ALMA Memo No. 495
Estimated Performance of the Water Vapour Radiometers

Richard Hills

Cavendish Laboratory
Cambridge, CB3 0HE, UK
richard@mrao.cam.ac.uk
14th June 2004

Abstract

The sensitivity and accuracy to be expected from the ALMA water vapour radiometers is estimated. Even with quite pessimistic assumptions about noise temperature and gain stability, it appears that the sensitivity of the present design is at least a factor of 2 better than the specifications. It is more difficult to predict the medium term stability and accuracy that will be achieved, but the indications are that these will not be a limiting factor in using the data from the radiometers. An estimate of the errors that would occur in correcting the single-dish pointing for the effects of water vapour gradients is also given.

Introduction

Sensitivity, stability and accuracy are the main parameters which will determine how well the ALMA water vapour radiometers fulfil their functions.

The sensitivity is set by the fluctuations in the output due to noise, on timescales of tenths of a second up to a few seconds. The radiometers have been designed so that these fluctuations should be dominated by the intrinsic thermal noise arising in the receivers and the atmosphere, but in principle there could also be contributions due to things like digitization errors, short-term gain fluctuations and power-supply transients.

By stability we mean spurious variations in the output on timescales of a few seconds to a few minutes, due to things like drifts in the temperature of the internal reference load and gain variations in the electronics. Note that because the radiometers are differential – i.e. they measure the difference between the sky and the reference load – a drift in the load temperature would show up directly as a spurious signal (but one which is common to all the output channels), while a gain variation multiplies only the difference between the sky and the load, so that there is no error when these temperatures are equal.

By accuracy we mean how well the value recorded by the radiometer (averaged over any noise fluctuations or instabilities) represents the true sky temperature. This involves things like the linearity of the detectors, the accuracy of the various internal temperature sensors, the frequency response of the filters and the optical coupling of the horns to the sky.

This note gives some initial estimates of the magnitude of these errors and shows how to convert them into errors in the quantities of interest, most importantly the corrections to be made to the path through the atmosphere. More accurate values will become available as a result of testing in the laboratory and at the ATF.

Sensitivity

The radiometers are designed to measure the brightness temperature of the atmosphere in 4 IF bands on either side of the water line frequency of 183.31GHz. In the prototypes, the frequencies (in GHz) chosen for these bands are as follows (table1) :

Filter	A	B	C	D
Centre Frequency	0.88	1.94	3.175	5.20
Full Width	0.16	0.75	1.25	2.50
Range	0.80–0.96	1.565–2.315	2.55–3.80	3.95–6.45

As already noted, the short term fluctuations in the outputs should be dominated by the intrinsic noise radiometer noise. Contributions from the A-D converters and the post-detection amplifiers are expected to be below 1% each. Excess noise due to gain fluctuations on timescales shorter than the switching rate should also be negligible, given that we are switching at 10Hz or faster. This means that the rms error in an integration time t should be

$$\Delta T = F\ T_{sys} / \sqrt{(B\ t)} \qquad (1),$$

where T_{sys} is the system temperature and B is the effective bandwidth of the filter (which, to be conservative, we take as ~0.9 of the full width quoted above). The switching factor F would be $\sqrt{2}$ in the ideal case, given that we are using either a correlation receiver or a Dicke-switched system with two mixers which look at the sky on alternate phases of the chop cycle. In practice both systems have some inefficiencies – in the Dicke system a small fraction of the integration time has to be blanked to remove the transitions of the chop cycle, and in the correlation system there are losses due to phase and gain mismatches. For this reason we use a factor of $F = 1.5$ instead of $\sqrt{2}$ in the calculations which follow.

The noise temperatures of the harmonic mixers plus uncooled IF amplifiers that we are using in the prototypes are ~1500K or better. (Note that this is the double sideband noise temperature, which is the relevant quantity here since we are measuring the average brightness of the atmosphere in the two sidebands. With the correlation receiver we can also measure the differences between the upper and lower sidebands, for which the errors would be twice as large.) To this must be added the losses in the waveguide components and the optics as well as the atmospheric emission. In order to ensure that the beam of the radiometer is well coupled to the sky there is an internal stop in the optics which will cut off the sidelobes and this will have a transmission of ~90%. In the Dicke system other losses should be small, so the system temperatures should be no worse than 2000K. In its present form, the correlation receiver contains waveguides and a hybrid with a total loss of at least 1dB giving T_{sys} ~2500K. To be conservative we take this latter value in what follows. The intrinsic noise for a 1-second integration time from equation (1) is then (table 2)

