1997) being as small as those in the NH summer and autumn (October 1995) is that, because of the more complex stratospheric wintertime dynamics, stratospheric analyses exhibit larger errors during the respective winter seasons. The comparison of the two RMS deviations would then suggest that the seasonally related variability in analysis uncertainty is actually larger than the uncertaintainty caused by data spareseness in the summertime stratospheric circulation. On the other hand, the statistical optimisation applied to radio occultation bending angles as well as the initialisation of the temperature profile at some high altitude require a priori assumptions, e.g. the use of a climatology. Marquardt et al. (2001) argued that the climatological a priori may not sufficiently represent the true variability of the wintertime stratosphere. This, too, would cause an increase in the standard deviations between radio occultation temperature soundings and analyses data on the respective winter hemisphere. Clearly, the issue cannot be decided without additional, independent auxiliary data against which both the retrievals and the NWP analyses have to be validated.

In a more recent study undertaken in the context of validation activities for CHAMP, Schöllhammer et al. (2003a,b) compared three different stratospheric analyses. Based on RMS differences between different analyses, the authors argue that, for present day stratospheric analyses, there is actually little difference in the agreement between the hemispheres in the lower stratosphere (up to 30 hPa) during both a NH winter and summer month. This is different at the 10 hPa level, where the enhanced variability of the stratospheric flow in the NH winter may even more than compensate issues related to the data sparesity on the SH. The finding for the lower stratosphere is consistent with Simmons et al. (2003), who concluded that the analyses error in the lower stratosphere of the ECMWF data assimilation system is below the error of radiosonde observations, and comparable to the accuracy of the analyses in the NH. Thus, for present NWP products, and therefore for the validation of GRAS products, the data quality of analyses in the lower stratosphere is not affected by the density of conventional meteorological observations. We note that the improvement of the quality in SH stratospheric analyses became possible through advances in the use of satellite data, which has improved significantly over the last few years. Thus, a seasonal hemispheric difference in the NWP product quality is probably likely to have contributed to the observed behaviour in the meridional distribution of the RMS deviations between GPS/MET and ECMWF temperatures.

The role of the use of a priori data in radio occultation retrievals of stratospheric temperatures is dicussed by Marquardt et al. (2003c), who showed how a change in its use may affect both mean and RMS deviations of CHAMP vs. ECMWF data. Furthermore, (Schöllhammer et al., 2003b) and Marquardt et al. (2003b) show statistics of CHAMP data vs. co–located radiosonde data for the months of January and May 2002, respectively. While the comparison was global, the majority of radiosonde stations is located in the NH. It therefore is effectively a comparison for the NH. Dry temperature retrievals were obtained from GFZ Potsdam. At the time when the comparison was performed, the setup of the GFZ retrieval was very similar with respect to using climatology to the GPS/MET retrieval. Thus, the comparison may have some relevance for the GPS/MET problem. The results in Schöllhammer et al. and Marquardt et al. show significantly larger standard deviations between CHAMP and radiosondes for all levels in the stratosphere. Temperature profiles obtained from a 1D–Var retrieval (which makes a statistical optimal use of a priori data) showed considerably less variability in the standard deviations during the two months. This suggests that the larger standard deviations of dry temperature CHAMP data - and possibly GPS/MET data as well – during NH winter are in fact related to the use of a priori in the retrieval. In contrast to the assumption of Rocken et al. (1997) and Steiner et al. (1999), therefore, it is very well possible that the accuracy of radio occultation retrievals does depend on hemisphere or season, simply due to the use of a priori. In order to reach this conclusion, however, validation against NWP data needed to be carried out during two seasons (to identify the problem), and comparisons with other datasets were required to find out what factors contribute to it.

Having these difficulties in mind, we do suggest that the validation of retrieved atmospheric parameters against NWP analyses and/or short range forecasts will play an important part in any validation exercise. However, statistics need to be calculated over sufficiently long periods (i.e., over more than just a few days), and should cover at least two different seasons. If possible, several different analyses should be used for intercomparison, if only to identify systematic differences between the analyses. If systematic deviations between retrieved geophysical parameters and analyses are found, additional comparisons against independent meteorological data sets are probably required in order to understand which differences have to be attributed to a particular analysis, to a priori used in the retrieval of the geophysical parameters, or to the raw measurements themselves. Atmospheric variability needs to be taken into account properly.

3.6.2 Comparisons with ground based and airborne co–located measurements

Similar to NWP products, routinely used meteorological measurement systems like radiosondes usually undergo periodic validation exercises themselves, and therefore have some traceability chain established. In contrast to NWP analyses, however, conventional meteorological observations are often only available over certain regions, or in few selected locations. One problem for a validation exercise using these data sets is to obtain a sufficiently large number of coincident measurements within a reasonably short period; otherwise, seasonal variations in atmospheric variability may mask seasonally dependent systematic problems in the parameters to be validated. This is especially important for activities within the commisioning phase for the GRAS instrument. The previously mentioned co–location / variability problem, of course, also applies.

We note that only those measurements which exhibit a significant overlap in terms of information content with the observations to be validated can contribute to the validation. For example, measurements of integrated water vapour (which is largely determined by boundary layer humidity) can hardly add useful information for the validation of radio occultation measurements, because the latter provide little or no information on boundary layer moisture abundance.

3.6.2.1 Radiosondes Radiosondes still form the backbone of the global meteorological observation network. They are subject to regular and extensive comparison tests, both coordinated by WMO (e.g., in 1984, 1985, 1989, 1995, and 2001) as well as undertaken by national meteorological services. Thus, radiosondes probably exhibit the best traceability chain for atmospheric observations that can reasonably be expected. Modern sensors like the Vaisalla RS80 radiosonde may provide temperature soundings with errors in the order of 0.5 K, even at stratospheric altitudes, and relative humidities with errors in the order of 5. . . 10% throughout the troposphere (Nash, pers. comm.).

Nevertheless, large geographical differences in the data quality of radiosondes do exist. This is mainly related to poor handling practices and to the use of out–of–date equipment and sensors. National meteorological services as well as the ECMWF maintain lists of reliable radiosonde stations with a proven record of high quality measurements, and monitor the performance of the radiosonde network. We recommend to restrict the validation of GRAS retrieval products to high quality radiosonde data; the required quality information is readily available from Met Services.

Because of the nearly global coverage of radiosonde observations, a reasonably large number of matches can be obtained within a sufficiently short period of time. The number of matches will, of course, depend on the collocation criteria. CHAMP, for example, regularly provides between 300 and 500 matches each month for distances below 300 km and 3 hours between occultation and radiosonde launch site. This is sufficient

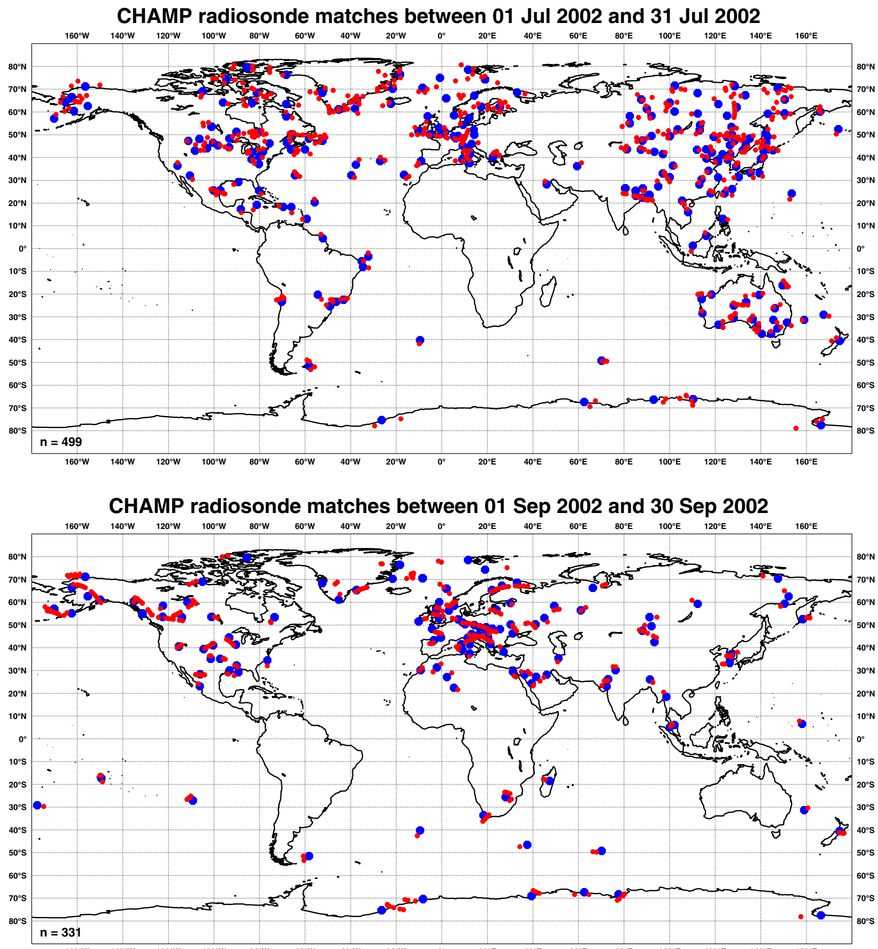

Fig. 8: Distribution of matches between radiosondes (blue) and CHAMP occultations (red) during July (left) and September 2002 (right).

to calculate monthly global statistics. Fig. 8 shows the distribution of matches during July and September 2002. Note that the details of the distribution will depend on the orbit geometry of the satellite. In case of CHAMP, because of its drifting orbit, the global distribution changes over time. In July 2002, a large number of matches was obtained over Asia and Australia, while there were only few co–located observations over Asia and none over Australia during September 2002. Also note that even in mid–latitudes where the data density of occultations is highest, an individual radiosonde site is usually not involved in more than 4 or 5 matches per month. In many cases, no matches occur at all. Thus, for a single satellite radio occultation mission like GRAS, the operational radiosonde network will provide sufficiently many matches to perform routine validation tasks. Validation measurements undertaken at a single site, however, will require long periods of data gathering before meaningful statistics can be calculated, and will therefore suffer from the seasonality problem mentioned earlier.

Radiosonde data obtained via the Global Telecommunication System (GTS) is available on a vertical grid of standard ("mandatory") pressure levels. Additionally data on “significant” levels, in principle, allows the reconstruction of somewhat higher resolved vertical profiles, and will also give information on the tropopause as seen in the radiosonde ascent. However, even significant levels will not provide the complete information as is available from a fully resolved radiosonde ascent, where the vertical resolution is in the order of 50 m. Some national meteorological services like the Met Office have recently started to archive the full resolution radiosonde data for ascents undertaken at their own sites. This type of data could be used to validate small vertical scale variation as caused by atmospheric gravity waves; signatures of the latter have been found in radio occultation data (e.g., Steiner and Kirchengast, 2000; Tsuda et al., 2000). Note, however, that the retrieval of atmospheric parameters from GRAS level 1b data needs to be sufficiently well resolved in order to represent these small scale types of atmospheric fluctuations. As we pointed out in section 3.5 (and will demonstrate in section 3.8), a 1D–Var using a standard NWP vertical level structure does not provide a sufficient resolution; the exploitation of fully resolved radiosonde data would therefore require a high resolution 1D–Var retrieval.

Also note that gravity waves occur predominatly within the tropics and subtropics. Thus, the validation of the small scale features of GRAS level1b data should primarily be undertaken using tropical and subtropical radiosonde data. Unfortunately, fully resolved radiosonde data is not usually available from these regions. Because of the small number of matches between radio occultations and radiosonde ascents for an individual site, however, we recommend against a dedicated tropical validation campaign for GRAS. Instead, a network of tropical and subtropical radiosonde stations willing to launch radiosondes on demand whenever a GRAS occultation is located near to their site might be established for the purpose of GRAS level1b validation, at least during the commissioning phase, provided that the full resolution radiosonde is passed back to EUMETSAT. This would require the ability to predict the locations of GRAS occultations several days ahead (which is probably possible). Data from tropical measurement campaigns run for other scientific reasons, if available, might also be utilised. It should be kept in mind, though, that the number of co–located measurements will still be small compared to the amount of comparison data obtained from the operational, but less well resolved radiosonde network.

3.6.2.2 Lidars Lidar (Light detection and ranging) instruments emit a beam of light (from a laser), and determine properties of the atmosphere by analysing the backscattered signal. Rayleigh lidars (e.g., Hauchecorne and Chanin, 1980) measure upper air temperature in the altitude range between (typically) 25 and 100 km. Temperature lidars exhibit theoretical temperature errors being below 1-2 K below 50 km, especially after temporal averaging; their vertical resolutions vary between 75 and 300 m, depending on operating modes. Lidar systems for tropospheric humidity (Raman lidars) as well as for trace gas species (ozone, aerosol) are also frequently used.

Lidar measurements have been successfully used for satellite validation in the past, although the number of co–located measurements is usually small (typically several tens). It should also be kept in mind that Lidars, as all remote sensing instruments, do not provide in situ measurements of geophysical quantities. Processing software may add errors and require some "optimisation" (e.g., Leblanc et al., 1998). Thus, only data from lidar sites that undergo regular validation and intercomparison efforts themselves should be considered for validation of GRAS products.

A group of sites that does untertake regular validations is the Network for the Detection of Stratospheric Change (NDSC). Currently, the NDSC network consists of 11 primary and complementary stations providing stratospheric temperature measurements by lidar. An additional mobile temperature lidar is available for measurement campaigns[6]. Note, however, that not all of these instruments are operated all year round. Similar to radiosondes, lidar measurements undertaken at a single site only will not produce many co–located measurements over a short period like a month. Lidars also only operate at nighttime and under clear sky conditions. This further limits the number of available co–located observations. Thus, it is possible that only a few co–located measurements can reasonably be expected over a period of a month. We therefore believe that the usefulness of Lidar observations for the validation of GRAS products is limited, especially during the commissioning phase. In the long term, Lidar measurements will provide an additional data set to compare GRAS data with, provided the problem of seasonally varying atmospheric conditions can be solved or shown to be negligible.

3.6.3 AMDAR

The AMDAR (Aircraft Meteorological Data Reporting) system consists of automated wind, temperature and pressure measurements along with the aircrafts position and a measure of turbulence, obtained from a large fleet of civilian commercial passenger aircrafts, predominantly from Europe, the US and Japan. Data is transmitted via satellite links to national Met Services and routinely used in today's data assimilation systems. Although the data processing involved it quite complex, errors in reported wind and temperatures are comparable with those of radiosonde systems. Thus, AMDAR observations can provide high quality single level data in cruise and detailed profile data up to cruise levels near airports. Because of the distribution of the major flight tracks (mainly over the US, Europe, and the Northern Atlantic), data coverage is restricted to these areas. The European AMDAR project (E-AMDAR) aims at providing ascent and descent measurements as a complement and potential substitute for radiosondes on the territory of European members at over 140 European airports, 35 of which should provide at least three profiles per hour during the day time. The programme also works on optimising the data collection process and addressing the need for humidity measurements as soon as practicable. Similar activities are undertaken in other countries. Unfortunately, it is not clear when large numbers of humidity measurements will become available on a routine basis due to the technical difficulties in developing a robust,

[6] Source: NDSC web page, `http://www.ndsc.ncep.noaa.gov`

reliable and accurate sensor[7].

Vertical AMDAR profiles would certainly provide a useful comparison set for GRAS tropospheric refractivity, temperature and moisture retrievals if a sufficiently large number of humidity measurements become available during the GRAS commissioning phase. The large number of the routinely undertaken ascent and descent measurements over major airports should be sufficient to gather several ten to hundred matches over periods of a few weeks or a month. It should be noted, however, that the data is restricted to altitudes below $\approx$ 200 hPa. If no, or only few humidity measurements, are available at the time, therefore, their value for the validation of GRAS retrievals might be limited as water vapour might contribute significantly to GRAS observations for the largest part of this altitude range.

3.6.3.1 MOZAIC MOZAIC (Measurement of Ozone and Water Vapour by Airbus In–service Aircraft, see Marenco et al., 1998; Helten et al., 1998) provide in-situ measurements of relative humidity and temperature (among others) along the flight path of five commercial aircraft. Of special interest for the validation GRAS are the vertical profiles of these quantities which are obtained during ascent and descent near about 200 cities. However, data is collected from the respective aircraft only once a month, and becomes available to the research community even later. Because of the small number of aircrafts involved, only few coincident measurements with GRAS occultations can be expected. We therefore do not expect that MOZAIC measurements will provide a useful source of validation data for GRAS.

3.6.3.2 GPS ground stations GPS ground stations measuring total zenith path delay or slanted path delay exploit the same signals as radio occultations do, but are mainly sensitive to humidity in the boundary layer of the Earth's atmosphere. Radio occultations, however, exhibit the largest errors in the lowest few kilometres of the atmosphere, and may not be able to provide information on the atmospheric boundary layer at all. Thus, the regions of the atmosphere which are covered by GPS ground and space based soundings have little or no overlap, and an intercomparison between these will not provide additional information for the validation of GRAS data.

3.6.4 Comparisons with other types of space based remote sensing instruments

Compared to, e.g., radiosondes, the number of vertical temperature and humidity profiles retrieved from remote sensing instruments like those from ATOVS package is huge. Therefore, many co–located measurements can potentially be compared in an validation exercise. However, remote sensing instruments do not provide true "measurements" of atmospheric parameters. Instead, retrieved atmospheric profiles contain

[7] Sources: EUMETNET web page, `http://www.eumetnet.eu.org`, and E-AMDAR web page, `http://www.metoffice.com/research/interproj/amdar/index.html`.

a priori data. An intercomparison of retrievals from different sources may therefore be, at least partially, a comparison of the a priori data used within these retrievals. They also differ in terms of resolution, and usually exhibit complicated error statistics. If retrievals of two different observing systems are to be compared with each other, these issues have to be taken into account.

