


Met Office

Progress towards the Implementation of Correlated Observation Errors in 4D-Var

Forecasting Research Technical Report 560

September 2011

Peter Weston



0 Introduction

Currently high resolution sounders, such as AIRS (the Atmospheric Infrared Sounder on the NASA satellite Aqua) and IASI (the Infrared Atmospheric Sounding Interferometer on the EUMETSAT satellite Metop-A), use diagonal fixed observation error covariance (\mathbf{R}) matrices within the Met Office assimilation scheme, assuming no correlation between channels. This is inadequate due to the presence of errors of representativeness, forward model error and errors associated with the pre-processing of satellite data.

Previous work at the University of Reading [9], the Met Office [3] and ECMWF [1, 2] has demonstrated that correlations exist in IASI data particularly for channels sensitive to water vapour. It is likely that a better description of the error correlations in 4D-Var will allow for improved use of the water vapour channels.

Using a generalised \mathbf{R} matrix complicates the processing within the assimilation scheme. In particular the inverse of \mathbf{R} needs to be computed at every iteration for each profile which could turn out to be too costly to implement. Representing the correlations as block diagonals or a matrix with an inverse that can be reconstructed efficiently may be the practical solution. This will hopefully allow for a more realistic \mathbf{R} matrix to be efficiently implemented into the operational assimilation system.

This report contains the following sections:

1. Introduction to AIRS and IASI
2. Data Assimilation
3. Desroziers Diagnostic
4. 1D-Var Results
5. 4D-Var Results
6. Timings
7. Conditioning
8. Conclusions
9. Future Work
10. Acknowledgements

§1 to §3 give an introduction to the instruments providing the data, some general theory and the methods used in this project. §4 & §5 show some results produced via the Desroziers diagnostic procedure. The remaining sections concentrate on what exact form of the matrix to use. The appendices focus on changes to the VAR code which will enable treatment of a full \mathbf{R} matrix.

1 Introduction to AIRS and IASI

AIRS and IASI are both infra-red sounders with high spectral resolution. They both have well over 1000 channels compared to the 20 channels of the previously used infra-red sounder HIRS (High resolution InfraRed Sounder). Currently only a small subset of ~ 100 channels from each instrument are assimilated into the NWP model. This is due to the cost of processing the data from each channel, unknown effects of ozone and solar reflectance on certain channels causing forward model error and the unknown nature of the correlations between channels [4]. It is hoped that the results of this work may allow the use of a larger channel selection.

1.1 AIRS

AIRS is a grating spectrometer which has been operating aboard the NASA satellite Aqua since its launch in May 2002. Aqua is in a sun-synchronous orbit 705km above the Earth's surface. AIRS has 2378 channels at a spectral resolution of 0.5cm^{-1} ranging from 650cm^{-1} to 2665cm^{-1} . Its spatial resolution at nadir is 13.5km. Due to the cost of processing the data from all the channels only 324 of these channels are received for use at the Met Office and just 142 of these are assimilated into the model. For this reason all of the diagnostic \mathbf{R} matrices I have produced for AIRS will have dimensions of 142×142 .




Figure 1: Sample spectrum of the 324 AIRS channels used at the Met Office

A sample spectrum of the 324 channels is shown in figure 1 with the coloured crosses indicating the 142 channels which are assimilated into the model. The red, green and blue crosses represent the temperature sounding channels, window channels and water vapour sensitive channels respectively.

1.2 IASI

IASI is a Michelson interferometer which has been operating aboard the EUMETSAT satellite Metop-A since its launch in October 2006. Metop-A is in a sun-synchronous orbit 817km above the Earth's surface. IASI has 8461 channels at a spectral resolution of 0.25cm^{-1} ranging from 645cm^{-1} to 2760cm^{-1} . Its spatial resolution at nadir is 12km. Only 314 of these channels are used at the Met Office and just 138 of these are assimilated into the model. However 183 channels are processed in the observation processing system (OPS), so the diagnostic R matrices I have produced for IASI in §4 will have dimensions 183×183 and those produced in §5 will have dimensions 138×138 .




Figure 2: Sample spectrum of the 314 IASI channels used at the Met Office

A sample spectrum of the 314 channels is shown in figure 2 with the coloured crosses indicating the 183 channels which are processed in the OPS. The 45 channels which are used in the OPS but are not assimilated are represented by the yellow crosses and are mostly high peaking temperature sounding channels. The remaining coloured crosses represent the same bands of channels as in figure 1.

Comparing figures 1 & 2 show that the wavelengths of the channels of AIRS and IASI are very similar. However, due to differences in the instrument design and different absorption lines at similar wavelengths, the characteristics of individual channels differ between the two instruments.

2 Data Assimilation

Data assimilation is the process of combining information from the previous forecast (background) and the observations to produce an analysis (the best estimate of the true state of the atmosphere). This is done by minimising a cost function of the form

$$J(x) = \frac{1}{2} (x - x_b)^T \mathbf{B}^{-1} (x - x_b) + \frac{1}{2} (y - H(x))^T \mathbf{R}^{-1} (y - H(x)) \quad (1)$$

where x is any model state vector, x^b is the background state vector produced from a previous short-range forecast, y is the vector of observations, \mathbf{B} is the background error covariance matrix, \mathbf{R} is the observation error covariance matrix and H is the observation operator.

The observation and background error covariance matrices provide information on the error characteristics of the observations and background and therefore how much weight is given to each source of data. It is the observation errors which are to be studied here.

2.1 Observation Errors

Each observation that is assimilated into the model has a corresponding error. The two main types of observation error are systematic and random errors. Systematic errors are biases in the data and the data assimilation scheme assumes that all of the data has no systematic data. This means that all satellite data which is assimilated into the model is put through a bias correction procedure. Deficiencies in the bias correction can lead to some systematic errors appearing in the assimilation, which can lead to some error correlations.

Random errors are the errors which are stored in the \mathbf{R} matrix and passed into the assimilation scheme in this form. The reasons for these errors specifically applying to satellite observations include:

- Instrument noise
- Errors in the radiative transfer model (often called forward model errors)
- Errors of representativeness, caused when the observations can resolve finer scales than the model
- Errors relating to pre-processing of satellite data; including processes like cloud detection.

