# ALMA Memo No. 568

# Optimization of the IF Filters for the ALMA Water Vapour Radiometers

Richard Hills
Cavendish Laboratory, Cambridge
6th August 2007

## Abstract

The specifications of the IF filters to be used in the production radiometers are derived in a way which should provide optimum performance over a range of conditions. In addition to finding the best set of centre frequencies and bandwidths, the tolerances that can be allowed are determined, taking account of the fact that the quantities that will be used for making the phase corrections are the differences between the readings of pairs of radiometers. Other key specifications of the radiometers, such the allowable range of sideband ratios in the 183GHz mixers, are also investigated in this memo.

## 1 Introduction

Thus far the design of the IF system for the ALMA water vapour radiometers has been based on what were essentially qualitative arguments. This was backed by a reasonably thorough analysis (ALMA memo 495) of what performance could be expected in measuring small fluctuations in the path due to water vapour along the line of sight. Before embarking on the full-scale production of the final radiometers it is worth carrying out a fuller investigation where the various parameters are adjusted to find a design which gives a generally optimum overall performance. To do this it was necessary to construct a detailed numerical model of the system and run a suitable optimisation procedure. The same model can then be used to investigate the effects of small variations in the parameters so that tolerances can be set.

## 2 Choice of Frequencies

For reasons discussed in earlier memos (e.g. 303 and 352) the prototype radiometers used four IF filters and for the moment we will continue to consider only designs with <u>four</u> filters. The parameters describing the filters used for the prototypes are given in table 1.

| Number > | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| Centre | 0.88 | 1.94 | 3.18 | 5.20 |
| Width | 0.16 | 0.75 | 1.25 | 2.50 |
| Low | 0.80 | 1.57 | 2.55 | 3.95 |
| High | 0.96 | 2.32 | 3.80 | 6.45 |

Table 1. IF Frequencies (GHz) of the filters in the prototype radiometers.

The radiometers use double-sideband mixers for the RF stage and the local oscillator frequency is set to be close to that of the water vapour line, i.e. 183.31GHz. This means that the frequency coverage at the RF input is as illustrated in the following figure (which was kindly provided by Bojan Nikolic).

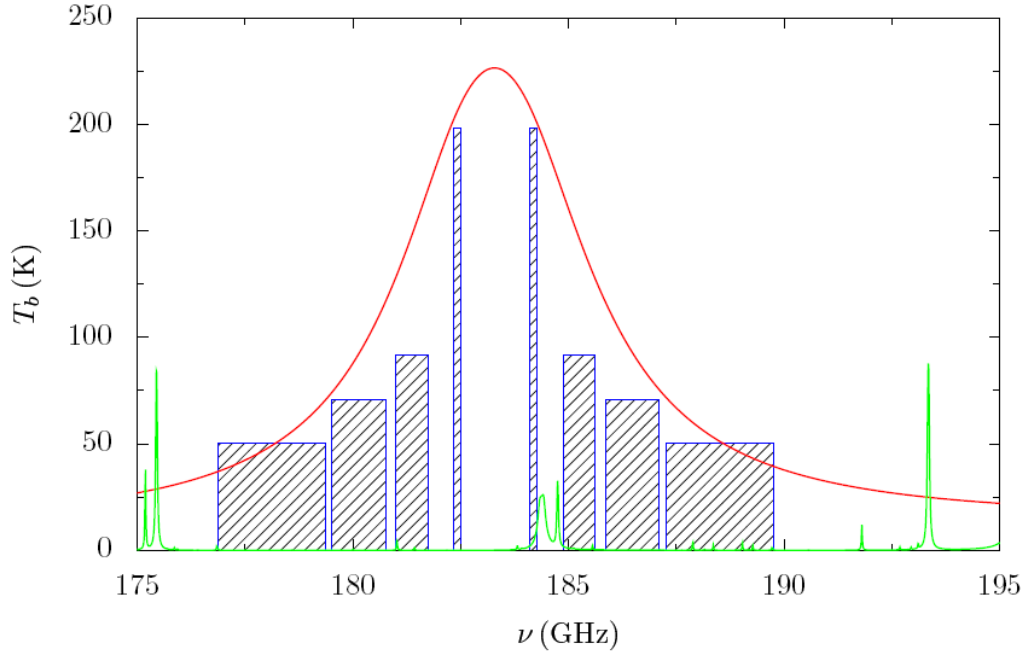


Figure 1. RF coverage produced by the filters in the prototype radiometers.

Here the red line shows the typical form of the 183GHz water vapour line, which is strongly broadened by collisions because most of the water lies at relatively low altitude. The green trace shows emission due to ozone, which is at much higher altitudes. The greatest sensitivity to fluctuations in the amount of water along the line of sight occurs when the opacity is of order unity, i.e where the brightness temperature is ~100K. The logic behind the choice of filters was to provide channels which would match this criterion over a wide range of amounts of water, with wider bandwidths in the weaker parts of the line where more sensitivity would be required. The first channel was made narrow to avoid the ozone feature at ~184.5GHz. This narrow width reduces the sensitivity of this channel so that in practice it is not often useful. Note that, for clarity, the opacity of the Ozone line has been exaggerated in figure 1. The performance of radiometers using this choice of filters was reported in memo 495.

## 2.1 Revised Choices

At the time of the PDR it was recognized that the choice of filters in the prototypes was not ideal, so some work was done to see how much improvement might be on offer with different choices. This was written up in a note on "Revised Noise Estimates" which was circulated internally but not put out as a memo, so for completeness we include that here.

| | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| Centre | 1.20 | 1.90 | 3.20 | 6.00 |
| Width | 0.40 | 0.80 | 1.50 | 3.00 |
| Low | 1.00 | 1.50 | 2.45 | 4.50 |
| High | 1.40 | 2.30 | 3.95 | 7.50 |

Table 2. Frequencies used in making Revised Noise Estimates[1].

[1] The centre frequency of the third channel was mistakenly given as 3.6GHz in the note circulated on this.

A plot of representing this set of filters is given next for comparison with later ones.

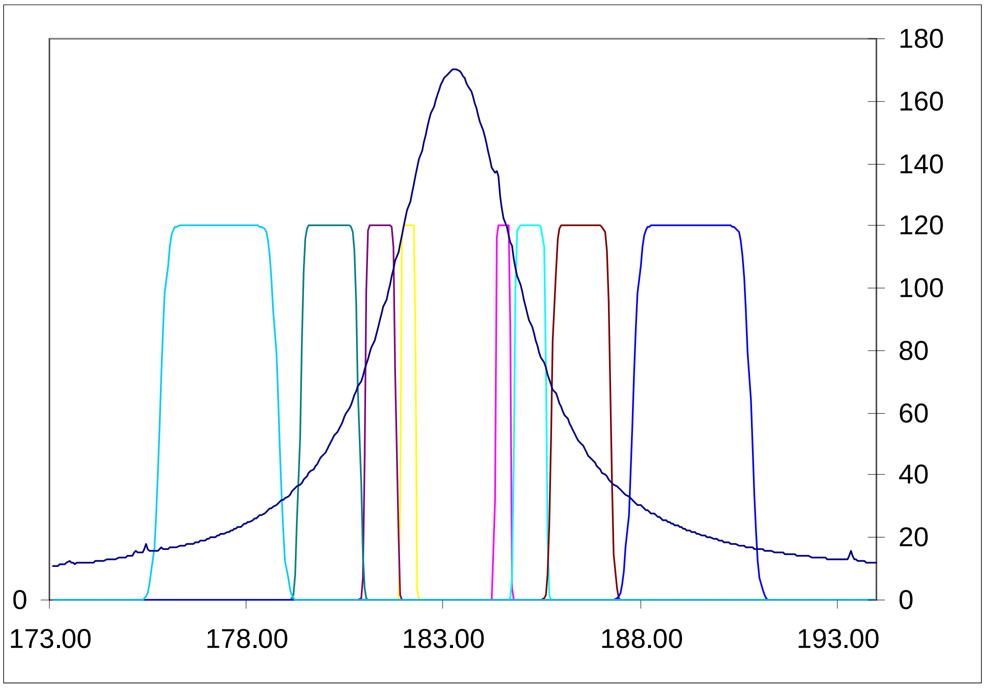


Figure 2. Filter choices used in the memo on revised noise estimates.

The dark blue curve is the total brightness temperature (right hand scale, K) versus frequency in GHz for rather dry conditions: ~0.5mm of water. The Ozone line can only just be seen as a blip to the right of the peak of the main (water) line.

The filters are represented in colour here and a simple function has been used represent their shapes, as discussed in more detail below. Note that the frequencies and widths of the filters have been chosen so that the IF band is more or less all filled and that the widths increase by roughly a factor of two from one filter to the next as one moves away from the line centre. The intention here was again to balance the fact that the water emission is weaker in the outer part of the line by having wider channels and therefore greater sensitivity.