A framework for a comparison of retrievals from different remote sensing instruments was given by Rodgers (2000) and Rodgers and Connor (2004), and is partially based on Eyre (1987). It is, of course, obvious that only those instruments can be reasonably compared with each other if they are sensitive to the same atmospheric variables, e.g. humidity and temperature. In the following, we initially assume that the state vectors used in both retrievals are identical. We will also assume that the same a priori information $\mathbf{x}_b$ has been used in both retrievals. A formal treatment for different a prioris, only partially overlapping state vectors, and for the comparison of derived quantities (like layer means, thicknesses or values interpolated from different state vectors) will be given in later sections. Another approach, namely to use the result of one retrieval as a priori in the other, will also be discussed briefly.

3.6.4.1 Comparing retrieved quantities We begin by writing down a linearised version of the retrieval procedure Eyre (1987); Rodgers and Connor (2004):

$$\hat{\mathbf{x}} - \mathbf{x}_b = \mathbf{A}(\mathbf{x} - \mathbf{x}_b) + \varepsilon \,. \tag{26}$$

Here, $\mathbf{x}_b$ and $\hat{\mathbf{x}}$ denote the background (or a priori) data used within the retrieval and its solution, respectively; ε describes the retrieval error originating from the measurement errors, and $\mathbf{x}$ represents the "true" state of the atmosphere. The linear operator $\mathbf{A}$ is usually called the "averaging kernel" (Backus and Gilbert, 1970) or "model resolution matrix" (Menke, 1989); "state resolution matrix" and "resolving kernel" are also common names.

In principle, any retrieval can be linearised and written in the above way; the same framework may therefore be applied to both variational and non–variational retrievals. In practice, however, the linearised forms of non–variational retrievals are not often available; the use of a priori data complicates the derivation of the appropriate linearised expression. Within the variational framework, however, averaging kernels are readily available, and given by

$$\mathbf{A} = \left(\mathbf{H}^T(\mathbf{E}+\mathbf{F})^{-1}\mathbf{H} + \mathbf{B}^{-1}\right)^{-1}\mathbf{H}^T(\mathbf{E}+\mathbf{F})^{-1}\mathbf{H} \,. \tag{27}$$

We therefore believe that the application of the framework developed in this and the following section will, in practice, be restricted to cases where the retrieved atmospheric products of both remote sensing instruments are obtained within a variational (1D–Var) context.

Once the averaging kernels $\mathbf{A}_1$ and $\mathbf{A}_2$ of the two observing systems are known, the difference of two retrievals $\hat{\mathbf{x}}_1$ and $\hat{\mathbf{x}}_1$ is given (in linear approximation) as

$$\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1 = (\mathbf{A}_2 - \mathbf{A}_1)(\mathbf{x} - \mathbf{x}_b) + \varepsilon_2 - \varepsilon_1 \ . \tag{28}$$

The corresponding covariance matrix of the difference is then given by

$$\mathbf{P}_{\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1} = (\mathbf{A}_2 - \mathbf{A}_1)^T (\mathbf{x} - \mathbf{x}_b)(\mathbf{A}_2 - \mathbf{A}_1) + \mathbf{P}_1 + \mathbf{P}_2 \ , \tag{29}$$

and the corresponding χ^2 by

$$\chi^2 = (\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1)^T \mathbf{P}_{\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1} (\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1) \ . \tag{30}$$

In theory, two retrievals from different observing systems, sharing the same a priori and state vectors, may be compared by calculating χ^2 as given by eq. (30), with the degrees of freedom being the number of elements of the state vector. If the two retrievals fail this test, it may be concluded the retrievals are inconsistent with each other.

In practice, however, the error estimates for both the background and the observations (and therefore the averaging kernels) may not be known sufficiently well. Similar to the use of 1D–Var diagnostics, some tuning of critical threshold values might be required.

3.6.4.2 Removing a priori In case that the two retrievals do not share the same a priori information, it is possible to formally "remove" the a priori information if the latter is known for each individual retrieval. If $\mathbf{x}_c$ is a state vector drawn from a "comparison ensemble" (which may well be identical with one of the two a prioris used in the retrievals), transformed retrievals may be calculated as

$$\tilde{\mathbf{x}}_i = \hat{\mathbf{x}}_i + (\mathbf{A}_i - \mathbf{I})(\mathbf{x}_{b,i} - \mathbf{x}_c) \ , \tag{31}$$

where $\mathbf{x}_{b,i}$ denote the individual retrieval's a priori state vectors. Inserting the linearised retrieval (26) yields

$$\tilde{\mathbf{x}}_i - \mathbf{x}_c = \mathbf{A}_i(\mathbf{x} - \mathbf{x}_c) + \varepsilon_i \ . \tag{32}$$

Thus, $\tilde{\mathbf{x}}_i$ is, within the validity of the linear approximation, the retrieval for observing system i, but corrected for the new $\mathbf{x}_c$ as (joint) a priori. The comparison of the two retrievals may then proceed as in the previous section, but using $\tilde{\mathbf{x}}_i$ instead of $\hat{\mathbf{x}}_i$, and $\mathbf{x}_c$ instead of the $\mathbf{x}_{b,i}$.

3.6.4.3 State vectors only overlapping partially If the state vectors of the observing systems overlap only partially, eq. (26) needs to be generalised slightly. If $\mathbf{s}$ denotes the full state vector, $\mathbf{x}$ the common and $\mathbf{e}$ the extra elements, the linearised retrieval equation can be written as

$$\hat{\mathbf{s}} - \mathbf{s}_b = \begin{pmatrix} \hat{\mathbf{x}} - \mathbf{x}_b \\ \hat{\mathbf{e}} - \mathbf{e}_b \end{pmatrix} = \begin{pmatrix} \mathbf{A}_{xx} & \mathbf{A}_{xe} \\ \mathbf{A}_{ex} & \mathbf{A}_{ee} \end{pmatrix} \begin{pmatrix} \mathbf{x} - \mathbf{x}_b \\ \mathbf{e} - \mathbf{e}_b \end{pmatrix} + \begin{pmatrix} \varepsilon_x \\ \varepsilon_e \end{pmatrix} \tag{33}$$

The state vector part which is to be compared with the retrieval from the second observing system, therefore, is

$$\hat{\mathbf{x}} - \mathbf{x}_b = \mathbf{A}_{xx}(\mathbf{x} - \mathbf{x}_b) + \mathbf{A}_{xe}(\mathbf{e} - \mathbf{e}_b) + \varepsilon_x \ . \tag{34}$$

Rearranging gives the error of $\hat{\mathbf{x}}$ as

$$\hat{\mathbf{x}} - \mathbf{x} = (\mathbf{A}_{xx} - \mathbf{I})(\mathbf{x} - \mathbf{x}_b) + \mathbf{A}_{xe}(\mathbf{e} - \mathbf{e}_b) + \varepsilon_x \ . \tag{35}$$

Thus, the error in $\hat{\mathbf{x}}$ now contains an additional part due the interrelation between the common part of the state vector and its extra terms. The covariance of the the common part of the state is therefore given by

$$\mathbf{P}_{\mathbf{x}} = \mathbf{A}_{xe}\mathbf{P}_{eb}\mathbf{A}_{xe}^T + \mathbf{P}_{\varepsilon_x} \tag{36}$$

where $\mathbf{P}_{eb}$ is the error covariance of the a priori of $\mathbf{e}$; Rodgers and Connor denote the contribution of the term $\mathbf{A}_x e(\mathbf{e} - \mathbf{e}_b)$ as "interference error". If the common part of the state vectors have been identified and their covariance matrices been corrected for the interference error, the comparison may proceed as described in section 3.6.4.1.

3.6.4.4 Derived quantities Let us assume we are interested in the estimate of a linear function of the true state vector $\mathbf{x}$, e.g. a layer mean, a column or a profile linearly interpolated onto a new set of levels. This may be written as

$$\mathbf{z} = \mathbf{z}_b + \mathbf{B}(\mathbf{x} - \mathbf{x}_b) \ . \tag{37}$$

Given we have an estimate $\hat{\mathbf{x}}$ of the true state vector - how can we obtain an optimal (in the sense of Rodgers) estimate of $\mathbf{z}$? Assuming that we formally know the probability of the state given the measurements, i.e. $P(\mathbf{x}|\mathbf{y})$, the expected value (i.e., the most likely estimate) of $\mathbf{z}$ is given by

$$\hat{\mathbf{z}} = \int P(\mathbf{x}|\mathbf{y})[\mathbf{z}_b + \mathbf{B}(\mathbf{x} - \mathbf{x}_b)]\, d\mathbf{x} \ . \tag{38}$$

If the linear transform does not depend on $\mathbf{x}$, the equation can be rearranged to give

$$\hat{\mathbf{z}} = \mathbf{z}_b + \mathbf{B}\int P(\mathbf{x}|\mathbf{y})(\mathbf{x} - \mathbf{x}_b)\, d\mathbf{x} \ . \tag{39}$$

Thus, if $\hat{\mathbf{x}}$ is an optimal retrieval, i.e. if $\hat{\mathbf{x}} = \int P(\mathbf{x}|\mathbf{y})\mathbf{x}\, d\mathbf{x}$, an optimal estimate of $\mathbf{z}$ can be obtained via

$$\hat{\mathbf{z}} = \mathbf{z}_b + \mathbf{B}(\hat{\mathbf{x}} - \mathbf{x}_b)\, d\mathbf{x} \ . \tag{40}$$

Clearly, the averaging kernel and error covariance matrix of $\hat{\mathbf{z}}$ are then given by

$$\mathbf{A}_{\mathbf{z}} = \mathbf{B}\mathbf{A} \tag{41}$$

and

$$\mathbf{P_z} = \mathbf{BPB}^T \ , \tag{42}$$

respectively.

If $\hat{\mathbf{x}}$ is not a statistically optimal retrieval, or was obtained using a different a priori, a procedure similar to the one outlined in section 3.6.4.2 may be applied. We refer to section 4.2 of Rodgers and Connor (2004) for details.

Again, after correcting averaging kernels and error covariance matrices for each retrieval, the intercomparison may proceed as described in section 3.6.4.1.

3.6.4.5 Var–in–Var A slightly simpler method might be to use a retrieval from one (say the second) observing system as a priori in the retrieval of the other (say the first). Because the retrieved profile of system 2 exhibits smaller errors than the a priori, the 1D–Var for the first instrument will be more constricted. Note that the higher resolution retrieval would be expected to reproduce the lower resolution measurement better than the other way round. With respect to the order of the retrievals, therefore, the method will not be symmetric.

If the Var–in–Var cannot be implemented, it is possible to simulate this approach if, as in the previous sections, the averaging kernels and a comparison ensemble for both retrievals are available. If $\hat{\mathbf{x}}_{12}$ denotes the first retrieval, using the second as a priori, the linearised retrieval gives

$$\hat{\mathbf{x}}_{12} = \mathbf{x}_c + \mathbf{A}_1(\hat{\mathbf{x}}_2 - \mathbf{x}_c) \tag{43}$$

and it’s averaging kernel is apparently

$$\mathbf{A}_{12} = \mathbf{A}_1\mathbf{A}_2 \ . \tag{44}$$

The difference between $\hat{\mathbf{x}}_{12}$ and the usual retrieval solution $\hat{\mathbf{x}}_1$ is given by

$$\delta_{12} = \hat{\mathbf{x}}_1 - \hat{\mathbf{x}}_{12} = (\mathbf{A}_1 - \mathbf{A}_1\mathbf{A}_2)(\mathbf{x} - \mathbf{x}_c) + \varepsilon_1 - \mathbf{A}_1\varepsilon_2 \tag{45}$$

with covariance

$$\mathbf{P}_{\delta_{12}} = (\mathbf{A}_1 - \mathbf{A}_1\mathbf{A}_2)\mathbf{P}_c(\mathbf{A}_1 - \mathbf{A}_1\mathbf{A}_2)^T + \mathbf{P}_1 - \mathbf{A}_1\mathbf{P}_2\mathbf{A}_1^T \ . \tag{46}$$

3.6.4.6 Practical application In practice, one might set up a 1D–Var retrieval for both GRAS bending angles and, say, co–located ATOVS soundings, using identical a prioris and the same vertical level structure (e.g., on the 43 RTTOV standard pressure levels). Because a priori and state vector are identical, the simplified procedure described in section 3.6.4.1 can be applied directly. This involves calculating the averaging kernels of both observing systems, of $\mathbf{P}_{\hat{\mathbf{x}}_2 - \hat{\mathbf{x}}_1}$, and finally of the χ^2 given by (30).

Comparing the actual value with the theoretically expected one finally gives an indication if the two measurements are consistent with each other (within their uncertainty, and within limits of the different measurement characteristics).

If retrievals from observing systems which are are not handled inhouse shall also be compared with GRAS data, data providers of the external retrievals need to provide their averaging kernels and any a priori along with the retrievals. Some data providers, e.g. for MIPAS, are aware of the problem and interested in comparing data within the above framework (Stiller et al., 2003). As the two state vectors of the retrievals are likely to differ, and may be based on different vertical level structures, a procedure as outlined in sections 3.6.4.3 and 3.6.4.4 will be required.

The above demonstrates that the intercomparison of retrievals from different remote sensing instruments is a complicated task. As a result, we get information if two retrieved profiles are consistent within the errors of the raw measurements, or not. We do not yet know which of the two retrievals is in error. It can be argued that the same information can already be obtained from the 1D–Var diagnostics: here, we get information whether a specific GRAS (or ATOVS) measurement is consistent with the NWP a priori. If the raw data from the second instrument to be compared with has been assimilated into the NWP product, little or no additional information can obviously be expected from a comparison of two retrievals.

We therefore recommend that the additional work associated with the intercomparison of different remote sensing instruments is only undertaken if an additional benefit can be expected. One example would be the comparison of GRAS retrievals with retrievals from a remote sensing instrument which provides a higher vertical resolution than NWP products. This might help in the validation of small scale vertical structures in the lower tropical stratosphere. Note, however, that this requires retrievals with a sufficiently high vertical resolution for both observing systems. There is no point in undertaking this comparison with retrievals on, say, the 43 pressure levels from the RTTOV setup.

3.7 Validation campaigns

In theory, dedicated validation campaigns for a remote sensing instrument like GRAS have several advantages: More co–located measurements from a multitude of different ground based and possibly airborne instruments may become available in a comparatively short amount of time than otherwise possible. Standard instruments (like radiosondes) may provide improved research quality data and higher vertical and temporal sampling than usual. Data from additional research instruments may also be available, and campaigns may involve experience from academia. Measurements taken from research aircraft will be able to circumvent the co–location problem, because flights can be directed towards the expected location of a GRAS occultation. Airborne radio occultation soundings could be of special interest, provided that “space–based equivalent bending angles” with error characteristics comparable or better than those obtained from space borne radio occultations can indeed be derived.

The main disadvantage of dedicated measurement campaigns was already pointed out in section 3.6.2.1: Ground based measurements at a single or only few sites will not provide enough matches with observations obtained from a single satellite radio occultation mission like GRAS within a short period of time. Both groundbased and airborne measurements are highly weather dependent. Validation campaigns are also expensive. If airborne radio occultation measurements should not be available with a bending angle quality comparable to space based radio occultations, we recommend against a dedicated validation campaign for GRAS. If they are available, it should be carefully considered if the small number of co–located bending angle profiles that can be obtained during a measurement campaign are indeed worth the effort. In particular, it should be considered which additional information can realistically be obtained from a validation campaign compared to what's already available from 1D–Var diagnostics and the comparison with conventional operational data sets.

3.8 Practical experience

For the purpose of illustration, we have applied some of the proposed methods to data obtained from a preliminary version of the Met Office's Radio Occultation Processing Package (ROPP). The system uses CHAMP excess phase delay data as available from GFZ Potsdam or UCAR in order to calculate ionospheric corrected bending angle profiles. Calculations are based on the standard geometrical optics approach (Vorob'ev and Krasil'nikova, 1994); no advanced retrieval methods like the Canonical Transform (Gorbunov, 2002b,c,a) or the Full Spectral Inversion (Jensen et al., 2003) are applied. Results of the error propagation for the error correlation structure were already presented in Fig. 5. In the current processing scheme, ionospheric corrected bending angles are downsampled to a 247 level vertical grid between the surface and 60 km altitude. The grid roughly samples one Fresnel diameter with four data points.

