

ALMA Memo 592 V1.0: The Power Spectrum of Atmospheric Path Fluctuations at the ALMA Site from Water Vapour Radiometer Observations

R. Bolton, B. Nikolic, J. Richer

Astrophysics Group, Cavendish Laboratory, Cambridge CB3 0HE, UK
email: rosie@mrao.cam.ac.uk

14 September 2011

ABSTRACT

This memo investigates the properties of atmospheric turbulence above the ALMA site using measurements from the Water Vapour Radiometers (WVRs), in particular testing the Kolmogorov turbulence model and searching for evidence of an outer length scale.

We fit model power-spectral density (PSD) curves to the measured power spectra for a number of observations and extract the slope, the amplitude and (when convincingly found) the break frequency for a sloping component and also the amplitude (or limit) on a white noise (thermal) component. We find that the majority of the power-spectra are consistent with a thick-screen Kolmogorov turbulence model, but find a transition from thick- to thin-screen behaviour in four of 18 observations. When we scale the frequency of these transitions to spatial scales using the measured (ground-level) wind speed, we find that these transitions occur at lengths in the range 40–400 m.

The WVRs are used to predict and correct for the line-of-sight path fluctuations caused by water vapour for each of the 54 12 m ALMA antennas. When noise is present, applying the WVR correction on baselines where the path fluctuations are dominated by this noise will *increase* the phase fluctuations rather than reducing them. We use the recovered PSD of fluctuations to estimate the lower limits on time scales and baselines on which application of the WVR data might be expected to help.

By combining several adjacent observations, we have assembled WVR path estimates over 3 hours. We find no evidence of an outer scale in the turbulence (which would present itself in a turnover in the PSD at long timescales or spatial frequencies), which suggests that in this case the outer scale was longer than approximately 16 km.

1 INTRODUCTION

One of the key challenges for astronomical observations with large interferometers operating at mm and sub-mm wavelengths are the fluctuations in the *effective* path to each of the antennas which are caused by the fluctuating quantity of water vapour in the troposphere along their line sight (see, e.g., [Evans et al. 2003](#), for summary of conditions at the ALMA site). The cause of the longer effective path (or, more precisely, of the delay) is the high refractive index of water vapour but the driver of the *fluctuations* is the natural turbulence in atmosphere. Understanding the spatial and temporal structure of the path fluctuations and the turbulence that drives them is useful for improving our techniques for correcting these path fluctuations and for quantifying the residual effect on the science data.

Most of the previous work, in the context of astronomy, on understanding this structure of the atmosphere has been done by measuring the resulting phase fluctuations in an astronomical interferometer. For example, [Lay \(1997\)](#) used data from the 12-GHz phase monitoring interferometer at the Owens Valley Radio Observatory to generate temporal phase power spectra and compare these predictions of models of a turbulent atmosphere. [Lay \(1997\)](#) found results consistent with [Kolmogorov \(1941\)](#) theory but did not fit

the slope of the power spectra. [Carilli & Holdaway \(1999\)](#) used the Very Large Array (VLA) interferometer in New Mexico to study the phase Root-Mean-Square (RMS) variation as a function of baseline length. They found that the root phase structure function exhibited the behaviour expected from [Kolmogorov \(1941\)](#) thick (relative to shortest baseline) screen turbulence models, and used these results to make predictions of required phase calibration cycle times for observations above about 20 GHz.

In this memo, we use data from ALMA 183 GHz Water Vapour Radiometers (WVRs) to study the structure of the atmospheric water vapour turbulence. The WVRs are installed on each of the 12-m diameter antennas in ALMA and are normally used to predict and correct the path fluctuations due to the water vapour. Here, we use the path fluctuation prediction from WVR observations to study the properties of the atmosphere and do not consider the measurements from the astronomical interferometer. An important advantage of using the WVR path fluctuation estimates rather than the phase of astronomical signal is that the WVR path fluctuation estimates are *not* differenced between a pair of antennas (i.e., on a baseline) but rather are estimates of the total path fluctuation seen by each antenna individually.

We characterise the structure of the atmospheric turbulence

by estimating the power spectra of the path fluctuations. Our observations intrinsically measure the temporal fluctuations in path and therefore we generally estimate the temporal power spectra. If the Taylor (1938) hypothesis of ‘frozen-screen’ turbulence is applicable, the power spectrum of temporal variations can be converted to spatial spectra and correlation functions using an estimate of the wind speed (illustrated schematically in Figure 1). We intend to investigate the validity of this hypothesis in future work but in this memo we assume that it is valid.