All of these bullet points can contribute to correlated errors between different channels which are the errors studied in this report. The \mathbf{R} matrix that is currently used is diagonal and so contains no information about how errors from different channels are related to or correlated with each other. However, it has been previously shown that these correlations exist for IASI and AIRS [1, 2, 3, 9] and so to account for them the currently operational 4D-Var diagonal \mathbf{R} matrix contains variances which have been inflated by an arbitrary amount. This is not an ideal way of modelling these correlations which is why this project is aimed at enabling the 4D-Var assimilation scheme to deal with a full \mathbf{R} matrix. Such a full \mathbf{R} matrix would contain the correlation information in its non-zero off diagonal elements.

3 Desroziers Diagnostic

To estimate the structure of the full \mathbf{R} matrix a diagnostic procedure introduced by Desroziers et al. [5] has been used. This uses observation minus background (O-B) and observation minus analysis (O-A)

statistics to produce observation error variances and covariances. The formula is:

$$\mathbf{R} = E \left[(y - H(x_a)) (y - H(x_b))^T \right] \quad (2)$$

where \mathbf{R} is the observation error covariance matrix, y is the observation, H is the observation operator which takes data in model space (variables such as temperature, humidity etc.) into observation space (variables such as brightness temperature, radiance etc.), x_a is the analysis which is the best estimate of the true state of the atmosphere and x_b is the model background which is the atmospheric state as produced by a short range forecast.

3.1 Derivation

The full derivation of (2) can be found in [5] & [9]. The main ideas behind the derivation are to use the best linear unbiased estimator $x_a = x_b + \mathbf{B}\mathbf{H}^T (\mathbf{R} + \mathbf{H}\mathbf{B}\mathbf{H}^T)^{-1} (y - H(x_b))$ and then make two main assumptions. The first is that observation errors and background errors are independent of one another. The second is that the \mathbf{R} and \mathbf{B} matrices used to produce the analysis are exactly correct.

It is the latter assumption which could cause some problems here. This is because in this project it is known that the \mathbf{R} matrix is not correct initially. Therefore the results should not be entirely trusted.

However it has been shown, in very simple examples, that iterating the Desroziers diagnostic after starting with incorrect errors can lead to convergence to the true errors, which is encouraging. This iterative procedure involves diagnosing some \mathbf{R} and \mathbf{B} matrices then running the assimilation again but with these diagnosed matrices used in the assimilation. Then some new diagnostic matrices are produced and this process can be repeated until the matrices converge to what should be the true matrices.

However, this only works if both the \mathbf{R} and \mathbf{B} matrices can be iterated. When using satellite data although a \mathbf{B} matrix can be diagnosed using a similar formula to (2) [5], the diagnosed matrix is in observation space and to use it in the assimilation it needs to be transformed back into (RTTOV) model space. This transformation is difficult so although the \mathbf{R} matrix can be iterated it will converge to something that is not quite the true matrix due to the inaccuracy of the \mathbf{B} matrix.

Nonetheless this procedure was run and the results can be seen in §4.

3.2 Technical Details

To apply the theory of the Desroziers diagnostic to real data the observation, background and analysis data required to calculate the O-A and O-B statistics had to be obtained. This was done in two ways which are explained in more detail in §4 & §5. How this data was used is explained here.

Firstly, the following formula was used to produce the raw observation error covariance matrix (\mathbf{R}):

$$\mathbf{R}(i, j) = \frac{1}{n} \sum_{k=1}^n [(O_i^k - A_i^k) (O_j^k - B_j^k)] - \left(\frac{1}{n} \sum_{k=1}^n (O_i^k - A_i^k) \right) \left(\frac{1}{n} \sum_{k=1}^n (O_j^k - B_j^k) \right) \quad (3)$$

where $\mathbf{R}(i, j)$ is the element of the \mathbf{R} matrix in the i th row and j th column representing the error covariance between channel i and channel j . O , A , and B are the brightness temperatures for the observation, analysis and background respectively. The subscripts i or j denote which channel and the superscript k denotes which profile that brightness temperature is from. Due to the product of brightness temperatures in the formula the units of the covariances and variances in the matrix are K^2 .

The first term comes from (2), whilst the second term is included to remove biases in the data.

3.3 Other Matrices

As well as the raw \mathbf{R} matrices some other matrices were also produced. The first of these is the observation error correlation matrix (\mathbf{C}). This gives a non-dimensional value between -1 and +1 to show how negatively or positively correlated two channels' errors are. The \mathbf{C} matrix is formed as follows:

$$\mathbf{C}(i, j) = \frac{\mathbf{R}(i, j)}{\sqrt{\mathbf{R}(i, i) \mathbf{R}(j, j)}} \quad (4)$$

where $\mathbf{R}(i, i)$ represents the error variance for channel i . It follows from the definition that the diagonal elements of \mathbf{C} are identically equal to 1.

The other matrices produced were as a result of looking at the first set of diagnostics (shown in §4). It turns out that the violation of the assumption about the true errors being used in the assimilation results in some unrealistic features of the diagnostic matrices. One of these features is that the diagnostic matrices are never perfectly symmetric. Obviously as the expected matrices are covariance and correlation matrices they must be symmetric by definition so in order to use them in the assimilation they would need to be made symmetric. This is quite a simple procedure which involves taking the mean of the original matrix and its transpose. So the symmetrised version of \mathbf{R} is $\mathbf{R}_{\text{sym}} = \frac{1}{2} \mathbf{R} + \mathbf{R}^T$ and similarly for \mathbf{C} .

Also, in order to visualise the asymmetry in the diagnostics some wholly asymmetric matrices were formed. This means that the values in these matrices visualise just how asymmetric the diagnostic matrices are. These matrices were formed by taking half of the difference of the original matrix and its transpose. So the antisymmetrised version of \mathbf{R} is $\mathbf{R}_{\text{asym}} = \frac{1}{2} \mathbf{R} - \mathbf{R}^T$ and similarly for \mathbf{C} .

4 1D-Var Results

To produce the diagnostic matrices found in this section the OPS was run for AIRS and IASI separately. The outputs from these runs included observation, background and analysis data, all in units of brightness temperature. The analysis brightness temperatures were produced by the 1D-Var retrieval that the OPS performs. There were individual brightness temperature values for each channel within every profile.

To avoid problems associated with the retrievals over land the data used was filtered to only come from observations over the sea. Also, only observations which had passed the quality control tests were used. Then the error covariances for each channel combination were calculated using (3). These covariances are shown in figures 3 to 6.