It was found that this choice of frequencies gave considerably better performance than the original choice. It is clear, however, that the process of choosing these is still a rather arbitrary. Ideally we would use the set of channels which gives us optimum performance in the measurement of the path fluctuations. To do this we need a model which describes the performance of the radiometer and allows us to alter the positions of the bands freely.

## 2.2 The New Model

This consists of the following components:

1) A model of the atmosphere which includes the wet and dry components, with Ozone included in the latter. The opacities are taken from outputs provided by Scott Paine's AM model and checked against Juan Pardo's ATM model, which gives very similar results. Instead of doing a full multi-layer computation I have just used two layers, the upper one containing the dry component and the lower the wet one. The amount of water in the lower layer is varied over the range of interest, here 0.2 to 5mm. Obviously the use of only two layers is a rather crude approximation, but the accuracy of the resulting brightness temperatures is good enough for our purposes and it makes the problem tractable in Excel.

2) A representation of the filters. These are described by a "top-hat" function convolved with a narrow Gaussian. In Excel this can be implemented using the "NORMDIST" function.

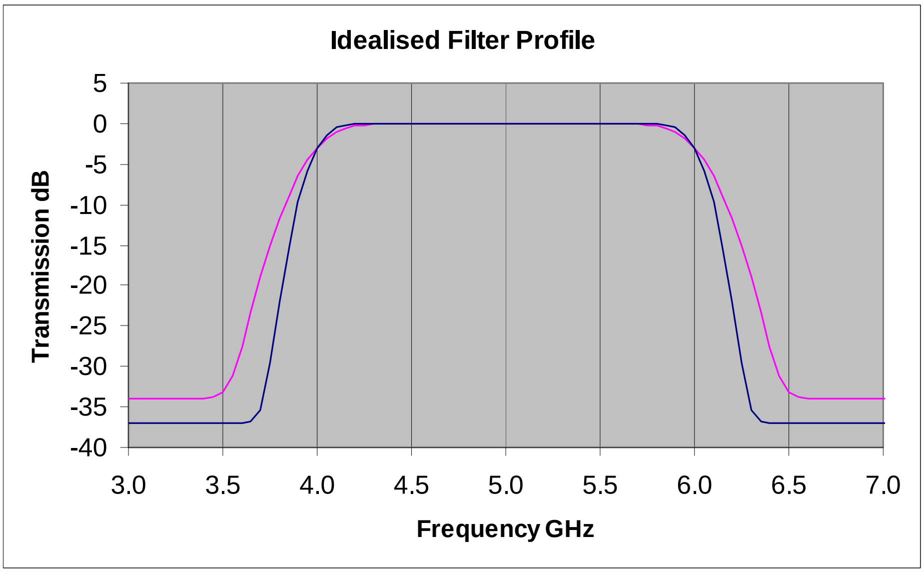


Figure 3. Illustrating the form of filter profile use in the model.

Here the nominal pass-band – the original the top-hat function – is from 4 to 6 GHz, and these become the 3dB points when it is convolved with a Gaussian. The steepness of the sides is defined by the ratio of the 3dB bandwidth to the standard deviation of the Gaussian, which I'll call the "sharpness". This was 25 for the blue curve and 15 for the pink. There is also a "floor" which is included to mimic leakage from outside the band. This was taken to be 0.02% of peak transmission for the blue curve and 0.04% for the pink one. Note that for a sharpness of 15 the 20dB width is 1.31 times the 3dB width, while for 25 it is ~1.20 times it.

3) The average brightness temperatures seen through the different filters are then calculated, initially for the upper and lower sidebands separately and then combined into a single number for each double-sideband channel, allowing for the possibility of a sideband imbalance. This produces a "curve of growth" showing how the brightness temperatures will vary with the amount of water. Here is an example showing the values for the four channels as the water along the line of sight varies over the range 0.2 to 5mm.

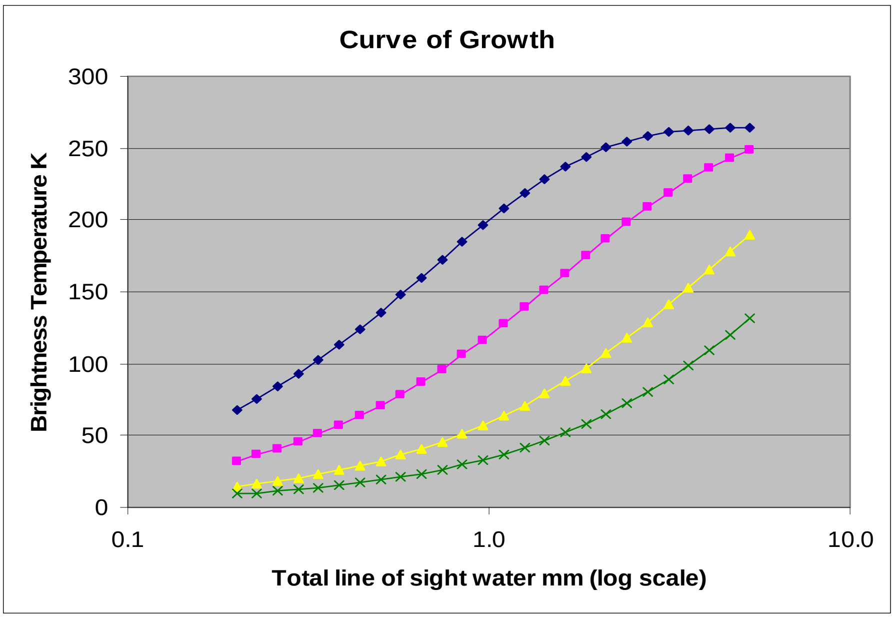


Figure 4. Growth of Brightness Temperature with line of sight water for filters in Table 2.

The changes in these values, as well as the extra path, due to a small change in the amount of water vapour are then calculated. This enables us to estimate the sensitivities of the channels in terms of Kelvins of brightness temperatures variation per millimetre of additional path.

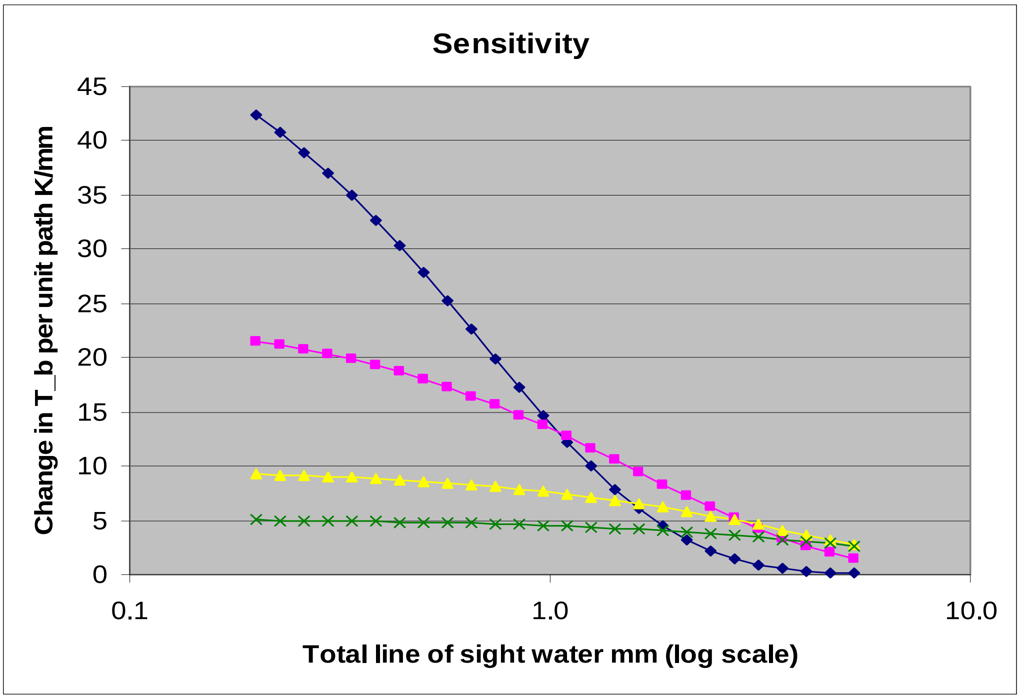


Figure 4. Sensitivity as a function of line of sight water for filters in Table 2.

4) An estimate of the noise in the instrument is produced by another element of the model describing the detailed performance of the radiometer. This takes account of input losses, the noise and gain fluctuations in the RF and IF stages and also post-detection noise. The parameters of this have been adjusted to fit the results obtained during the testing of the prototype radiometers. These include gain fluctuations at the level of 1.5 $10^{-4}$ and post-detection noise equivalent to 0.015K at the input. The results here assume a single-channel Dicke-switched receiver with noise temperature of 1100K, an overall coupling to the sky of 0.82 (which includes an allowance for an RF filter) and a switching efficiency of 0.94.