The Kolmogorov (1941) theory predicts (see e.g., Tatarskii 1971) that power spectrum of a turbulent field is a power-law with an intrinsic index of $-5/3 \sim -1.67$. However, because the measured path fluctuations are in fact the integral along the line of sight through a turbulent layer of finite thickness (see Figure 1 again) a steeping by -1 in the index of the power law is expected at time- (or equivalently length-) scales which are shorter than the thickness of the layer. In summary the theory predicts a spectrum with broken power-law shape, with $-8/3 \sim -2.67$ slope in the thick-screen (short timescale) regime and $-5/3 \sim -1.67$ in the thin-screen (long timescale) regime. As indicated in Figure 1, the transition between thick- and thin-screen regimes occurs at a timescale of $t = D/w$, where D is the thickness of the turbulent layer and w is the effective wind speed. The predicted spectrum for *phase* fluctuations is in general different as these are due to the *difference* of paths fluctuations on the baseline used in the measurement. However, in the case when baseline is longer than the ‘outer’ scale of the turbulence, that is the atmospheres above the two antennas are uncorrelated, the expected shape of the spectrum again has slopes $-5/3$ and $-8/3$. Therefore, the model shown in Figure 5 of Lay (1997) is applicable for long baselines as well as for the un-differenced measurements presented here.

The plan of this memo is as follows. In Section 2 we describe our observations, the algorithms used to fit models to the observed spectra the results of this fitting. In Section 3 we interpret these fitting results in the context of models of turbulence, and in Section 4 we summarise our findings.

2 MEASURING POWER SPECTRA OF ATMOSPHERIC PATH FLUCTUATIONS

This study is based on data from the ALMA WVRs. The WVRs are recording measurement of sky brightness along the bore-sight of 12 m diameter the antennas¹ during all ALMA observations. In this study, we re-use data from some observations of quasars that were originally intended for tests the efficacy of phase correction based on WVR data. These observations are suitable because they are reasonably long (~ 20 minutes or more) and the antennas are continuously tracking a single object, so that there are no complications associated with changing line of sight between science and calibration sources that are typical of most science observing.

During these observations the WVR sky brightness measurements from each of the antennas are internally integrated to 1.152 seconds and sampled and recorded at this same interval. These measurement of sky brightness are then converted to estimates of path fluctuations by the `wvrgcal`: the primary purpose of this program is correction of phase errors due to the path fluctuations using its CASA calibration table output, but it also outputs the computed

path fluctuations that can be used for the present analysis (these are output to a HDF5-format file if these libraries are detected at compile time). The internals of the `wvrgcal` program are discussed by Nikolic (2011), but the relevant workings can be summarised as follows. First, a physical model for emission from water vapour in the atmosphere is fit to the absolute sky brightness recorded at one instant in time. This model is then used to predict the coefficient which relates the sky brightness fluctuation to the fluctuation in effective path for each channel of the radiometers. Finally, the sky fluctuations are multiplied by these coefficients and the four WVR channels are combined together to give estimates of path fluctuations to each antenna.

Because the antennas were tracking a celestial source and we wish to make comparisons between different observations, we rescale the estimated paths by $\sin(\theta)$, where θ is the antenna elevation, to get an estimate of the path length fluctuations toward zenith (i.e. we make a simple airmass correction). In practice the effect of changing airmass and its subsequent correction are relatively small: for example, during a 20 minute observation starting with the source at an elevation of 50 degrees, the end elevation will be at least 45 degrees which corresponds to a change in airmass of 8%. The effect of the elevation changes on the resulting spectra (or incorrect corrections to these) would be to increase the power on long time-scales, steepening the slope. For longer runs, any errors will become more significant and therefore the proposed static WVR, which will continuously take readings in the zenith direction would be useful for further studies of atmospheric turbulence on longer timescales than examined here.

We reviewed the available single quasar observations and identified 18 that are useful (i.e., are of suitable length, there was a minimum number of antennas available and there were no hardware or software problems) for the present study. The majority of these are ~ 20 minutes long and were made during a two-week period in April 2011, but we also consider six longer observation of ~ 1 hour each, 4 from December 2010 and 2 from June 2011. These observations include between 4 and 12 antennas with working WVRs (median 8). We show the temporal power spectra for all these observations in Figures 9 to 11. The power spectra are computed from the full length of each run, but considering only about 20 minutes at a time (so that the short and long data sets are processed consistently).

We calculate the power spectral density using Welch’s average periodogram by using the `python(2.5)` function `psd` in `matplotlib.mlab` (see http://matplotlib.sourceforge.net/api/mlab_api.html#matplotlib.mlab.psd). We use the same parameters for all runs (specifically: `NNFT=512`, `scale_by_freq=True`, `noverlap=0`, `detrend=none`, and `window=mlab.window.hanning`) The resulting frequency resolution in the PSDs is 0.0017 Hz. In Section 3, we also briefly examine power spectra at slightly higher frequency resolution available from longer runs.

Within each of our observations, data are recorded in 10-minute long ‘scans’ separated by gaps that are about 30 seconds long. We take this into account in computation of the spectra by padding replacing missing data during gaps by path fluctuations estimated by linear interpolation between the values at the start and end of the gap. This will cause some ripple in the power spectra, but much less than if zero-padding was used. The ripple is more pronounced in the spectra computed from shorter 20-minute observations than from the longer 1-hour observations, see for example Figure 10(b).

