

GRAS SAF Report 06
Ref: SAF/GRAS/METO/REP/GSR/006
Web: www.grassaf.org
Date: 4 February 2008

The EUMETSAT
Network of
Satellite Application
Facilities



GRAS SAF Report 06

Levenberg-Marquardt minimisation in ROPP

Huw Lewis

Met Office, UK

Document Author Table

	Name	Function	Date	Comments
Prepared by:	H. Lewis	GRAS SAF Project Team	4 February 2008	
Reviewed by:	D. Offiler	GRAS SAF Project Team	4 February 2008	
Approved by:	K.B. Lauritsen	GRAS SAF Project Manager	17 October 2008	

Document Change Record

Issue/Revision	Date	By	Description
Version 1.0	4 Feb 2008	HL	First version

GRAS SAF Project

The GRAS SAF is a EUMETSAT-funded project responsible for operational processing of GRAS radio occultation data from the Metop satellites. The GRAS SAF delivers bending angle, refractivity, temperature, pressure, and humidity profiles in near-real time and offline for NWP and climate users. The offline profiles are further processed into climate products consisting of gridded monthly zonal means of bending angle, refractivity, temperature, humidity, and geopotential heights together with error descriptions.

The GRAS SAF also maintains the Radio Occultation Processing Package (ROPP) which contains software modules that will aid users wishing to process, quality-control and assimilate radio occultation data from any radio occultation mission into NWP and other models.

The GRAS SAF Leading Entity is the Danish Meteorological Institute (DMI), with Cooperating Entities: i) European Centre for Medium-Range Weather Forecasts (ECMWF) in Reading, United Kingdom, ii) Institut D'Estudis Espacials de Catalunya (IEEC) in Barcelona, Spain, and iii) Met Office in Exeter, United Kingdom. To get access to our products or to read more about the project please go to <http://www.grassaf.org>.

1 Background

The ROPP 1dVar retrieval aims to provide profiles of pressure, temperature and humidity using the refractivity or bending angle profile measured from a GPS radio occultation. This is achieved through the minimisation of a quadratic cost function J . This routine calculates

$$J = \frac{1}{2} \langle x - x_b | B^{-1} | x - x_b \rangle + \frac{1}{2} \langle y_o - H(x) | O^{-1} | y_o - H(x) \rangle \quad (1.1)$$

where the background state x_b is given by the state vector, y_o is the observation vector, H is the forward model and B and O are the background and the observation covariance matrices respectively.

In ROPP v1.0 the minimisation of J is performed in the subroutine `ropp_1dvar_solve` by calling an external minimiser. In ROPP v1.0 this is conducted using `M1QN3`, one of the INRIA limited memory Quasi-Newton codes designed to solve large-scale unconstrained minimisation problems (1). A new ROPP-specific minimiser `minROPP` has recently been developed to replicate the performance of `M1QN3`, but avoid any potential licencing issues arising from the use of `M1QN3`. Further details of the `minROPP` and `M1QN3` minimiser algorithms is provided by (2). The tests conducted by (2) demonstrated that the performance of `minROPP` is equivalent to `M1QN3`. The cost function itself is calculated by the `ropp_1dvar_cost` subroutine called from `ropp_1dvar_solve` before proceeding to the minimisation problem.

The ROPP software has been recently developed to enable Met Office background data to be processed in addition to ECMWF data. As part of this work, tests have shown significant differences between the convergence rates obtained using `minROPP` compared with that achieved by the operational Met Office 1dVar code. The latter uses the Levenberg-Marquardt method (3)(4) to solve the 1dVar minimisation problem. A new subroutine has been written to implement the Levenberg-Marquardt method (`ropp_1dvar_LevMarq`) within ROPP. This document provides an overview of `LevMarq` and demonstrates its performance compared with `minROPP`.

2 Levenberg-Marquardt method

The Levenberg-Marquardt method, developed by (4) from an earlier suggestion by (3), is a standard routine for solving a nonlinear least-squares problem. It is implemented in the operational Observation Processing System at the Met Office (OPS) to solve the 1dVar minimisation problem for example.

The minimum of the cost function J is found by the iterative solution of

$$J'' \times (x^{n+1} - x^n) = -J' \quad (2.1)$$

where x^n and x^{n+1} are the n and $(n+1)$ approximations of the state vector x and J'' and J' are the second and first derivative of the cost function with respect to x .

From Equation 1.1,

$$J' = B^{-1}(x^n - x_b) - H'(x^n)^T O^{-1}(y_o - H(x^n)) \quad (2.2)$$

and

$$J'' = B^{-1} + H'(x^n)^T O^{-1} H'(x^n) \quad (2.3)$$

where H' is the derivative of the forward model with respect to the state vector.

In the Levenberg-Marquardt method the diagonal elements of J'' are modified on each iteration towards a solution as

$$J'' = J''(1 + \lambda) \quad (2.4)$$

where λ is a positive scalar value. The Levenberg-Marquardt method therefore continues as follows for each iteration n until convergence is reached.

- Compute cost function J using Equation 1.1 for the current value of x .
- Compute the gradient of the forward model $H'(x^n)$ using routine `ropp_fm_refrac_1d_tl`.
- Compute J' using Equation 2.2 for the current value of x .
- Compute J'' using Equation 2.3 for the current value of x .
- Modify diagonal elements of J'' by scaling with $(1 + \lambda)$, where $\lambda = 10^{-4}$ on the first iteration.
- Solve Equation 2.1 to obtain $dx = (x^{n+1} - x^n)$ and the next value for $x = x^{n+1}$
- Recompute the cost function J using the new state vector x^{n+1}
- If J has reduced from the previous iteration then reduce λ by a factor of 10 and recompute $H'(x^n)$, J' and J'' for the new state vector x^n .
- If J has increased from the previous iteration then increase λ by a factor of 10 and solve Equation 2.1 again to obtain a new estimate of x^{n+1} .

The iteration continues until the convergence criteria are satisfied. In ROPP, convergence is assumed if either the state vector does not change by more than a set value between iterations

$$\frac{|x^{n+1} - x^n|}{\sqrt{B}} < \text{max_delta_state} \quad (2.5)$$

or the cost function does not change by more than a set value between iterations

$$|J^{n+1} - J^n| < \text{max_delta_J} \quad (2.6)$$

Either of these conditions need to be met for at least `conv_check_n_previous` successive iterations for convergence to be assumed. Parameters `max_delta_state`, `max_delta_J` and `conv_check_n_previous` can be set in a ROPP 1dVar configuration file. Default values of `max_delta_state=0.1`, `max_delta_J=0.1` and `conv_check_n_previous=2` are defined in `ropp_1dvar_types`.

The Levenberg-Marquardt algorithm has been implemented in the ROPP software in the subroutine `ropp_1dvar_levmarq`. This replaces the call to `ropp_1dvar_solve`, used when solving the minimisation problem using `minROPP`.