Filter	A	B	C	D
Effective Bandwidth	0.144	0.675	1.125	2.25
RMS error (K)	0.313	0.144	0.112	0.079

To relate these errors in the apparent sky brightness temperature to errors in the path correction we need to know the conversion factor between these quantities. This depends on the frequency of the channel and the amount of water vapour in the atmosphere – the lower IF frequencies are more sensitive when the atmosphere is very dry but they become saturated and therefore loose sensitivity, when there is a lot of water in the path. The conversion factors have been calculated using the "ATM" numerical model as discussed in ALMA memo 496. Taking the rms errors from table 2 and converting to path gives the following values (table 3 – columns 2 to 5):

Water (mm)	Path Errors (microns)				Weights				Weighted Error
	A	B	C	D	A	B	C	D	
0	5.4	4.6	6.4	9.1	0.294	0.396	0.207	0.103	2.92
0.2	7.1	5.4	7.0	9.5	0.230	0.401	0.239	0.130	3.42
0.4	9.4	6.3	7.6	9.9	0.175	0.396	0.269	0.159	3.94
0.6	12.5	7.3	8.3	10.3	0.130	0.382	0.297	0.191	4.50
0.8	16.5	8.5	9.0	10.7	0.095	0.361	0.320	0.224	5.08
1	21.8	9.8	9.8	11.2	0.068	0.335	0.339	0.258	5.69
1.2	28.9	11.4	10.6	11.7	0.048	0.306	0.353	0.293	6.32
1.4	38.2	13.3	11.6	12.2	0.033	0.276	0.364	0.327	6.97
1.6	50.6	15.4	12.6	12.7	0.023	0.246	0.370	0.362	7.64
1.8	66.9	17.9	13.7	13.3	0.016	0.217	0.372	0.396	8.34
2	88.5	20.8	14.9	13.8	0.010	0.189	0.371	0.429	9.06
2.2	117.2	24.2	16.2	14.4	0.007	0.164	0.367	0.462	9.80
2.4	155.0	28.1	17.6	15.0	0.005	0.141	0.360	0.494	10.56
2.6	205.1	32.6	19.1	15.7	0.003	0.121	0.352	0.524	11.35
2.8	271.4	37.9	20.8	16.3	0.002	0.103	0.342	0.553	12.16
3	359.1	44.0	22.6	17.1	0.001	0.087	0.330	0.582	13.00
3.4	628.6	59.5	26.8	18.5	0.001	0.062	0.304	0.634	14.76
3.8	1100.5	80.3	31.7	20.2	0.000	0.043	0.276	0.681	16.64
4.2	1926.6	108.3	37.5	21.9	0.000	0.030	0.248	0.723	18.65
4.6	3372.9	146.2	44.3	23.9	0.000	0.020	0.220	0.760	20.80
5	5904.8	197.4	52.4	26.0	0.000	0.014	0.194	0.792	23.10

The noise in the four channels will be independent, so the estimates of the path correction from the four individual channels can be combined with an appropriately chosen weighting (inversely proportional to the square of the errors). These optimum weights vary with the water content, as shown above (columns 6 to 9). The resulting error in the weighted average for the path correction is given in the right hand column.

When the radiometer is balanced, i.e. when it is measuring an atmospheric temperature which is close to that of its internal reference load, then gain fluctuations are not important. When there is an imbalance, however, there will be an additional error which has the form

$$\Delta T = (T_{atm} - T_{ref}) \Delta g/g \qquad (2)$$

where $\Delta g/g$ is the fractional gain fluctuation over the integration time.

Note that, unlike the intrinsic noise fluctuations, the gain fluctuations will be strongly correlated between the four channels (but independent of the intrinsic noise).

This means that the correct treatment is to convert these errors into variations in path, multiply them by the weights for the channels and add them linearly (rather than quadratically) before combining them with the errors due to the noise, this time as root sum squares. The initial indications are that the gain variations on short timescales are small – probably no more than a few parts in 10^4. If this is confirmed then their contribution is negligible. For illustration, the plot below (figure 1) shows the results if the gain fluctuations are 1 part in 10^3 and the reference load is at 140K.