The estimated path fluctuations and their power spectra generally show similar features:

¹ The beams of astronomical receivers are arranged around the WVR beam so that the offsets are between 3 and 10 arcminutes – the effect of this offset has been discussed by Nikolic et al. (2007)

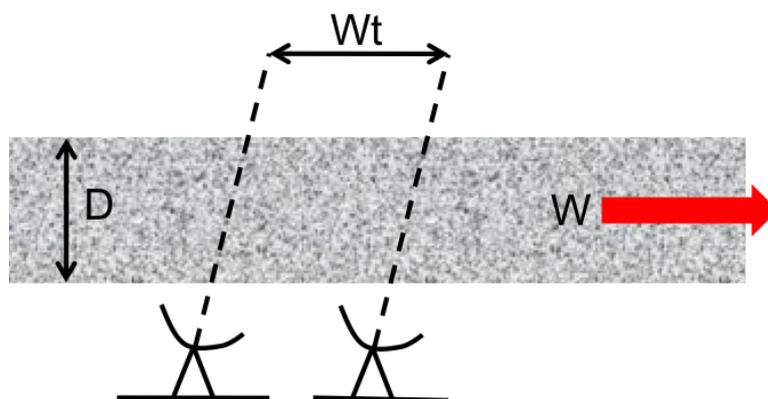


Figure 1. A sketch showing a phase screen blowing past an antenna: if there is no time evolution of the screen itself then we can imagine that the portion of this screen seen by the antenna changes as it moves. After time $t = D/W$ the sampled portions of the screen are further apart than the total thickness of the screen and we move from the thick-screen (three dimensional) turbulence model into the thin-screen (two dimensional) mode.

(i) At the highest frequencies (i.e., shortest time separations) the spectra become flat, i.e., independent of frequency, suggesting that the path fluctuation is dominated by uncorrelated noise component

(ii) At intermediate frequencies the spectrum appears to follow a simple power law

(iii) In some cases, a break at low frequencies can be seen where the slope of the spectrum flattens, i.e., the slope is shallower at the lower frequency side compared to the high frequency side

(iv) In many observations it is clear that one antenna sees quite different fluctuations compared to those seen by the other antennas: in December 2010 this is antenna DV09 and in April 2011 this is antenna DV01 (see, e.g., Figure 2). Figure 8 shows the geometry the antennas in April 2011, showing that one of the antennas was on a long baseline where the atmosphere can be expected to be different.

These features are broadly consistent with fluctuations from a turbulent random field as described in the Introduction, together with additional short term uncorrelated noise as expected from the internal intrinsic noise in the WVR mixers and amplifiers.

2.1 Fitting accuracy

We interpret the measured spectra by fitting models separately to each observation. The fitting provides estimates of the model parameter, that is, the slope of the PSD, the amplitude of the white noise component and the presence of any break in the power law consistent with a thick-thin screen transition (we search for a frequency where the slope flattens by 1.0, but do *not* force the value of the slope to be -2.67 in the thick-screen regime). We check the accuracy with which the estimates can be made from observations by analysing simulated path fluctuations with known spectral properties.

The simulated data are generated to consist of a pure power law component (with a specified slope, which we refer to as the ‘coloured’ component) and an uncorrelated Gaussian (white) noise component. The total power in each of these two components can be controlled to represent the range of observed conditions at the site.

We have generated several hundred such time series and used our fitting algorithms to estimate the slopes. We scale the ‘coloured’ component so that its PSD has a fixed target value at 1Hz and add white noise components with amplitudes corresponding to 0, 5 and

20 microns of path fluctuation. Because the coloured components are normalised to have the same power on timescales of 1 second, the total RMS values over 20 minutes vary as a function of slope. With a slope of -2.0, the coloured components have an amplitude corresponding to an RMS value of around 160 microns, which is within (but towards the low end of) the range seen in the data observed at the ALMA site. Typical values for the ‘white noise’ component seen in real data are between 5 and 15 microns, so the accuracies and any biases that we recover from these simulated data sets should reflect the real situation well.

In Table 1 we show mean estimated slope values for three sets of simulations with different input slope values and with three different white noise levels. In each case 100 time series data were simulated. The results show that the typical accuracy with which we estimate the slope of a single observation is about 0.15. The accuracy decreases if more white noise is added to the simulations, as expected since the range of frequencies over which the coloured component dominates the PSD is reduced. The results, however, also appear to show that the estimated slopes are slightly, but systematically, biased steeper as more noise is added. These biases are well within one standard deviation of individual observation in all cases². Therefore the presence of a white noise component in the real observations should not significantly bias estimated spectral slopes, but we should not attempt to interpret these slopes to an accuracy better than about 0.15.

2.2 Observations

We compute power spectra for each antenna and observation and then calculate the geometric (i.e. log-space) mean of the spectra across antennas to obtain best estimates of the spectra (shown in Figures 9 to 11) during each observation. As mentioned in relation to the simulated data, because we use finite segments of data, we are susceptible to statistical variance. To take this into account we make an initial fit to the antenna-averaged PSD from each observation

² We believe we have traced this apparent bias to the way the data are simulated rather than the fitting. We have not taken the step of re-simulating the data as the effect is smaller than the statistical uncertainty.