5) Finally, following the methods outlined in ALMA memo 495, the sensitivities and noise estimates are used to determine the optimum weights which are to be used in combining the data from the four channels. The expected errors in the estimating the path fluctuations due to water are then calculated. Two cases are considered, in the first (referred to as "Ideal") it is assumed that there is no additional source of fluctuation signal other than the variations in water, and in the second ("Cloud") it is supposed that there is also variable broad-band opacity due to thin cloud. In the second case, the weights are chosen such that frequency independent emission will have no effect on the estimate of the path fluctuations, which are assumed to be all due to water vapour. (Emission and scattering from fine droplets would in fact be expected to show a roughly frequency-squared behaviour, but because the receivers are double-sideband this will not seen by the production radiometers so the slope of the continuum has not been included in this model.)

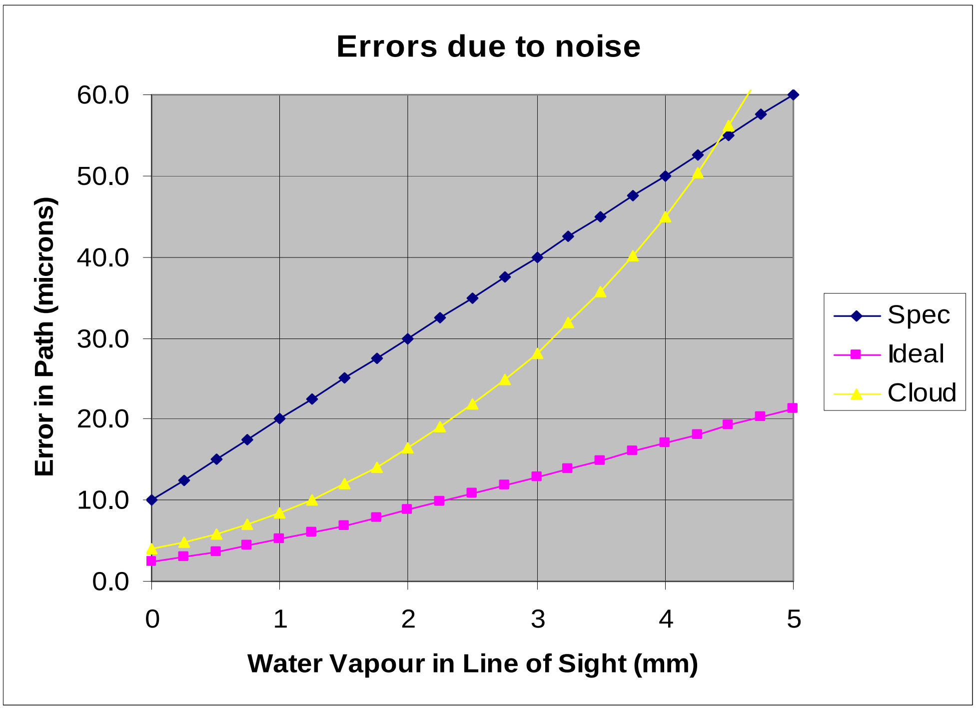


Figure 5. Estimated errors in path using filters in Table 2 and Figure 2.

Note that these results are somewhat different from those given in the note on revised noise estimates, partly because of the small differences in the modelling of the atmosphere and the instrument in the new calculations, but more importantly because I have now corrected a blunder in the calculation of the weights for the "cloud" case. This new calculation shows that at the top end of the range of water we would be slightly out of specification with a single channel Dicke system and this choice of filter frequencies. This adds to the importance of optimising the frequencies and widths of the filters.

## 2.3 Optimum Filter Choice

The first point is that we should be using all of the IF bandwidth that is provided by the IF amplifier. For the moment we take this to be 0.5 to 8GHz. To use all the information available, the filters should cross at roughly their 3dB points. The choice of bands therefore comes down to fixing the three frequencies where these crossover points should occur. To do this we need to decide what quantity to optimize. In order to cover the full range of conditions I have taken the rms value the quantity (error estimated / specification) for 21 values of line of sight water covering the range 0.2 to 5mm in logarithmic steps. One can do the optimisation for either the "ideal" case or the "cloud" case or some combination of the two. The results are as follows:

| | Bottom | Cross 1 | Cross 2 | Cross 3 | Top |
|---|---|---|---|---|---|
| Ideal | 0.50 | **1.87** | **3.52** | **5.52** | 8.00 |
| Cloud | 0.50 | **1.35** | **5.80** | **7.08** | 8.00 |

Table 3. Optimum filter frequencies, assuming that the whole IF frequency range from 0.5 to 8GHz is to be used, for the cases where only water vapour contributes (ideal) and where a thin continuum needs to be rejected (cloud).

Here are plots to show how the resulting filters look, plotted with the models of the atmospheric emission with 0.5mm (black) and 4mm (blue) of water in the line of sight.

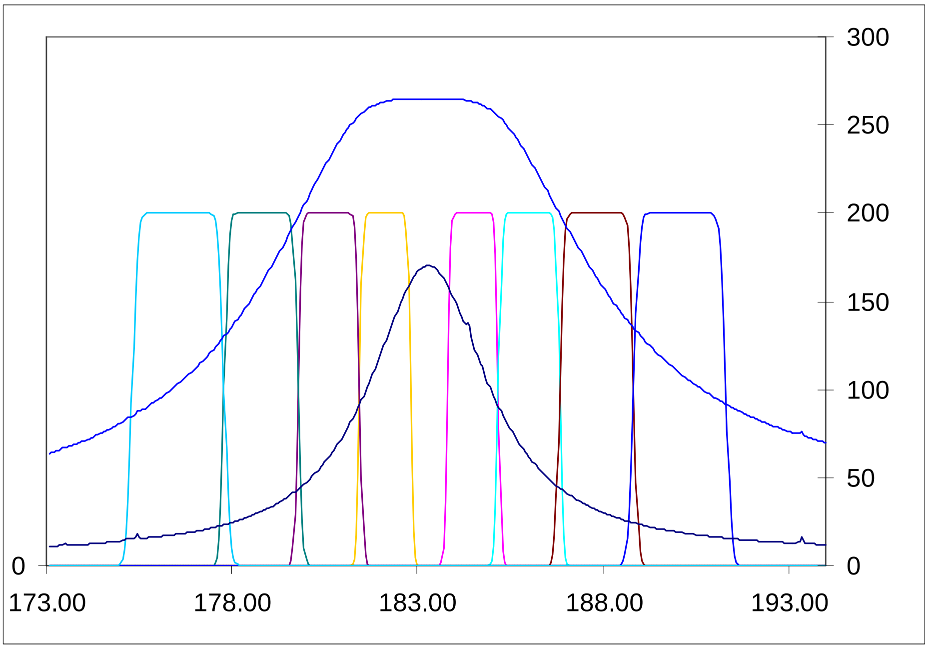


Figure 6. Filters optimized for "ideal" case – i.e. clear sky.

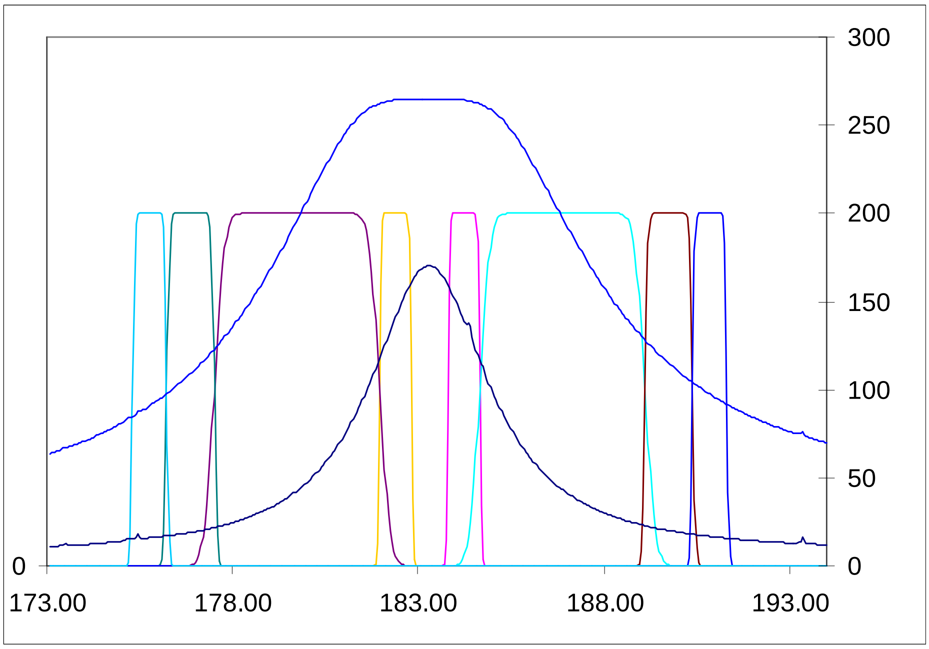


Figure 7. Filters optimized for "cloud" case – i.e. discriminates against continuum.