3.8.1 Statistical methods for noise estimation

Fig. 9 shows the global statistics of GCV estimates of the CHAMP L1 excess phase delay noise during September 2002 in a similar fashion as in the Monte–Carlo simulated Fig. 6. The qualitative behaviour (i.e., an increase in the estimated noise levels towards longer window lengths) is the same as in the simulated case. The transition towards higher noise estimates (again as function of window length) is less well pronounced in the real data, probably due to the fact that the 50 Hz tracking of the GPS signal at the beginning of an occultation occures at a different altitude for each event in the real data. A surpring effect is the large variability of estimated excess phase delay noise, indicated by the large standard deviations. This suggests that individual profiles may vary significantly with respect to their noise level. Fig. 10 shows the meridional distribution of the estimated excess phase noise of the L1 channel as well as in the standard deviation of the estimate, grouped into 10° wide latitude bins for three different months. The figures suggest that there seems to be a generally enhanced noise level in

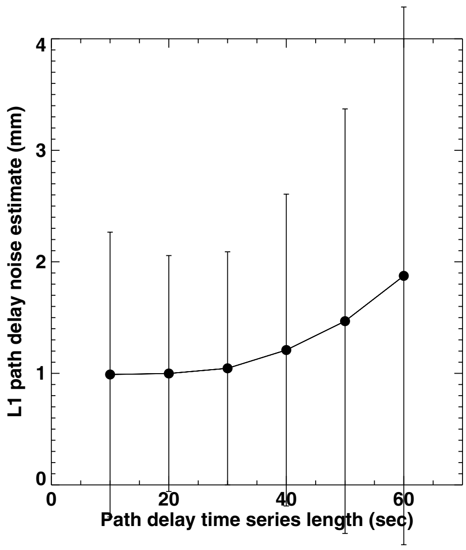

Fig. 9: Mean L1 phase delay noise estimated by GCV as function of data window length subject to the GCV analysis (measured from the beginning of the occultation) for all CHAMP data during September 2002. The error bars denote the 1 σ standard deviation of the estimated values.

tropical and subtropical occultations. Variability is largest in the mid latitudes of the NH, although a secondary maximum appears in the SH subtropics. To our knowledge, this is the first indication of geographically varying noise levels in (relatively) little processed radio occultation data; it is currently not known what causes the meridional variability in the mean noise levels, or their seasonal and longitudinal variability. One possible reason is ionospheric activity seen on the excess phase delay data; different error characteristics of ground station data used in the double differencing applied for clock correction is another. Thermal changes of the spacecraft due to day/night effects might also be relevant. More research is certainly required.

CHAMP data is known to exhibit 1 and 2 Hz "spikes" in the raw phase data (Wickert, pers. comm.); in double differenced data, these become visible in a smoothed form with the naked eye in ionospheric corrected excess phases. Despite the usually applied smoothing of the excess phase delays, these spikes show up in ionospheric corrected bending angles as wave–like structures with vertical wavelengths of a few km. Amplitudes can reach several few microradians. Applying a GCV noise estimator for uncorrelated noise to a signal with such highly correlated noise fails (in that the estimated noise is too small)[8]; a Sokolovskiy and Hunt (1996) like approach systematically estimates larger noise levels. These are typically in the order of several microradians, which is by a factor of two or more larger than which is expected from the error propagation

[8]Similarly, we assume that spike-related structures in the excess phases are not taken into account when the noise is estimated by GCV . This conclusion is based on the width of the smoothing windows generated by the GCV algorithm as well as on our analysis of the GCV residuals. As a consequence, the meridional distribution of the noise levels in Fig. 10 is likely not to be related to the spikes.

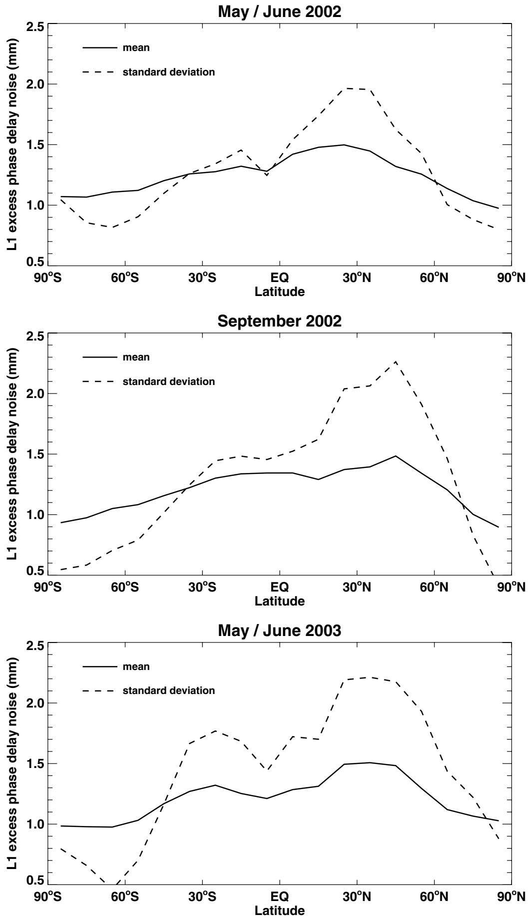

Fig. 10: Estimated L1 excess phase delay noise as function of latitude for three 4 week periods during May/June 2002, September 2002, and May/June 2003.

presented in section 3.2. We note that this example illustrates how the inconsistency between noise estimates of the real data and theoretically expected noise levels may indicate a potential problem in the data (in this case the existence of the 1 Hz and 2 Hz spikes).

3.8.2 1D–Var statistics and forward modelling

The thinned bending angle data are used in a prototype implementation of the bending angle 1D–Var (see Appendix A). The main difference of the 1D–Var implemented in the Met Office processing scheme is that it uses ECMWF's operational 60 vertical levels rather than the fixed pressure levels as in the prototype software. As background data, short range forecasts from ECMWF have been interpolated from a 3 hourly / $0.5^\circ \times 0.5^\circ$ resolution onto the occultation's position.

For the purpose of this section, we have chosen to let observational errors follow a simple model which is similar to a refractivity error model originally proposed by Kursinski et al. (1997b). In the lower stratosphere, bending angle errors are set to a fixed relative

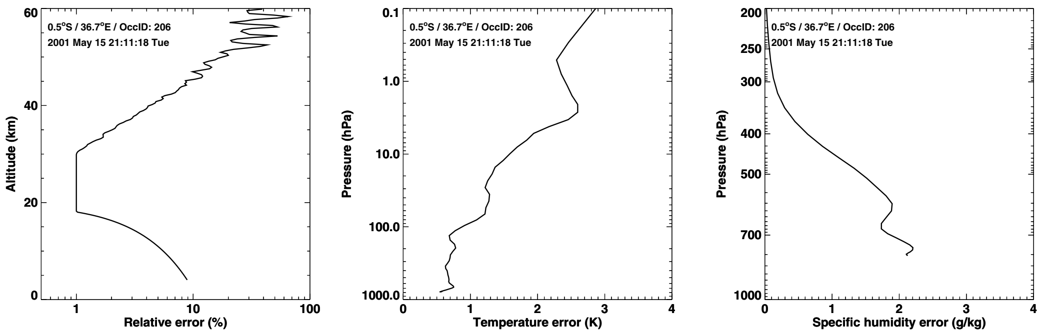

Fig. 11: Assumed errors in bending angle (relative, in %; left), background temperature (in K, middle) and specific humidity (in g/kg; right) for the CHAMP profile shown in Fig. 14. Note the different vertical scales.

value of 1%; below the tropopause, the error increases linearly to 10% at the surface. Note that the tropospheric error estimate uses some background information on the location of the tropopause. In the mid and upper stratosphere, the minimum value of the bending angle error is set to 3μrad, a heuristically set value which takes into account our results from the bending angle noise estimation discussed above. For simplicity, the errors are assumed to be uncorrelated in the vertical. The vertical profile of the relative bending angle error for a tropical observation taken by CHAMP in May 2001 is shown in Fig. 11. We note that this particular CHAMP observation is close to the radiosonde station of Nairobi, and will be used for a number of illustrations in the following.

Background errors are based on the operationally available error estimates provided by ECMWF (Fisher and Courtier, 1995; Andersson et al., 2000) for the initial state of a forecast. Errors are inflated according to the length of the forecast period assuming an exponential error growth with an error doubling time of 1.5 days (for error growth in NWP forecasts, see, e.g., Savijarvi, 1995). The corresponding vertical error profiles for temperature and specific humidity are also shown in Fig. 11. Error covariance matrices are constructed using globally averaged vertical error correlations for temperature and humidity kindly provided by M. Fisher (ECMWF).

In the Met Office's implementation of the bending angle 1D–Var, thinned bending angle values undergo a background quality control (BGQC); data points deviating by more then 10 expected standard deviations from the background are not used in the retrieval (see eq. (5) in Section 2). The BGQC typically removes the data points in the lowest few hundred meters at the lower end of the observed profiles. This is related to the well known systematic bias of bending angles for the phase–only retrieval. The BGQC here simply acts as limiting the amount of bias that is accepted by the system. In rare cases, individual data points in the mid and upper stratosphere are rejected from the bending bending angle profile. These are related to peaks of larger oscillatory perturbations in the ionospheric corrected bending angle profiles related to the spikes in the raw CHAMP data. Complete profiles are rejected if more than 50% of the data points are rejected.

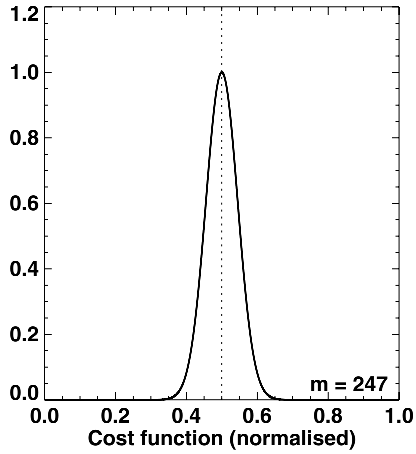

Fig. 12: Theoretical frequency distribution of a normalised cost function for 247 degrees of freedom. Note that maximum of the distribution has been normalised to 1.

The theoretical distribution of the cost function values for 247 degrees of freedom (i.e., the number of observations used in a typical 1D–Var retrieval) is shown in Fig. 12. It exhibits its peak at a value of 0.5, and a half width of less than 0.1. A histogram of the actual distribution of normalised cost function values J_n for extratropical CHAMP profiles during September 2002 is shown in Fig. 13. As can be seen immediately, both the mean location (around 0.84) and the width of the cost function distribution are larger than theoretically expected. This suggests that the error specifications for either background or observations (or both) are too optimistic; they need revision. In the current context (where we knowingly have assumed a crude, imperfect error model for the observations), this is not surprising. If we would not have known about the imperfection of our error estimates, however, the J_n statistics would have told us so immediately.

So even if the error characterisation is not perfect, cost function values provide useful information. Another example is the meridional distribution of the cost function, also shown in Fig. 13. Note the overall increase of cost function values in the subtropics and tropics; the same regions are also characterised by a larger number of profiles exhibiting unusally 'large' cost function values compared to the extratropics.

We already showed in Section 2 how the monitoring of values of J_n may help to identify profiles exhibiting gross errors in the observations. The cost function can also be used to implement an automated quality control. This does, of course, assume that gross errors do indeed occur only rarely, so that the distribution of J_n is reasonably well behaved. A heuristic approach to quality control could be to calculate a threshold value for J_n based on its actual distribution. For example, those 5% of the data with the largest J_n values might flagged as being of poor quality. The solid line in Fig. 13 denotes the critical threshold value as function of latitude if this procedure is applied to data in individual 10° wide latitude bins; the threshold curve has been made symmetric with respect to the equator. By construction, this method rejects about 5% of the entire data set. In an

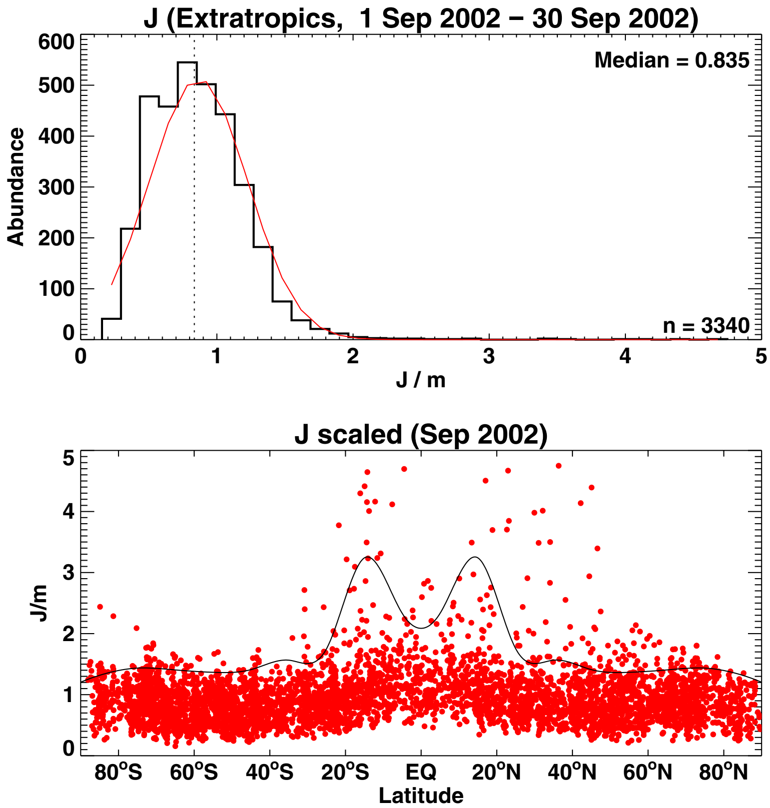

Fig. 13: Top: Frequency distribution of normalised cost function values for extratropical (poleward of 30°) CHAMP bending angles during September 2002 (black). The solid red curve shows a normal distribution fitted to the actual frequency distribution. Bottom: Meridional distribution of all cost function values during the same period. The solid line denotes a heuristic critical value for a cost function based objective quality control (see text for details).

operational setup, rejection limits might be based on, e.g., the last 30 days of data.

The meridional costfunction distribution shown in Fig. 13 suggests that the fit between observations and background is generally poorer in the subtropics and tropics compared to the extratropics. A possible explanation could be that the well known problems in the retrieval of bending angles in the lower and mid troposphere due to atmospheric multipath surface in the cost function value distribution. This is, however, not the case. Fig. 14 shows retrieved temperature and specific humidity profiles for the occultation near Nairobi, along with the vertical profiles used as background. Also shown are the assumed background errors of humidity. Clearly, the retrieval is well within the error limits of the background; at least in this case, atmospheric multipath does not cause an increase in the cost function value.

As a matter of fact, the value of the normalised cost function for this particular case is 0.55. With the overall characteristics of the cost function values in the tropics and subtropics in mind, this does not signal the profile in question as in any way problem-

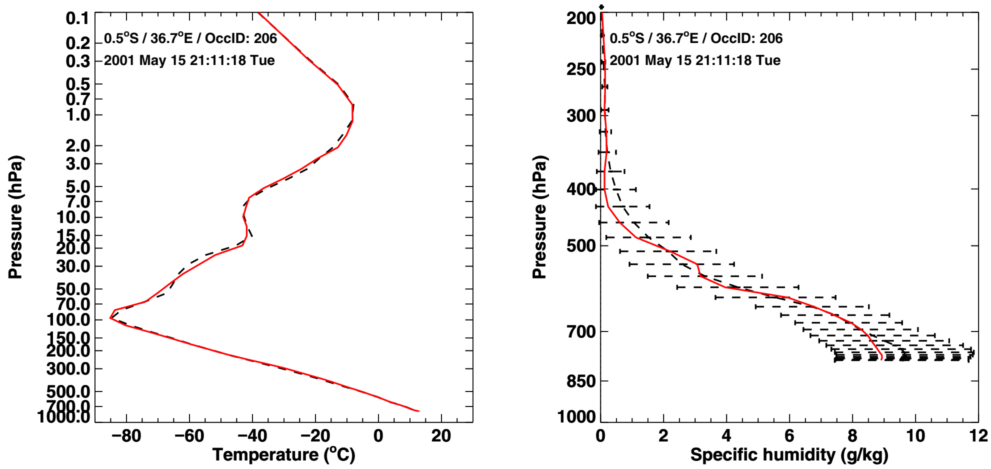

Fig. 14: Retrieved (solid, red) and background (dashed) temperature (left) and specific humidity profiles (right) for a CHAMP occultation on 15 May 2001. Note the different vertical scales. The error bars in the specific humidity plot denote the a priori errors assumed in the retrieval.

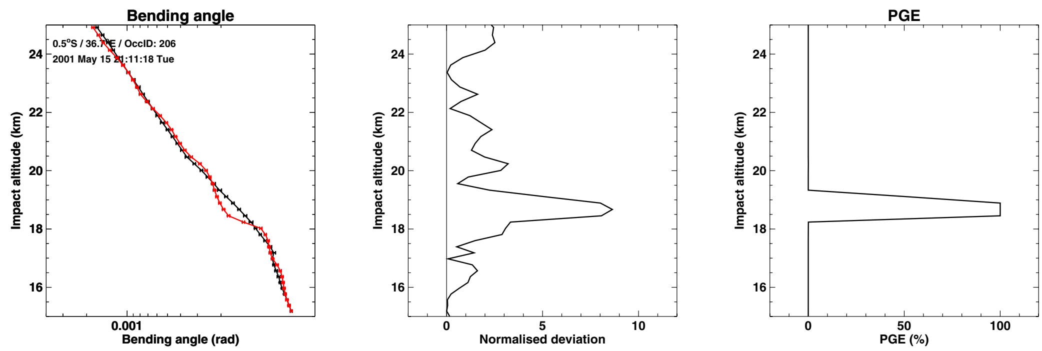

Fig. 15: Observed (red) and forward modelled (black) bending angles, normalised deviation, and PGE (from left to right) for the CHAMP occultation shown in Fig. 14. The error bars in the leftmost figure denote the assumed observation and background errors, the latter forward modelled into bending angle space.

atic. Nevertheless, the values of J_n restricted to tropospheric and stratospheric altitude ranges, being 0.38 and 0.65, respectively, indicate that the tropospheric observational error assumptions might actually be overly pessimistic. A potential problem, however, might exist in the stratosphere. In this case, the value of the cost function helps to narrow down the region where a possible problem in the data exists. We note that the increase of the stratospheric contribution to the cost at convergence is also a generic feature of the CHAMP data (and probably of radio occultation data in general).