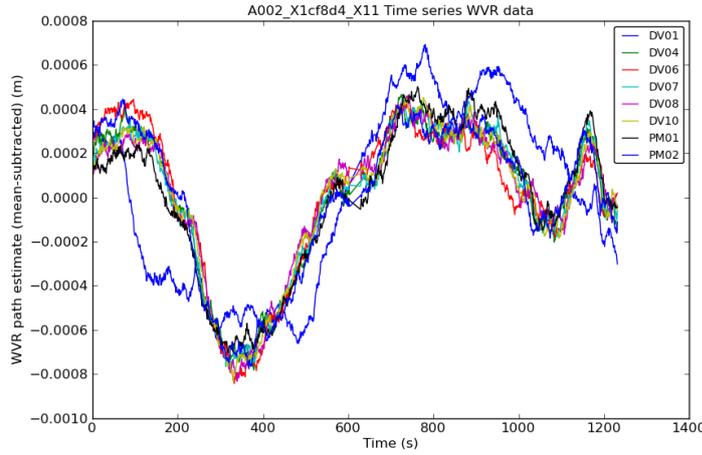


Figure 2. A typical example of (mean-subtracted) path fluctuations estimated from WVR data from one 20-minute observation (A002_X1cf8d4_X11), during which antenna DV01 was ~ 600 m from the main cluster of antennas.

Table 1. Estimated spectral slopes for different combinations of input slopes and Gaussian noise amplitudes in simulated data.

PSD slope used to generate date	Amplitude of white noise component, microns	Mean recovered psd slope	Standard deviation
-1.5	0	-1.50	0.15
-1.5	5	-1.52	0.17
-1.5	20	-1.59	0.26
-2	0	-2.02	0.15
-2	5	-2.02	0.12
-2	20	-2.02	0.17
-3	0	-3.02	0.11
-3	5	-3.06	0.12
-3	20	-3.08	0.13

and then use the resulting best-fitting model to generate a set of 100 simulated observation with similar properties. We then compute the power spectra of the simulated observations and use the variance in these to estimate the statistical variance in the real data. The variances estimated in this way are usually somewhat larger than the variance estimated simply by comparing the PSD from the different antennas. We then use the real data and the uncertainty estimates from the simulations to fit final models to the observations.

The fitting of models is based on the `bnmin1` Bayesian inference library (see Nikolic 2009). We fit two alternative models:

(i) ‘Single Power Law’ (SPL), where the model takes three parameters: the slope and the amplitude of the power law component and the amplitude of the white noise component.

(ii) ‘Broken Power Law’ (BPL), which is like the above model, but the slope changes by +1 at a break frequency which is a free parameter. Therefore, the BPL model fits spectra where the low frequency slope is shallower (by 1.0) than the intermediate frequency slope. As with the SPL model, we also estimate the amplitude of the white noise component. The BPL model therefore searches for Kolmogorov thick to thin screen transitions

We use the Bayesian ‘evidence’ (see for example Jaynes 2003; Sivia & Skilling 2006) to decide if the more complex BPL model explains the observations significantly better than the SPL model. We base the

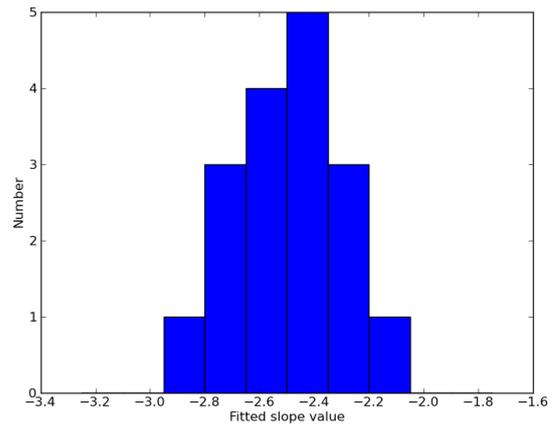


Figure 3. Histogram of recovered slopes (With the ...X191 data set slope set to -2.61 rather than -3.61). Mean -2.50. Standard Deviation 0.18.

Table 2. Parameters and search ranges used for the SPL and BPL models.

Parameter name	Search range used for model fits
Power Law Comp. Slope	-4.0 to -1.0
Power Law Comp. Amplitude	0.1 to 1000 microns
Power Law Comp. Break Freq.	10^{-3} to $10^{-0.5}$ (not used in SPL fit)
White noise Amplitude	0.1 to 1000 microns

threshold for selecting a more complex model on the Jeffreys’ Scale (Jeffreys 1961), and only choose the BPL model if the ratio of the Bayesian evidences is higher than 148 (i.e. $\log \epsilon > 5$), corresponding to a ‘decisive’ favouring of the BPL model over the SPL model.

The ranges of model parameters considered in the fitting procedure (i.e., the ‘prior’ ranges) are summarised Table 2. The estimates of model parameters resulting from the fitting are shown in Table 5. The results show a remarkably similar slope value for almost all fits: the range is -2.81 to -2.18 , the mean is -2.50 and standard deviation is 0.18 (see Figure 3).

Data run	Thick-thin break distance scale (m)
A002 X1ce560 Xc6	41 +/- 8
A002 X1d1cbf X3b	100 +/- 40
A002 X20c6c7 X4b	390 +/- 160
A002 X20c6c7 X71	225 +/- 37

Table 3. Thick- to thin- screen Kolmogorov scales for the 4 runs where such a break was found.