The error predictions for these two cases look like this.

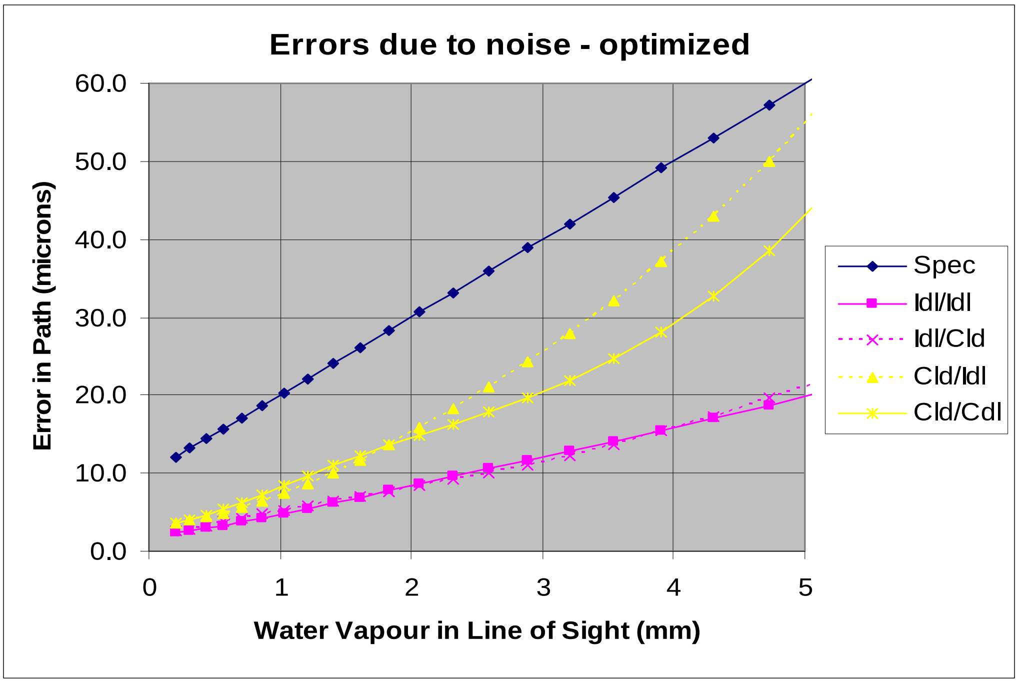


Figure 8. Estimated errors in path for the two sets of filters specified by Table 3.

Here the pink lines give the estimated error for the Ideal case and the yellow lines shows the Cloud case, and in each case the solid line is where the filters have been optimised for that case and the dotted line is where the filters have been chosen for the other case. It is seen that the performance in the Ideal case is not very sensitive to the exact choice of frequencies so long as the whole band is used. This is not surprising since one is basically measuring the integrated strength of the line. The Cloud case is more sensitive because it relies on taking the differences between the brightness temperatures in the different channels. From the plot on the previous page it can be seen that the optimization has resulted in a choice of one wide channel in the middle, a narrow one at low frequencies and two narrow ones at high frequencies. It turns out that there is an alternative configuration with two narrow channels at low frequencies and one at the high frequency end that works almost as well.

At first sight it would appear that the obvious thing to do is adopt the set optimized for the "Cloud" case. There are however at least two reasons for not doing this. First, we haven't really established that this process of subtracting the continuum emission works well in practice. We have some data from the SMA tests when there was cloud present that shows that the correction is better when the weights are adjusted to subtract the emission, but this has not been investigated thoroughly. Secondly, the optimisation is somewhat sensitive to what assumptions are made about the performance of the radiometer – e.g. how much of the noise is due to gain fluctuations or post-detection noise, relative to that arising in the RF and IF sections. We should perhaps also think about how well the radiometer would perform with one channel out of action. Clearly the cloud-optimized case puts very strong emphasis on the one broad channel.

I therefore recommend that we take an intermediate choice, which give good performance in both cases. Rounding-off the cross-over points leads to them being at 2.0, 4.5 and 6.5GHz. The filter specifications are then as given in table 4 and illustrated in figure 9 with the resulting performance in figure 10.

| | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| Centre | 1.25 | 3.25 | 5.50 | 7.25 |
| Width | 1.50 | 2.50 | 2.00 | 1.50 |
| Low | 0.50 | 2.00 | 4.50 | 6.50 |
| High | 2.00 | 4.50 | 6.50 | 8.00 |

Table 4. Recommended IF Filter Frequencies.

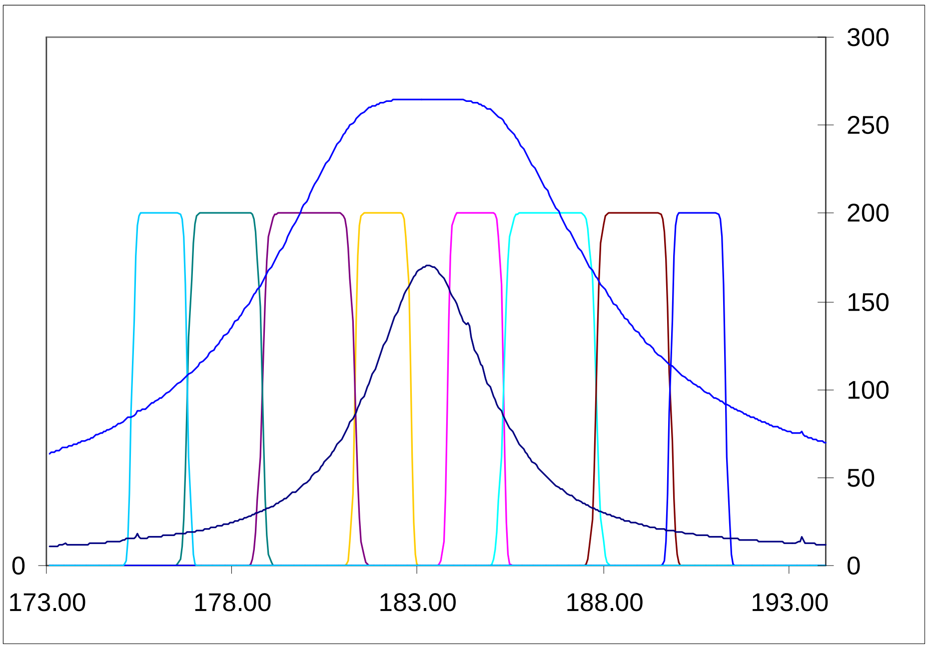


Figure 9. RF frequency coverage for filters in Table 4 together with the line emission profile for 0.5 and 4mm of water vapour.

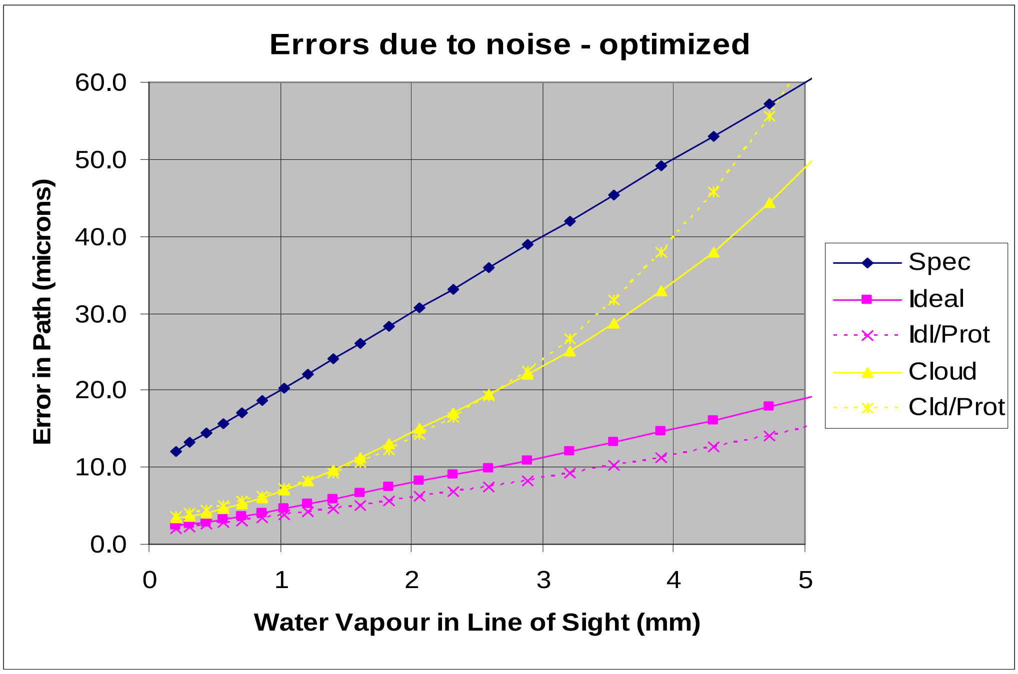


Figure 10. Production radiometer performance with equivalent for prototype for comparison.