Turning to $\mathbf{O} - \mathbf{B}$ statistics (Fig. 15), we find that the largest discrepancies between observations and background occur shortly above the tropopause, where retrieved and forward modelled bending angles differ by several standard deviations. The same information can also be drawn from the normalised deviation between the two (which is,

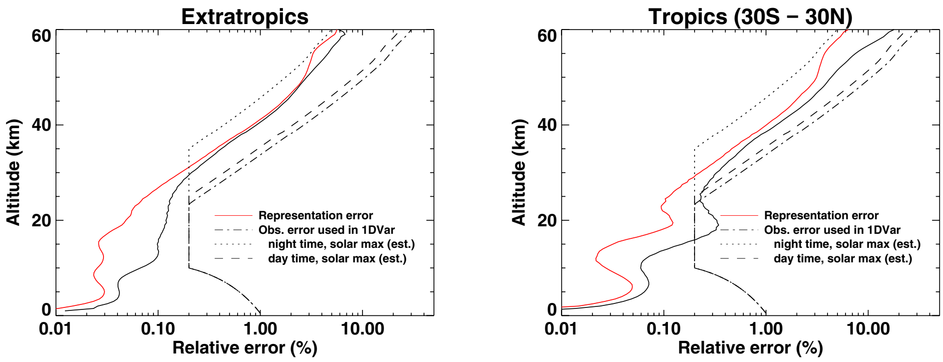

Fig. 16: Estimated relative representation error for 60 (solid) and 90 vertical levels (red) in refractivity. The dotted, dashed and dash-dotted lines denote relative refractivity observation errors, loosely modelled following Kursinski et al. (1997b). The estimate is based on CHAMP data during May and June 2001.

up to a factor of $\sqrt{2}$, identical to the quantity u in Section 2), defined as

$$u^* = \sqrt{\frac{(\alpha_o - \alpha_b)^2}{\sigma_o^2 + \sigma_b^2}} \;,$$

or from the a priori probability of gross error (PGE; Fig. 15). Both the normalised deviation and the PGE are largest in the altitude range shown, and never exceed these values in other altitude regions for this specific profile.

We note that the use of the PGE for an automated level wise quality control in the retrieval (as discussed in Section 2) would have led to the rejection of all data points between 18 and 19.5 km altitude. If the data in question is indeed subject to gross errors, this would be reasonable; if, however, the large PGE is related to a representation problem of the a priori (as will be shown in the following subsection), the levelwise QC would have rejected valid observational data.

We therefore recommend that a levelwise QC based on the probability of gross error is only activated if the error estimates of both observations and a priori are sufficiently well understood and properly tuned. For the validation of GRAS level1b data, and specifically during the commissioning phase, the feature should not be used. However, we do recommend to monitor the diagnostically calculated PGE which may give hints where systematic problems in the error models of the observations or the background occur.

3.8.3 Validation of retrieved parameters

In the above example, it is not clear if the large deviations between CHAMP and ECMWF bending angles in the lower stratosphere are due to gross errors of the observation, or if an issue exists with respect to the representativity of the a priori. Fig. 16 shows an estimate of the representativity error of a 60 level state vector compared to a 90 level version. The figure is based on a refractivity 1D–Var; refractivity errors modelled after Kursinski et al. (1997b) are also shown. In the tropics, the mean representation error in the lower stratosphere (around 20 km altitude) is indeed larger than the observational error. The same problem apparently does not exist in the extratropics. This suggests that atmospheric processes specific to the tropics cannot be represented well in the 60 level version (and therefore not by the 60 level background data). Bending angle variations like the one shown in Fig. 15 may therefore represent true atmospheric structures.

Fig. 17 shows a retrieval for the same occultation, but using a refractivity 1D–Var with 90 vertical levels. Also shown are the temperature sounding from a nearby radiosonde station in Nairobi (in a nominal distance of 30 km, launched about two hours after the occultation took place). The better resolved retrieval shows a wave like structure just above the tropopause which is in excellent agreement with the radiosonde data, suggesting that the deviation between CHAMP and ECMWF bending angles in Fig. 15 does indeed represent a true atmospheric structure. At the same time, the agreement between the ECMWF background and the radiosonde is poor.

Fig 17 also shows a physical dry temperature retrieval. The comparison with the radiosonde is as good as that of the 1D–Var in the lower stratosphere, but larger deviations exist above the 50 hPa level. Because dry temperature retrievals usually do not provide error estimates, they are not as well suited in the validation context. We therefore recommend against their use, as 1D–Var retrievals are likely to be able to provide a comparable vertical resolution, but having the benefit of better diagnostics.

We note that the solution of the problem required the comparison of a retrieval with other auxiliary data than the a priori; we also stress that a sufficiently well resolved 1D–Var had to be developed.

4 Conclusions and recommendations

We have proposed a validation strategy which aims at establishing a traceability chain for GRAS level 1b (and possibly retrieved) data. This is achieved by a stepwise approach, in which the actual and theoretically expected error characteristics are validated for each processing step.

During the pre–launch phase of the GRAS instruments, emphasis is put on deriving a full error characterisation of calibrated and pre–processed GRAS level 1b data, and on testing the implemented software for consistency with the derived error characteristics. Further numerical simulations might be used to assess various aspects of the error characteristics, like the applicability of the proposed statistical methods for the noise estimation in GRAS level 1b data, or the role of horizontal gradients and ionospheric perturbations on the error characteristics.

For the validation during the commissioning and operational phase of the GRAS instrument, the error characteristics of excess phase delays and bending angles should be

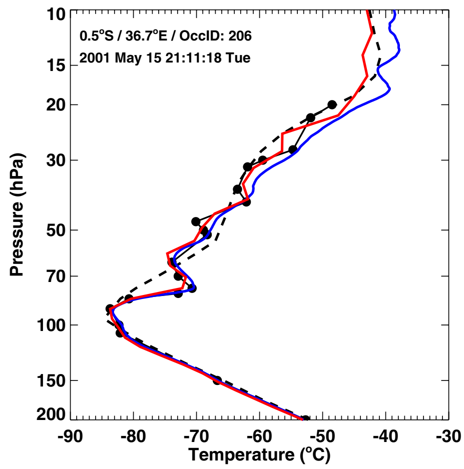

Fig. 17: Experimental 1D–Var 90 level retrieval, based on a low resolution (L60) background. Shown are the retrieved temperature profile (red), background temperature (dashed), data from a radiosonde ascent in Nairobi (solid with black bullets), and a physical dry temperature retrieval (blue). See text for details.

assessed by statistical methods, and compared to the expected error characteristics. If consistency cannot be achieved, the theoretical error models require revision. For monitoring purposes, we recommend that these methods are applied throughout the lifetime of the instrument.

Co–located data from other spaceborne radio occultation missions like CHAMP, SAC-C, COSMIC and EQUARS, while not allowing an estimate of the true error of GRAS level 1b products, will be useful in assessing the instrumental noise characteristics. Such comparisons may act as independent validation of the statistical noise estimates of level1b data. They will be especially important during the commissioning phase, where a rapid understanding of the actual noise characteristics of the GRAS instrument is desirable. We recommend that the raw radio occultation data from other satellite missions are processed by the same calibration and preprocessing system (as far as possible). Additional random–like errors introduced by the differences in the various processing system would thereby be avoided. Once the applicability and reliability of the statistical monitoring algorithms is established, i.e. after the commissioning phase, comparisons with level 1b data from other radio occultation instruments will not be required on a routine basis.

Airborne radio occultation measurements could play a role similar to spaceborne radio occultations, provided that the feasibility of conducting such measurements can be established in time. If this is the case, measurements are likely to take place within the framework of a dedicated measurement campaign. If airborne radio occultations turn out not to be feasible, or if they should have poorer error characteristics then GRAS radio occultations, we recommend against a validation campaign. The main reason is that the number of co–located measurements that can be taken during the short period of time a validation campaign lasts is probably too small to be of significant use, at least for a single satellite radio occultation mission like GRAS.

A central part in the ongoing validation of GRAS level 1b is the application of a bending angle based 1D–Var and its diagnostics. During the commissioning phase, the diagnostic output will be used to improve the observation's error characteristics and derive heuristic thresholds for the later application of the 1D–Var's diagnostic for quality control purposes. The validated error statistics should be made available to data users, e.g.in form of bending angle error covariance matrices. After the commissioning phase, the 1D–Var will mainly act as a monitoring and possibly objective quality control tool for GRAS level 1b data. We have demonstrated the potential usefulness of 1D–Var diagnostics (as well as of the proposed statistical methods for noise estimation) using CHAMP data.

We note that the assimilation of GRAS data into a modern data assimilation system will also provide useful information on the data quality of GRAS observations. For example, the data monitoring usually undertaken at meteorological services will aim at verifying if the assumed error characteristics are met by the actual data. Studying the impact of GRAS data in data assimilation trial runs might hint towards sytematic problems in either the NWP data or the observations. Thus, NWP centres should be involved in the validation of GRAS data as early as possible.

Additional information on the quality of GRAS level 1b data products may be obtained from regular intercomparisons of retrieved atmospheric temperature and humidity profiles with NWP and radiosonde data. We recommend to use data from several NWP analyses, at least during the commissioning phase. Radiosonde data should be restricted to data from a set of reliable radiosonde stations. Both activities should be continued after the commissioning phase for monitoring purposes.

A specific aspect of radio occultation sounding is the high vertical resolution. Validating small scale vertical fluctuations in, e.g., stratospheric temperature profiles may require high resolution data. In the tropics, where such fluctuations are regularly present due to the strong gravity wave activity, such data is not readily available. A possible solution is the organisation of a network of tropical and subtropical radiosonde stations which are willing to launch dedicated radiosondes on demand, i.e. if a GRAS occultation occurs close by. The full radiosonde data should then be provided by the radiosonde sites. It should be kept in mind, though, that the number of matches will still be restricted. The exploitation of high resolution radiosonde data also requires a high resolution retrieval; a 1D–Var based on a standard NWP background (i.e., 60 or less vertical levels) is not able to provide gravity wave information.

With respect to other ground based measurement techniques (like Lidars), the main problem is the small number of co–located measurements that can be achieved in a short amount of time (like a month). Experience with CHAMP data suggests that not more than 4 or 5 monthly matches can be obtained for a single site, if at all. Therefore, such data will not play an important role in the validation of GRAS data during the commissioning phase. For the same reason, we advise against any dedicated validation campaign that mainly relies on ground based measurements. In the long term, however, comparisons of GRAS data with these measurements might be useful for monitoring purposes.

Space based remote sensing systems may provide a significant higher number of co–located measurements than ground based observing systems. The main problem here is that all remote sensing measurements require a retrieval, which includes the use of a priori data. Measurement characteristics also differ between the instruments. Within the variational framework, it is possible to take these issues into account; we have outlined the necessary steps. Apart from being complicated and tedious, the procedure uses the same information (error characteristics of observations and background) as a 1D–Var. The diagnostics of a 1D–Var also provides similar information for each individual profile (i.e., the consistency of the data with the assumed error characteristics, or otherwise). We therefore believe that only little additional information can be gained by intercomparing retrievals of different remote sensing instruments, especially if these are already assimilated into the NWP products which are used as a priori for the 1D-Var.

Ackowledgements

We would like to thank William Bell, Steven English, John Eyre, John Nash, Dave Offiler, and Roger Saunders from the Met Office for discussions and inputs for various aspects of the validation problem. We would also like to thank J. Wickert and T. Schmidt from GFZ Potsdam for providing CHAMP excess path delay data, and Andrew Collard from the Met Office for providing the matrix inversion and decomposition routines for the prototype software.

References

E. Andersson, F. Fisher, R. Munro, and A. NcNally. Diagnosis of background errors for radiances and other observable quantities in a variational data assimilation scheme, and the explanation of a case of poor convergence. *Quart. J. Roy. Meteorol. Soc.*, 126:1455–1472, 2000.

E. Andersson and H. Jarvinen. Variational quality control. ECMWF Technical Memorandum 250, ECMWF, 1998.

G.E. Backus and J.F. Gilbert. Uniqueness in the inversion of inaccurate gross earth data. *Phil. Trans. R. Soc. Lond.*, 266:123, 1970.

P. Craven and G. Wahba. Smoothing noisy data with spline functions. *Numer. Math.*, 31:377–403, 1979.

A. Dudhia and N. Livesey. Validation of temperature measurements from the improved stratospheric and mesospheric sounder. *J. Geophys. Res.*, 101:9795–9809, 1996.

GRAS Calibration and Validation Plan. EUMETSAT, Am Kavalleriesand 31, D-64295 Darmstadt, Germany, 2004.

J.R. Eyre. On systematic errors in satellite sounding products and their climatological mean values. *Quart. J. Roy. Meteorol. Soc.*, 113:279–292, 1987.

E.F. Fishbein, R.E. Cofield, L. Froidevaux, R.F. Jarnot, T. Lungu, W.G. Read, S. Shippony, J.W. Waters, I.S. McDermid, T.J. McGee, U. Singh, M. Gross, A. Hauchecorne, P. Keckhut, M.E. Gelman, and R.M. Nagatani. Validation of UARS Microwave Limb Sounder temperature and pressure measurements. *J. Geophys. Res.*, 101:9983–10016, 1996.

M. Fisher and P. Courtier. Estimating the covariance matrices of analyses and forecast error in variational data assimilation. Technical memorandum 220, Europen Cetre for Medium–Range Weather Forecasts, Reading, 1995.

J.C. Gille, P.L. Baily, S.T. Massie, L.V. Lyjak, D.P. Edwards, A.E. Roche, J.B. Kumer, J.L. Mergenthaler, M.R. Gross, A. Hauchecorne, P. Keckut, T.J. McGee, I.S. McDermid, A.J. Miller, and U. Singh. Accuracy and precision of cryogenic limb array etalon spectrometer (CLAES) temperature retrievals. *J. Geophys. Res.*, 101:9583–9601, 1996.

M.E. Gorbunov. Canonical transform method for processing radio occultation data in the lower troposphere. *Radio Sci.*, 37:10.1029/2000RS002592, 2002a.

M.E. Gorbunov. Radio-holographic analysis and validation of Microlab-1 radio occultation data in the lower troposphere. *J. Geophys. Res.*, 107:10.1029/2001JD000889, 2002b.

M.E. Gorbunov. Radioholographic analysis of radio occultation data in multipath zones. *Radio Sci.*, 37:10.1029/2000RS002577, 2002c.

A. Hauchecorne and M.-L. Chanin. Density and temperature profiles obtained by Lidar between 35 and 70 km. *Geophys. Res. Lett.*, 7:565–568, 1980.

S.B. Healy. Radio occultation bending angle and impact parameter errors caused by horizontal refractive index gradients in the troposphere: A simulation study. *J. Geophys. Res.*, 106:11875–11889, 2001.

S.B. Healy and J. Eyre. A forward model for the assimilation of GPS radio occultation bending angle measurements. *Quart. J. Roy. Meteorol. Soc.*, 2003. submitted.

S.B. Healy and J.R. Eyre. 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.

M. Helten, H.G.J. Smit, W. Sträter, D. Kley, P. Nedelec, M. Zöger, and R. Busen. Calibration and performance of automatic compact instrumentation for the measurement of relative humidity from passenger aircraft. *J. Geophys. Res.*, 103:25643–25652, 1998.

K. Ide, P. Courtier, M. Ghil, and A.C. Lorenc. Unified notation for data assimilation: Operational, sequential and variational. *J. Met. Soc. Jap.*, 75:181–189, 1997.

N.B. Ingleby and A.C. Lorenc. Bayesian quality control using multivariate normal distributions. *Quart. J. Roy. Meteorol. Soc.*, 119:1195–1225, 1993.

A.S. Jensen, M.S. Lohmann, H.-H. Benzon, and A.S. Nielsen. Full Spectrum Inversion of radio occultation signals. *Radio Sci.*, 38:1040, doi:10.1029/2002RS002763, 2003.

E.R. Kursinski, G.A. Hajj, W.I. Bertiger, S.S. Leroy, T.K. Meehan, L.J. Romans, J.T. Schofield, D.J. McCleese, W.G. Melbourne, C.L. Thornton, T.P. Yunck, J.R. Eyre, and R.N. Nagatani. Initial results of radio occultation observations of earth's atmosphere using the Global Positioning System. *Science*, 271:1107–1110, 1996.

E.R. Kursinski, G.A. Hajj, J.T. Schofield, R. Linfield, and K.R. Hardy. Observing earth's atmosphere with radio occultation measurements using gps. *Journal of Geophysical Research*, 102:23429–23465, 1997a.

E.R. Kursinski, G.A. Hajj, J.T. Schofield, R.P. Linfield, and K.R. Hardy. Observing earth's atmosphere with radio accultation measurements using the Global Positioning System. *J. Geophys. Res.*, 102:23.429–23.465, 1997b.

T. Leblanc, I.S. McDermid, A. Hauchecorne, and P. Keckhut. Evaluation and optimization of lidar temperature analysis algorithms using simulated data. *J. Geophys. Res.*, 103:6177–6187, 1998.

A.C. Lorenc. Analysis methods for numerical weather prediction. *Quart. J. Roy. Meteorol. Soc.*, 112:1177–1194, 1986.

A. Marenco, V. Thouret, P. Nedelec, H. Smit, M. Helten, D. Kley, F. Karcher, P. Simon, K. Law, J. Pyle, G. Poschmann, R. von Wrede, C. Hume, and T. Cook. Measurement of ozone and water vapor by Airbus in–service aircraft: The MOZAIC airborne program, An overview. *J. Geophys. Res.*, 103:25631–25642, 1998.

C. Marquardt, J. Eyre, S. Healy, A. Jupp, and D. Offiler. Use of GPS radio occultation data in meteorological services. NWP Forecasting Research Technical Report 403, Met Office, 2003a.