3 CHARACTERISTIC TIMESCALES AND IMPLICATIONS FOR WVR PHASE CORRECTION

We find evidence for a break in the slope of the power spectrum in five of the 18 data sets analysed. In for four of these five observations in which the BPL model is preferred, the slope at higher frequencies is -2.71 ± 0.08 (changing by constraint to -1.71 at the lower frequencies), which is in excellent agreement with the predictions of the Kolmogorov theory of turbulence and a thick-screen to thin-screen transition (this predicts -2.67 at higher frequencies).

If we assume that the wind speed aloft is the same as the wind speed at 8m height (this is likely to be an underestimate, see e.g. Holdaway (1995)), we can recast our temporal scales as spatial scales. The equivalent spatial scales for the inferred thick-screen to thin-screen transition is shown in Table 3: in the model these approximately correspond to the thickness of the water vapour layer. We note that the last two observations (...X4b and ...X71) were made consecutively and that their break scales are mutually consistent. Observation A002_X1ce560_Xc6 was made during very high levels of PWV (7mm) when the weather was ‘cloudy’.

For all our other observations we do not see a break from thick to thin screen behaviour. This suggests that the turbulent screen is thicker than 100 m or so (see Figure 6). We note that Holdaway (1995) show indirect evidence to suggest that the turbulent layer is thicker than 300 m or so.

The remaining observation for which the BPL model was preferred appears to be anomalous. For this observation, A002_X1849a5_X191, the slope of the spectrum at high frequencies is -3.61 (changing by constraint to -2.61 at lower frequencies). The estimated break timescale is 11 s, which is a similar timescale to the breaks in two of the other four observations for which the BPL model was preferred. However, the wind speed during the run A002_X1849a5_X191 was very low, and the 11 s timescale corresponds to a spatial scale of 20 ± 4 m. Because the WVR measurement are made with a beam with finite width as it passes through the water vapour layer, we do expect to see the amplitude of the sloping component in the PSD drop rapidly when spatial scales comparable to or smaller than the beam size are sampled. The WVR beam is between about 15-20m across, depending on height (see Nikolic et al. 2007). The fact that the observation with the lowest wind speed, which therefore samples the highest spatial frequencies of all observations, is the only one where this break is detected hints that this beam-averaging might be responsible.

We have also examined the observed power spectra for evidence of turn-over (or flattening to zero slope) of the spectrum at long timescales. According to the theory of turbulence, there must be such an ‘outer’ scale of turbulence on which the energy is injected. On scales longer than this, the fluctuations become entirely uncorrelated. We see no evidence for such turn over in the 20-minute observations that we analysed. The lowest frequency we measure for these observations is approximately 10 minutes, therefore we only

constrain the outer scale of turbulence to be greater than $\sim 300s \times w$ where w is the effective wind speed.

Besides the 20 minute observations, we also have a few longer observations, several of which were made consecutively. We combine these into one 2-hour and one 3-hour data set. We show the estimates of path fluctuations and the resulting spectra of the combined data in Figure 4. The plots corresponding to the two data sets show substantially different behaviour: the path estimates plotted in Figure 4(a) clearly shows variation on long time scales. Some of the overall trend in the path could be due to errors in the airmass correction (elevation at the end of observation was only 45 degrees), but in any case the differences in the PSD with or without the airmass correction are well inside the error bars even at the lowest frequencies. The spectrum does not show convincing evidence of a flattening on the longest time scales shown here which are about 1 hour. The wind speed (measured on the ground, not in the water vapour layer) was around 4.5 m s^{-1} , so one hour corresponds to a spatial scale of ~ 16 km.

The second long observation (~ 2 hours in total, Figure 4b) shows much higher power on timescales up to 30 seconds, and a nicely defined break where the slope transitions from ~ -2.6 to ~ -1.6 — see the fitting plots for the individual 1-hour runs in Figure 11. Again here there is no obvious turn-over in the spectrum, out to frequencies of 10^{-3} Hz or 17-minute times scales. The ground-level wind speed was high, at 10.5 m s^{-1} , so 17 minutes corresponds to 10.5 km of transverse movement.

A thorough analysis of the outer scale of the turbulence clearly requires longer runs during which the WVRs are pointed at a fixed direction. The proposed ALMA fixed monitoring WVR will be very useful for constraining the outer scale of turbulence, which will become increasingly relevant as ALMA begins to use longer baselines. The results of this preliminary analysis suggest that the outer scale of the turbulence is as at least as long as the longest baselines that are anticipated to be used by ALMA. This suggests that the magnitude of phase fluctuations will increase beyond what is measured on the baselines used so far. Consequently, effective correction of these phase fluctuations will become more important, and this in turn will depend increasingly on accurately predicted coefficients that convert sky brightness fluctuations in path fluctuations.