3 ROPP 1dVar results

The performance of the ROPP 1dVar retrieval using `LevMarq` and `minROPP` minimisers has been compared. The stand-alone tools `ropp_1dvar_bangle` and `ropp_1dvar_refrac` have been applied to retrieve temperature, humidity and pressure profiles from bending angle and refractivity observations.

A selection of observations and corresponding background data have been used for testing. Several different radio occultation scenarios are included in the ROPP test procedure IT-1DVAR-03 together with co-located background temperature, humidity and surface pressure data defined on the ECMWF NWP model grid levels. More recent observed occultations are also considered with the corresponding background pressure and humidity data from the operational Met Office NWP model.

3.1 Synthetic data: IT-1dVAR-03

The IT-1dVAR-03 test data have been generated by modifying the atmospheric parameters corresponding to a given refractivity and bending angle profile to create synthetic background profiles with which to test the 1dVar. Both temperature and humidity profiles are modified with a sine curve while the surface pressure is modified with a bias of 0.5%. Figure 3.1 illustrates the difference between the background data and the 1dVar solution for temperature, humidity and pressure when the 1dVar retrieval is performed relative to refractivity observations using the `minROPP` and `LevMarq` minimisers. The corresponding results for the 1dVar solution using bending angle observations are plotted in Figure 3.2.

Figures 3.1 and 3.2 indicate that the choice of minimiser results in differences in the retrieved temperatures of less than 0.2K, in specific humidity of less than 0.1 g/kg and in pressure of less than 0.5 hPa for this case. Interestingly, these results indicate that the pressure solution deviates from the background by less than 0.02 hPa when solving the 1dVar problem using the `LevMarq` minimiser. This apparent insensitivity of pressure data is to be investigated. Tests using data from the other IT-1DVAR-03 profiles show very similar results.

Tables 3.1 and 3.2 list the cost function values at each iteration towards the 1dVar solution for Profile 1 of IT-1DVAR-03 using `minROPP` and `LevMarq`. Results are listed for the 1dVar retrieval using refractivity and bending angle observations. The maximum relative change in the state vector is also listed for each iteration. This change is computed as a fraction of the background error covariance. Convergence is assumed to occur in the ROPP software when the maximum fractional change detected between iterations or the cost function change between iterations is less than 0.1 for two successive iterations. The processing time for each 1dVar retrieval is listed for each case.

It is clear from Tables 3.1 and 3.2 that while `LevMarq` requires very few iterations to reach a solution, the processing time required to reach that solution is at least 3 times longer than the time taken by `minROPP`. The cost function reaches a value close to that at convergence

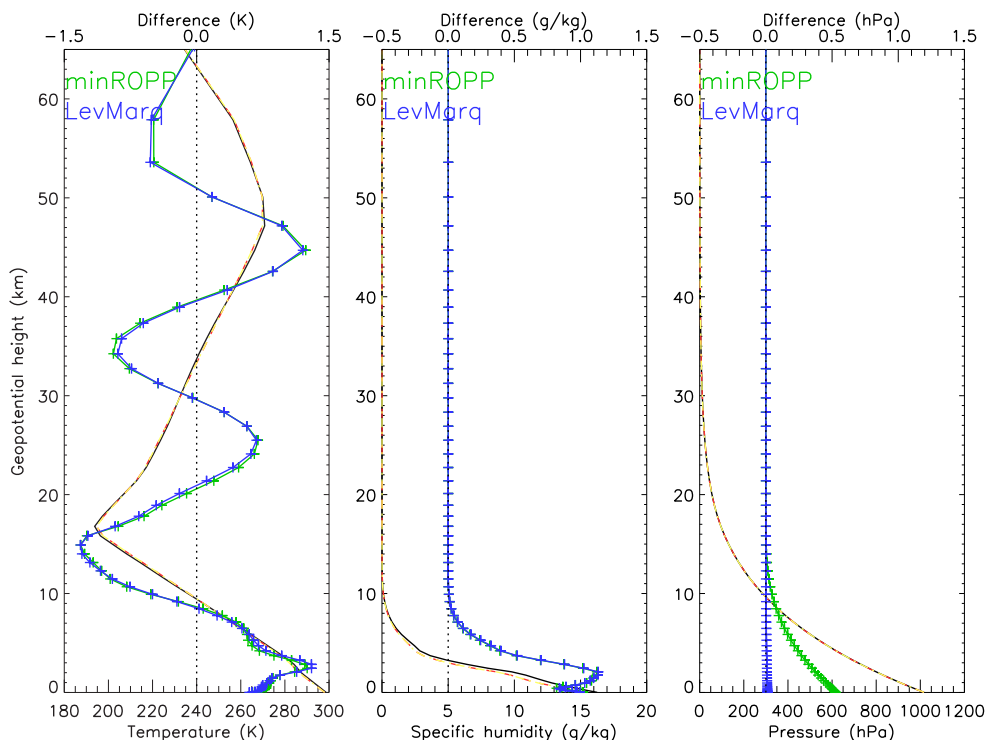


Figure 3.1: Comparison of background (black) and retrieved (red minROPP, yellow LevMarq) profiles of temperature, humidity and pressure for refractivity observations of Profile 1 in the ROPP 1dVar module test IT-1DVAR-03. The difference in profiles (background - solution) is plotted in green for 1dVar using minROPP and in blue for LevMarq minimiser.

after just one iteration using LevMarq while the minROPP minimiser results in a far more gradual change towards convergence. The significant additional processing time taken by LevMarq can be attributed to the need to calculate all elements of $H'(x^n)$ each iteration, requiring $nstate$ calls to the tangent linear calculations in the ROPP software on each iteration. In contrast, implementation of minROPP only requires a single call to the adjoint calculations each iteration. Although further developments may lead to a more efficient implementation of LevMarq, such as using the convergence criteria implemented in the Met Office operational OPS, this analysis indicates that the minROPP algorithm provides the more efficient minimiser for the 1dVar problem.

3.2 COSMIC occultation data and Met Office background data

The ROPP 1dVar software has also been applied to recent COSMIC occultation data and co-located background data derived from the Met Office operational NWP model. This provides a more useful test of the minROPP and LevMarq minimisers, applied to real problems where the differences between the observed and background temperature or humidity fields are not smoothly varying functions for example.