Here, for interest, I have plotted as dotted curves the sensitivities of the prototype radiometers, assuming the same parameters for noise temperature, loss and so on, but allowing for the additional root 2 of sensitivity arising from the fact that the prototypes were dual-channel instruments, and also using 150K for the temperature of the internal reference. For the new case, I have assumed only one channel and a reference temperature of 223K, such as we might hope to achieve with Peltier cooling. We see that the new choice of IF frequencies compensates for much of the root 2 loss of sensitivity in the “Ideal” case and actually does better than the dual-channel prototypes in the “Cloud” case.

## 2.4 Other Issues Relating to the Frequency Choice

As mentioned, the original choice of filters was partly determined by the desire to avoid the Ozone line. The new recommendation includes this line in channel 1. We therefore need to look at this issue again. A preliminary check shows that even a change of a factor of two in the strength of the Ozone line has no significant effect (~0.2% change in sensitivity). By contrast, raising the lower limit of the range of frequencies available to around 1.25GHz, to be sure of avoiding this line, produces a loss of sensitivity of around 30% under dry conditions (~1mm). It therefore seems rather clear that we should choose the bottom limit to be as low as practical, depending on things like the performance of available LO noise amplifiers and the size and cost of the IF filters. Moving it up to say 0.75GHz produces about a 10% penalty under dry conditions which might be acceptable if it produces significant savings.

One might also ask whether the upper end of the band should be extended. Moving it out to 10GHz produces about a 5% improvement in the clear-sky case for ~3mm of water, but does nearly 20% better in the Cloud case. To see whether this is worthwhile one would have to look at whether IF amplifiers to cover this range have poorer noise performance. I suspect that the 5% gain in sensitivity would soon be eaten up.

The other point regarding the parameters of the filters is the question of how steep their responses need to be. As explained above, in this model the shape is described by two parameters – the “sharpness”, which is the ratio of the nominal filter width to the standard deviation of the Gaussian function describing the edges, and the “leakage”, which is assumed to be constant across all the other IF frequencies.

In all the models above a value of 20 has been taken for the sharpness. It was found that there was no measurable difference in performance for higher values than 20 and that even reduce it to 10 had very little effect. At 8 one was starting to see changes at about the 1% level. (It is worth noting that, if the filters are constructed as a series of high-pass / low-pass crossover networks instead of a set of band-pass filters, then a more relevant ratio is that of the crossover frequency to the width of the smoothing function. This alternative will be modelled if that seems to be an appropriate design concept. See section 3.4 for more discussion.)

The “leakage” was set at 0.05% (–33dB) in the modelling above. It was found that setting this to zero had no effect on performance and that raising it to 0.5% (–23dB) produced changes in performance of one to a few percent depending on the water vapour content and the case. This appears to suggest that setting a specification of around -30dB for the out-of-band rejection would be adequate (but see discussion of tolerances below).

Clearly the basic conclusion here is that the demands in terms of filter design, on both steepness and out of band rejection, are very modest. The real issue here is, however, the effect of differences between the filters in different radiometers, rather than the performance of the system when the radiometers are assumed to be identical. That is discussed in the next section.

# 3 Sensitivity to Errors in Filter Characteristics

With the new model of the whole system it is relatively easy to estimate the effects of changes in the various parameters that describe the IF bands, and also the sideband ratio of the mixers, on the measured path. The critical question here is whether or not we can make all the radiometers in the production run sufficiently similar that we do not need to use separate calibration parameters to describe each instrument. To quantify the requirement here we need to think about how the radiometers are to be used. The total path associated with 1mm of water is about 6600 microns. (This value is inversely proportional to the temperature, but this is a typical figure.) This means that our target accuracy of 10(1+w) microns amounts to between 0.5 and 0.2% of the total path. This does not however mean that we have to make two independent radiometers agree to such high accuracy. The radiometers are only to be used to measure and correction for <u>fluctuations</u> in the path that take place within a certain length of time and over a certain angular distance on the sky. The relevant time is that between observations of a reference source and the angular distance is that from the object being observed to the reference source. In both cases it is unlikely that that the change in path is more that one tenth of the total. This obviously provides significant relief in the accuracy required. On the other hand we do not want to use up all out error budget on this item and there are quite a large number of independent variables to take into account. As a starting point we will therefore look at what tolerance is allowed on each individual parameter such that it produces an error on the total path that is no greater than the target of 10(1+w) microns.

To do this, a version of the model has been made with two separate representations of the radiometers, both looking at the same atmospheric emission profile but with slightly different parameters for the filters, etc. The differences in the resulting brightness temperatures are then processed as if they were real differences in the atmosphere and the resulting difference in the estimated path found. As an illustration the following diagrams show the effects of the 0.1% change in the centre frequency of each filter. The colours represent the four channels – 1 blue, 2 pink, 3 yellow and 4 green.

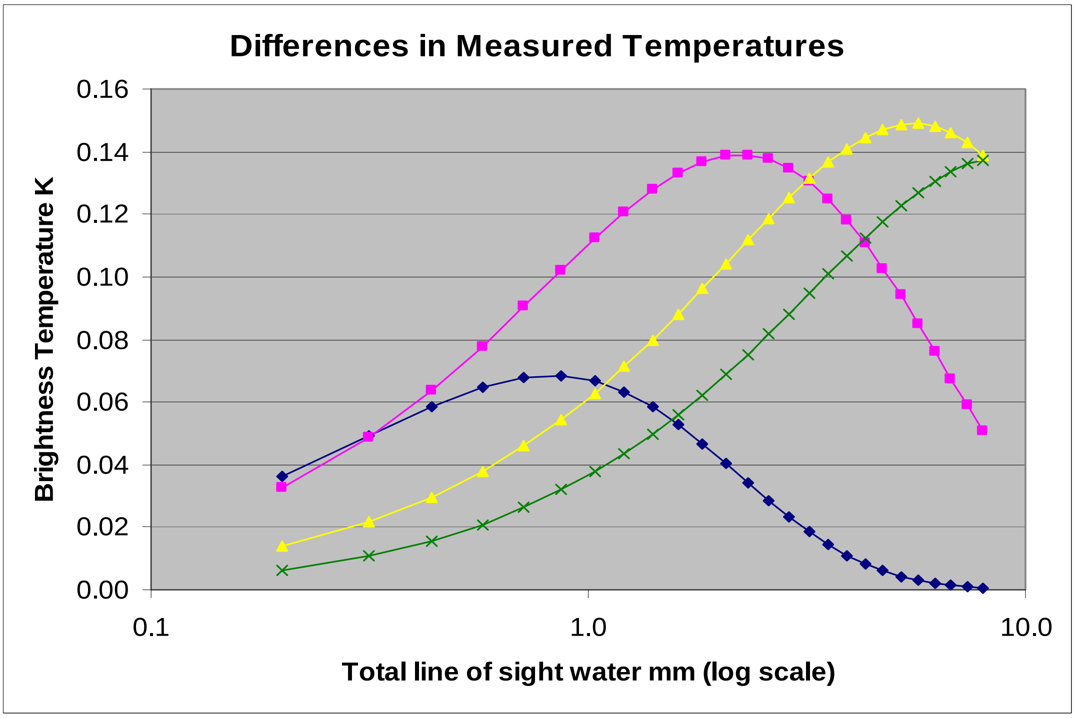


Figure 11. Effect of 0.1% change in the centre frequencies of the IF filters.

As before we deal with two cases, one where the data from the four channels are combined to give the best sensitivity assuming that water and dry air are the only sources of emission and a second where the weights are adjusted to remove the effects of broad-band emission from small particles.  The purple line is the target accuracy of 10(1+w) microns.