Figure 3: Diagnostic IASI R matrix




Figure 4: Diagnostic AIRS R matrix




Figure 5: Diagonal of the diagnostic IASI \mathbf{R} matrix




Figure 6: Diagonal of the diagnostic AIRS \mathbf{R} matrix

By comparing figures 3 & 5 and 4 & 6 (taking note of the scales) it can be seen that in general the covariances between different channels are much smaller than variances for individual channels. This is because there are only vertical errors of representativeness in 1D-Var, due to the fact that the 1D-Var assimilation is performed on a column of atmosphere at that profile's location. Therefore in this case the columns of retrieved values (analysis vectors) and the observations have the same horizontal resolution, hence no horizontal errors of representativeness. The errors that are being picked up are the instrument noise on the diagonal and elements of forward model error and vertical errors of representativeness away from the diagonal.

Other aspects of the diagnostics that are very noticeable are some of the unrealistic features. Firstly there are some channels for which the diagonal values of the diagnostic \mathbf{R} matrix, which represent error variances, are negative. By definition variances have to be positive so this is an unrealistic result. Secondly there is a large amount of asymmetry in the matrices. This is highlighted in figure 3 by the block of negative covariances below the diagonal but not above for the window channels indexed between 120 and 150. These unrealistic results are presumably due to the violated assumption that the correct \mathbf{B} and \mathbf{R} are being used in the assimilation when it is pretty clear that the correct matrices are not actually being used.




Figure 7: Difference between the original diagnostic IASI \mathbf{R} matrix and its transpose

One way of visualising the asymmetry from the diagnostic matrices is by looking at the matrices produced by taking the difference of the original matrix and its transpose as described in §3.3. An example of this type of matrix for the IASI data is shown in figure 7 where any non-zero value shows asymmetry.

The largest values appear in the covariances between window channels as previously mentioned. This potentially shows that the current diagonal \mathbf{R} matrix is furthest away from the truth for these channels.

In an ideal assimilation scheme where the correct \mathbf{B} and \mathbf{R} matrices were being used these difference matrices would be identically equal to zero and the diagnostic \mathbf{R} matrix would identically match the \mathbf{R} matrix used in the assimilation. In order to get closer to this perfect setup an iteration of the Desroziers diagnostic was run where the diagnostic \mathbf{R} matrix was used in the assimilation and then new diagnostics were produced.

To run this iteration the diagnostic matrices had to be edited. This was because none of the ones initially produced were positive definite or symmetric. To make them symmetric the mean of the matrix and its transpose was taken as described in §3.3. To make them positive definite the eigenvectors and eigenvalues were found, some of which were negative. Then these negative eigenvalues were set to small positive values and the matrix was reformed by using the opposite of an eigendecomposition as shown by the following equation:

$$\mathbf{R}_{\text{new}} = \mathbf{Q}\mathbf{D}\mathbf{Q}^{-1} \quad (5)$$

where \mathbf{R}_{new} is the new positive definite version of the diagnostic \mathbf{R} matrix, \mathbf{Q} is the matrix of eigenvectors and \mathbf{D} is a diagonal matrix with the new (all positive) eigenvalues along the diagonal.

The assumption leading to these unrealistic results also referred to the wrongly specified \mathbf{B} matrix as a possible cause. The operational \mathbf{B} matrix used in the 1D-Var retrievals is almost 10 years old and dates back to a time when 3D-Var was operational at the Met Office. For these reasons it is almost certainly wrongly specified at the moment. To counter this Marek Wlasak and Fiona Hilton produced a new \mathbf{B} matrix by sampling the behaviour of the much more accurate 4D-Var \mathbf{B} matrix currently used at the Met Office.

The main differences in the diagnostics when using the iterated \mathbf{R} matrix as well as the improved \mathbf{B} matrix are that generally the values for the variances and covariances for all channels are significantly larger. This can be seen by the redder appearance of the matrix in figure 8 when compared to figure 3 and the higher values of the diagonal in figure 9 when compared to figure 5. The larger variances mean that none of them are negative this time which is obviously an improvement.

Comparing figures 7 & 10 shows that using a full \mathbf{R} matrix and improved \mathbf{B} reduce the asymmetry as shown by the paler colours in the latter. This result can also be seen in the diagnostics produced using just the iterated \mathbf{R} matrix along with the old \mathbf{B} matrix (not shown here). These two main results are encouraging and show that the full \mathbf{R} matrix produced by the Desroziers diagnostic is probably closer to the true \mathbf{R} matrix than the original diagonal matrix.




Figure 8: Diagnostic IASI **R** matrix using an iterated **R** and improved **B** in the 1D-Var retrieval




Figure 9: Diagonal of the diagnostic IASI **R** matrix using more realistic **R** and **B** matrices in 1D-Var




Figure 10: Difference between the diagnostic IASI \mathbf{R} matrix from figure 8 and its transpose

As an extension this iteration could be run repeatedly until the diagnostic \mathbf{R} matrix converged to what should be the true matrix. However, as explained in §3.1 the \mathbf{B} matrix would need to be iterated as well and this is difficult.

Figure 11 compares the diagnosed and currently used operational standard deviations with the instrument noise (in standard deviation space). The instrument noise data has been derived from principal component residuals by EUMETSAT [7]. This gives a more accurate specification than the instrument noise which came from the pre-launch tests carried out by CNES (Centre National d'Etudes Spatiales), the French organisation which led the development of IASI in association with EUMETSAT.

The currently operational standard deviations were calculated by adding a constant forward model error contribution of 0.04K^2 to the pre-launch noise in variance space and then taking the square root. This is why the operational standard deviations are larger than the instrument noise for every channel.

The diagnosed standard deviations are similar to the instrument noise and slightly smaller than the operational standard deviations for most channels. However, the diagnosed standard deviations are significantly larger for some of the water vapour sensitive channels. This must be due to larger vertical errors of representativeness and forward model error affecting these channels.




Figure 11: Currently operational and diagnostic error standard deviations compared with IASI instrument noise

Due to the absence of horizontal representativeness errors in the 1D-Var assimilation the next section focuses on performing the Desroziers diagnostic procedure on data which has come from the 4D-Var assimilation system.