C. Marquardt, J. Eyre, S. Healy, A. Jupp, and D. Offiler. Use of GPS radio occultation data in meteorological services. In *OIST-4 Proceedings*, pages 261–268. Danish Meteorological Institute, Copenhagen, 2003b.

C. Marquardt, K. Labitzke, Ch. Reigber, T. Schmidt, and J. Wickert. An assessment of the quality of GPS/MET radio limb soundings during February 1997. *Phys. Chem. Earth*, 26:125–130, 2001.

C. Marquardt, K. Schöllhammer, G. Beyerle, T. Schmidt, J. Wickert, and Ch. Reigber. Validation and data quality of CHAMP radio occultation data. In *First CHAMP mission results for gravity, magnetic and atmospheric studies*, pages 384–396. Springer, 2003c.

W. Menke. *Geophysical data analysis – discrete inverse theory*. Academic Press, 1989.

P.I. Palmer, J.J. Barnett, J.E. Eyre, and S.B. Healy. A nonlinear optimal estimation inverse method for radio occultation measurements of temperature, humidity, and surface pressure. *J. Geophys. Res.*, 105:17.513–17.526, 2000.

D.F. Parrish and J.C. Derber. The National Meteorological Center's spectral statistical interpolation analysis system. *Mon. Wea. Rev.*, 120:1747–1763, 1992.

W.H Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical recipes in Fortran – The art of scientific computing*. Cambridge University Press, Cambridge, New York, 2nd edition, 1992.

C.M. Reinsch. Smoothing by spline functions. *Numer. Math.*, 10:177–183, 1967.

C.M. Reinsch. Smoothing by spline functions, II. *Numer. Math.*, 16:451–454, 1971.

C. Rocken, R. Anthes, M. Exner, D. Hunt, S. Sokolovsky, R. Ware, M. Gorbunov, W. Schreiner, D. Feng, B. Herman, Y.-H. Kuo, and X. Zou. Analysis and validation of GPS/MET data in the neutral atmosphere. *J. Geophys. Res.*, 102:29.849–29.866, 1997.

C.D. Rodgers. Retrieval of atmospheric temperature and composition from remote sounding measurements of thermal radiation. *Rev. Geophys. Space Phys.*, 14:609–624, 1976.

C.D. Rodgers. Characterization and error analysis of profiles retrieved from remote sounding measurements. *J. Geophys. Res.*, 95:5587–5595, 1990.

C.D. Rodgers. *Inverse methods for atmospheric sounding: Theory and practice*. World Scientific Publishing, Singapore, New Jersey, London, Hong Kong, 2000.

C.D. Rodgers and B.J. Connor. Intercomparison of remote sounding instruments. *J. Geophys. Res.*, 2004. in print.

R. Saunders, M. Matricardi, and P. Brunel. An improved fast radiative transfer model for assimilation of satellite radiance observations. *Quart. J. Roy. Meteorol. Soc.*, 125: 1407–1425, 1999.

H. Savijarvi. Error growth in a large numerical forecast system. *Mon. Wea. Rev.*, 123: 212–221, 1995.

K. Schöllhammer, C. Marquardt, and K. Labitzke. Comparison of three different meteorological datasets (ECMWF, Met Office and NCEP). In *First CHAMP mission results for gravity, magnetic and atmospheric studies*, pages 528–535. Springer, 2003a.

K. Schöllhammer, C. Marquardt, and K. Labitzke. Comparison of three different meteorological datasets (ECMWF, Met Office and NCEP) and CHAMP temperature measurements. In *OIST-4 Proceedings*, pages 255–260. Danish Meteorological Institute, Copenhagen, 2003b.

A. Simmons, M. Hortal, G. Kelly, A. McNally, A. Untch, and S. Uppala. Analyses and forecasts of stratospheric winter polar vortex break-up: September 2002 in the Southern Hemisphere and related events from ECMWF operations and ERA-40. ERA-40 Project report series 5, European Centre for Medium–Range Weather Forecasts, 2003.

A.J. Simmons and A. Hollingsworth. Some aspects of the improvement in skill of numerical weather prediction. *Quart. J. Roy. Meteorol. Soc.*, 128:647–677, 2002.

S. Sokolovskiy and D. Hunt. Statistical optimization approach for GPS/MET data inversions. URSI GPS/MET Workshop, Tucson, Arizona, 1996.

A. Steiner and G. Kirchengast. Gravity wave spectra from GPS/MET occultation observations. *J. Atmos. Ocean. Tech.*, 17:495–503, 2000.

A.K. Steiner, G. Kirchengast, and H.P. Ladreiter. Inversion, error analysis, and validation of GPS/MET occultation data. *Ann. Geophys.*, 17:122–138, 1999.

G.P. Stiller, T. Steck, M. Milz, T. von Clarmann, U. Grabowski, and H. Fischer. Approach to the cross–validation of MIPAS and CHAMP temperature and water vapour profiles. In *First CHAMP mission results for gravity, magnetic and atmospheric studies*, pages 551–556. Springer, 2003.

S. Syndergaard. Retrieval analyses and methodologies in atmospheric limb sounding using the GNSS radio occultation technique. Scientific Report 99–6, Danish Meteorological Institute, Copenhagen, 1999.

A.T. Tarantola and B. Valette. Generalized non-linear inverse problems solved using the least squares criterion. *Reviews of Geophysics and Space Physics*, 20(2):219–232, 1982.

T. Tsuda and K. Hocke. Vertical wave number spectrum of temperature fluctuations in the stratosphere using GPS occultation data. *J. Met. Soc. Jap.*, 80:925–938, 2002.

T. Tsuda, M. Nishida, C. Rocken, and R.H. Ware. A global morphology of gravity wave activity in the stratosphere revealed by the GPS occultation data (GPS/MET). *J. Geophys. Res.*, 105:7257–7273, 2000.

V.V. Vorob'ev and T.G. Krasil'nikova. Estimation of the accuracy of the atmospheric refractive index recovery from doppler shift measurements at frequencies used in the NAVSTAR system. *USSR Phys. Atmos. Ocean, Engl. Transl.*, 29:602–609, 1994.

G. Wahba. Smoothing noisy data with spline functions. *Numer. Math.*, 24:383–393, 1975.

Y. Wang. Smoothing spline models with correlated errors. *J. Am. Stat. Soc.*, 93:341–348, 1998.

C. Zuffada, G.A. Hajj, and E.R. Kursinski. A novel approach to atmospheric profiling with a mountain–based or airborne GPS receiver. *J. Geophys. Res.*, 104:24.435–24.447, 1999.

A Bending Angle 1D–Var S/W Overview and User Guide

A.1 Introduction

The purpose of this Appendix is to present a simple overview and 'user guide' for a one-dimensional variational (1D-Var) retrieval of temperature, humidity and surface pressure from radio occultation (RO) bending angle profiles. It provides an introduction to the numerical methods that have been employed, specifies the data and interface requirements of the main routine and gives an outline of the program structure.

In Section A.2, the variational retrieval technique is outlined, including descriptions of the minimisation method and forward model that have been employed in this implementation. The method used to estimate the probability of gross error is outlined in section A.3. An outline of the program structure as well as the input requirements and user interface of the 1dvar routine are explained Section A.4. Section A.5 includes the instructions for installation and compilation, followed by a description of two simple test programs which are part of the distribution. A Reference Guide to all subroutines is provided in Section A.6.

The bending angle 1D-Var software described here is available from the authors on request.

A.2 The variational retrieval method

In a variational retrieval the most probable state, $\hat{\mathbf{x}}$, is calculated by combining the *a priori* or background information, $\mathbf{x_b}$, and the measurements/observations, $\mathbf{y^o}$, in a statistically optimal way. The approach has been described in detail by many authors (e.g., Rodgers, 1976; Tarantola and Valette, 1982; Lorenc, 1986). The method requires the solution of the forward problem, mapping the state vector information into the measurement/observation space. It is essentially a sophisticated least squares calculation. In simple terms, the solution is found by adjusting the state vector elements (for example, temperatures on fixed pressure levels) in a way that is consistent with the estimated errors in the background information, in order to produce simulated measurements that fit the observations to within their expected errors.

It can be shown that in the case of unbiased Gaussian error distributions, obtaining the most probable state is equivalent to minimizing a cost function given by:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x}-\mathbf{x_b})^{\mathbf{T}}\mathbf{B}^{-\mathbf{1}}(\mathbf{x}-\mathbf{x_b}) + \frac{1}{2}(\mathbf{y^o}-H(\mathbf{x}))^T(\mathbf{E}+\mathbf{F})^{-\mathbf{1}}(\mathbf{y^o}-H(\mathbf{x})) \qquad (47)$$

adopting the notation outlined in Ide et al. (1997), where $\mathbf{B}$ is the expected background error covariance matrix; $H(\mathbf{x})$ is the forward operator, mapping the atmospheric information into measurement (bending angle) space; $\mathbf{E}$ and $\mathbf{F}$ are the expected error covariance matrices of measurements and forward model respectively. The superscripts $\mathbf{T}$ and $-\mathbf{1}$ denote matrix transpose and inverse.

A.2.1 Minimizing the cost function

In this work we map the state vector information to a co-ordinate system where the background errors are uncorrelated, to improve the conditioning of the problem. This is achieved with the transform,

$$\mathbf{v} = \mathbf{B}^{-\mathbf{1/2}}\mathbf{x} \tag{48}$$

The Marquardt-Levenberg approach (Press et al., 1992) is employed to minimise the cost function. In general, the minimum can be found by the iterative solution of

$$\mathbf{J}''.(\mathbf{v}^{n+1} - \mathbf{v}^{n}) = -\mathbf{J}' = \mathbf{0} \tag{49}$$

where $\mathbf{v}^n$ and $\mathbf{v}^{n+1}$ are the *nth* and *(n+1)th* approximation of $\mathbf{v}$. $\mathbf{J}'$ and $\mathbf{J}''$ are the first and second derivatives of the cost function with respect to $\mathbf{v^n}$. These are given by

$$\mathbf{J}' = (\mathbf{v}^n - \mathbf{v_b}) - \mathbf{H}'(\mathbf{v}^n)^{\mathbf{T}}(\mathbf{E}+\mathbf{F})^{-1}(\mathbf{y^o} - H(\mathbf{v}^n)) \tag{50}$$

and (in the linear limit)

$$\mathbf{J}'' = \mathbf{I} + \mathbf{H}'(\mathbf{v}^n)^T(\mathbf{E}+\mathbf{F})^{-1}\mathbf{H}'(\mathbf{v}^n) \tag{51}$$

where $\mathbf{I}$ is the identity matrix. In the Marquardt-Levenberg approach the diagonal values of $\mathbf{J}''$ are modified with $J''_{ii} = J''_{ii}(1+\lambda)$, where λ is a positive scalar value. If $\mathbf{v}^{n+1}$ reduces the penalty function value, λ is reduced for the next iteration. Conversely, if it is found that $\mathbf{v}^{n+1}$ increases the cost function value, λ is increased and the increment is recalculated. This procedure is repeated in until the cost function value falls.

When we have found $\hat{\mathbf{v}}$, the solution that minimizes the cost function, we map back to physical space with

$$\hat{\mathbf{x}} = \mathbf{B}^{\mathbf{1/2}}\hat{\mathbf{v}} \tag{52}$$

A.2.2 The forward model

This maps the atmospheric state vector $\mathbf{x}$ into measurement space. Since the forward model operates on the physical state variables $\mathbf{x}$, but the minimisation uses the control variable $\mathbf{v}$, we have to perform $\mathbf{x} = \mathbf{B}^{\mathbf{1/2}}\mathbf{v}$ before actually calling the forward model. In this problem the measurement vector, $\mathbf{y^o}$, is bending angle as a function of impact parameter and we assume that the state vector is composed of a surface pressure estimate (hPa) together with temperature (K) and (natural) log(specific humidity(g/Kg)) values on a set of fixed pressures.

The bending angle forward operator is composed of three main steps (3 subroutines):

- *Evaluate the geopotential height and refractivity on fixed pressure levels*

The virtual temperature T_i^v on the fixed pressure levels is calculated using the relationship,

$$T_i^v = T_i(1 + 0.608 \times Q_i) \tag{53}$$

where T_i and Q_i are the temperature and specific humidity values respectively on the *i-th* pressure level. The geopotential height of the *j-th (j=i+1)* level, Z_j, is given by

$$Z_j = Z_i + \frac{R(T_j^v + T_i^v)}{2g} \log_e \left(\frac{P_i}{P_j} \right) \tag{54}$$

The refractivity on each fixed pressure level is given by

$$N_j = \frac{c_1 P_j}{T_j} \left(1 + \frac{k Q_j}{T_j} \right) \tag{55}$$

given, $k = c_2/(c_1 \varepsilon)$ where c_1 and c_2 are known constants (Kursinski et al., 1997b) and ε is the ratio of the molecular weights of water-vapour and dry air.

- *Calculate 'nr' on the fixed pressure levels*

  The geopotential heights are mapped to geometric heights using the transformation given by List (1984). The geometric height of the *jth* level is given by,

$$H_j = \frac{Z_j}{\left(\frac{g_{lat}}{g_0} - \frac{Z_j}{R_e} \right)} \tag{56}$$

  where g_{lat} is the value of gravity at Mean Sea Level for the given latitude and R_e 'effective radius of the earth' (see List for a detailed discussion). The radius of the *jth* pressure level is then,

$$r_j = R + H_j + \Delta \tag{57}$$

  where R is the radius of curvature value and Δ[9] is the difference between the geoid and WGS-84 ellipsoid radius value at the location of the measurement (known as the undulation). It is useful to evaluate and store the product of the refractive index and radius values for each fixed pressure level,

$$x_j = (1 + 10^{-6} N_j) r_j \tag{58}$$

  noting that x is the conventional notation for this product, and should not be confused with the 1D–Var state vector.

- *Calculate the bending angle value*

[9] Note that we actually add the undulation to the radius of curvature outside the 1D–Var routine. We then pass the "effective" radius of curvature to the 1D–Var routine.

This requires evaluating integrals of the form,

$$\alpha(a) = -2a \int_a^\infty \frac{d\ln n/_{dx}}{(x^2-a^2)^{1/2}} dx \tag{59}$$

where a denotes the impact parameter. This is simplified using the following approximations:

$$d\ln n/_{dx} \approx 10^{-6} dN/_{dx}$$

which is valid because the refractivity is small, and,

$$(x^2-a^2)^{1/2} \approx \sqrt{2a(x-a)}$$

This is valid because the refractivity scale height is small compared to the radius of the earth. We also assume that N varies exponentially with $x = nr$ between the fixed levels. The gradient between the *jth* and *(j+1)th* levels is then

$$dN/_{dx} = -k_j N_j \exp(-k_j(x-x_j)) \tag{60}$$

The bending between the *jth* and *(j+1)th* levels is then

$$\Delta\alpha_j = 10^{-6} k_j N_j \exp(k_j(x_j-a))\sqrt{2a} \int_{x_j}^{x_{j+1}} \frac{\exp(-k_j(x-a))}{\sqrt{x-a}} dx \tag{61}$$

which can be evaluated in terms of the 'error function'. The total bending angle is found by summing contributions of this form.

A.3 Screening of observations and the probability of gross error (PGE)

The 1D–Var code can have considerable difficultly finding the optimal solution if the observation vector contains gross-errors. We define gross errors as large errors that are not consistent with the assumed combined observation/forward model errors, (***E+F***). Ideally, gross errors should be 'screened out' before the observation vector is presented to the 1D–Var. For example, the data could be rejected if the 'O-B' difference exceeds 10 times the combined O and B errors. The 1D–Var package contains the subroutine `set_pge` which provides an estimate of the probability of gross error (PGE) for each observed bending angle. The PGE is based on the O-B difference, following the approach outlined by Ingleby and Lorenc (1993) and Andersson and Jarvinen (1998). The PGE for the *ith* bending angle is given by,

$$PGE_i = 1 - \frac{1}{\gamma \exp(u) + 1} \tag{62}$$

Subroutines	**Functions**
OneDvar_Solve	Penfunc
SatMatInv	erf
cholesky	svp
eval_derivs	e_rad
alpha_op	g_lat
refrac_levs	
calc_nr	
calc_alpha	
humid_check	
alpha_op_and_K	
refrac_levsK	
calc_nrK	
calc_alphaK	
check_input_alpha	
set_pge	

Tab. 17: Subroutines and functions in the `libalpha_1dvar.a` library.

where,

$$u = \frac{(y_o^i - y_b^i)^2}{2({\sigma_o^i}^2 + {\sigma_b^i}^2)} \tag{63}$$

and γ is set within refrac_info, using equation 11 of Andersson and Jarvinen (1998). We can also define a "QC weighting factor" given by,

$$w_{qc}^i = 1 - PGE_i \tag{64}$$

This factor can be used to reduce the gradient of the cost function with respect to the *ith* bending angle. The qc weighting is implemented if qc_on = .TRUE. in the OneDvar_Solve argument list.

We would strongly advise that qc_on = .FALSE. **for validation applications**.