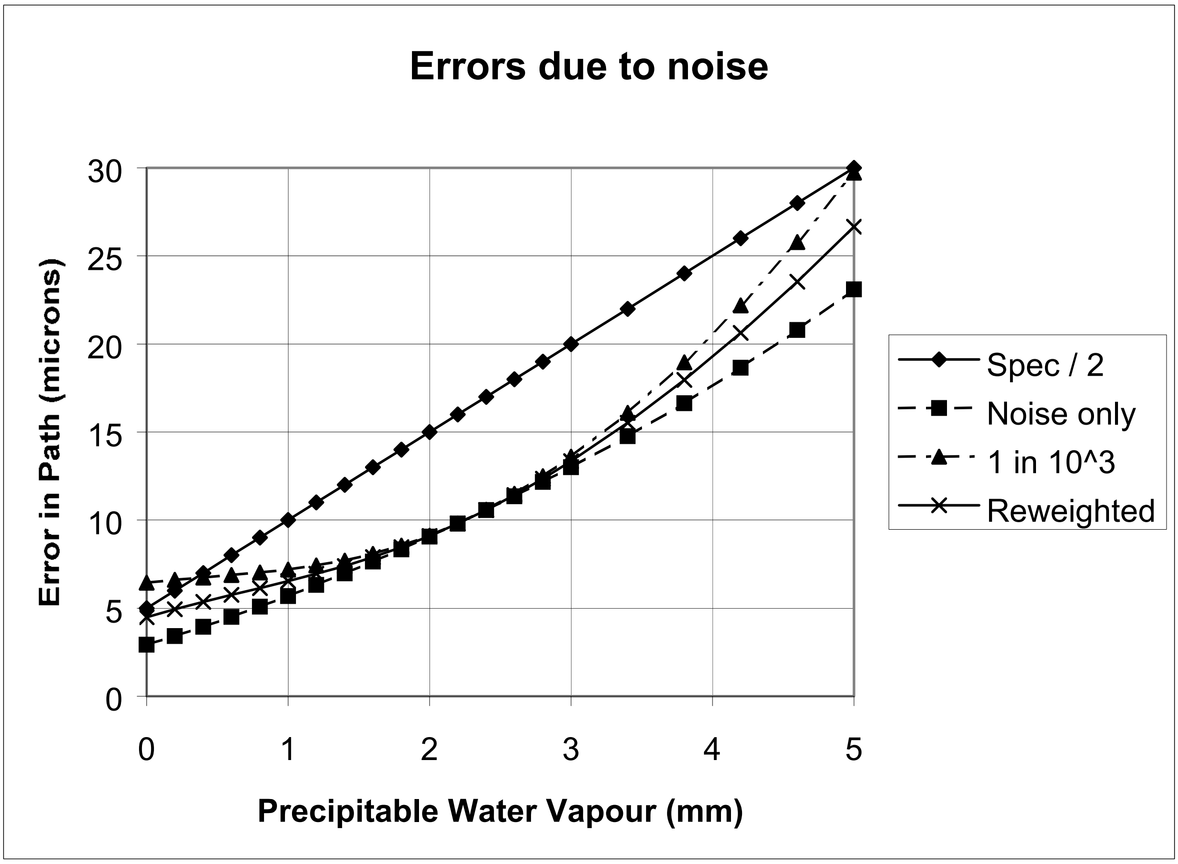

Figure 1.

Here the top curve is half the nominal specification, i.e. 5*(1+*w*) microns, with *w* being the precipitable water in mm's, and the lowest curve is for the radiometer noise only, while the curve "1 in 10^3" includes gain fluctuations (equation 2). It is also possible to adjust the weights to minimize the overall error and the result of this, for the same level of gain fluctuations, is shown in the curve "Reweighted". It can be seen that this gives a small advantage, but that this is unlikely to be significant. (See memo 496 for more discussion of these weighting factors.)

Recalling that the estimates used were quite pessimistic, it is clear that the sensitivity of the radiometers should be at least a factor of 2 better that the nominal specification over the whole range of water vapour values likely to be of relevance to ALMA, and that they should be a factor of 3 better for the range 1 and 4mm. It is also worth noting that, because of its narrow bandwidth (chosen to avoid an Ozone line), channel A does not contribute much to the sensitivity. If the Ozone emission to be unimportant we would obtain some increase in sensitivity under dry conditions by using a larger bandwidth for this channel.