5 4D-Var Results

To produce the diagnostic matrices found in this section the Met Office's variational assimilation system, VAR, was run for AIRS and IASI separately. Gordon Inverarity and I developed code that would output the innovations (O-B) and residual (O-A) statistics for all observations assimilated into the model. Specifically for the satellite observations these statistics were in units of brightness temperature. From these statistics the Desroziers diagnostics shown in figures 12 to 17 were produced.




Figure 12: Diagnostic IASI R matrix




Figure 13: Diagnostic IASI C matrix




Figure 14: Diagnostic AIRS R matrix




Figure 15: Diagnostic AIRS C matrix

Note that when comparing the 4D-Var diagnostics with the 1D-Var diagnostics the scales are usually different. Figures 12 & 14 show that the covariance values for the window and water vapour sensitive channels are much larger in the 4D-Var diagnostics than the 1D-Var ones. This is due to the much larger horizontal representativeness errors in 4D-Var. This shows that the representativeness errors, and hence the error correlations, are largest between water vapour sensitive channels.

These representativeness errors can also be seen in figures 16 & 17 where the diagonal of the matrix from 4D-Var is significantly larger than the diagonal coming from 1D-Var for these channels.

Figures 13 & 15 are the first appearances of observation error correlation (C) matrices. The figures show that the strongest positive error correlations are between water vapour sensitive channels and window channels for both IASI and AIRS. These are shown by the red blocks surrounding the diagonal towards the bottom right corner of the matrices. Also apparent are the weaker error correlations between temperature sounding channels shown by the paler colours in the top left corner of the matrix.

The advantage of looking at these C matrices is that the scale is always the same. This is because the correlations themselves are dimensionless quantities always varying between -1 and +1.




Figure 16: Diagonal of the diagnostic IASI R matrix




Figure 17: Diagonal of the diagnostic AIRS \mathbf{R} matrix

Figure 18 shows that the diagnosed standard deviations are much smaller than the operational ones. This is because of the artificial inflation of the operational standard deviations to account for the inter-channel error correlations that are not modelled directly. This idea was introduced in §2.1. The inflation is larger for the water vapour sensitive channels as this is where the correlations were thought to be largest and my results have confirmed this.

Figure 19 also shows the instrument noise for comparison. Where the diagnosed standard deviations are close to the instrument noise it shows that the errors for these channels are barely affected by representativeness errors. Conversely channels where the diagnosed standard deviations are much larger than the instrument noise, have errors which are heavily affected by representativeness errors. This shows that for the temperature sounding channels instrument noise dominates the errors. For the window channels there are slight representativeness errors but instrument noise is more influential. For most of the water vapour sensitive channels representativeness errors are large and instrument noise is comparatively small.

This backs up the initial ideas that the representativeness errors causing correlations mainly affecting the water vapour sensitive channels. Similar figures for AIRS are not shown here but they do show the same patterns. Instrument noise dominates the diagnosed errors in the temperature sounding channels and representativeness errors are more important in the window and water vapour sensitive channels.




Figure 18: Currently operational and diagnostic IASI standard deviations




Figure 19: As figure 18 but zoomed in and compared with IASI instrument noise as well

6 Timings

To test the difference in the timings of running 4D-Var with or without the full \mathbf{R} matrices for AIRS and IASI, some VAR jobs were run using my new code, as explained in Appendix A. A basic 'AnalysePF' job was run which produces an analysis increment from a set of observations and background data. Jobs of this type also require the \mathbf{R} and \mathbf{B} matrices. The \mathbf{R} matrix is specified for each instrument individually and the contributions to the penalty function are built up by calculating the contributions from all of the observations of each observation type separately and then adding all of these together. This is done because it is assumed that different observation types do not have errors correlated with each other and it makes the code more efficient and easier to edit for individual observation types.

The new code is written so that it can deal with either a full or diagonal matrix. This meant that the only change needed to enable the jobs to run with a full \mathbf{R} matrix was to replace the diagonal \mathbf{R} matrix in the coefficients directory with a full one of the correct format. As when the job using the full \mathbf{R} matrices in 1D-Var was run, the matrices provided in 4D-Var had to be positive definite and symmetric. So the matrices were symmetrised as before and it turned out that all of the diagnostic matrices produced from 4D-Var data were positive definite. This meant that the manipulation of the eigenvalues explained in §4 didn't have to be repeated. The matrices were then put into files of the correct format and the experiments were ready to be run.

As well as the timings of different jobs using different combinations of full matrices the value of the cost function when using these different matrices is also included. As shown in (1) the cost function contains the inverses of both the \mathbf{R} and \mathbf{B} matrices so the smaller the errors in the matrices the larger the value of the cost function. As figure 18 showed, the full diagnostic matrices which I will be using have much smaller errors on the diagonal than the original diagonal matrices. This should result in an increase in the cost function.

If the observation errors used in the assimilation are correctly specified then the cost function should be proportional to the total number of observations divided by 2 [6, 10]. In the standard jobs I have been running there are ~ 1.9 million separate observations assimilated of which ~ 0.5 million are from each of AIRS and IASI. This suggests that if the observation errors are correctly specified the total cost function value should be ~ 950000 with contributions of ~ 250000 from each of IASI and AIRS.

6.1 Standard 4D-Var assimilation test

The control experiment is the run where diagonal \mathbf{R} matrices were used for all observation types and the results of this are shown in the top row of table 1. The other rows show the results of runs using a full \mathbf{R} matrix for either AIRS only, IASI only or both AIRS & IASI. The first result which jumps out from the table is the vast increase in overall time required to run the assimilation when running with any of the full matrices. For instance running with full matrices for both AIRS and IASI results in an increase in time of 127%. However what is also noticeable is the increase in iterations of the minimisation algorithm which in the case above equates to an increase of 128%. This suggests that the increase in time is wholly due

to the increase in iterations. This increase in iterations means that it is taking the minimisation longer to converge when using these full matrices, which possibly implies that the minimisation itself becomes unstable. This issue shall be returned to in more detail in §7.

Full \mathbf{R} used for	Number of Iterations	Overall Time	Final Value of J
No observation types	39	3605	319346.9
IASI	49	4514	528740.4
AIRS	79	7276	523894.5
AIRS & IASI	89	8184	732296.5

Table 1: The number of iterations, timings and cost function values resulting from using diagonal and full \mathbf{R} matrices in 4D-Var

The final column in table 1 shows that the value of the cost function does indeed shoot up when using the full matrices for IASI and AIRS. This may also help to explain the large increase in iterations as most of the convergence criteria in the minimisation are based on the value of the cost function. Referring back to the expected values of the cost function it can be seen that the value in the first line of the table where the diagonal matrices are used is around 3 times smaller than the expected value of ~ 950000 . In the same job the contributions to the cost function from IASI and AIRS are ~ 39000 and ~ 69000 respectively. These are both much lower than the expected contributions of ~ 250000 . This suggests that the inflated diagonal values in the IASI and AIRS matrices are *too large*.