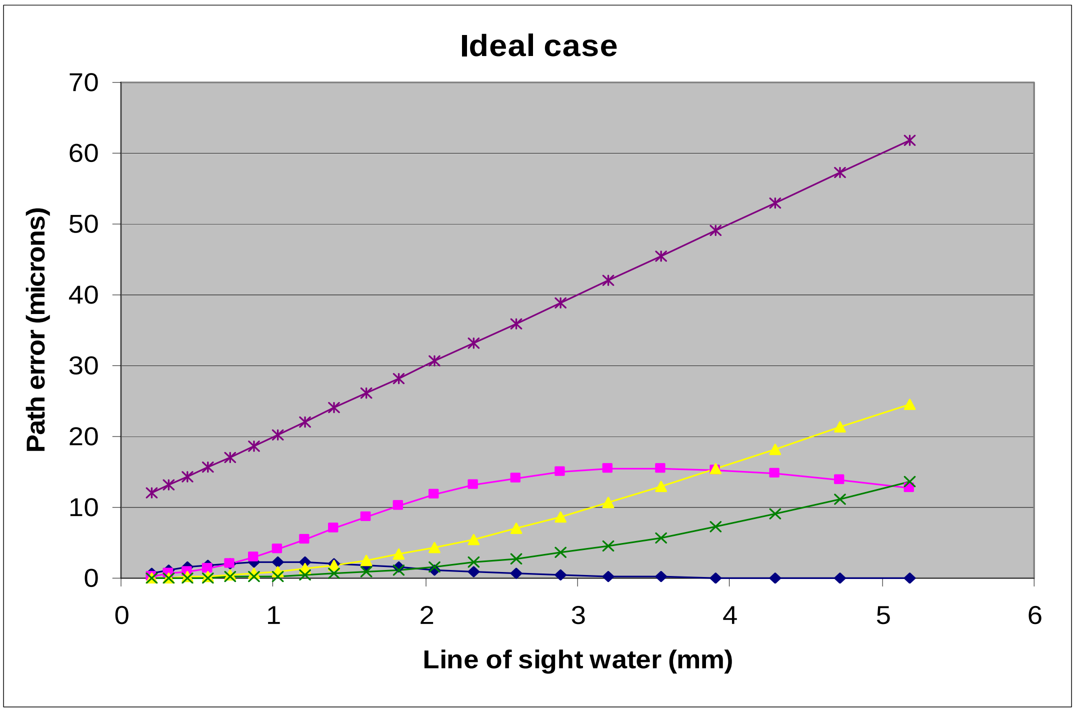


Figure 12.  Effect of 0.1% change in each centre frequency on estimated path – clear sky case.

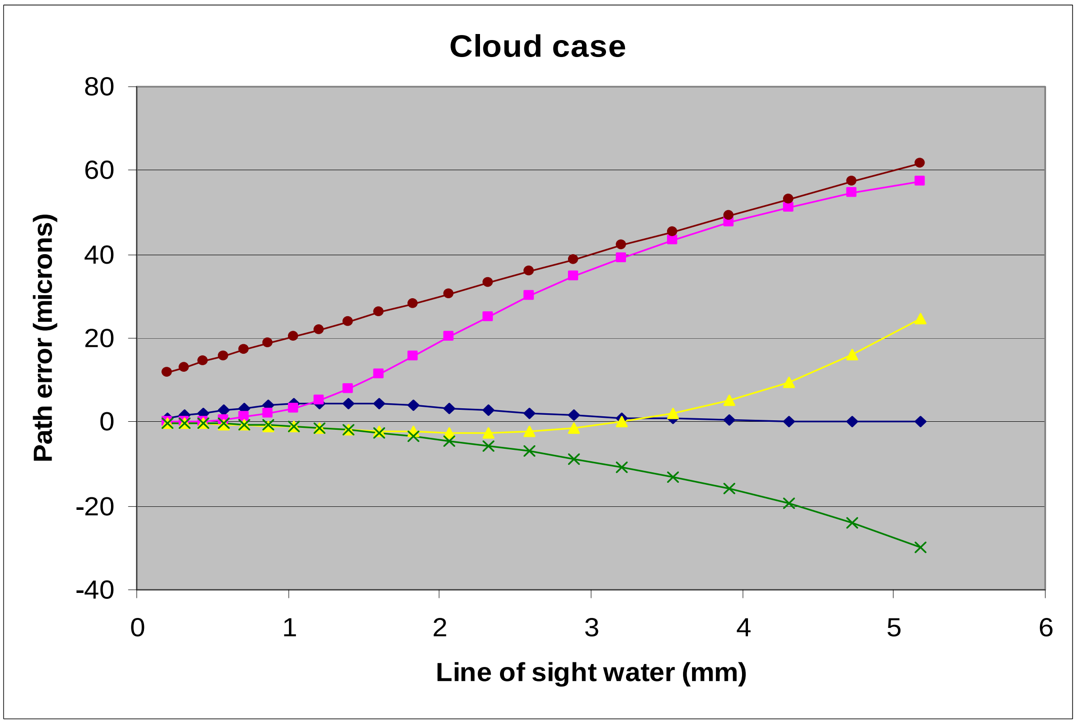


Figure 13.  As for figure 12 but for the “cloud” case.

This immediately shows that the most critical requirement is on the accuracy of filter 2 and that the cloud case is the more difficult one.

Carrying this process through leads to the following list of errors on the filter parameters in MHz, each of which results in the target error on the total path[2]:

| | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| Centre freq | -6.0 | -3.3 | -13.0 | 14.0 |
| Bandwidth | -85.0 | -55.0 | 320.0 | -500.0 |

Table 5. Changes in filter parameters (MHz) that individually cause the allowed error in path.

It is clear the filter width is very non-critical, but the requirements on the centre frequency are tight. As an alternative we can describe the channels in terms of their low and high frequency edges, which leads to the following tolerances, again in MHz.

| | | | | |
|---|---|---|---|---|
| Low cut | -13.0 | -7.0 | -25.0 | 30.0 |
| High cut | -11.0 | -6.0 | -30.0 | 30.0 |

Table 6. Same as table 5 but in terms of edge frequencies of filters.

We can also check the effects of other factors such as out of band leakage and "sharpness". If one radiometer has an average leakage of 0.05% (–33dB) on all of the filters, and the other has none, then the resulting path errors look like this:

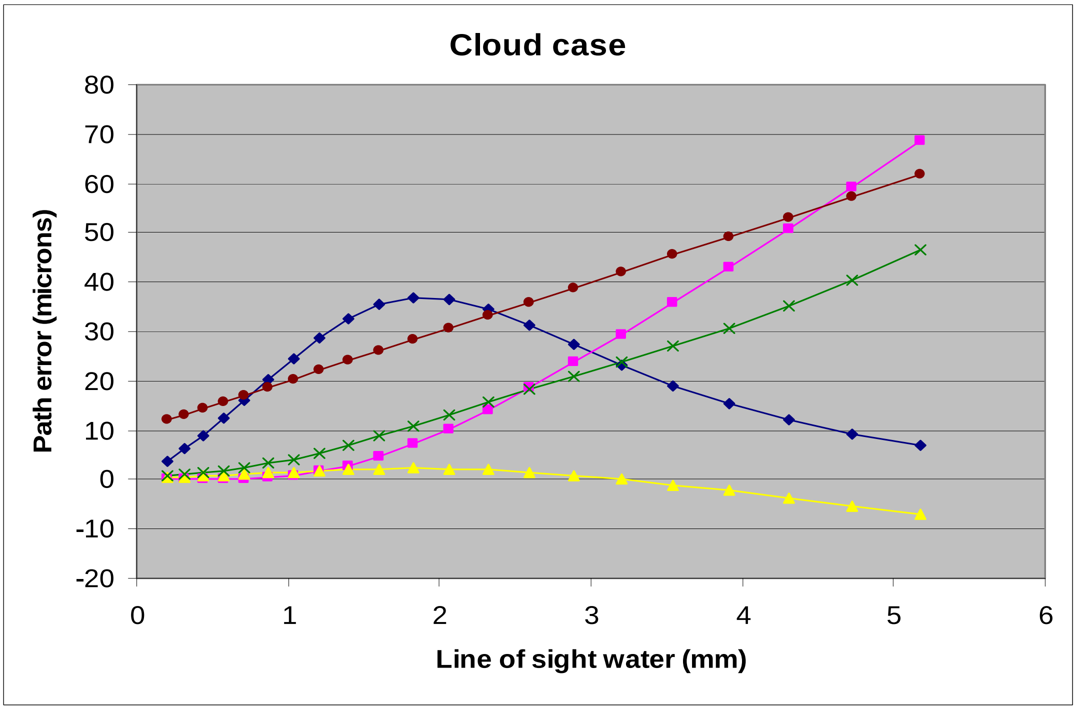


Figure 14. Effect of "leakage" of –33dB on average on path error.

Given that –33dB is not a hard specification to meet, it does not make much sense to require that the leakage patterns of the filters are matched. It would be better to set the leakage requirement rather lower, e.g. –40dB for channel 1 and 2 and –37dB for channels 3 and 4. This would ensure that the effects of any differences are negligible.