A.4 Implementation

A.4.1 Code organisation

The 1D–Var source code, that implements the theory outlined in Sections 2 and 3, is written in Fortran 90. After compilation of the package (see section A.5), several subroutines and functions are available in a library `libalpha_1dvar.a` and can be linked against other programs written in Fortran 90. Table 1 lists the user callable routines. The calling tree, beginning at the main user interface `OneDVar_Solve`, is

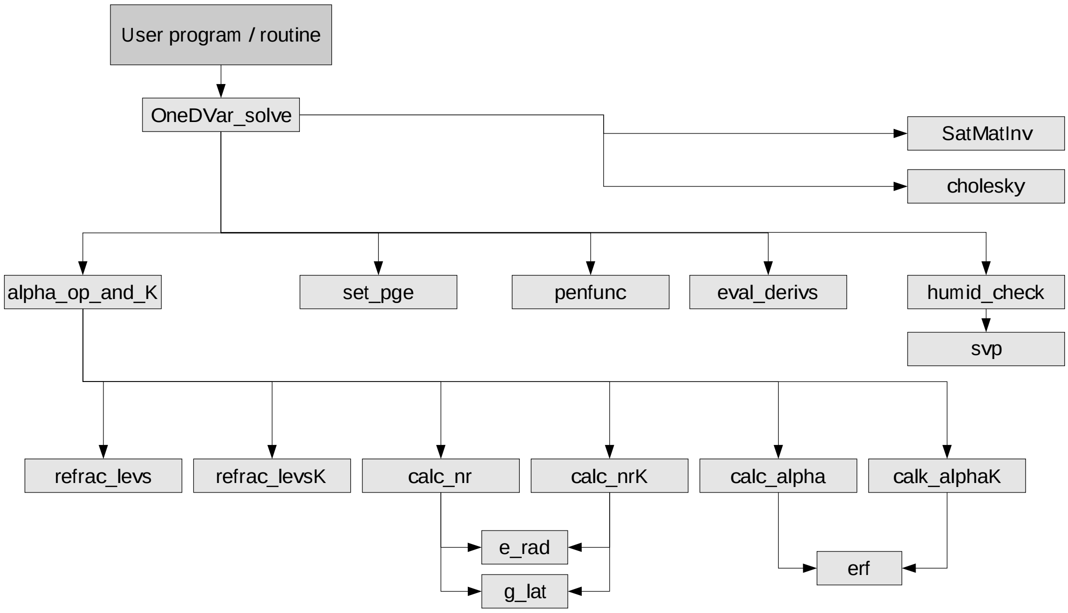

Fig. 18: Calling tree of the 1DVar software.

shown in Fig. 18. A detailed reference guide of all routines and call interfaces can be found in Section A.6.

Internally, we use one module refrac_info.f90, which contains several parameters used by the library. User programs written using the `libalpha_1dvar.a` library may benefit from `USE`'ing it. Details of this module and the derived types contained in it can be found in the reference section. However, users calling routines of the library do not need to use this module, as long as they do not want to change default settings of some variables used by the retrieval routines.

To ensure that input values are within physically reasonable ranges, the library further provides the subroutine `check_input_alpha` which checks the input data for plausibility; it should always be called prior to `OneDVar_Solve`. The package also contains two example programs (`text_champ_data` and `test_sim_data`) that illustrate the use of these routines. Several utility routines that are used for the setup of data are documented in the reference part of this Appendix; they include read routines to the example data (including error covariances, see below) which are also provided as part of the distribution.

A.4.2 The user interface: OneDVar_Solve

Subroutine OneDVar_Solve is the user-callable interface and main routine of the 1D–Var calculation. If the 1D–Var code is to be used as a 'black box', ensuring that the inputs to OneDVar_Solve are of the correct format and type should be sufficient to enable the calculation to be completed successfully. This subroutine is called from the user's application with:

```
call OneDVar_Solve(nstate,     &
                   nlev,       &
                   nwet,       &
                   nobs,       &
                   roc,        &
                   pz0,        &
                   lat,        &
                   press,      &
                   xb,         &
                   Bhalf,      &
                   B_min_half, &
                   a,          &
                   yobs,       &
                   OM1,        &
                   Osigma,     &
                   qc_on,      &
                   it,         &
                   J_pen,      &
                   x,          &
                   pge,        &
                   yb,         &
                   ycalc,      &
                   Amat,       &
                   Kmat,       &
                   converged,  &
                   error)
```

The meaning of each variable is described in Table 2, and the the subroutine argument types, dimensions and intent are shown in Table 3. Additional information can be found in the reference guide section.

A.4.3 The inverted covariance matrices: Bhalf, B_min_half & OM1

In general, the user must provide inverted background and observation error covariance matrices (BM1 and OM1) for subroutine `OneDVar_Solve` to be executed successfully. We have provided a data file `new_atovs_bmat_tot.dat` which contains $\mathbf{B}^{-\mathbf{1}}$ and $\mathbf{B}^{-\mathbf{1}/2}$ on the RTTOVS fixed pressure levels, which has been used in the testing of the 1D–Var. Clearly, $\mathbf{B}^{\mathbf{1}/2}$ can be found by inverting $\mathbf{B}^{-\mathbf{1}/2}$ (with SatMatInv). The example data file can be read with the utility routine `read_binv`. Note also, more accurate error covariance matrices should be obtained from the provider of the *a priori* information whenever possible (ECMWF in this case).

The observation errors are currently estimated in the testing of the 1D–Var code as follows. We assume 10% at the surface falling to 1% at 10km (i.e., impact parameter minus radius of curvature = 10 km). The percentage error is assumed to be constant above 10km, but we use a lower absolute limit of 3μrad[10]. Correlations are not currently modelled, so the matrix is assumed to be diagonal.

nstate	Size of the state vector
nlev	Number of 'dry' pressure levels
nwet	Number of 'wet' pressure levels (nwet$\leq$nlev)
nobs	Number of observations in profile (i.e, number of bending angle values)
roc	Radius of curvature of the earth **+ UNDULATION** (in metres)
pz0	Surface elevation value (in metres) at the observation location, appropriate for surface pressure value, xb(nstate)
lat	Latitude of measurement (in degrees)
press	The vector of fixed pressure levels on which the temperature and humidity profiles are specified. The pressure values are in hPa and are in descending order of value. E.g. pres(1)=1013.25hPa, pres(2)=1005.43hPa and pres(nlev)=0.1hPa (Currently the RT-TOVS pressure levels).
xb	The background or *a priori* state vector. This is composed of nlev temperature values (K) on fixed pressure levels, followed by nwet $\log_e$ (specific humidity(g/Kg)) values and a surface pressure (hPa). E.g., xb(1) = T(at 1013.25hPa), xb(2) = T(at 1005.43hPa), xb(nlev+1) = log(Q(at 1013.15hPa)), xb(nlev+2) = log(Q(at 1005.43hPa)) and xb(nstate) = surface pressure(hPa)
Bhalf	The square-root of the background error covariance matrix = $\mathbf{B^{1/2}}$
B_min_half	The inverse of the square-root of the background error covariance matrix = $\mathbf{B^{-1/2}}$
a	The vector of observed impact parameter values, in ascending order (in metres)
yobs	The observed bending angle values (in radians) on 'a' levels.
OM1	The inverse of the observation and forward model $(\mathbf{E}+\mathbf{F})$ error covariance matrix
Osigma	The standard deviation of the observation errors
qc_on	Logical switch for applying the qc weighting based on the pge value. If qc_on=.FALSE. the qc weighting is not applied.

Tab. 18: The meaning of arguments to `OneDVar_Solve` *(cont'd on next page)*.

[10] In fact, we believe that 6 μrad may be more appropriate for CHAMP data.

it	The number of iterations that were required for convergence. If it=ITMAX then convergence has not been obtained within the maximum number of iterations. By default, ITMAX=20. If this is the case the solution may be questionable, particularly if the cost-at-convergence is high. We have found that the convergence is typically found in 6 iterations, when OM1 and Bhalf are good estimates of the actual errors
J_pen	The penalty function at convergence. This is useful for quality control. J_pen should be of order nobs/2, for a reasonable retrieval
x	The solution vector
pge	The probability of gross error vector
yb	The simulated bending angle values calculated with the first guess profile
ycalc	The simulated bending angle values calculated with the solution profile
Amat	The solution error covariance matrix
Kmat	The gradient of the forward model with respect to the state vector elements
converged	Logical, TRUE if the solution converged within ITMAX
error	Logical, TRUE if there was an error in the 1D–Var retrieval

Tab. 18: The meaning of arguments to `OneDVar_Solve`.

A.5 Installation

A.5.1 The package

The 1dvar ‘alpha_1dvar_package’ is supplied as ’.tar.gz’ or ’.zip’ file, which can be unpacked with the usual `gunzip`, `tar` or `unzip` commands under Linux/Unix, or with the `WinZip` application under Windows. After unpacking, the source code is available in the subdirectory `alpha_1dvar-n.n`, where `n.n` specifies the version of the package. Refer to the files `README` for general information, and the files `README.unix` and `README.cygwin` for detailed information on the installation of the software under Linux/Unix and Windows environments, respectively. In general, however, a sequence of commands like

```
> ./configure
> make
```

should be sufficient. If a Fortran 90 compiler is not immediately found by the configuration system, the name of the compiler can be specified as follows (in a `ksh` or `bash` environment:

Variable	Type	Size	Usage
nstate	Integer	Scalar	In
nlev	Integer	Scalar	In
nwet	Integer	Scalar	In
nobs	Integer	Scalar	In
roc	Real	Scalar	In
pz0	Real	Scalar	In
lat	Real	Scalar	In
pres	Real	Array(nlev)	In
xb	Real	Array(nstate)	In
Bhalf	Real	Array(nstate,nstate)	In
B_min_half	Real	Array(nstate,nstate)	In
a	Real	Array(nobs)	In
yobs	Real	Array(nobs)	In
OM1	Real	Array(nobs,nobs)	In
Osigma	Real	Array(nobs)	In
qc_on	Logical	Scalar	In
it	Integer	Scalar	Out
J_pen	Real	Scalar	Out
x	Real	Scalar	Out
pge	Real	Array(nobs)	Out
yb	Real	Array(nobs)	Out
ycalc	Real	Array(nobs)	Out
amat	Real	Array(nstate,nstate)	Out
Kmat	Real	Array(nobs, nstate)	Out
Converged	Logical	Scalar	Out
error	Logical	Scalar	Out

Tab. 18: Type and intent of arguments to the subroutine `OneDVar_Solve`.

```
> FC=<compiler> ./configure
> make
```

For example, the Intel V7.x compilers under Linux might require

```
> FC=ifc ./configure
> make
```

Additional options to the compiler can be specified by means of the `FCFLAGS` environment variable. We note that we have used the \ character within a strings which are used for describing the usage of the two example programs (in the files `test_sim_data.f90` and `test_champ_data.f90`. Some compilers are able to treat \ as an escape character (or do that by default), but may issue a warning that such escape characters are an extension of standard Fortran 90. These warning messages can safely be ignored.

A.5.2 A program using simulated data

After successful compilation, a simple test program that runs the 1D–Var retrieval with simulated data is available in the `drivers` subdirectory. The main steps of the program `test_sim_data` are:

- Read the B matrix information from the file `new_atovs_bmat_tot.dat`.
- Read the background profile information (on RTTOVS levels) from the file `Background_MidLatWin_Corr.dat`.
- Open each 'true' file e.g, `Profile_MidLatWin_001.dat`.
- Calculate the 'true' bending angles that would be produced with the true state, `Profile_MidLatWin_001.dat`.
- Set up the inverse observation error covariance matrix OM1.
- Add Gaussian random noise to the true bending angles to obtain simulated observation.
- Check the input prior to calling 1D–Var with check_input_alpha.
- Solve the 1D–Var problem using `OneDVar_Solve`.
- Output the results to the file `sim_out.dat`.

To run the sample program from within the `drivers` subdirectory, type

```
> test_sim_data -d ../data
```

If all goes well, this will produce an output file `sim_out.dat` which should be compared with the file `MO_sim_out.dat` in the `data` subdirectory of the distribution. The file `sim_out.dat` contains:

- The 'true' filename;
- it, the number of iterations for convergence;
- J_pen, the 'cost at convergence';
- Followed by the 1D–Var solution vector, $\hat{\mathbf{x}}$ and $\mathbf{x_b}$ on the fixed pressure levels. The output is of the form

```
TEMPERATURE (K)
   P(hPa) soln.  back
1 1013.25 272.05 272.07
2 1005.43 271.84 271.85
 ...
42 0.29 257.84 257.55
43 0.10 241.66 241.63

HUMIDITY (g/kg)
44 1013.2 3.2258 2.6618
45 1005.4 3.1034 2.6248
46  985.9 2.8672 2.5336
 ...
68  143.8 0.0029 0.0030
69  122.0 0.0029 0.0029

SURFACE PRESSURE
70 1014.3 1014.3
```

The state vector information is then followed by the bending angle values. Each line contains:

- An integer (the position in array).
- The observed impact parameter.
- The observed bending angle value.
- The bending angle calculated with the 1D–Var solution, $\hat{\mathbf{x}}$.
- The bending angle calculated with the background vector, $\mathbf{x_b}$.

Hence, in sim_out contains output of the form:

```
imp param observed solution background
1 6373500.0 0.23145E-01 0.22406E-01 0.22545E-01
2 6374000.0 0.19336E-01 0.20486E-01 0.20605E-01
3 6374500.0 0.23431E-01 0.21305E-01 0.18955E-01
4 6375000.0 0.19047E-01 0.18453E-01 0.17349E-01
5 6375500.0 0.15044E-01 0.15448E-01 0.15831E-01
```

At the end of the file we output the average (over the 10 simulated retrievals) cost function value and number of iterations required for convergence.

```
AVERAGE J_PEN = 59.39983
AVERAGE IT    = 3.300000
```

Ideally, the average J_pen should be 60, because we use 120 simulated bending angle measurements.

Note that additional information on the command line arguments of `test_sim_data` can be obtained from the reference section. When comparing the results obtained with the example results in the file `data/MO_sim_out.dat`, it should also be kept in mind that small numerical deviations (usually limited to the last digit in each floating point number) may occur because of compiler differences or round off errors on different platforms.

A.5.3 A program using a CHAMP measurement

After successful compilation, the `drivers` subdirectory also contains a simple test program (`test_champ_data`) that runs the 1D–Var retrieval with a CHAMP measurement. The main steps of the programs are:

- Read the B matrix information from the file `new_atovs_bmat_tot.dat`.
- Read the RTTOVS fixed pressure levels (from the file `Background_MidLatWin_Corr.dat`).
- Read the observed CHAMP measurement using the subroutine `read_ob`.
- Read the co-located ECMWF data using the subroutine `read_bg`.
- Interpolate the ECMWF data to fixed pressure levels and set up $\mathbf{x_b}$.
- Set the inverse of the observation error covariance matrix, `OM1`.
- Solve the 1D–Var problem using `OneDVar_Solve`.
- Output the results to `1dvar_out.dat`.

The program can be run in the `drivers` subdirectory by typing

```
> test_champ_data -d ../data <obs_file> <bg_file>
```

where <`obs_file`> and <`bg_file`> are the names of a bending angle and a background data file. We have provided one example data file of each in the `data` subdirectory of the source code distribution; the files `ob_data.dat` and `bg_data.dat` in the same subdirectory are copies of these to ease the task of typing. Thus, to test the installation, it should be sufficient to type

```
> test_champ_data -d ../data ../data/ob_file.dat \
>                            ../data/bg_file.dat
```

which will produce an output file `1dvar_out.dat`. The result file should be compared with the file `MO_champ_out.dat` in the `data` subdirectory of the distribution. The general format of the output is similar to the one used in `sim_out.dat`, as described above. Information on other command line arguments of this program can be found in the reference section.

As for the simulated data, we note that small numerical differences, typically limited to the last printed digit of each floating point number, may occur due to compiler differences or round off errors on different platforms.

A.6 Reference guide to all subroutines

This section has been generated from the source code, and included in the LaTeX–source of this Appendix. To regenerate the text included here, do a `make refsec.tex` in the `doc` subdirectory of the source code distribution. Note that a standalone version of the reference guide of the library (`libalpha_1dvar.pdf`) is also available in the same subdirectory. An experimental html version of the reference documentation is provided in the `doc/html` subdirectory of the distribution; point your browser to the file `index.html` to view it. The standalone reference documentation (both as `pdf` and as `html`) can be updated by running `make doc` in the `doc` subdirectory of the documentation. Note that recreating the documentation requires properly installed versions of `RoboDoc` and `pdfLaTeX`

A.7 Examples/test_champ_data

NAME

```
test_champ_data - Test bending angle 1DVar using radio occultation data
                  from the German CHAMP satellite.
```

SYNOPSIS

```
test_champ_data [-d <data_dir>] [-c <bg_corr_file>]     \
                [-l <bg_levels>] [-k <bg_ml_coeff>]     \
                [-o <output_file>] <obs_file> <bg_file>
```

DESCRIPTION

```
This program tests the bending angle 1DVar solver for CHAMP bending angle
profiles. The main steps of the program are:

   o Read the B matrix from the file new_atovs_bmat_tot.dat
   o Read the vertical level structure (for RTTOVS levels) from the file
     Background_MidLat_Win_Corr.dat
   o Read the observed CHAMP bending angle profile
```

- o Set up a (simple) error covariance matrix for the CHAMP data
- o Read co-located ECMWF background data
- o Interpolate ECMWF data to RTTOVS pressure levels and set up the state vector
- o Solve the 1DVar problem
- o Write results

ARGUMENTS

The two command line arguments are mandatory:

<obs_file>	file w/ observed bending angle data (including path).
<bg_file>	file w/ background temperature, humidity and surface pressure data (including path).

OPTIONS

All options are optional, i.e. default values (as used above) are provided.

-h	give some help.
-p	apply PGE based automated quality control.
-d <data_dir>	directory where system files are stored.
-c <bg_corr_file>	file w/ background error correlations.
-k <bg_ml_levels>	file w/ the background’s model level coefficients.
-l <p_levels>	file w/ pressure levels of the retrieval.
-o <output_file>	output file (including path).