However, when the full matrices are introduced the cost function does increase to a value much closer to the expected value. In fact the contributions to the cost function from IASI and AIRS increase to ~ 240000 and ~ 270000 respectively. These values are comparable to the expected value of ~ 250000 and this gives another indication that these full matrices are close to the true matrices.

6.2 Fixed iterations 4D-Var assimilation test

To isolate the effect of the extra processing required to invert the full matrices on the timings, the same jobs were run again but with a fixed number of iterations. This meant that any time increase would solely be due to the difference in processing the diagonal and full matrices for IASI and AIRS.

The time differences between using the full and diagonal matrices are minimal, table 2 shows. In fact when using just the full IASI \mathbf{R} matrix the overall time reduces by 0.75%! When using full matrices for both IASI and AIRS the time only increases by 0.19%. These results are very encouraging showing that because 4D-Var is performing so many other computationally expensive operations the change to processing a full \mathbf{R} matrix barely has any effect on the overall timings.

Full R used for	Number of Iterations	Overall Time	% Change in Time
No observation types	39	3605	0
IASI	39	3578	-0.75
AIRS	39	3640	+0.97
AIRS & IASI	39	3612	+0.19

Table 2: The number of iterations, timings and cost function values resulting from using diagonal and full **R** matrices in 4D-Var with a fixed number of iterations

6.3 Conjugate gradient 4D-Var assimilation test

The jobs used to run the timing tests above were standard 4D-Var assimilation jobs running a limited memory quasi-Newton minimisation known as the L-BFGS (after its founders Broyden, Fletcher, Goldfarb and Shanno). However since November 2010 at the Met Office the operational 4D-Var assimilation is run using a conjugate gradient method of minimisation. For more information on these minimisation methods see [8]. The upshot of this is that there are two VAR runs: the first runs for a fixed number of 30 iterations at N108 resolution; the second runs for another fixed 30 iterations but at the higher N216 resolution. This is an efficient and cheap way of producing a high resolution analysis. Now, because of these fixed iterations the problem of increasing iterations due to using the full **R** matrices should go away. To show this the operational configuration was run with diagonal and full matrices.

However, when running the operational configuration using full **R** matrices for both IASI and AIRS there was a problem. In the first low resolution run the minimisation got to the 16th iteration and then failed. The error that caused the job to fail was that the value of the cost function was increasing between iterations. Since the whole aim of the 4D-Var assimilation is to minimise the cost function this is a major issue.

This provided further evidence that the minimisation was becoming unstable when using these full matrices. It has been suggested that this problem could be caused by the matrices being ill-conditioned. In §7 the conditioning of the matrices used is investigated and some methods of improving it are suggested.

6.4 OPS test

The difference between using the diagonal and full matrices for the 1D-Var within the OPS was also tested.

In table 3 there are three matrices used for each instrument. As well as the diagonal and full matrices used there are also rows containing timings for 'sparse' matrices. The difference between the sparse matrices and the full matrices are that the full matrices contain non-zero off diagonal elements for all ~300 channels (324 for AIRS and 314 for IASI) whereas the sparse matrices only contain non-zero off diagonal elements for the ~150 channels (142 for AIRS and 183 for IASI) processed in the OPS.

Instrument	Matrix used	Overall Time	% Change in Time
IASI	Diagonal	578	0
	Sparse	700	+21.1
	Full	639	+10.6
AIRS	Diagonal	472	0
	Sparse	530	+12.3
	Full	489	+3.6

Table 3: The timings from using diagonal and full \mathbf{R} matrices in 1D-Var

The timings show that using either the sparse or full matrices compared to the diagonal ones result in an increase in the time required to run the OPS. However the surprising result is that using the full matrices results in a smaller increase in time than using the sparse matrices. This may be due to the consistency of the errors in the full matrices. It is also interesting to note that using the full or sparse matrix for IASI results in a larger time increase than using the corresponding matrix for AIRS.

As expected these time increases are larger than in the 4D-Var system and this is due to the OPS being much less computationally expensive. This means that the same increase in processing time involved in inverting the full matrices results in a larger time increase proportionally in the 1D-Var.

6.5 Summary

The difference in timings when using the full \mathbf{R} matrices for AIRS and IASI in both VAR (running 4D-Var) and the OPS (running 1D-Var) have been investigated. The results are that the extra processing time involved in inverting the matrices results in a very small, almost negligible, increase in the overall time taken to run the 4D-Var assimilation. The increase in time is slightly larger in running the 1D-Var assimilation. This means that it may be acceptable to use these full matrices if it can be demonstrated that using them in 1D-Var is beneficial. This would avoid the need to use an approximation to the full matrices when they are implemented operationally.

7 Conditioning

As referred to in the previous section one of the possible causes of the vast increase in cost function value and the instability in the minimisation when using the full \mathbf{R} matrices is the conditioning of the matrices used. The conditioning of a matrix measures the ‘distance’ between the matrix and the set of singular or non-invertible matrices. It is quantified by the condition number which is defined in (6):

$$\kappa = \frac{\lambda_{\max}}{\lambda_{\min}} \quad (6)$$

where κ is the condition number, λ_{\max} is the largest eigenvalue and λ_{\min} is the smallest eigenvalue.

The larger this value is the more poorly conditioned the matrix and the ‘closer’ it is to being singular. An example of the effect poor conditioning can have is if one is solving the linear equation $Ax = b$, the larger the condition number of A , the more sensitive the solution x is to changes in b . Using an ill-conditioned matrix in the minimisation in data assimilation can result in the minimisation becoming unstable.

In this section the condition numbers of the various matrices that have been used in the assimilation tests will be calculated. Also different methods of improving the condition numbers and the results of using reconditioned matrices will be investigated.