3.1 Spatial power spectra

The observations we analyse here were made under different prevalent wind speeds which makes it difficult to fairly compare the temporal power spectra between the observations. We can make a better comparison by converting these spectra into spatial power spectra and overlying the results of all observations, as shown in Figure 5. In this calculation we used the average ground-level wind speed for each observation and shifted all data by $\log_{10}(w)$ in the left and up direction. (The units are m^{-1} rather than radians/m - conversion to true ‘ k ’ values just leads to another shift in x and y , but by the same amount for all runs.) The spread of power at a particular frequency seen in Figure 5 corresponds to almost 4 orders of magnitude in power, equivalent to a factor of 100 in the fluctuations. We investigate the relationship between the total Precipitable Water Vapour (PWV) and amplitude of fluctuations at a fixed frequency by plotting in Figure 6 the power at frequency of $1/40 \text{ m}^{-1}$ versus elevation correction precipitable water vapour at the middle of the observation. It can be seen there is a general correlation between the two quantities, which is as expected if we assume that the fractional changes in water vapour concentration caused by the turbulence are about the same in all observations.

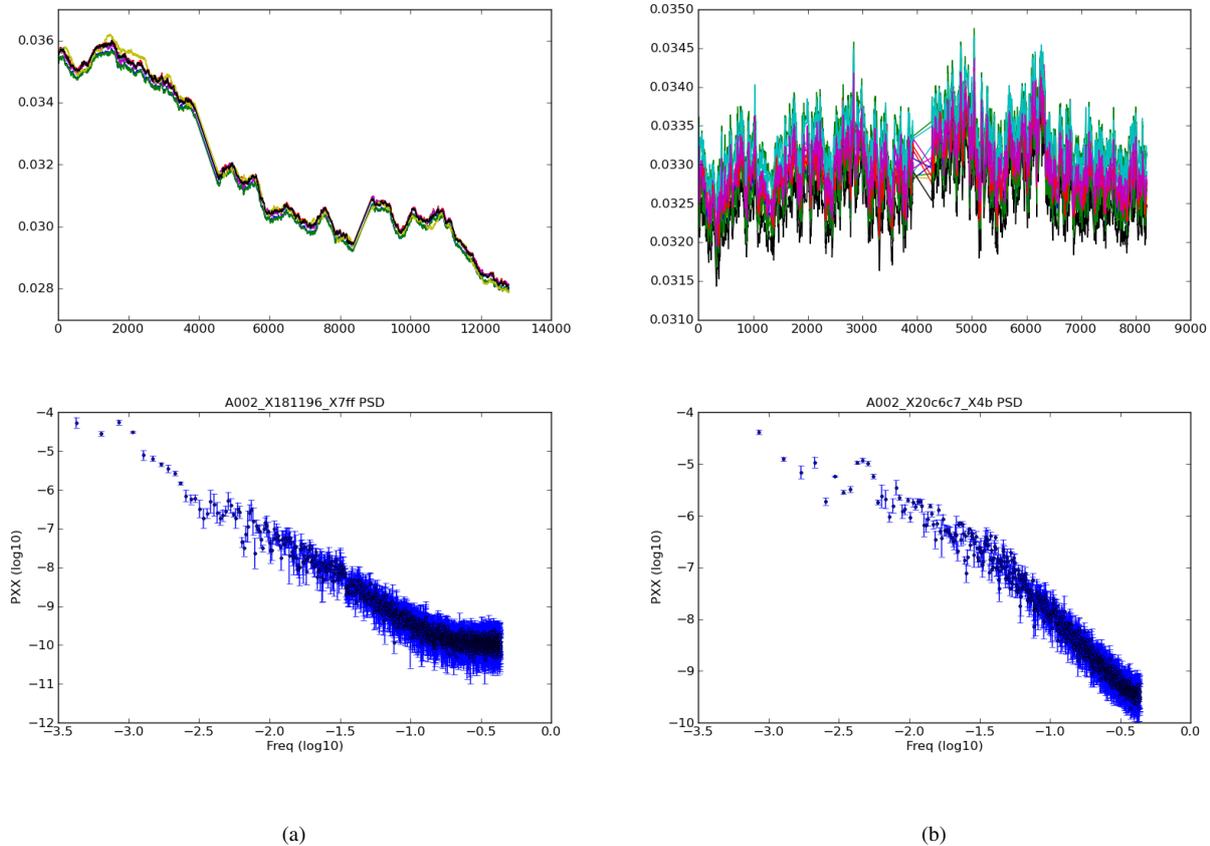


Figure 4. (a): Path and PSD from approximately 3 hours of WVR data taken on 2010-12-20, between 01:00 and 04:00 local time. The total PWV was 3.0 +/- 0.2 mm. (b): Path and PSD from approximately 2 hours of WVR data taken on 2011-06-05 between 04:15 - 06:30 local time. The PWV was 2.6-2.7 mm.