The -d <data_dir> options defines the path for all ’system’ data files, i.e. to the filenames that can be changed by means of the -c, -l and -k options. The default is path the current directory. Note, however, that the installation procedure will install examples of such data files in the directory $(prefix)/share/alpha_1DVar.

Input (<obs_file> and <bg_file>) as well as the output file can contain a full path specification if they do not reside in the current directory.

By default, the PGE based automated QC of the 1DVar is not enabled.

OUTPUT

The file 1dvar_out.dat is generated in the local directory; it contains results of the retrieval run. For details of the format, see the documentation of output_res.

USES

```
The following routines from the libalpha_1DVar library are used:
   satmatinv
   set_xb
   check_input_alpha
   OneDVar_solve

Additional routines:
   read_bg
   read_ob
   read_binv
   nread_tovs_data
   output_res
```

NOTES

```
The program assumes that the background profiles are given on ECMWF's
hybrid vertical levels; the ak and bk level coefficients that allow
the calculation of the model level's actual pressure are expected in
a different file (which can be specified via the -k option). Both the
example background data files as well as the model level coefficient
file are thos of the currently (early 2004) operational L60 version
of the ECMWF system. The program then interpolates the profile given
on model levels onto the pressure levels defined in the level
structure data file as specified via the -l option. The example data
file replicates the RTTOVS set of fixed pressure levels.
```

SEE ALSO

```
test_sim_data
```

A.8 Examples/test_sim_data

NAME

```
test_sim_data - Test bending angle 1DVar using simulated radio occultation
                bending angle data.
```

SYNOPSIS

```
test_sim_data [-d <data_dir>] [-b <bg_file>] [-c <bg_corr_file>] \
              [-o <output_file>] [<list_file>]
```

DESCRIPTION

```
This program tests the bending angle 1DVar solver for multiple simulated
bending angle profiles. The main steps of the program are:

   o Read the B matrix from the file new_atovs_bmat_tot.dat
```

```
o Read a background profile (on RTTOVS levels) from the file
  Background_MidLat_Win_Corr.dat
o Loop over several 'true' atmospheric profiles, where the names
  of the test cases are contained in the file true_files.lst, and
   - Read each 'true' profile
   - Calculate the 'true' bending angle profile
   - Set up the inverse observation error covariance for
     uncorrelated errors
   - Add uncorrelated random gaussian noise to the 'true' bending
     angles to simulate the 'measurement'
   - Check the simulated 'measurement' for reasonable data ranges
   - Solve the 1DVar problem
   - Write results to an output file
o Add diagnostic information on the mean value of the cost function
  and the average number of required iterations.
```

ARGUMENTS

```
<list_file>         Name of file containing a list of sample profiles
                       used for simulating bending angle measurements.
```

OPTIONS

```
All options are optional, i.e. default values (as used above) are
provided by the program.

   -d <data_dir>       Data directory.
   -b <bg_file>        Background profile (including 1DVar level structure).
   -c <bg_corr_file>   Background error covariance (on 1DVar levels).
   -o <output_file>    Output file name.

The -d <data_dir> options defines the path for all input data files;
the location of the output data file still needs to be specified
explicitely.
```

OUTPUT

```
The file sim_out.dat is generated in the local directory; it contains
results of all retrievals run.
```

USES

```
The following routines from the libalpha_1DVar library are used:
   satmatinv
   alpha_op
   check_input_alpha
   OneDVar_solve

Additional routines:
   read_binv
   nread_tovs_data
   gasdev
```

NOTES

```
The list of true profiles contained in the source code distribution
points to ten example profiles contained in files named like
Profile_MidLatWin_<nnn>.dat, where <nnn> refers to the example
number.
```

SEE ALSO

```
test_champ_data
```

A.9 Examples/Tools

DESCRIPTION

```
As part of the drivers subdirectory, a number of useful subroutines
including some to read the sample data in the data subdirectory of the
alpha_1DVar library are provided. They are not required to use the
actual bending angle 1DVar, but may serve as example on how to
prepare data for the use with the 1DVar.
```

SEE ALSO

```
Data input and output:
   read_bg           Read background data given on ECMWF model levels.
   read_binv         Read the inverse background covariance matrix.
   read_ob           Read bending angle observations.
   nread_tovs_data   Read RTTOV vertical level structure.
   output_res        Write retrieval results to a data file.

Random number generator:
   ran1              Minimal randon number generator.
   gasdev            Normally distributed random numbers.

Other:
   nag_interfaces    Interfaces to the NAG f90/f95 f90_unix_* routines.
```

A.9.1 Tools/gasdev

NAME

```
gasdev - Normally distributed random numbers.
```

SYNOPSIS

```
value = gasdev(idum)
```

DESCRIPTION

```
This function, repetetively called, returns a series of normally
distributed pseudo random numbers with mean zero and unit variance.
```

INPUTS

```
idum    Seed; to initialise, call with a negative value; afterwards.
           don't change (see Press at al. for details).
```

OUTPUT

```
value   Pseudo random number.
```

REFERENCES

```
This function was copied manually from

   Press, W., S. Teukolsky, W. vetterling and B. Flannery, Numerical
      Recipes in Fortran, 2nd Ed., Cambridge University Press, 1988.
```

A.9.2 Tools/nag_interfaces

NAME

```
nag_interfaces - Interfaces to the NAG f90/f95 f90_unix_* routines.
```

SYNOPSIS

```
When Fortran 90 routines call standard Unix system calls but are
compiled with the NAGWare Fortran 90/95 compiler, link the resulting
object code with a compiled version of this file.
```

DESCRIPTION

```
The NAG f90 / f95 compiler provides access to most standard Unix
system routines via a set of modules named f90_unix_*. While this
is a clear and well defined interface to system routines, it leaves
source code developed for the NAG compilers incompatible with Other
Fortran compilers (which usually provide access to the same routines
as part of their standard library), as the corresponding NAG modules
are missing in Other compiler distributions. This causes a lot of
unnecessary recoding work.

As an alternative, this file collects (part of) the system routines
addressed by the various f90_unix_* modules and provides an interface
to them. Instead of adding the appropriate use f90_unix_* entries
in each affected source file, this file can be compiled with the NAG
compiler and linked to the Other routines.
```

NOTES

```
These interfaces are only included in the libalpha_1DVar library if
the NAG Fortran 90/95 compiler is used; it is ignored otherwise.

Not all routines are currently implemented; I will add them if they
will be requested (or I have a need for them).

Some of the NAG routines do have optional arguments, like ALARM.
Obviously, this is not implemented. Also, all parameter and type
definitions available from the NAG modules have been omitted.
```

SEE ALSO

```
For the NAG compiler:
   f90_unix
   f90_unix_dir
   f90_unix_dirent
   f90_unix_env
   f90_unix_file
   f90_unix_proc
```

A.9.3 Tools/nread_tovs_data

NAME

```
nread_tovs_data - Read RTTOV vertical level structure and a vertical
                     background profile from a data file.
```

SYNOPSIS

```
call nread_tovs_data(file, nstate, nlev, nwet, lat, pz0, press, xb)
```

DESCRIPTION

```
This subroutine reads the RTTOV vertical level structure and a
vertical background profile from a data file.
```

INPUTS

```
file    Name of data file.
nstate  Number of elements in the state vector.
nlev    Number of levels for temperature.
nwet    Number of levels for humidity.
```

OUTPUT

```
lat     Latitude (in degree).
pz0     Surface elevation (in m).
press   Pressure of RTTOV levels (in hPa).
xb      Example state vector.
```

NOTES

The routine is intended to be used as read routine for the file Background_MidLatWin_Corr.dat which is part of the alpha_1DVar package. This file contains a background profile on the standard set of RTTOV fixed pressure levels. Contents of this file may be used for simulation purposes (as in the test_sim_data example program which is part of the alpha_1DVar package), or simply as a way to define the RTTOV pressure levels (as in the test_champ_data example program which is part of the alpha_1DVar package).

A.9.4 Tools/output_res

NAME

output_res - Write retrieval results to a data file.

SYNOPSIS

```
call output_res(outfile, nstate, nlev, nwet, &
                it, j_pen, press, x, xb, ob)
```

DESCRIPTION

This subroutine writes the result of a 1DVar retrieval based on bending angles to an ASCII data file.

INPUTS

```
outfile  Name of the output data file.
nstate   Number of elements in the state vector.
nlev     Number of levels for temperature.
nwet     Number of levels for humidity.
it       Number of iterations required for convergence.
J_pen    Value of cost / penalty function at convergence.
press    Pressure levels of state vector.
x        Solution state vector.
xb       Background state vector.
ob       Observation structure, as defined in refrac_info.
```

USES

refrac_info

NOTES

For details of the format implemented, see the source code of this routine.

A.9.5 Tools/ran1

NAME

```
ran1 - Minimal randon number generator.
```

SYNOPSIS

```
value = ran1(idum)
```

DESCRIPTION

```
This function generates a uniform series of pseudo random numbers.
```

INPUTS

```
idum    Seed; to initialise, call with a negative value; afterwards.
           don’t change (see Press at al. for details).
```

OUTPUT

```
value   Pseudo random number.
```

REFERENCES

```
This function was copied manually from

   Press, W., S. Teukolsky, W. vetterling and B. Flannery, Numerical
      Recipes in Fortran, 2nd Ed., Cambridge University Press, 1988.
```

A.9.6 Tools/read_bg

NAME

```
read_bg - Read background data given on ECMWF model levels.
```

SYNOPSIS

```
call read_bg(bg_file, bk_file, bg)
```

DESCRIPTION

```
This subroutine reads background data given on ECMWF model levels,
calculates pressure on the full levels, and copies everything into
the background structure. The information on the level coefficients
is read from the second data file.
```

INPUTS

```
bg_file      Name of the file containing the actual background data
                on ECMWF model levels.
bk_file      Name of the file containing the ECMWF model level
                definition coefficients.
```

OUTPUT

```
bg           Background data structure, as defined in refrac_info.
```

USES

```
refrac_info
```

NOTES

```
In the example programs provided in the alpha_1DVar package, the
model level background information is interpolated onto a set of
fixed pressure levels; this calculation is done in the routine
set_xb.

The routine is intended to be used as read routine for one of the
example background data files which are part of the alpha_1DVar
package. Other background data may be used as well, provided they
are in the same ASCII format of the above mentioned files. For
details of the format, see the source code of this routine.
```

SEE ALSO

```
set_xb
```

A.9.7 Tools/read_binv

NAME

```
read_binv - Read the inverse background covariance matrix.
```

SYNOPSIS

```
call read_binv(file, nstate, BM1, BMHALF, Bdiag)
```

DESCRIPTION

```
This subroutine reads the inverse of a background error covariance
matrix, along with the inverse of (a) symmetric square root of that
matrix, and the diagonal elements of the non-inverted matrix.
```

INPUTS

```
file      File name
nstate    Number of elements in the state vector.
```

OUTPUT

```
BM1       Inverse of the background covariance matrix.
BMHALF    Inverse of a square root of the background error covariance
             matrix.
Bdiag     Diagonal elements of the background error covariance matrix.
```

NOTES

```
The routine is intended to be used as read routine for the file
new_atovs_bmat_tot.dat which is part of the alpha_1DVar package.
Other covariance matrices may be used as well, provided they are
in the same ASCII format of the above mentioned file. For details
of the format, see the source code of this routine.
```

A.9.8 Tools/read_ob

NAME

```
read_ob - Read bending angle observations.
```

SYNOPSIS

```
call read_ob(ob_file, ob)
```

DESCRIPTION

```
This subroutine reads observed bending angle data and copies it into
the bending angle observation structure.
```

INPUTS

```
ob_file   Name of the bending angle observation data file.
```

OUTPUT

```
ob        Bending angle observation structure, as defined in
             the module refrac_info.
```

USES

```
refrac_info
```

NOTES

```
The routine is intended to be used as read routine for one of the
example bending angle data files which are part of the alpha_1DVar
package. Other observations may be used as well, provided they are
in the same ASCII format of the above mentioned files. For details
of the format, see the source code of this routine.
```

A.10 libalpha_1dvar/1DVar

DESCRIPTION

```
These functions form the main part of the bending angle 1DVar in that
they actually calculate and minimise the cost / penalty function.
```

SEE ALSO

```
eval_derivs     Evaluate 1st and 2nd derivatives of the penalty function.
OneDVar_solve   Solve the bending angle 1DVar problem.
penfunc         Evaluate the penalty function.
```

A.10.1 1DVar/eval_derivs

NAME

```
eval_derivs - Evaluate 1st and 2nd derivatives of the penalty function.
```

SYNOPSIS

```
call eval_derivs(nstate, nobs, v, vb, yobs, ycalc, OM1, qcwt, Kmat, &
                 dJ_dv, d2J_dv2, diag_d2J)
```

DESCRIPTION

```
This subroutine calculates the first and second derivatives of the
cost/penalty function with respect to the state vector.
```

INPUTS

```
nstate      Number ofelements in the state vector.
nobs        Number of observations.
v           Current estimate of solution.
vb          Background vector.
yobs        Observation vector.
ycalc       y(x)
OM1         Inverse of observation + forward model error covariance matrix.
qcwt        QC weighting.
Kmat        Gradient matrix.
```

OUTPUT

```
dJ_dv       Negative(!) of first derivative of the cost function.
d2J_dv2     Second derivative of the cost function.
diag_d2J    Vector containing the diagonal values of the matrix above.
```

USES

```
refrac_info
```

A.10.2 1DVar/OneDVar_solve

NAME

```
OneDVar_solve - Solve the bending angle 1DVar problem.
```

SYNOPSIS

```
call OneDVar_solve(nstate, nlev, nwet, nobs,                    &
                   roc, pz0, lat, press, xb, Bhalf, B_min_half, &
                   a, yobs, OM1, Osigma, qc_on,                 &
                   it, J_pen, x, pge, yb, ycalc, Amat, Kmat,    &
                   converged, error)
```

DESCRIPTION

```
This subroutine is the main user interface to the bending angle 1DVar.
```

INPUTS

```
nstate      Number of elements in the state vector.
nlev        Number of levels for temperature.
nwet        Number of levels for humidity.
nobs        Number of observations.
roc         Radius of curvature + geoid undulation (im m).
pz0         Surface elevation abobe Mean Sea Level (in m).
lat         Latitude (in degrees).
press       Pressure levels for the state vector (in hPa).
xb          Background state vector.
Bhalf       The square root of the background error covariance matrix.
B_min_half  The inverse of Bhalf.
a           Impact parameters for the bending angle observations (in m).
yobs        Bending angle observations (in rad).
OM1         The inverse of the observation and forward model (E + F)
               covariance matrix.
Osigna      The standard deviations of the observation errors (i.e., the
               square roots of the diagonal elements of (E + F)).
qc_on       If .false., PGE based QC will not be applied; otherwise
               if .true.
```

OUTPUT

```
it          Number or iterations required for convergence.
J_pen       Penalty / cost function value at convergence.
x           Solution state vector.
pge         Probability of Gross Error vector.
yb          Simulated bending angle calculated from the background state.
ycalc       Simulated bending angle calculated from the solution state.
Amat        Error covariance matrix of the solution.
```

```
Kmat        Weighting functions / Gradient of the forward model with
               respect to the state vector.
converged   .True. if solution converged within ITMAX iterations (as
               defined in refrac_info).
error       .True. if an error occured in the 1DVar retrieval.
```

USES

```
refrac_info
```

A.10.3 1DVar/penfunc

NAME

```
penfunc - Evaluate the penalty function.
```

SYNOPSIS

```
J_pen = penfunc(qc_on, nstate, nobs, &
                v, vb, yobs, ycalc, qcwt, OM1)
```

DESCRIPTION

```
This function evaluates the value of the cost/penalty function.
```

INPUTS

```
qc_on       QC check on or noff.
Nstate      Number of elements of the state vector.
Nobs        Number of elements of the observation vector.
v           Current estimate of the solution.
vb          Background.
yobs        Observed values.
ycalc       y(x).
qcwt        Weight from QC.
OM1         Inverse of observation and forward model covariance matrix.
```

OUTPUT

```
J_pen       Value of the cost / penalty function.
```

USES

```
refrac_info
```

A.11 libalpha_1dvar/Diagnostics

DESCRIPTION

```
The alpha_1DVar library contains two diagnostic routines.
```

SEE ALSO

```
check_input_alpha   Check input data for plausibility.
set_pge             Calculate the Probability of Gross Error (PGE).
```

A.11.1 Diagnostics/check_input_alpha

NAME

```
check_input_alpha - Check input data data for plausibility.
```

SYNOPSIS

```
call check_input_alpha(nstate, nlev, nwet, nobs, &
                       roc, pz0, xb, a, yobs, input_error)
```

DESCRIPTION

```
This subroutine checks both the bending angle values and background
data for physical plausibility.
```

INPUTS

```
nstate        Number of elements in the state vector.
nlev          Number of pressure levels for temperature.
nwet          Number of pressure levels for humidity.
nobs          Number of observations.
roc           Local radius of curvature (in m).
pz0           Surface elevation (in m).
xb            State vector
a             Impact parameters (in m).
yobs          bending angle observations (in rad).
```

OUTPUT

```
input_error   .True. if the data appears to be implausible or
                 inconsistent.
```

USES

```
refrac_info
```

A.11.2 Diagnostics/set_pge

NAME

```
set_pge - Set Probability of Gross Error (PGE).
```

SYNOPSIS

```
call set_pge(nobs, nstate, yobs, ycalc, Osigma, kmat, pge, qcwt)
```