Instrument	Matrix	Condition Number
IASI	Diagonal 1D-Var	2844.4
	Full 1D-Var	50472.5
	Diagonal 4D-Var	1600.0
	Full 4D-Var	26668.3
AIRS	Diagonal 1D-Var	73.7
	Full 1D-Var	1.766×10^8
	Diagonal 4D-Var	256.0
	Full 4D-Var	12337.9

Table 4: Condition numbers of the AIRS and IASI full R matrices used in the timing tests

Apart from the 1D-Var AIRS R matrix these condition numbers don’t seem too large when one takes into account the size of the matrices. However, using any matrix with a condition number which is a number of orders of magnitude larger than 1 can lead to the problems of minimisation instability, hence these matrices will be reconditioned and retested in the assimilation runs.

The condition numbers of the diagonal matrices are also included in table 4 for reference.

7.1 Methods of Reconditioning

Three different methods were used to recondition the matrices. All of the methods involve setting a threshold maximum value for the condition number to be. Then the eigenvalues are changed to satisfy this condition number and the matrices are reconstructed using the reverse eigendecomposition technique from (5) in §4. The three methods used are now explained in more detail.

The first method involves decreasing the condition number by adding a small increment to the eigenvalues. The increment used is equal to the current smallest eigenvalue so once the increment has been added the smallest eigenvalue is doubled and so the condition number is approximately halved (as the small increment has a very small effect on the largest eigenvalue especially when the condition number is $\gg 1$). This technique is repeated until the condition number is below the predetermined maximum

threshold value. This technique does preserve the approximate structure of the eigenvalues but does not guarantee that the new condition number is exactly the same as the threshold value

The second method is similar to the first but instead of incrementing by the smallest eigenvalue, the eigenvalues are incremented by an exact amount calculated to ensure the matrix would have exactly the threshold condition number. The formula used to calculate the increment was

$$\kappa_{\max} = \frac{\lambda_{\max} + \lambda_{\text{inc}}}{\lambda_{\min} + \lambda_{\text{inc}}} \quad (7)$$

and rearranging this gives

$$\lambda_{\text{inc}} = \frac{\lambda_{\max} - \lambda_{\min} \kappa_{\max}}{\kappa_{\max} - 1} \quad (8)$$

where κ_{\max} is the threshold condition number and λ_{inc} is the exact increment used. This method preserves the overall structure of the eigenvalues and ensures that the condition number is exactly the threshold value.

The third method involves setting a threshold minimum value for the eigenvalues by dividing the maximum eigenvalue by the required condition number. Then any eigenvalues which are smaller than this threshold value are changed to be equal to it. This results in a matrix with a condition number of exactly the threshold value set but with a different eigenvalue structure.

7.2 Results of Reconditioning

AIRS & IASI R Matrices used	Condition Number Threshold	Method used	Number of Iterations	Overall Time	Final Value of J
Diagonal	N/A	N/A	39	3605	319346.9
Full	N/A	N/A	89	8184	732296.5
Full	1000	1	69	6348	519978.7
Full	100	1	69	6326.3	344716.1
Full	1000	2	79	7255.4	537180.2
Full	100	2	60	5509.9	354831.0
Full	100	3	69	6317.2	422012.5
Full	25	3	49	4481.7	301932.3

Table 5: The number of iterations, timings and cost function values resulting from using reconditioned full R matrices in 4D-Var

A range of different matrices using different condition number thresholds and the three different methods explained above were then produced. Using these reconditioned matrices the same jobs were run as for

the timing tests and the outputs were compared with the jobs which were run with the initial raw diagnostic matrices.

In general table 5 shows that using the reconditioned matrices does indeed reduce the number of iterations required for convergence. This shows that the high condition numbers of the original matrices were at least partly responsible for the instability in the minimisation. The other effect of improving the conditioning is that the value of the cost function decreases. As discussed in §6 this is suggesting that the observation errors are not correctly specified. This means that although reconditioning results in an improvement in the stability of the minimisation, it also results in the errors being larger than the actual true errors. The overall choice of matrix will have to be a compromise between using the correct errors and preserving the stability of the minimisation.

Lowering the condition numbers of the matrices to 25 results in only 10 extra iterations compared to using the diagonal matrix and a smaller cost function value. This smaller cost function value suggests that this matrix contains errors which are too large which will lead to the AIRS and IASI having less impact in the analysis. Also when the condition number is reduced to 25 the matrices bear hardly any resemblance to the original diagnostics and are almost diagonal in appearance. For these reasons trying to use a matrix with the condition number edited to be that low would not model the correlations accurately enough.

Comparing method 1 with either method 2 or 3 is difficult because although the condition number threshold can be the same, due to the difference in the methods, the actual condition number of the matrices will be smaller for matrices produced using the first method. This is the reason for the smaller cost function value when using the matrices generated by method 1 with a condition number threshold of 100. In fact the condition number of the AIRS and IASI R matrices used in that job were 97.4 & 53.1 respectively.

If the condition number is to be limited to a specific value this suggests that either methods 2 or 3 are preferable. Additionally if the overall structure of the eigenvalue distribution is to be preserved rather than setting all of the smallest eigenvalues to some constant reference value, then method 2 is preferable. Also, for the matrices with a condition number threshold of 100, the second method results in 9 fewer iterations than the third method.

Having used the reconditioned matrices in the standard assimilation they were then tried out in the operational configuration as explained in §6.3. Recall that when the full raw diagnostic matrices were tried in the operational configuration there was a problem in the minimisation. When running with the reconditioned matrices there were no such problems, even with the condition number threshold set at 1000. Also, the analysis increments from the job using the diagonal matrices and the job using the reconditioned matrices were compared and they looked very similar. This suggests that the two sets of 30 iterations are enough for the analysis to converge to the correct solution, even when using the full reconditioned matrices.

7.3 Best Matrix to Use

Ideally a matrix which models the correlations as accurately as possible whilst not upsetting the stability of the minimisation will be used. As the operational configuration has a limited number of iterations the increase in iterations should not be a problem when using the full matrices. This is true providing that the analysis converges fully within the two sets of 30 iterations. This appears to be the case for the reconditioned matrices so the best matrix to use is the one which has been reconditioned via method 2 with a condition number of 100 in the initial tests. If it turns out that this matrix does not perform well then a tighter constraint could be enforced on the condition number. Conversely if this matrix performs well then a matrix with a larger condition number, and hence more accurate error correlations, could be used.