We can also compare the power of path fluctuations we measure with previous measurements with the 300 m baseline site-testing interferometer described by Radford et al. (1996). The results of the measurement by the interferometer are summarised by Evans et al. (2003). We take the values for rms delay fluctuations (in microns) on a 300m baseline at 10, 25, 50 and 75 percentiles from table 1 of Evans et al. (2003) and plot the square of these values as the black diamonds on the spatial PSD in Figure 5. The spatial PSDs shown are generally a little higher than the measured phase fluctuations: for comparison, observation A002_X1cd467_X45 has a PWV close to the median value of 1.3, but its PSD would not pass through the 50th percentile diamond (second from top) in the plot if extrapolated linearly. However, underestimation of the wind speed in the turbulent layer by a factor of 2 would shift the PSD values to the left by 0.3 dec (increasing all the spatial scales) and up by 0.3 dec, making it broadly consistent, and any break in the PSD to the thin-screen regime would help. Scatter in the amplitude of the PSD for fixed PWV (in Figure 6 means that although ...X45 sees median PWV levels, it may not see median path rms values.

3.2 Estimation of ‘Thermal’ Noise

Some of the apparent short-time scale fluctuations of path estimated from WVR data may due to the intrinsic, uncorrelated, noise in the WVR mixers and amplifiers (‘thermal’ noise). If the real path

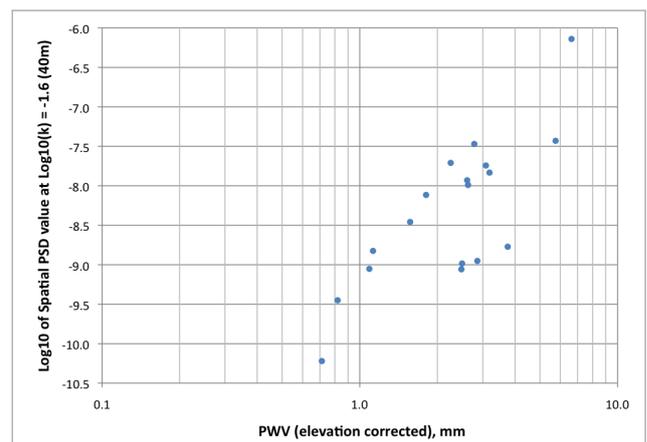


Figure 6. Spatial PSD value at a spatial frequency of $10^{-1.6}$ plotted against the elevation-corrected PWV value for the mid point of the run.

fluctuation between two antennas is smaller than the combination of the noise path fluctuations in the two WVRs, then applying phase correction based on the differenced WVR signals would be counter productive and make the resulting phase fluctuations on this baseline worse. The `wvr_gcal` program predicts the magnitude of path

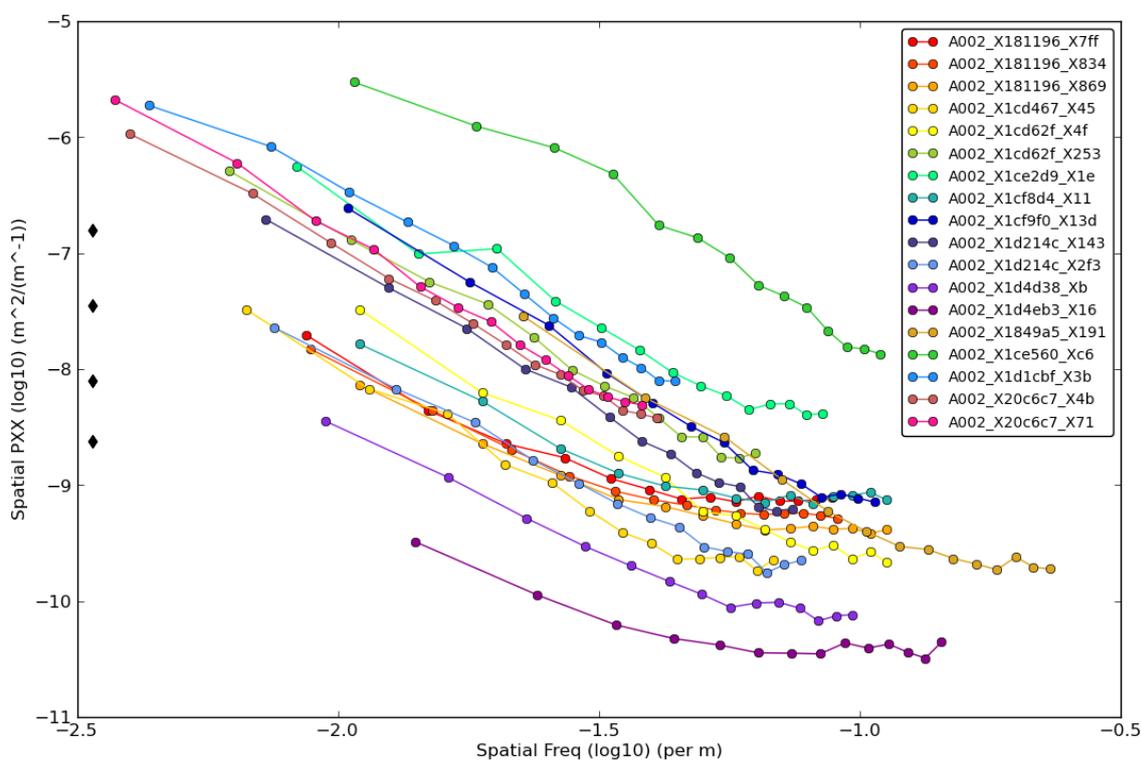


Figure 5. Spatial PSDs for all runs. Note that the 5 BPL runs are plotted last. ...X191 is the run with the lowest wind speed (and the anomalous BPL fit). The black diamonds show the RMS values of 49, 89, 187 and 394 microns measured using the 300m baseline with the NRAO site-testing phase monitor described by Radford et al. (1996).