Figures 3.3 and 3.4 show the background and 1dVar temperature, humidity and pressure profiles for a sample COSMIC occultation (200801260405,CO01UCAR). As in Figure 3.1, the difference between background and 1dVar solution is shown for the 1dVar performed

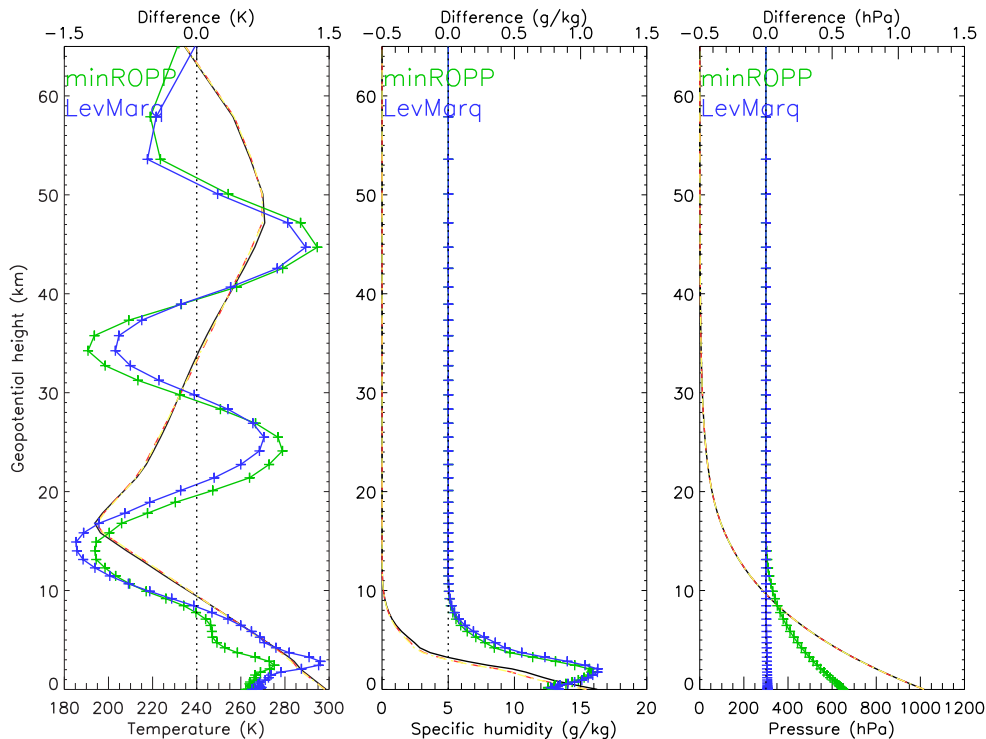


Figure 3.2: As in Figure 3.1 for 1dVar using bending angle observations.

using the `minROPP` and `LevMarq` minimisers. Note that results using ROPP with `LevMarq` are identical to the solutions obtained by the operational 1dVar in the Met Office Observation Processing System (OPS), which uses a Levenberg-Marquardt algorithm, for the corresponding profiles. Tables 3.3 and 3.4 list the cost function values at each iteration for this case.

The temperature, humidity and pressure profiles in Figure 3.3 all indicate that the 1dVar solution relative to the refractivity observations deviates more from the background profiles when using `LevMarq` than for `minROPP`. This suggests that the convergence criteria have been reached using `minROPP` before the best 1dVar solution has been obtained in this case. This is evident by comparing the cost function values on convergence listed in Table 3.3. While `minROPP` is clearly more efficient than `LevMarq`, these tests suggest that `LevMarq` can provide a better solution (greater deviation from the background and a lower cost function value on convergence) for a given set of convergence criteria.

Further examples showing the 1dVar solution obtained using `minROPP` being closer to the original background than the `LevMarq` results are provided in Figures 3.5 to 3.10. Other cases, including the bending angle 1dVar solutions in Figure 3.4, provide examples where the `minROPP` and `LevMarq` solutions agree very closely. This indicates that `minROPP` provides similar results to `LevMarq` when the solution has reached an adequate convergence.

This evidence suggests that users may wish to implement ROPP using the `LevMarq` minimisation algorithm rather than `minROPP`. Alternatively, more appropriate default convergence criteria need to be assigned to ensure that the efficient `minROPP` algorithm can be used to obtain the best possible solution to the 1dVar problem.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	68.425	-	1	68.425	-
2	55.937	.18441	2	9.4923	1.5649
3	21.387	.95170	3	9.4517	.13663E-01
4	16.274	.25741	4	9.4432	.23069E-02
5	10.738	.48096			
6	9.1313	.26927			
7	8.7589	.68130E-01			
8	8.5903	.47234E-01			
CPU time taken: 0.310 s			CPU time taken: 0.931 s		

Table 3.1: Comparison of cost function values at each iteration applying the ROPP 1dVar to refractivity observations and Profile 1 in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	227.16	-	1	227.16	-
2	118.36	.43317	2	12.717	1.6895
3	69.583	.48928	3	10.942	.68452E-01
4	55.905	.21387	4	10.888	.68884E-02
5	46.414	.16197			
6	36.617	.21414			
7	31.084	.19231			
8	26.452	.13777			
9	22.927	.12356			
10	16.780	.27773			
11	17.149	.19858			
12	14.201	.90892E-01			
13	13.834	.18152E-01			
CPU time taken: 0.473 s			CPU time taken: 1.901 s		

Table 3.2: Comparison of cost function values at each iteration applying the ROPP 1dVar to bending angle observations and Profile 1 in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

3.3 minROPP convergence criteria

Comparison of the 1dVar solutions generated using the minROPP and LevMarq has highlighted cases where the solution obtained using minROPP when the default ROPP convergence criteria have been met differs considerably from the LevMarq solution. In all