DESCRIPTION

```
This subroutine calculates the a priori Probability of Gross Error
for a bending angle profile as described in Healy and Marquardt (2004).
```

INPUTS

```
nobs               Number of elements in the observation vector.
nstate             Number of elements in the state vector.
yobs(nobs)         Array containing observations
ycalc(nobs)        Forward modelled observations.
Osigma(nobs)       Error in observations.
kmat(nobs, nstate) Weighting functions (i.e., linearised forward model).
```

OUTPUT

```
pge    Probability of Gross Error (0...1).
qcwt   QC weighting factor, equals 1 - pge.
```

USES

```
refrac_info
```

NOTES

```
It is assumed that the weighting functions are given in the control
variable where the problem has been preconditioned with the square
root of the background covariance matrix, i.e. that the transformed B
matrix is 1.
```

A.12 libalpha_1dvar/Forward_model

DESCRIPTION

```
These routines form the forward model that calculates bending angles
as function of impact parameter from a given state vectore, as well as
the gradient of the forward model with respect to the state vector
elements.
```

SEE ALSO

```
alpha_op          Forward operator for bending angles.
alpha_op_and_K    Forward operator and its gradient for bending angles.
calc_alpha        Calculate bending angles for observation impact
                    parameters.
calc_alphaK       Gradient of calc_alpha.
calc_nr           Calculate the 'refractivity x radius' product on
                    geopotential levels.
calc_nrK          Gradient of calc_nr.
refrac_levs       Calculate refractivity and geopotential heights.
refrac_levsK      Gradient of the refrac_levs.
```

A.12.1 Forward_model/alpha_op

NAME

```
alpha_op - Forward operator for bending angles.
```

SYNOPSIS

```
call alpha_op(nstate, nlev, nwet, nobs, roc, pz0, lat, pres, x, a, alpha)
```

DESCRIPTION

```
This subroutine provides a forward operator to calculate bending angles
on a given set of observation impact parameters.
```

INPUTS

```
nstate   Number of elements in the state vector.
nlev     Number of levels for temperature.
nwet     Number of levels for humidity.
nobs     Number of observations.
roc      Radius of curvature (in m).
pz0      Surface elevation above Mean Sea Level (in m).
lat      Latitude.
press    Pressure levels for state vector.
x        State vector.
a        Impact parameter of observations (in m).
```

OUTPUT

```
alpha    Forward modelled bending angles on observation heights (in rad).
```

USES

```
refrac_levs
calc_nr
calc_alpha
```

A.12.2 Forward_model/alpha_op_and_K

NAME

```
alpha_op_and_K - Forward operator and its gradient for bending angles.
```

SYNOPSIS

```
call alpha_op_and_K(nstate, nlev, nwet, nobs, roc, pz0, lat, press, x, a, &
                    alpha, Kmat, ErrorCode)
```

DESCRIPTION

```
This subroutine provides a forward operator to calculate bending angles
on a given set of observation impact parameters. The gradient of the
operator with respect to the state vector is also calculated.
```

INPUTS

```
nstate   Number of elements in the state vector.
nlev     Number of levels for temperature.
nwet     Number of levels for humidity.
nobs     Number of observations.
roc      Radius of curvature (in m).
pz0      Surface elevation above Mean Sea Level (in m).
lat      Latitude.
press    Pressure levels for state vector.
x        State vector.
a        Impact parameter of observations (in m).
```

OUTPUT

```
alpha    Forward modelled bending angles on observation heights (in rad).
Kmat     Gradient of the forward operator.
```

USES

```
refrac_levs
refrac_levsK
calc_nr
calc_nrK
calc_alpha
calc_alphaK
```

A.12.3 Forward_model/calc_alpha

NAME

```
calc_alpha - Calculate bending angles for observation impact parameters.
```

SYNOPSIS

```
call calc_alpha(nobs, nlev, a, refrac, nr, alpha)
```

DESCRIPTION

```
This subroutine calculates bending angles from forward modelled refractivity
and impact parameter ('refracitivity x radius' product) and interpolates
them onto a set of given impact parameters.
```

INPUTS

```
nobs     Number of observations.
nlev     Number of levels in the nr and refrac forward modelled profiles.
a        Impact parameter of the bending angle observations (in m).
refrac   Refracivity
nr       'Refractivity x radius' product.
```

OUTPUT

```
alpha    Forward modelled bending angles at impact parameters a (in rad).
```

USES

```
refrac_info
```

A.12.4 Forward_model/calc_alphaK

NAME

```
calc_alphaK - Gradient of calc_alpha.
```

SYNOPSIS

```
call calc_alphaK (nobs, nlev, a, refrac, nr, Kmat_ref, Kmat_nr)
```

DESCRIPTION

```
This subroutine calculates the gradient of calc_alpha with respect t
refractivity and the 'refractivity x radius' product.
```

INPUTS

```
nobs      Number of observations.
nlev      Number of levels in the nr and refrac forward modelled profiles.
a         Impact parameter of the bending angle observations (in m).
refrac    Refracivity
nr        'Refractivity x radius' product.
```

OUTPUT

```
Kmat_ref  Gradient of calculated bending angles with respect to refractivity.
K_mat_nr  Gradient of calculated bending angles with respect to nr.
```

USES

```
refrac_info
```

A.12.5 Forward_model/calc_nr

NAME

```
calc_nr - Calculate the 'refractivity x radius' product on geopotential
             levels.
```

SYNOPSIS

```
call calc_nr(nlev, roc, lat, zg, refrac, nr)
```

DESCRIPTION

```
This subroutine calculates the product of refractive index and radius (i.e.,
distance to the center of curvature).
```

INPUTS

```
nlev       Number of vertical levels.
roc        Radius of curvature (in m.)
lat        Latitude (in degrees).
zg         Geopotential height (in gpm).
refrac     Refractivity.
```

OUTPUT

```
nr         Refractivity x radius product.
```

USES

```
refrac_info
e_rad
g_lat
```

A.12.6 Forward_model/calc_nrK

NAME

```
calc_nrK - Gradient of calc_nr.
```

SYNOPSIS

```
call calc_nrK(nlev, roc, lat, zg, refrac, dnr_dzg, dnr_dref)
```

DESCRIPTION

```
This subroutine calculates the gradient of calc_nr with respect to
geopotential heights and refractivity.
```

INPUTS

```
nlev       Number of vertical levels.
roc        Radius of curvature (in m.)
lat        Latitude (in degrees).
zg         Geopotential height (in gpm).
refrac     Refractivity.
```

OUTPUT

```
dnr_dzg    Gradient of nr with respect to geopotential height.
dnr_dref   Gradient of nr with respect to refractivity.
```

USES

```
refrac_info
e_rad
g_lat
```

A.12.7 Forward_model/refrac_levs

NAME

```
refrac_levs - Calculate refractivity and geopotential heights.
```

SYNOPSIS

```
call refrac_levs(nstate, nlev, nwet, pz0, press, x, zg, refrac)
```

DESCRIPTION

This subroutines calculates refractivity values at a given set of observation heights.

INPUTS

```
nstate  Number of elements in the state vector.
nlev    Number of levels for temperature.
nwet    Number of levels for humidity.
pz0     Surface elevation (in m).
press   Pressure levels for the state vector.
x       State vector.
```

OUTPUT

```
zg      Observation heights (in gpm).
refrac  Refractivity on observation heights.
```

USES

```
refrac_info
```

A.12.8 Forward_model/refrac_levsK

NAME

```
refrac_levsK - Gradient of refrac_levs.
```

SYNOPSIS

```
call refrac_levsK(nstate, nlev, nwet, pz0, pres, x, dzg_dx, dref_dx)
```

DESCRIPTION

This subroutine calculates the gradient of refrac_levs with respect to the elements of the state vector.

INPUTS

```
nstate  Number of elements in the state vector.
nlev    Number of levels for temperature.
nwet    Number of levels for humidity.
pz0     Surface elevation (in m).
press   Pressure levels for the state vector.
x       State vector.
```

OUTPUT

```
dzg_dx  Gradient of geopotential height with respect to state vector.
dref_dx Gradient of refractivity with respect to state vector.
```

USES

```
refrac_info
```

A.13 libalpha_1dvar/Math

DESCRIPTION

```
The alpha_1DVar library contains a small number of mathematical
routines.
```

SEE ALSO

```
Cholesky    Cholesky decomposition of a real symmetric matrix.
erf         Error function.
satmatinv   Invert a positive definite matrix.
```

A.13.1 Math/Cholesky

NAME

```
Cholesky - Cholesky decomposition of a real symmetric matrix.
```

SYNOPSIS

```
call subroutine Cholesky(U, v, N, q, ErrorCode)
```

DESCRIPTION

```
This subroutine solves the linear equation Uq = v for q by means of a
Cholesky decomposition, where U is a NxN symmetric positive definite
matrix, and V and q are vectors of length N.
```

INPUTS

```
U(N, N)    Positive definite symmetric matrix (only upper half is used).
v(N)       Array.
N          Number of elements.
```

OUTPUT

```
q(N)       Solution.
errorcode  0: Ok, 1: V is not positive definite.
```

NOTES

```
If U is not positive definite this will be detected by the program
and flagged as an error.  U is assumed to be symmetric as only the
upper triangle is in fact used.
```

A.13.2 Math/erf

NAME

```
erf - Error function.
```

SYNOPSIS

```
y = erf(x)
```

DESCRIPTION

```
This function calculates the value of the error function.
```

INPUTS

```
x    real number.
```

OUTPUT

```
y    value of the error function at x.
```

A.13.3 Math/satmatinv

NAME

```
satmatinv - Invert a positive definite real matrix.
```

SYNOPSIS

```
call satmatinv(n, m, A, status)
```

DESCRIPTION

```
This subroutine calculates the inverse of a real symmetric positive
definitive matrix using a Cholesky decomposition
```

INPUTS

```
 N:                Size of the matrix being inverted
 M:                If MATRIX is not present this is the same as N, else
                      this is the Other dimension of MATRIX.
 A:                Real matrix (assumed square and symmetrical)
                   overwritten by its inverse if MATRIX is not
                   present.
```

OUTPUT

```
A:              Real matrix (assumed square and symmetrical)
                overwritten by its inverse if MATRIX is not
                present.
Status:         0: ok, 1: A is not positive definite.
```

NOTES

```
Cholesky decomposition solves the Linear equation UQ = V for Q
where U is a symmetric positive definite matrix and U and Q
are vectors of length N.

The method follows that in Golub and Van Loan although this is
pretty standard.

If U is not positive definite this will be detected by the program
and flagged as an error.  U is assumed to be symmetric as only the
upper triangle is in fact used.
```

A.14 libalpha_1dvar/Other

DESCRIPTION

```
The library contains a small number of otherwise useful routines and one
type declaration.
```

SEE ALSO

```
Derived types and programming utilities:
   refrac_info   Derived types and constants for RO observations.
   set_xb        Fill the state vector with background data values.

Geodesy:
   e_rad         Effective radius of Earth (for gravity and geopotential
                    height calculations).
   g_lat         Gravity at Mean Sea Level.

Thermodynamics:
   humid_check   Check that humidity is below saturation (and correct it
                    if not).
   svp           Water vapour saturation pressure.
```

A.14.1 Other/e_rad

NAME

```
e_rad - Effective radius of Earth (for gravity and geopotential height
          calculations).
```

SYNOPSIS

```
R_eff = e_rad(lat)
```

DESCRIPTION

This function calculates the effective radius of Earth, which is intended to be used in the calculation of geopotential height. The function implements the method described by List (1985).

INPUTS

```
lat     Latitude (in degrees).
```

OUTPUT

```
R_eff   Effective Earth radius (in m).
```

USES

```
g_lat
```

REFERENCES

List, R.J., Smithsonian Meteorological Tables, 6th ed., Smithsonian Institution Press, Washington, 1985.

A.14.2 Other/g_lat

NAME

```
g_lat - Gravity at Mean Sea Level
```

SYNOPSIS

```
g = g_lat(lat)
```

DESCRIPTION

This function calculates gravity at Mean Sea Level as function of latitude. The formulation follows List (1985).

INPUTS

```
lat   Latitude (in degrees).
```

OUTPUT

```
g     Gravitational acceleration (in m/s^2).
```

REFERENCES

List, R.J., Smithsonian Meteorological Tables, 6th ed., Smithsonian Institution Press, Washington, 1985.

A.14.3 Other/humid_check

NAME

```
humid_check - Check that humidity is below saturation (and correct it
                 if not).
```

SYNOPSIS

```
call humid_check(nstate, lev, nwet, press, x)
```

DESCRIPTION

```
This subroutine checks that the humidity values in the state vector
are below supersaturation (i.e., below 100 % relative humidity). If
they are, the humidity value is corrected to reflect a 100 % relative
humidity value.
```

INPUTS

```
nstate    Number of elements in the state vector.
nlev      Number of levels used for temperature.
nwet      Number of levels used for humidity.
press     Pressure levels used in the state vector (hPa).
x         State vector.
```

OUTPUT

```
x         State vector (elements might be changed by this routine).
```

SEE ALSO

```
set_xb    for the definition of the state vector.
```

A.14.4 Other/refrac_info

NAME

```
refrac_info - Derived types and constants for RO observations.
```

SYNOPSIS

```
use refrac_info
```

DESCRIPTION

```
This module provides derived types and physical constants useful
for radio occultation data.
```

SOURCE

```
Useful constants:

 logical, parameter :: check_hum = .FALSE.
 integer, parameter :: ITMAX = 20

 real,    parameter :: delta     = 0.01
 real,    parameter :: R         = 287.05
 real,    parameter :: CP        = 1005.0
 real,    parameter :: kappa     = R/CP
 real,    parameter :: Pref      = 1000.0
 real,    parameter :: Epsilon   = 0.62198
 real,    parameter :: C_virtual = 1.0E-3*(1.0/Epsilon - 1.0)
 real,    parameter :: aval      = 77.6
 real,    parameter :: bval      = 3.73E5
 real,    parameter :: g         = 9.80665
 real,    parameter :: RMDI      = -9999.0
 real,    parameter :: RMDItol   = -(1.0E-6*RMDI)
 real,    parameter :: Rog       = R/g
 real,    parameter :: pi        = 3.14159
 real,    parameter :: root_pi   = 1.77245

Constants for PGE QC:

 real,    parameter :: big_diff  = 0.05
 real,    parameter :: qcaval    = 0.001
 real,    parameter :: qcdval    = 10.0
 real,    parameter :: gamma     = 1.253314*qcaval/((1.0-qcaval)*qcdval)

Constant for check_input:

 real,    parameter :: ref_min   = 1.0E-3
 real,    parameter :: ref_max   = 500.0
 real,    parameter :: alpha_min = -1.0E-3
 real,    parameter :: alpha_max = 0.1
 real,    parameter :: a_min     = 6.2E6
 real,    parameter :: a_max     = 6.5E6
 real,    parameter :: zg_min    = 0.0
 real,    parameter :: zg_max    = 1.0E5
 real,    parameter :: T_min     = 150.0
 real,    parameter :: T_max     = 350.0
 real,    parameter :: lnQ_min   = -25.0
 real,    parameter :: lnQ_max   =  4.0

Derived type for (bending angle) observations):

 type ob_type
    real            :: lat
```

```
    real             :: lon
    integer          :: year
    integer          :: month
    integer          :: day
    integer          :: hour
    integer          :: min
    integer          :: sec
    integer          :: doy
    integer          :: ocnum
    real             :: roc
    integer          :: nobs
    real,   pointer :: a(:)
    real,   pointer :: alpha(:)
    real,   pointer :: alpha_b(:)
    real,   pointer :: alpha_s(:)
    real,   pointer :: pge(:)
 end type ob_type

Derived type for background data:

 type bg_type
    real             :: lat
    real             :: lon
    integer          :: year
    integer          :: month
    integer          :: day
    integer          :: hour
    integer          :: min
    integer          :: sec
    integer          :: doy
    integer          :: ocnum
    real             :: pz0
    real             :: psurf
    integer          :: nlev
    real,   pointer :: pres(:)
    real,   pointer :: temp(:)
    real,   pointer :: qval(:)
 end type bg_type
```

A.14.5 Other/set_xb

NAME

```
set_xb - Fill the state vector with values.
```

SYNOPSIS

```
call set_xb(nstate, nlev, nwet, press, bg, xb)
```

DESCRIPTION

This subroutines sets the values of a state vector array from the background data provided in the bg structure.

INPUTS

```
nstate    Number of elements in the state vector.
nlev      Number of pressure levels (for temperature).
nwet      Number of pressure levels (for humidity).
press     Pressure levels to be used by the state vector (nlev elements).
bg        Structure of type(bg) as defined in refrac_info.
```

OUTPUT

```
xb        State vector (nstate elements).
```

USES

```
refrac_info
```

NOTES

```
The routine interpolates the background profile as given in the bg
structure onto the pressure levels press of the state vector. The
state vector has the following elements:
   1 ... nlev:              temperature
   nlev+1 ... nlev+nwet:    log(specific humidity)
   nstate:                  surface pressure
```

SEE ALSO

```
refrac_info
```

A.14.6 Other/svp

NAME

```
svp - Calculate water vapour saturation pressure.
```

SYNOPSIS

```
p = svp(T)
```

DESCRIPTION

```
This function calculates water vapour saturation pressure in hPa,
given the temperature in K.
```

INPUTS

```
real :: T    Temperature in K.
```

OUTPUT

```
real :: p    Water vapour saturation pressure in hPa
```

NOTES

```
svp uses a look-up array es_table and interpolates to the required
temperature. The returned value corresponds to svp over ice for
T <= 8 deg C, to water for T >= 5 deg C, and to transitional values
in between 5 & 8 deg C.
```