Stability

For integration times that are somewhat longer or shorter than one second the predicted errors would scale as the inverse square root of the integration time. (The control system has been designed to allow meaningful readings to be taken at least as often as 10 times per second, but this is unlikely to be necessary for phase correction.) On timescales longer than about ten seconds, we expect gain variations and other drifts in the electronics to become more significant relative to the noise and, because these effects are likely to behave as something like 1/f, we expect that the errors will stop coming down and eventually come up again. As an indication of the type of performance that would be expected, here is a curve showing how the errors might scale with integration time.

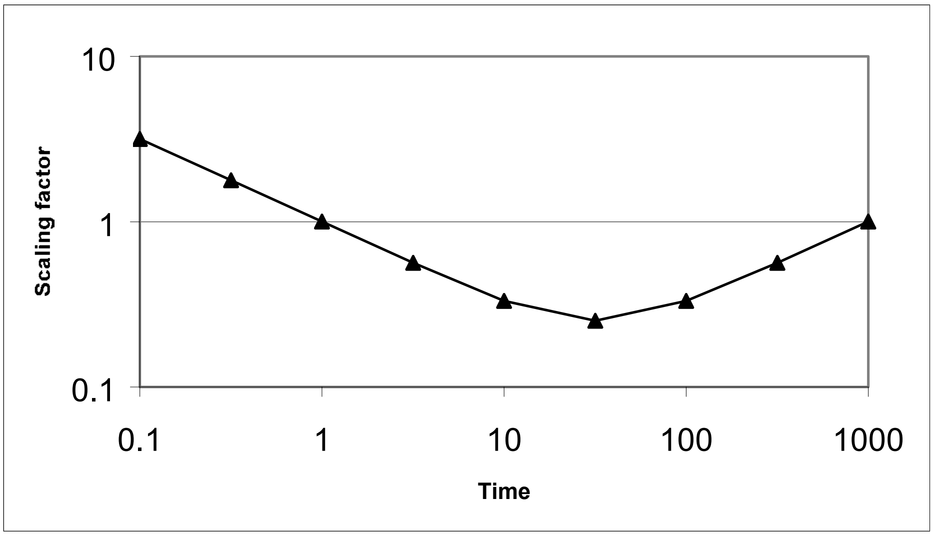

Figure 2

Here a term inversely proportional to root time has simply been added quadratically to one rising proportional to root time, scaled to reach unity at 1000 seconds. The preliminary experimental data indicate that any such turn-around occurs at longer times than shown here but, until firm data become available, it would be reasonable to use this curve to scale the values in table 3 if this is required for simulation purposes.

In fact, as can be seen in figure 1., the effect of gain variations depends on the total amount of water in the path. The errors are significant for very dry and very wet conditions, but less important for 2 to 3 mm's of water, when the gain errors largely cancel out. If the gain variations were to reach 3 parts in 10^3, the path errors would be approaching the nominal specification for ~0.5 and ~5mm, but still more than a factor of two below the specification in the middle part of the range. The data so far indicate that the gain variations are actually well below 1 part in 10^3 at 100 seconds. Gain variations on timescales longer than 100 seconds or so will be removed by the internal calibration process of the radiometer, which involves switching the reference beam to an ambient temperature load for about one second.

Similarly a drift of ~0.2K in the temperature of the cold load would produce a path error which is roughly equal to the specification. The test data on the Stirling cycle cooler indicate, however, that the temperature variations should be well below 0.01K, so this contribution should again be insignificant.

Accuracy

Here the question is what the systematic errors will be in the radiometer's measurements of the brightness temperatures, and how these will affect the estimates of the water vapour in the path. The sources of such errors will be things like inaccuracies in the monitoring devices used to read the temperatures of the cold and ambient reference loads, the linearity of the detectors and the accuracy with which the coupling of the radiometer beam to the sky is known. There are three cases to consider:

1) when the water vapour fluctuations are very large, so that scaling errors affect our ability to correct them;
2) when we move the antenna from the source to a reference object, so that there is a step in the water value and we wish to use the radiometer to apply the phase correction across this step; and
3) when we want to measure the total amount of water in the path, rather than just the fluctuation in it, so that it can be used as part of the amplitude calibration.

In all three cases what we are interested in is the percentage error in the amount of water derived resulting from the errors in the instrument. (See Memo 496 for discussion of the errors due to uncertainties in the atmospheric models.) We can get a flavour for the quantities involved using the same modelling of the radiometer as in the previous sections.