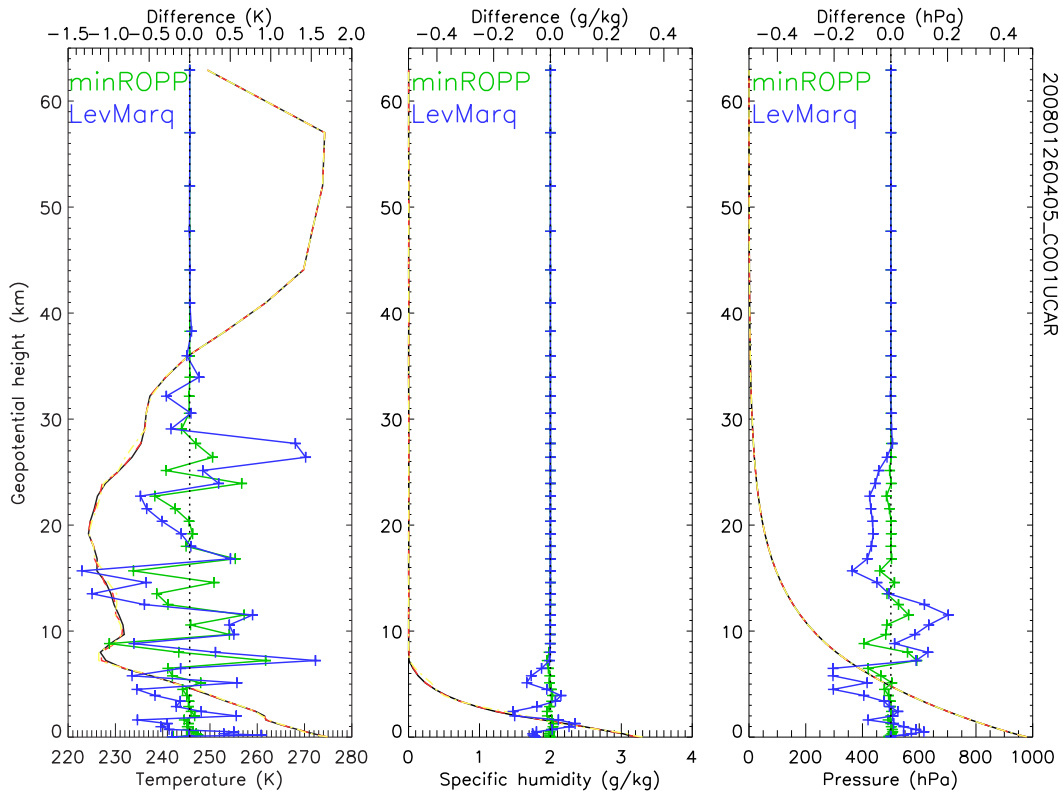


Figure 3.3: Comparison of background (black) and retrieved (red minROPP, yellow LevMarq) profiles of temperature, humidity and pressure for refractivity observations from a COSMIC occultation (200801260405,CO01UCAR) and Met Office background data using the (a) minROPP and (b) LevMarq minimisers. The difference in profiles (background - solution) is plotted in green for 1dVar using bending angle observations and in blue for refractivity observation.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	16.720	-	1	16.720	-
2	14.180	.27709E-01	2	9.2503	.33998
3	11.364	.69604E-01	3	9.3085	.13906
			4	9.3085	.13906
CPU time taken: 0.149 s			CPU time taken: 0.596 s		

Table 3.3: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC refractivity observations (200801260405,CO01UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

such cases the cost function J is smaller for LevMarq than minROPP, suggesting that the minROPP solution is not robust. Figure 3.9 shows an example of a particularly poor agreement between the minROPP and LevMarq solutions.

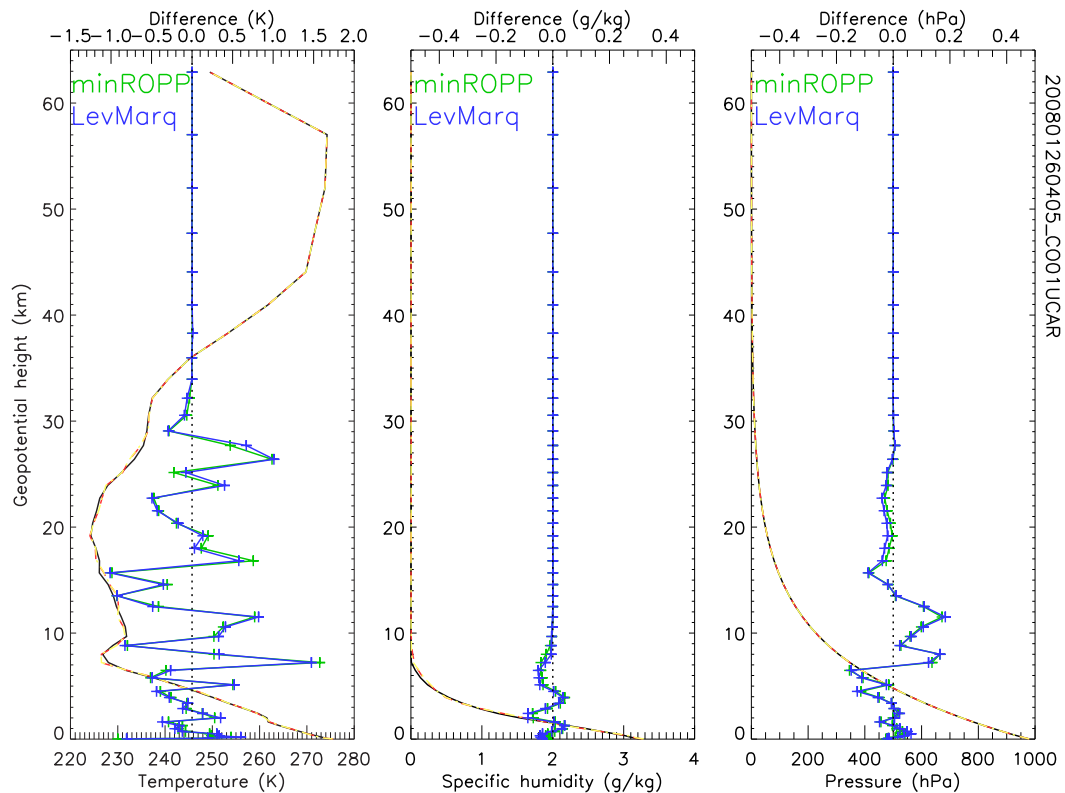


Figure 3.4: As in Figure 3.3 for 1dVar using bending angle observations.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	3.7215	-	1	3.7215	-
2	3.1823	.64210E-01	2	2.7533	.22247
3	2.8325	.10669	3	2.7629	.59030-01
4	2.7673	.48660E-01	4	2.7497	.29048-01
5	2.7371	.42951E-01			
CPU time taken: 0.195 s			CPU time taken: 1.389 s		

Table 3.4: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC bending angle observations (200801260405,CO01UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

Figure 3.11 shows the corresponding results when requiring that more strict convergence criteria are met before assuming that convergence has been achieved. In this case the convergence parameters `max_delta_state` and `max_delta_J` are both reduced from their default values of 0.1 to 0.01. Table 3.11 lists the cost function values at each iteration. The LevMarq solution still converges after 4 iterations using the new convergence criteria, so that the solution is unaltered from that in Figure 3.9.

Agreement between the minROPP and LevMarq profiles in Figure 3.11 is better than that