8 Conclusions

Estimates have been produced of the structure of the correlated errors in IASI and AIRS channels using the Desroziers diagnostic procedure in §4 & §5. There are slight differences between the structures diagnosed from the 1D-Var and 4D-Var output. There are larger variances and covariances in the window and particularly the water vapour sensitive channels in the diagnostics from 4D-Var output. This is due to larger horizontal errors of representativeness.

The overall structure of the diagnostic matrices suggests that the temperature sounding channels are very weakly correlated with each other, the window channels are moderately strongly correlated with one another and the water vapour sensitive channels are strongly correlated with one another. This structure appears to be accurate for both AIRS and IASI.

Code has been developed to allow full \mathbf{R} matrices to be used in the 4D-Var assimilation scheme at the Met Office in Appendix A. The diagnostic matrices have been tested and edited to improve their performance in the assimilation in §6 & §7. A reconditioned matrix has been proposed, which should be used in the initial testing of the treatment of correlated errors in VAR.

Finally, further changes to the VAR code for the consistent and error-proof handling of the full \mathbf{R} matrices have been proposed in Appendix B. This has hopefully been done in such a way that all types of VAR jobs which use \mathbf{R} matrices will run with either diagonal or full matrices. This will have to be tested thoroughly.

After consultation based on this report the required code changes will be made, tested and reviewed. Then several acceptance tests will be run with diagonal and full matrices making sure the code is doing what it is supposed to. The new matrices will be formed for input into both the OPS and VAR. Then these matrices will be tested in an assimilation trial, the impact of using the matrices will be gauged and, if positive, the system should be implemented into the operational configuration.

9 Future Work

There are a few ways this work could be extended:

- Perform an iteration of the Desroziers diagnostic in 4D-Var similar to the one done in §4 and analyse the results.
- Investigate whether there are significant correlated errors in other instruments, such as the existing microwave sounders ATOVS and SSMIS as well as the new instruments ATMS and CrIS aboard the NPP platform.
- Research the implications that using correlated observation errors for IASI and AIRS have on the optimum channel selection.
- Use other techniques to estimate the inter-channel error correlations and compare the results to those found using the Desroziers diagnostic.
- Use the Desroziers diagnostic to estimate spatial correlations and investigate ways of providing this information into the assimilation.
- Derive possible approximations or compact representations of \mathbf{R} for future use. This might become necessary when using more channels from existing instruments or when introducing new hyper-spectral instruments. These changes could result in the cost of the inversion becoming too large when using the full matrices in their current form.

10 Acknowledgements

I would like to thank a number of people who have helped me throughout this project. Firstly Brett Candy for introducing the project and providing an overall plan to follow. Also James Cameron and Laura Stewart, whose work I am following on from, for their continued support and guidance. Thanks to Fiona Hilton who provided technical guidance particularly in manipulating the matrices. Thanks to Gordon Inverarity for the code changes that allowed the required statistics to be outputted from the 4D-Var assimilation. Thanks to Nigel Atkinson for information on the IASI instrument noise estimates. I would also like to thank John Eyre and Ed Pavelin for their inputs into the direction the project was heading in at several key stages, and their advice on the content of this report.

References

- [1] Niels Bormann and Peter Bauer. Estimates of spatial and interchannel observation-error characteristics for current sounder radiances for numerical weather prediction. I: Methods and application to ATOVS data. *Q. J. R. Meteorol. Soc.*, 136:1036–1050, 2010.

- [2] Niels Bormann, Andrew Collard, and Peter Bauer. Estimates of spatial and interchannel observation-error characteristics for current sounder radiances for numerical weather prediction. II: Application to AIRS and IASI data. *Q. J. R. Meteorol. Soc.*, 136:1051–1063, 2010.
- [3] James Cameron. Options for improving the treatment of observation error used in the assimilation of hyperspectral infrared sounder data. Technical report, Met Office, 2010.
- [4] A. D. Collard. Selection of IASI channels for use in numerical weather prediction. *Q. J. R. Meteorol. Soc.*, 133:1977–1991, 2007.
- [5] G. Desroziers, L. Berre, B. Chapnik, and P. Poli. Diagnosis of observation, background and analysis-error statistics in observation space. *Q. J. R. Meteorol. Soc.*, 131:3385–3396, 2005.
- [6] Gerald Desroziers and Serguei Ivanov. Diagnosis and adaptive tuning of observation-error parameters in a variational assimilation. *Q. J. R. Meteorol. Soc.*, 127:1433–1452, 2001.
- [7] EUMETSAT. *EPS Product Validation Report: IASI L1 PCC PPF*, August 2010.
- [8] T. J. Payne. *VAR Scientific Documentation Paper 5*. Met Office, January 2007.
- [9] Laura M. Stewart. *Correlated observation errors in data assimilation*. PhD thesis, The University of Reading, 2009.
- [10] O. Talagrand. A posteriori evaluation and verification of analysis and assimilation algorithms. In *Proceedings of ECMWF Workshop on Diagnosis of data assimilation systems*, 1999.

Appendices

A Basic Code Changes

Although the OPS code had previously been written to deal with full **R** matrices, the VAR code had not. Hence in order to run experiments on the timings of using the full **R** matrices the code had to be edited to allow for the treatment of these full matrices.

The routine which sets up and reads in the SatRad **R** matrices is `Var_SatRad_Setup`. A diagonal, band diagonal or full **R** matrix can then be read in by `Ops_SatRad_GetRmatrix`. This meant it would just be the processing which followed the reading that would need editing.

Now, due to varying channel selections the inversion of the full **R** matrix couldn't be performed here as the channels have not been selected at the point when this routine is called. Therefore the only change made to `Var_SatRad_Setup` was to insert an if statement which only nullified the full matrix and inverted the diagonal if the matrix read in is indeed diagonal. If the matrix read in is full then no further processing is by this routine. For now any new code that would be required to deal with a band diagonal matrix has been ignored. However this could be added in the future if a band diagonal matrix is the best way of representing the correlations.

The only routine which then uses the SatRad **R** matrix is Var_SatRad2PenAndGrad which is where the cost function for the SatRad observation types is calculated. This routine contains lots of processing specific to calculating observation sensitivities as well as code to allow for variable observation errors. The changes to these sections of code are explained in Appendix B.