Similarly, the model shows that, with the sharpness set to the nominal value of 20 on one radiometer, then having a values ranging from 15 to say 40 in the other radiometer produces errors of up to about half of the target. If we wanted to be sure that differences in sharpness will have no effect we could increase it to be say 25, but this is probably overkill.

[2] Note that the signs on these errors are not important. They were just chosen to make the errors positive.

Now we consider the effects of differences in the sideband ratios of the mixers. These have an effect on the derived path because the line profile is not symmetric and because there is a slope on the continuum emission produced by the atmosphere. We take the same criteria as before – i.e. we match the accuracy target of 10(1+w) microns on the whole path for each variable – and we assume that one radiometer has an exactly balanced response – i.e. that the response are 0.5 and 0.5 – while the other has the sideband responses of $0.5 + \delta$ and $0.5 - \delta$. If we consider the results for the 4 IF bands separately, the allowed values of $\delta$ are 0.2, 0.1, 0.07 and 0.035. As would be expected the requirement for the outer frequencies are the most severe because the slope in the continuum has the most effect there.

Finally we can check that there is no effect from the fact that we intend to use slightly different LO frequencies in the different radiometers in order to avoid introducing coherent signals into the astronomical signal channels. The answer is that the difference in LO frequency needs to be more than 100 MHz before the error target is approached. This means that errors from the planned offsets of only a few MHz can be ignored.

## 3.1 Combined Errors

Having found the sensitivity of the various parameters in this way, we can select a set of tolerances chosen in such a way that the combination of likely errors is acceptable but the difficulty of meeting them is minimized. It is clear from the above that the most important parameters are the mixer sideband ratio and the centre frequencies of the channels. If one assumes that the errors in all these are independent then one can take the individual estimates of the path length errors and add them quadratically to produce an overall error estimate. As an example, we take the tolerances on the filter centre frequencies and widths given in table 7.

| | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| Centre | 5.0 | 10.0 | 18.0 | 25.0 |
| Width | 15.0 | 40.0 | 50.0 | 50.0 |

Table 7. Assumed errors (MHz) in the filter parameters for case illustrated in figure 15.

We also allow the tolerance on sideband ratio to be $\delta < 0.1$, i.e. the ratio is no worse than 60:40. This gives the errors on the path shown in figure 15.

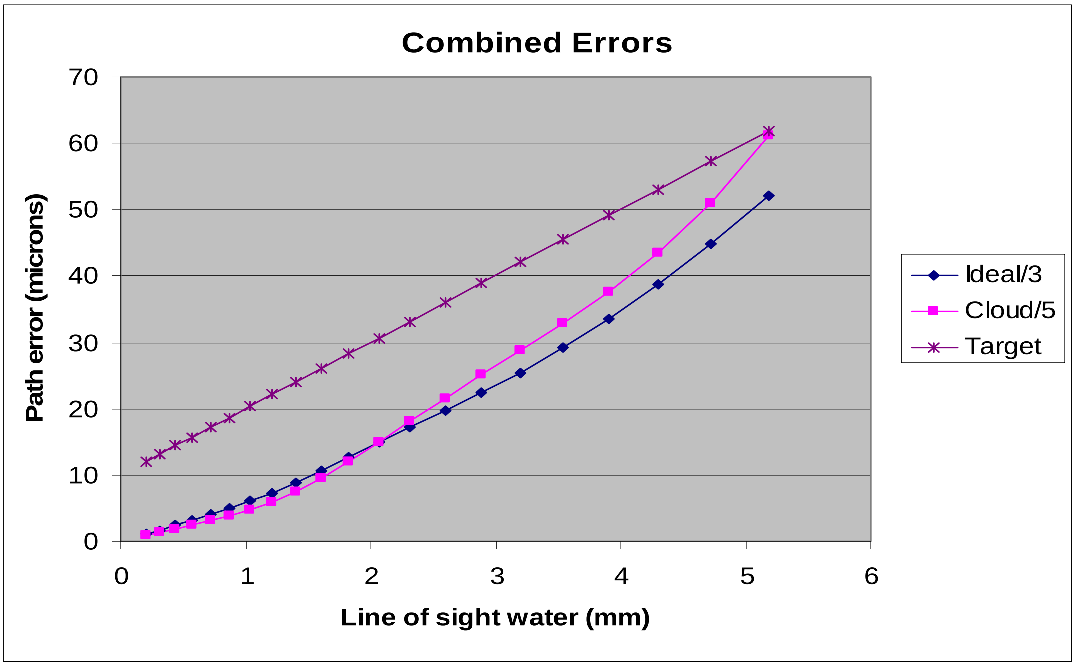


Figure 15. Results of combination of filter and sideband errors (divided by factors of 3 and 5).

Here the error estimates have been divided by arbitrary factors of 3 and 5, as indicated in the legend, in order to make them comparable to the standard target of 10*(1+w) microns. The point here is that what has been calculated are the estimated errors in measuring the total path, whereas we are actually only concerned with the accuracy in measuring changes in this on timescales of a few minutes or angles of a few degrees at most. As already discussed, we are not likely to see changes larger than a tenth of the total path, so if the predicted errors are only 3 or 5 times the target this means that we would still have some margin left, if we are able to achieve the tolerances proposed above.

## 3.2 Other Filter and System Properties

The other important parameters describing the response of the RF/IF system are the ripple and slope. The main effect of ripple is to reduce the effective bandwidth, which in turn reduces the sensitivity. For a system with frequency response $T(\upsilon)$, the effective bandwidth is given by $( \int T\, d\upsilon )^2 / \int T^2\, d\upsilon$. For a ripple of +/- 1dB this gives ~3% loss of effective bandwidth or ~1.5% loss of sensitivity, which is quite small. The effect is quadratic with the magnitude of the ripple so for +/– 0.5dB the loss of sensitivity is certainly negligible, at under 0.5%. It would be reasonably to specify +/– 0.5dB for the filters (or perhaps even better) and to require +/– 1dB for the response of the whole system including amplifiers, detectors, etc.

The main effect of any slope across the band is to shift the effective frequency. A slope from –1 to +1dB from one side of the band to the other would produce a shift in the effective frequency of ~ 4.3% of the bandwidth. On channel 2, which has a width of 2.5GHz, this amounts to more than 100MHz. This is large compared to the allowed tolerance of ~10MHz. The shift is linear with the amount of slope, so we would have to reduce it to about +/–0.1dB to be sure that the effective centre frequency is within the tolerance. Although it is probably possible to achieve this for the filter itself, it will clearly be difficult to get a response as flat as this from the system overall when the amplifiers, cables and detectors are included. I believe that this requirement – to get very low slope for the complete system across each of the channels – is going to prove the most difficult to meet.

## 3.3 Number of Channels

In all of the above discussion it has been assumed that the number of channels is fixed at four. It would be rather a lot of work to increase the number of filters in the model, but it is easy to reduce it to three by simply defining one of the present four to have essentially zero width. It turns out that the optimization process then chooses to have a roughly 1GHz-wide filter at the low end of the band and another one of similar width at the top, with one wide filter in between. The overall performance is then less than 10% worse than that of the optimum design with four filters. This strongly indicates that with four filters we are already getting into diminishing returns. The cost advantage of going down to three is probably not significant and some redundancy is certainly useful.

The only case using more channels that might be worth considering is one with say eight equally spaced channels. It might be possible to implement this by using a second down-conversion followed by a set of identical filters. Such a scheme might make it easier to achieve the high accuracy required for the effective frequencies of the bands, but note that a simple double-sideband down-conversion scheme looses root(2) in sensitivity. The additional complexity of using single-sideband conversion does however make this option unattractive.

## 3.4 Alternative Form for the "Sharpness" Factor

It was noted earlier that describing the edges of the filters with a parameter relating the widths of the filters to the sharpness of the edge may not be appropriate. I therefore looked at an alternative case where the width of Gaussian that models the transition – the "edge width" – is taken to be a constant fraction of the frequency of the edge. (Note that this is the IF frequency.) The following figures illustrate different cases.

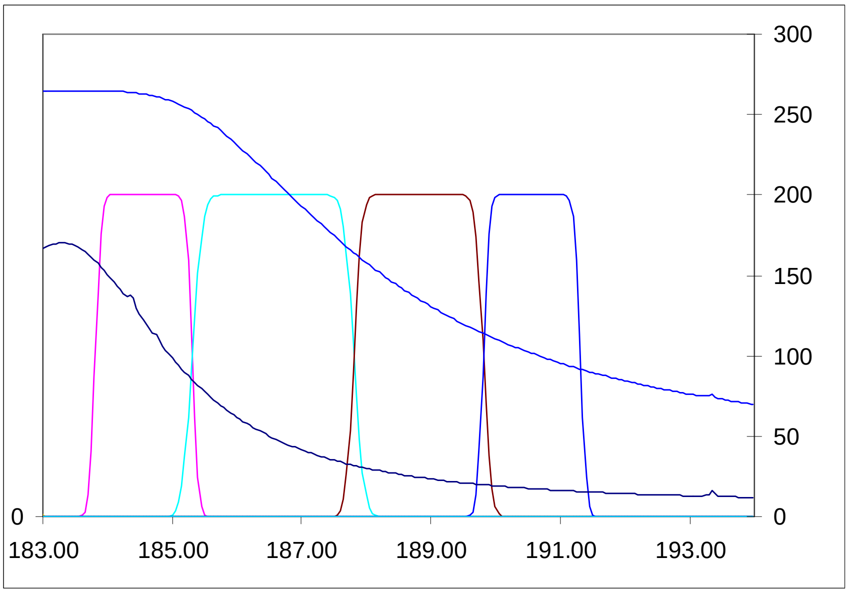


Figure 16. Filters with edge widths equal to bandwidth divided by 20.