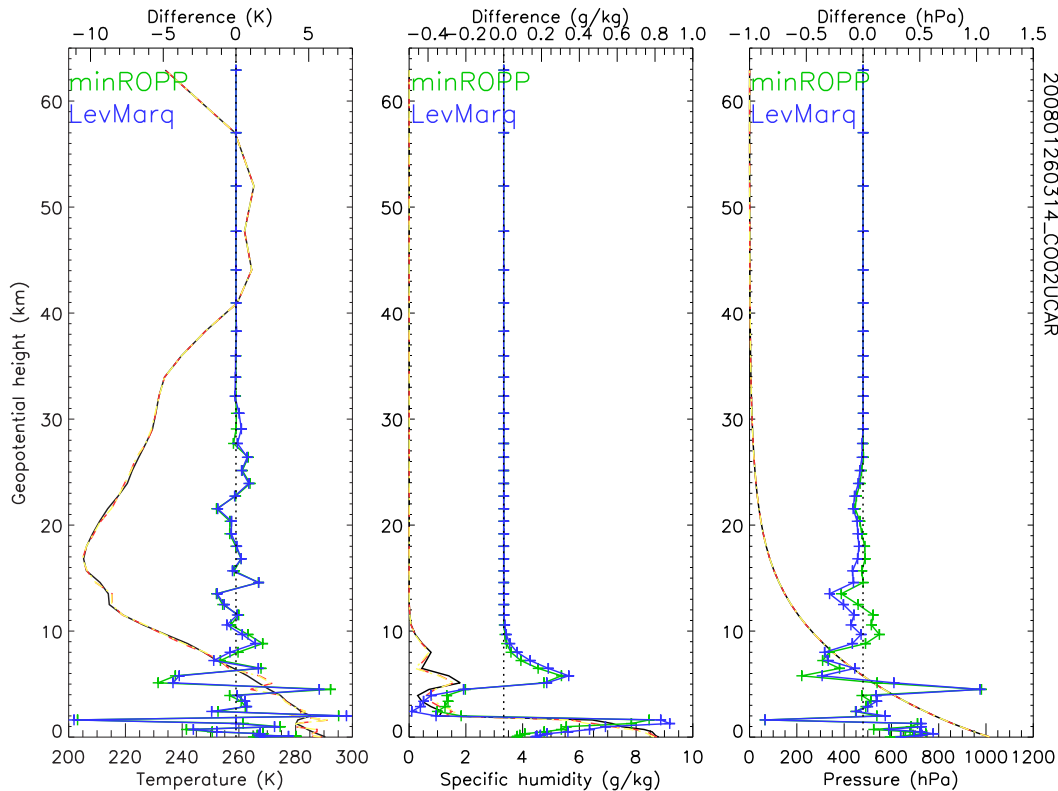


Figure 3.5: Comparison of background (black) and retrieved (red minROPP, yellow LevMarq) profiles of temperature, humidity and pressure for refractivity observations from a COSMIC occultation (200801260314,CO02UCAR) and Met Office background data using the (a) minROPP and (b) LevMarq minimisers. The difference in profiles (background - solution) is plotted in green for 1dVar using bending angle observations and in blue for refractivity observation.

in Figure 3.9 using the default ROPP convergence criteria. This can be quantified by comparing the standard deviation of differences between the two solutions. The standard deviation for temperature, humidity and pressure are all reduced by applying stricter convergence criteria from 1.06 to 0.59 K, 0.023 to 0.020 g/kg and from 0.18 to 0.12 hPa respectively. The maximum difference between the minROPP and LevMarq temperature, humidity and pressure profiles is also reduced from 2.2 to 1.2 K, 0.084 to 0.071 g/kg and from 0.63 to 0.33 hPa respectively by applying the new convergence criteria. Very similar results are obtained if applying the convergence criteria implemented in the Met Office OPS (17 iterations, $J=23.865$).

The minROPP solution approaches LevMarq very closely if parameters `max_delta_state` and `max_delta_J` are both reduced further to 0.001 (standard deviations 0.08 K, 0.0050 g/kg, 0.012 hPa). The final cost function value of 22.699 obtained in this case still exceeds the LevMarq value, and was reached after 65 iterations (2.3 s processing time). This is clearly not a viable approach for use in operational software.

This analysis suggests that the ROPP 1dVar solution can be improved by using default convergence criteria parameters of `max_delta_state=0.01` and `max_delta_J=0.01`. This option provides the best compromise between quantitative accuracy and a reasonable number of iterations and processing time required to obtain convergence.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	137.46	-	1	137.46	-
2	117.15	.99471E-01	2	56.552	1.1418
3	96.023	.23972	3	56.267	.30433
4	82.505	.25211	4	56.266	.10077E-01
5	71.260	.24048		56.266	.91790E-03
6	62.549	.25159			
7	59.320	.12409			
8	59.087	.10126			
9	57.993	.46692E-01			
10	57.817	.12896E-01			
CPU time taken: 0.354 s			CPU time taken: 1.375 s		

Table 3.5: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC refractivity observations (200801260314,CO2UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	49.149	-	1	49.149	-
2	40.752	.12353	2	17.179	.85548
3	25.815	.40339	3	13.635	.65206
4	18.048	.44401	4	13.515	.97565E-01
5	17.950	.23504	5	13.462	.94660E-01
6	16.620	.19844			
7	15.591	.44387			
8	15.400	.18712			
9	14.728	.71108E-01			
10	13.390	.14145			
11	13.883	.19429			
12	13.419	.15925			
13	13.418	.35035E-03			
14	13.415	.17342E-02			
CPU time taken: 0.411 s			CPU time taken: 1.764 s		

Table 3.6: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC bending angle observations (200801260314,CO2UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

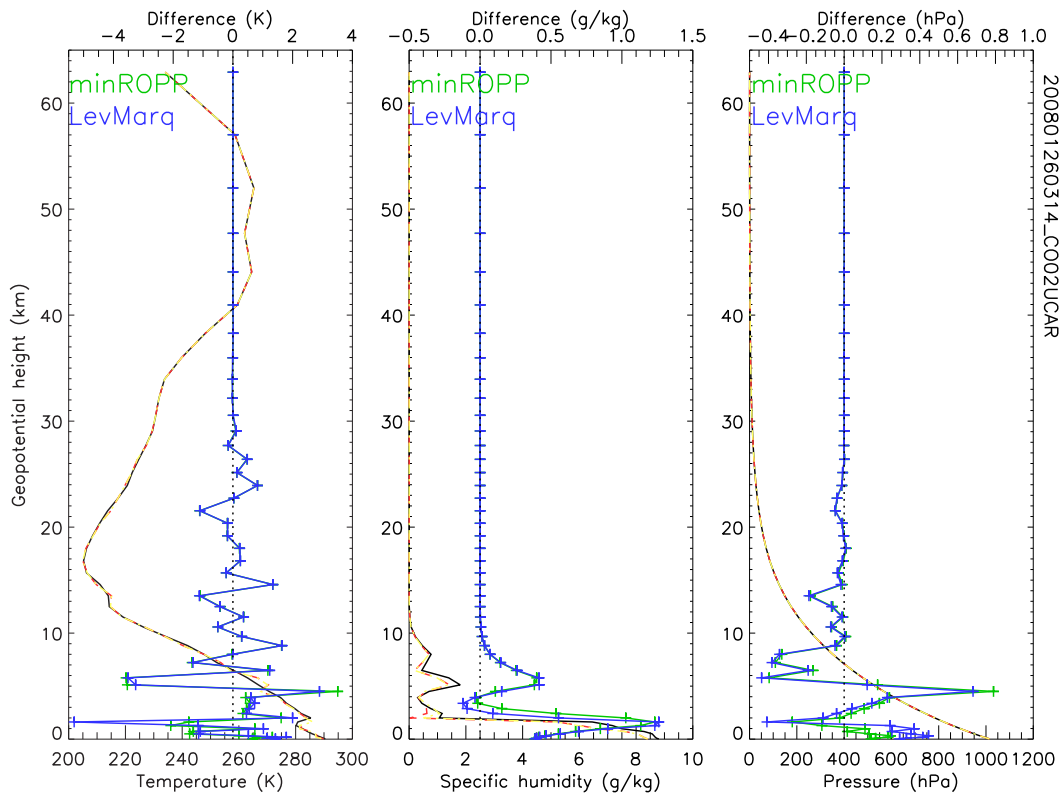


Figure 3.6: As in Figure 3.5 for 1dVar using bending angle observations.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	63.602	-	1	63.602	-
2	52.848	.63582E-01	2	35.514	.92236
3	39.392	.20787	3	35.513	.10659E-01
4	37.794	.91648E-01	4	35.513	.53296E-02
5	36.972	.72245E-01			
CPU time taken: 0.214 s			CPU time taken: 1.046 s		