For example, a 1K error in the temperature of the cold reference load would give an error of a little over 1% in the estimate of the total water over the range w from 1 to 4mm. (A more precise expression would be that the error in w is given by $0.006 + 0.01w$.) The same error in the ambient load temperature would give a much smaller error.

It is possible to obtain temperature monitors with an accuracy of better than 1K, but because there are various offsets likely to be present in the system (window losses, beam truncation, etc.) it is likely that the effective temperatures would have errors of at least 1K in any case. The radiometers will therefore be calibrated after construction and "burn-in" by means of external black-body loads at ambient and liquid-nitrogen temperatures. This calibration will provide corrections to the load temperatures which will be stored internally and applied in the data processing. This means that the requirement is more on the long term stability of the monitoring and control systems than the absolute accuracy. Here errors of well below 1K should be achievable quite easily.

The measurements will also be effected by errors in the effective frequencies of the IF channels. The reason for this is that the conversion factors will be incorrect if the channels are not measuring the part of the spectral line that they are supposed to. The sensitivity here is quite high: a 1% error in the centre frequency of each channel produces about 0.5% error in the value obtained for the amount of water vapour. (This is assuming that the errors in the frequencies of the 4 channels are uncorrelated.) The filters themselves can have accuracies in their effective frequencies of considerably less than 1%. Those obtained for the prototypes were generally no worse than 0.5%. It will however be difficult to maintain this level of accuracy when the slopes of other components in the system are taken into account. Again it will be possible to measure the effective frequencies of the channels for each completed radiometer by laboratory measurements (probably using a swept coherent source). This will however complicate the data processing somewhat and we will therefore try to avoid it.

It should also be possible to perform various checks on the calibration when the radiometers are installed, for example by performing "sky dips" when conditions are clear and stable and by comparing the "step" in water measured by the radiometers on different antennas when they are moved through the same range of angles.

We will not be able to provide firm estimates for the systematic errors until more testing has been done, but in general terms an instrumental accuracy of about 2% should be achieved and 1% seems possible with careful calibration.

Correcting Single-Dish Pointing

An option which the project is still carrying is to provide the means of measuring the gradient of the water vapour across the aperture and using this to correct the resulting single-dish pointing errors. We expect that this will be done by modifying the optics so that the radiometer illuminates an area of about 5m diameter on the dish. The pick-off mirror will be motorized so that it can move the beam to each of the four un-obstructed quadrants of the dish in turn. The differences in the water vapour amounts in these four positions then give the gradients in the Elevation and cross-Elevation directions. Since the pointing errors will be small and we can use a fairly rapid switching, the only errors that are likely to be significant here are those due to the intrinsic noise in the radiometers.

To give an indication of the performance in this mode, suppose that the beam spends 0.192 sec in each position and takes 0.096 sec to move between them. (These numbers are chosen to be integer multiples of the 48msec tick.) A complete cycle is therefore completed in 1.152 seconds. The integration time going into each sample, s, is 0.192 secs, so the errors due to noise on these will be ~2.28 times those given in table 3 above. A good estimate of the path errors on the samples is 6.7*(1+w) microns, for the range $0.5 < w < 2.5$. We then form a mean and difference of the samples from the 4 sectors, e.g. $md = (s_1 + s_2 - s_3 - s_4)/2$, and obtain the gradient by dividing by the spacing between the beams, which we take to be half the dish diameter, or 6m. The error on md is the same as that on an individual sample, so we find that the error in each angular coordinate is 0.23*(1+ w) arcsec. For the 50th percentile value of the water w = 1.275 this corresponds to 0.52 arcsec, which can be compared to the goal of 0.6 arcsec for the antenna pointing.

Remembering again that we have made quite conservative assumptions about the noise performance of the radiometers, we see that it should be possible to make a useful correction to the pointing on timescales of about 1 second. If the pointing fluctuations vary more slowly than this, the intrinsic errors could be reduced by smoothing the data. To do pointing corrections on timescales that are significantly faster than once per second would require more sensitive receivers.

Conclusions

Although the radiometers use un-cooled receivers they should have considerably better sensitivity for phase correction than is required by the specification. We also expect that their stability will more than adequate. Achieving an absolute accuracy of 1% will be more difficult and will probably require careful calibration. Correction of single-dish pointing errors due to water vapour should be possible.