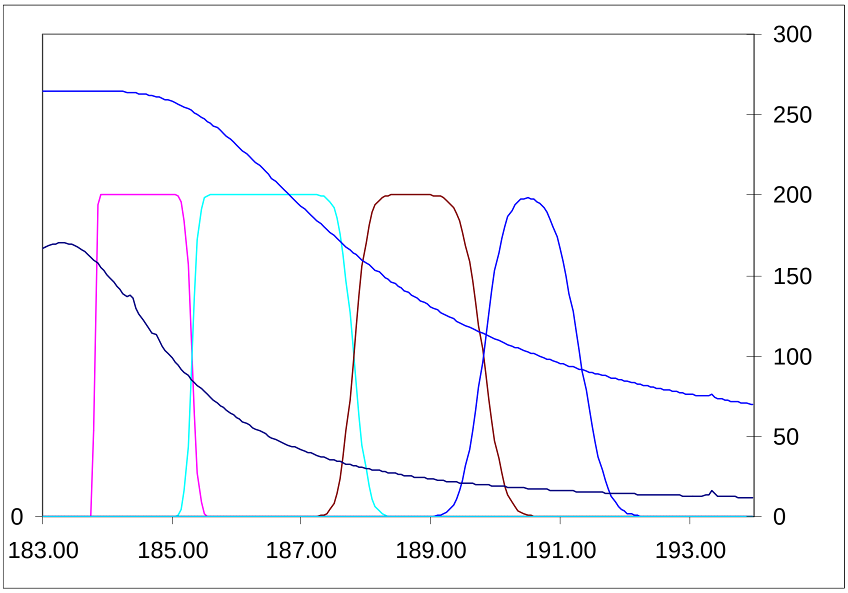


Figure 17. Filters with edge widths equal to IF frequency divided by 25.

Running this case through the model showed that the overall performance is hardly affected at all by this rather big change in the form of the filters. This simply reflects the fact that the water line has a very smooth shape, due to the pressure broadening.

Specifying the form of the filters as in figure 17 is, however, probably not very realistic, since the conventional design techniques for bandpass filters will, for a given number of filter sections, give rise to something closer to figure 16 than figure 17. In deriving the final requirements, therefore, a compromise form was adopted where the edge width was set by the bandwidth divided by 30 plus the IF frequency divided by 60, as shown in figure 18.

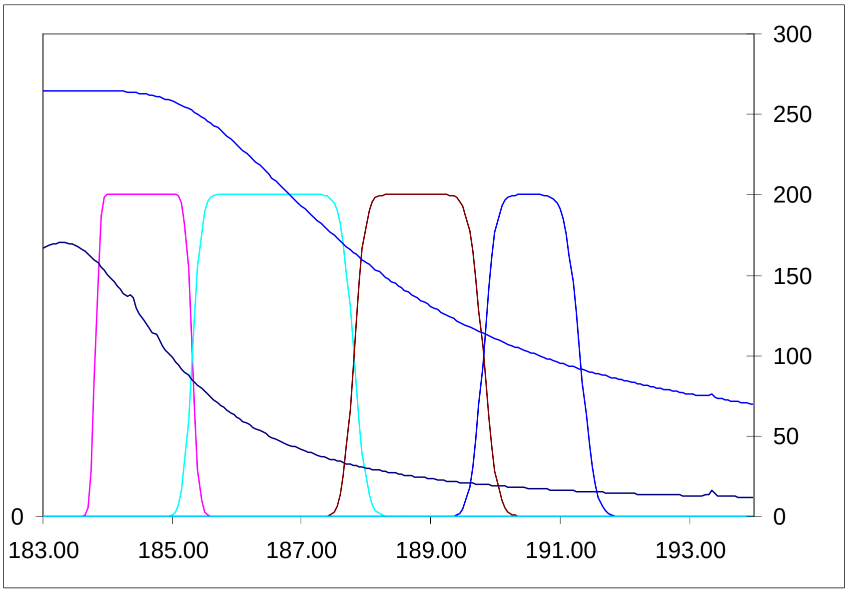


Figure 18. Compromise choice of “edge widths”.

It should again be made clear that there is nothing critical about this question of how to set the sharpness of the edges. The important thing is that the behaviour of the radiometers on different antennas matches to high accuracy.

# 4 Summary and Conclusions

We have shown that a significant gain in performance can be achieved by optimizing the frequencies of the filters. We have also shown that the requirements for the steepness of the filters, out of band leakage, etc., are relatively modest. The realization of such a set of filters therefore seems quite straight-forward.

We have also investigated how similar we would have to make the production radiometers in order to be able to treat them as being completely identical. This analysis shows that the effective centre frequencies would need to be reproduced very accurately – e.g. only a little over 1 part in a 1000 in the case of filter 2. This, together with the demanding numbers found for the slopes across the bands and the mixer sideband ratios, suggests that it is not realistic to require that the radiometers be effectively identical in this way. It therefore seems more appropriate to assume that the data reduction system will need to hold two or three parameters for each channel of each radiometer. Shifts in the effective frequency of a channel would for example be covered by adjusting the parameter in the model which gives the opacity, averaged across the band, per mm of water. Other parameters might be the coupling efficiency to the sky and a small offset to be applied to all the temperatures. (The latter was found necessary in the prototype radiometers because of reflections from the calibration loads.) One can imagine two ways of deriving these parameters: 1) by laboratory calibration, for example by using an FTS, to measure the effective frequencies; 2) observations of the atmosphere made after installing the radiometers on the telescopes, e.g. comparing the response of each radiometer to a standard when doing “sky-dips” under stable conditions.

From experience with the prototypes I am confident that the second of these would work, but it would take a good deal of time and trouble. The feasibility of achieving sufficient accuracy by the first method – laboratory measurements – needs to be examined. In either case, one would certainly not want to repeat the measurements often. One would instead set the requirements on stability of the performance of the filters, amplifiers, etc., such that one would only need to establish these parameters once, when the radiometer is first commissioned, and then expect to use them for at least several years without further alteration. This means that the final specifications should be in the form of a value, an allowed manufacturing tolerance on that value and a required stability. (In principle the stability should include thermal drifts as well as ageing but it is presumed that the temperature control is good enough to eliminate significant thermal errors.)

We also need to define the sharpness of the filters in a more conventional way than by using the convolution of a top-hat function with a Gaussian, which was convenient in the model but not as a specification. To do this the widths at 20 and 40dB rejection were calculated for the shapes shown in figure 18. This leads to the following set of requirements[3]:

| Effective frequency $f_{eff}$ and Tolerance [GHz] | Bandwidth@-3dB and Tolerance [GHz] | Bandwidth @-20dB Max [GHz] | Bandwidth @ -40dB Max [GHz] |
|---|---|---|---|
| 1.250±0.015 | 1.500±0.025 | 1.83 | 2.28 |
| 3.250±0.030 | 2.500±0.040 | 3.14 | 4.01 |
| 5.500±0.050 | 2.000±0.050 | 2.74 | 3.74 |
| 7.250±0.070 | 1.500±0.050 | 2.3 | 3.38 |

Table 8: Proposed Filter Specifications

In addition it is proposed that the stability of the effective frequency should be set at one third of the values given for the manufacturing tolerance given in the left hand column. The requirements on the widths are sufficiently loose that there seems no point in defining tolerance and stability separately. Finally the mean out-of-band leakage (outside the –40dB widths) should be below –40dB for channels 1 and 2 and –37dB for channels 3 and 4 and the ripple and slope across each band should be set at +/– 0.5dB.

I am indebted to Darrel Emerson for his helpful comments and corrections.

Richard Hills original 19th Nov 2006 revised 6th August 2007

[3] Note that these widths differ somewhat from those in the currently approved version of the specifications. Those were set before the options in section 3.4 of this note had been considered, while these reflect the "compromise" described there. It is likely that the formal specifications of the production radiometers will be revised after further investigations of the practicalities of implementing the filters.