Table 3.7: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC refractivity observations (200801260323,CO04UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

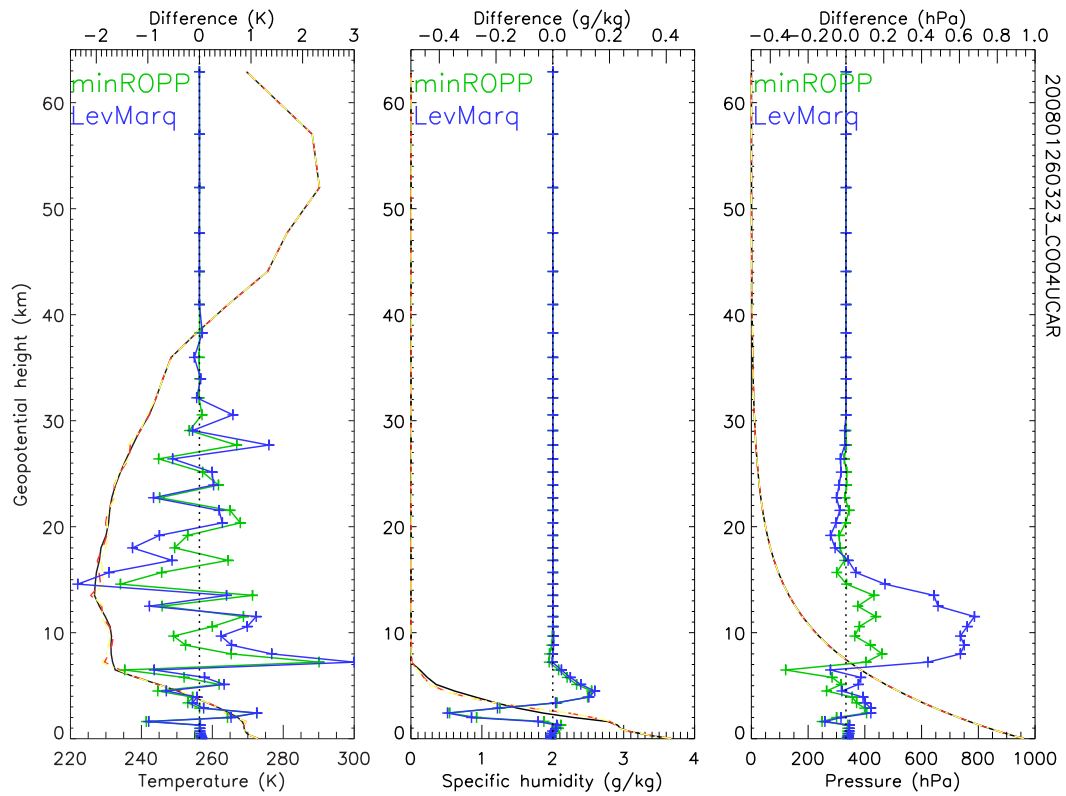


Figure 3.7: Comparison of background (black) and retrieved (red minROPP, yellow LevMarq) profiles of temperature, humidity and pressure for refractivity observations from a COSMIC occultation (200801260323,CO04UCAR) and Met Office background data using the (a) minROPP and (b) LevMarq minimisers. The difference in profiles (background - solution) is plotted in green for 1dVar using bending angle observations and in blue for refractivity observation.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	9.6230	-	1	9.6320	-
2	7.9671	.10183	2	5.6202	.61098
3	5.8030	.30020	3	5.6195	.18434E-01
4	5.6978	.10181	4	5.6203	.25676E-01
5	5.6465	.56619E-01			
6	5.6178	.72095E-01			
CPU time taken: 0.218 s			CPU time taken: 1.356 s		

Table 3.8: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC bending angle observations (200801260323,CO04UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

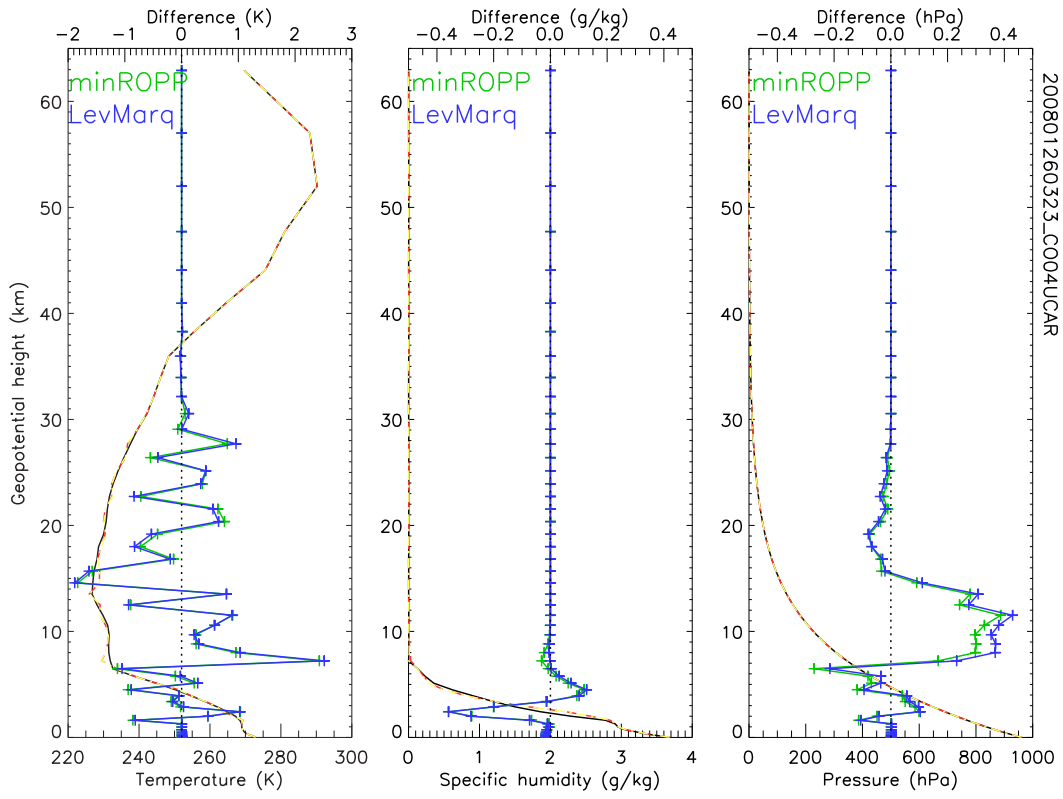


Figure 3.8: As in Figure 3.7 for 1dVar using bending angle observations.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	37.204	-	1	37.204	-
2	36.772	.88225E-01	2	22.681	.63974
3	34.042	.33127E-01	3	22.681	.57482E-02
			4	22.681	.23299E-04
CPU time taken: 0.147 s			CPU time taken: 1.094 s		