After ignoring those sections of code the only reference to the **R** matrix left was in one line where part of the cost function was calculated for each profile individually. This line of code assumed that the **R** matrix was diagonal. The changes made to Var_SatRad2PenAndGrad were:

- Insert an if statement based on whether **R** is diagonal or not.
- If it is diagonal then the existing code runs.
- If it is full then:
 - Copy the **R** matrix for the current channel selection into a new correctly sized array.
 - Call Ops_SatRad_Cholesky which performs the inversion and matrix multiplication between the inverted matrix and the 'BTdiff' array (this array contains the difference between the brightness temperature observation and the current (i.e. on this iteration) brightness temperature model value).
 - The output from Ops_SatRad_Cholesky is put in the correctly named array which will eventually get fed into the cost function.
 - If there is an error with the inversion then a fatal error is spawned.

When this code gets implemented operationally a fatal error would be a bit extreme so code could be put in to reject the observation but this is discussed in more detail in Appendix B.

The new code was then built and tested and now runs error free with both diagonal and full matrices. This meant that experiments could be run using different versions and combinations of full matrices for AIRS, IASI or both. The results of these tests can be seen in §6.

B Advanced Code Changes

In addition to the code changes explained, in Appendix A, code specific to performing observation sensitivity calculations and also the code which will, in the near future, process variable observation errors had to be edited. One final issue was error handling which will be discussed towards the end of this section.

B.1 Observation Sensitivity code

As well as being used in the calculation of the cost function in Var_SatRad2PenAndGrad, the **R** matrix is also used in the observation sensitivity code. It is used in two different contexts for each of the three different types of brightness temperature observations which this routine deals with.

Firstly, when using diagonal errors the observation error values themselves are stored in an observation error variable ready to be written out into the sensitivity output files. This will be left much the same as there is no accurate way to summarise all of the information from the full matrix for each observation without just providing the diagonal value for the relevant channel. This is also what is done with GPSRO errors when a full **R** matrix is used in the assimilation. There is a slight difference when using the full matrix however. If the matrix is diagonal it is inverted in `Var_SatRad_SetUp` and the inverted diagonal overwrites the original diagonal in the same array. When using the full matrix, the original diagonal values remain in the array.

Therefore a switch will be inserted so when the diagonal matrix is used the observation error is equal to the reciprocal of the diagonal array (as before). Whereas when the full matrix is used the observation error is equal to the diagonal array itself.

Secondly, the sensitivity values themselves are equal to the inverse of the **R** matrix multiplied by the 'BTdiff' array at the final iteration (this is equal to the residual or O-A statistic). When using the diagonal matrix this is a simple element by element multiplication of the reciprocal of the diagonal variances by the 'BTdiff' values for each channel. However when using the full matrix the matrix will need to be inverted and matrix multiplied with the 'BTdiff' array. This is very similar to the change made in the cost function calculation stage within this same routine. Namely using `Ops_SatRad_Cholesky` to do the inversion and multiplication. Then the outputted array is used to form the cost function value, or in this case the sensitivity value.

Due to the similarity to the change made previously, this calculation can just be done once and then the outputted value can be used wherever it is needed. Putting this into practice was slightly more complicated than first thought as the 'BTdiff' arrays are populated separately for the three different brightness temperature observation types in the routine. This meant some of the code had to be slightly reordered so that the 'BTdiff' and the 'Rmatrix' arrays are all populated before performing the inversion and multiplication. Finally another local variable, 'Sensitivity', was required to store the result of the multiplication of the inverse matrix and 'BTdiff' array.

Overall these changes seem pretty simple but they still need to be tested and checked over by people responsible for the observation sensitivity code.

B.2 Variable Observation Error code

Code has been developed by Brett Candy to allow for the treatment of variable observation errors depending on certain variables within individual observations. Examples of the parameters with which error values vary are surface to space transmittance and scan position. In order to implement a full variable **R** matrix it will have to be formed by multiplying the correlation matrix which is read in from the 'Rmatrix file' by the individual variable errors for each observation. This will potentially result in a different matrix for every different observation. The extra storage and processing required to deal with this will have to be investigated thoroughly before this is implemented. At present (September 2011) variable observation

errors are not used operationally so this is not a problem which has to be completely resolved at present but will need to be considered in the near future.

The code which was originally written by Brett assumes the variable error matrices are diagonal and are stored within the observation data structure under the variable name 'BriTempVarError'. The error data for this variable is calculated in the OPS and passed to VAR within the 'varobs' files. Currently the code is set up so that if variable errors are present in the observations the normal 'Rmatrix file' is ignored and the remainder of the assimilation just uses the diagonal variable errors.

To use both correlated and variable errors the structure of the code in Var_SatRad.SetUp will have to be changed. It will need to be possible to read in the **C** matrix from the 'Rmatrix file' and multiply this by the variable standard deviations for each observation instead of multiplying by the fixed standard deviations in the file. This could be done by writing a new routine or editing Ops_SatRad_GetRmatrix to do the reading and multiplying to form a separate **R** matrix for each observation. In this case the 'Rmatrix' variable would need 3 dimensions. The usual two, rows and columns of the matrix, and the third dimension would be for the observation number that the matrices are valid for. Then whenever the matrix is used it would have to include the index for the observation being processed at that time. Also the diagonal values usually stored in the structure 'Rmatrix % Diagonal' would have to be ignored and the 'BriTempVarError' variable would have to be used instead.

This will involve quite a lot of new code so I will leave this work to be continued once the variable observation error work becomes operational. At that point I will liaise with the person responsible for that work and we shall have to work out the most efficient way of implementing the changes required. The ideas above can act as a starting point for these changes.

B.3 Error handling

As explained in Appendix A a fatal error was inserted if there was a problem with the inversion of the full **R** matrix within Ops_SatRad.Cholesky. This could become a problem once the change goes operational as a fatal error would cause the operational assimilation to fail. This possibility must be avoided at all costs. One way of handling a possible error in the inversion is just to reject the observation which is being processed at that time. This would be a suitable error handling procedure as long as it didn't happen regularly. So far in all of my tests this error has not occurred at all so the chances of it occurring regularly are very small as long as each matrix which is used is tested thoroughly.

B.4 Summary

These appendices will make up most of the documentation required to accompany the code changes. This report can therefore be linked from a VSDP or OSDP document to explain the code changes.

Met Office

FitzRoy Road, Exeter
Devon, EX1 3PB
UK

Tel: 0870 900 0100

Fax: 0870 900 5050

enquiries@metoffice.gov.uk

www.metoffice.gov.uk