Table 3.9: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC refractivity observations (200801260306,CO05UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

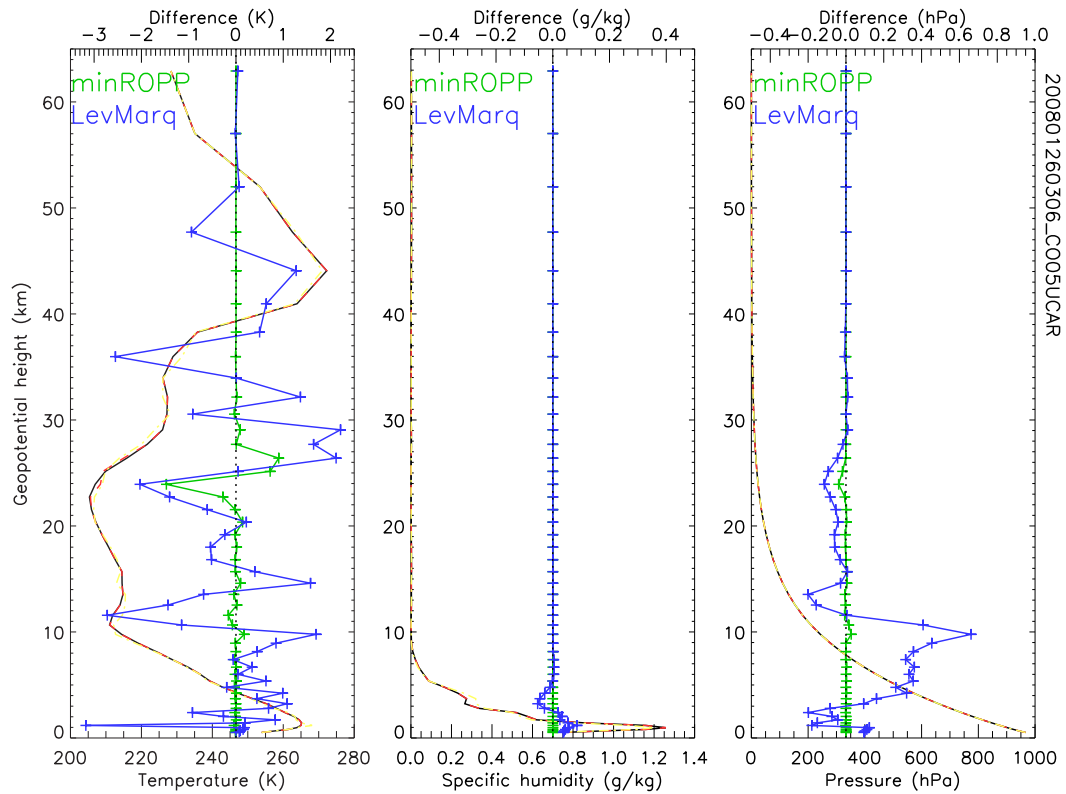


Figure 3.9: Comparison of background (black) and retrieved (red minROPP, yellow LevMarq) profiles of temperature, humidity and pressure for refractivity observations from a COSMIC occultation (200801260306,CO05UCAR) and Met Office background data using the (a) minROPP and (b) LevMarq minimisers. The difference in profiles (background - solution) is plotted in green for 1dVar using bending angle observations and in blue for refractivity observation.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	10.720	-	1	10.720	-
2	10.882	.17201	2	8.4790	.44415
3	9.9699	.85004E-01	3	8.4782	.17580E-01
4	9.8369	.13455E-01	4	8.4782	.18517E-02
CPU time taken: 0.170 s			CPU time taken: 1.391 s		

Table 3.10: Comparison of cost function values at each iteration applying the ROPP 1dVar to COSMIC bending angle observations (200801260306,CO05UCAR) and Met Office background data in the ROPP 1dVar module test IT-1DVAR-03 using the minROPP and LevMarq minimisers.

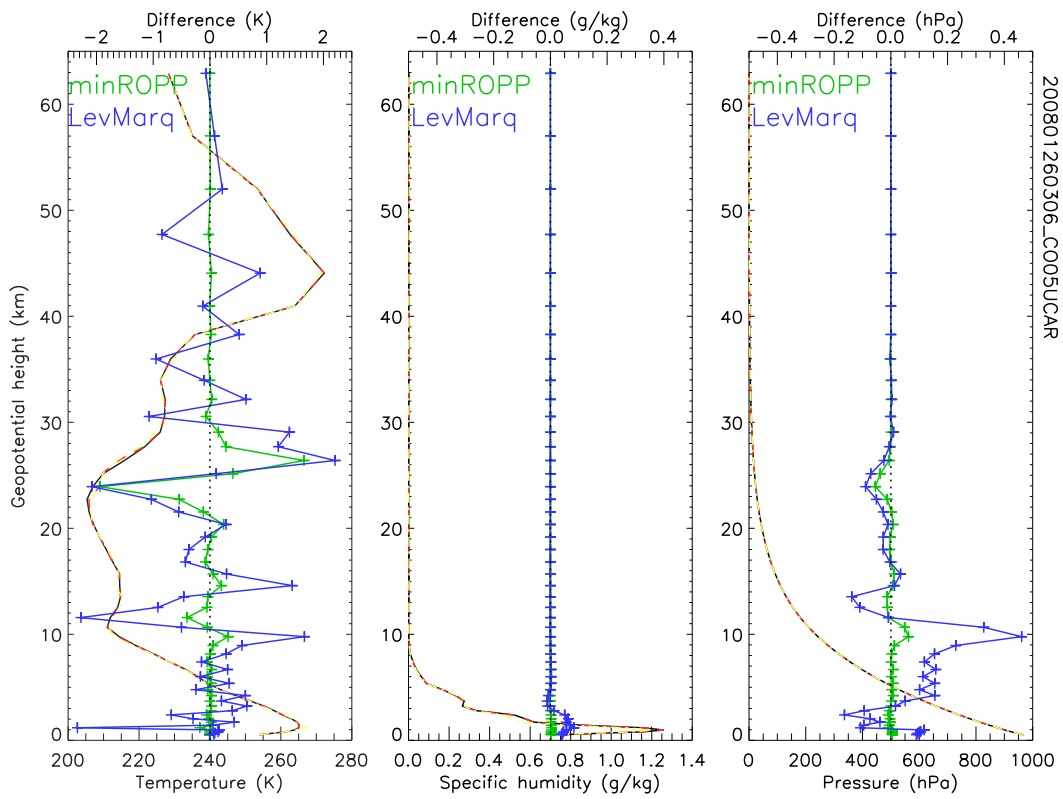


Figure 3.10: As in Figure 3.9 for 1dVar using bending angle observations.

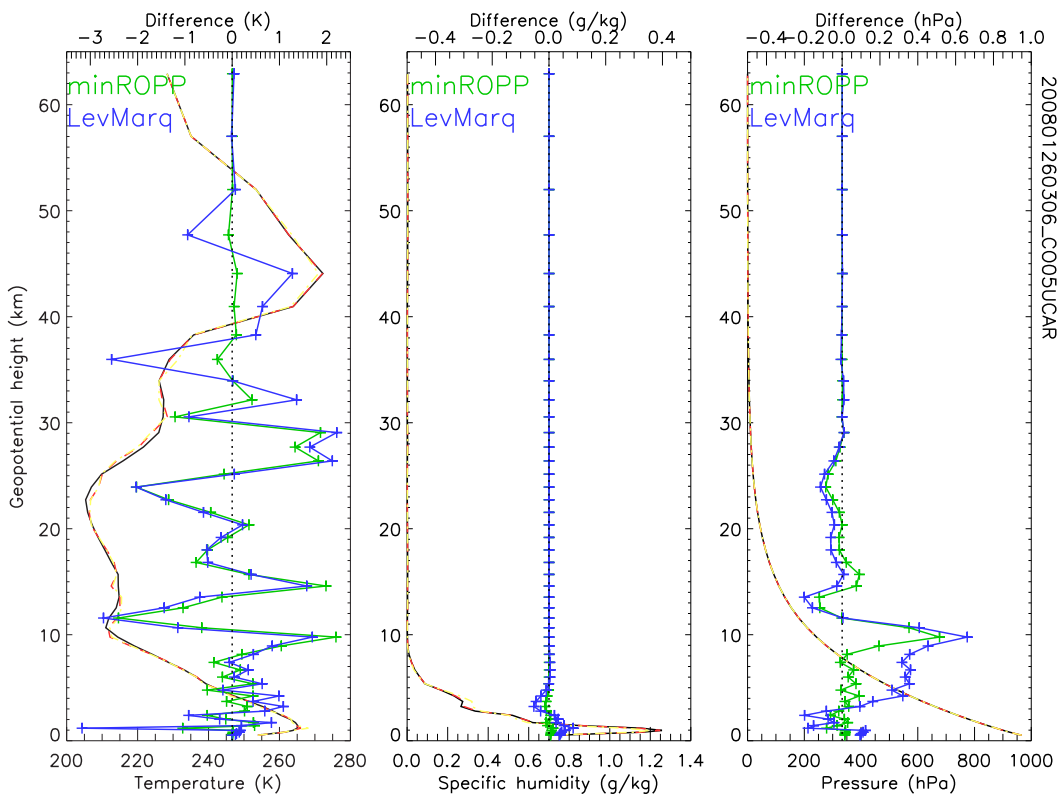


Figure 3.11: As in Figure 3.9 for 1dVar using ROPP convergence criteria with $\text{max_delta_state}=0.01$ and $\text{max_delta_J} = 0.01$.

minROPP			LevMarq		
n_iter	J	max. relative change state	n_iter	J	max. relative change state
1	37.204	-	1	37.204	-
2	36.772	.88225E-01	2	22.681	.63974
3	34.042	.33127E-01	3	22.681	.57482E-02
4	33.336	.82770E-02	4	22.681	.23299E-04
5	32.163	.26963E-01			
6	29.662	.77014E-01			
7	27.757	.69551E-01			
8	27.454	.12573E-01			
9	25.577	.16622			
10	24.944	.42325E-01			
11	24.428	.67347E-01			
12	24.212	.26251E-01			
13	24.154	.29177E-02			
14	24.067	.85837E-02			
CPU time taken: 0.486 s			CPU time taken: 1.094 s		

Table 3.11: As in Table 3.9 for 1dVar using ROPP convergence criteria with $\text{max_delta_state}=0.01$ and $\text{max_delta_J} = 0.01$.

4 Summary

These tests allow a comparison between ROPP 1dVar solutions obtained using the standard ROPP minimiser `minROPP` and a Levenberg-Marquardt algorithm `LevMarq`. All tests demonstrate that `minROPP` takes typically 3 times less processing time to converge to a solution than `LevMarq` for a given set of convergence criteria. This is because `LevMarq` requires all elements of the forward model gradient matrix to be computed on each iteration.

While several cases show the solutions obtained using `minROPP` and `LevMarq` to match very closely, a significant proportion of the cases tested show large differences. In all such cases, the `minROPP` solution is apparently not as robust as that obtained using `LevMarq`. The `minROPP` cost function at convergence is larger than for `LevMarq` and the temperature, humidity and pressure solutions deviate more from the background profiles for `LevMarq`. Tests have indicated that ROPP 1dVar solutions can be improved by using default values for the ROPP convergence criteria parameters of `max_delta_state=0.01` and `max_delta_J=0.01` without requiring excessive additional processing. The `LevMarq` algorithm is generally considered to produce more reliable solutions than `minROPP` however. Note that this conclusion also holds for comparisons between `LevMarq` and the third-party `M1QN3` algorithm previously implemented in ROPP.

This analysis suggests that the `minROPP` and `LevMarq` minimisers both have their strengths. It is therefore recommended that both are provided and supported as part of ROPP. Users requiring a very efficient algorithm should use `minROPP`, tuning the convergence parameters as required. Users requiring very robust 1dVar solutions may prefer to implement `LevMarq` at the cost of additional processing to reach that solution. Some further development work should be conducted to optimise the `LevMarq` algorithm before it is provided to users.

Bibliography

- [1] Gilbert, J. C. and Lemaréchal, C. M1QN3. A solver of large-scale unconstrained minimization problems, <http://www-rocq.inria.fr/estime/modulopt/optimization-routines/m1qn3/m1qn3.html>.
- [2] Lewis, H. (2007) ROPP minimiser - minROPP. GRAS SAF Report Series www.grassaf.org
- [3] Levenberg, K. (1944) A method for the solution of certain problems in least squares. *Quarterly of Applied Mathematics* **2** 164-168
- [4] Marquardt, D. (1963) An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics* **11** 431-441

GRAS SAF Reports

SAF/GRAS/METO/REP/GSR/001	Mono-dimensional thinning for GPS Radio Occultation
SAF/GRAS/METO/REP/GSR/002	Geodesy calculations in ROPP
SAF/GRAS/METO/REP/GSR/003	ROPP minimiser - minROPP
SAF/GRAS/METO/REP/GSR/004	Error function calculation in ROPP
SAF/GRAS/METO/REP/GSR/005	Refractivity calculations in ROPP
SAF/GRAS/METO/REP/GSR/006	Levenberg-Marquardt minimisation in ROPP
SAF/GRAS/METO/REP/GSR/007	Abel integral calculations in ROPP
SAF/GRAS/METO/REP/GSR/008	ROPP thinner algorithm
SAF/GRAS/METO/REP/GSR/009	Refractivity coefficients used in the assimilation of GPS radio occultation measurements
SAF/GRAS/METO/REP/GSR/010	Latitudinal Binning and Area-Weighted Averaging of Irregularly Distributed Radio Occultation Data
SAF/GRAS/METO/REP/GSR/011	ROPP 1dVar validation

GRAS SAF Reports are accessible via the GRAS SAF website <http://www.grassaf.org>.