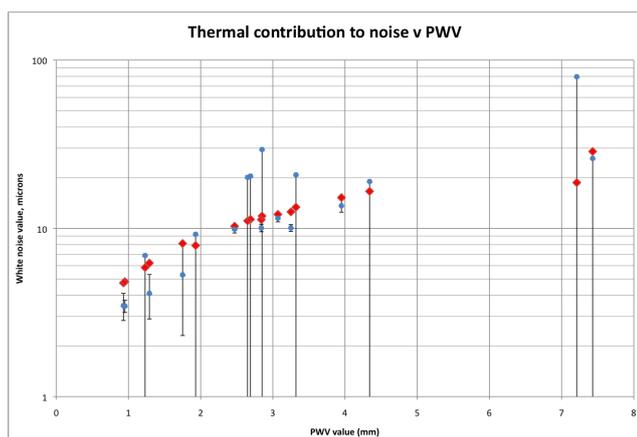


Figure 7. Estimates of thermal noise versus PWV, showing (red diamonds) noise predictions made by `wrgcal` and (blue circles): (i) The white noise component measured from the model fits (where a good fit was made) – these points have two-sided error bars; and, (ii) Upper limit on noise calculated from the lowest value of the PSD (points with negative error bar only).

fluctuation due to thermal noise and in Figure 7 we compare these predictions with estimates of the white noise component extracted from the fitting observed power spectra. We plot the estimated white noise component versus the (non-elevation corrected) PWV. Higher

values of PWV correspond to greater noise (approximately scaling as white noise amp \propto PWV^{0.8}). The plot shows that, in general, the `wrgcal` predictions are in line with the recovered parameters.

Estimates of the effect of noise on the path fluctuations and of the amplitude and slope of the power-law component can be used to compute the timescale on which the thermal noise dominates over the real path fluctuations. Using the wind speed, this can be converted to a minimum baseline length on which WVR phase estimates should be applied without smoothing. Below we use the estimates for the noise path from fluctuation from `wrgcal` as they are available for all of the observations. The results are shown in Column J of Table 5: the minimum baseline length to apply WVR phase correction can be as long as 40m (in low winds and with low PWV), down to as low as 4m (in poor weather, with 7mm PWV). These estimates in terms of length-scales can be converted back into a suggested smoothing timescale by dividing column J of Table 5 by the wind speed in column F; these timescales range from 1 to 14 seconds. We plan to implement and test the effects of such smoothing in `wrgcal` in continuing work.

4 SUMMARY AND FURTHER WORK

We have analysed estimates of path fluctuations from WVR measurements in 18 observations made between December 2010 and June 2011. These observations cover a wide range of observing

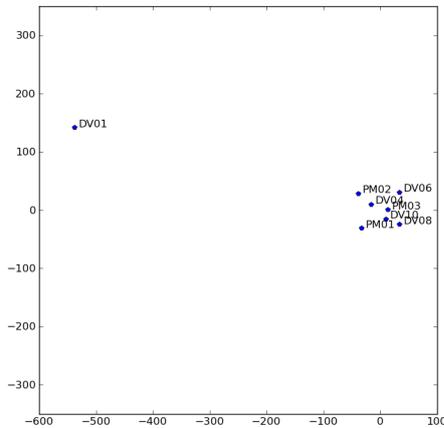


Figure 8. Sample antenna positions from run A002 X1cd62f X253, from Apr 2011. Positions are in m offset from the median value. The x axis is eastwards, y is north.

conditions, with the zenith PWV varying between 0.7 and 6.6 mm. We note the following conclusions:

(i) The estimated slopes of the spectra of path fluctuations have a mean slope of -2.5 and a standard deviation of 0.18 .

(ii) By fitting models to simulated data, we have shown that the typical error on the derived PSD slopes is 0.15 , in close agreement with the observationally estimated standard deviation.

(iii) The Kolmogorov thick-screen model of turbulence predicts a slope of -2.67 . The slopes we estimate are consistent with this model

(iv) In four of the 18 observations we detect a break in the slope from -2.71 ± 0.08 at higher frequencies, in a way consistent with transition from Kolmogorov thick screen to Kolmogorov thin-screen turbulence. The thickness of these screens (corresponding to the spatial scale of the change in slope) is in the range 40 to ~ 400 m.

(v) We see no evidence for an outer scale to the turbulence in the data, which constrains any outer scale to be greater than ~ 16 km. This suggests that application of WVR data is likely to increasingly important (and challenging) as ALMA starts using longer baselines.

(vi) One of the observations shows some evidence of an inner scale, which is likely caused by the averaging effects of the WVR beam.

(vii) The amplitude of the spatial power spectrum at fixed frequency (which is dominated by the sloping component) is as expected correlated with total water vapour.

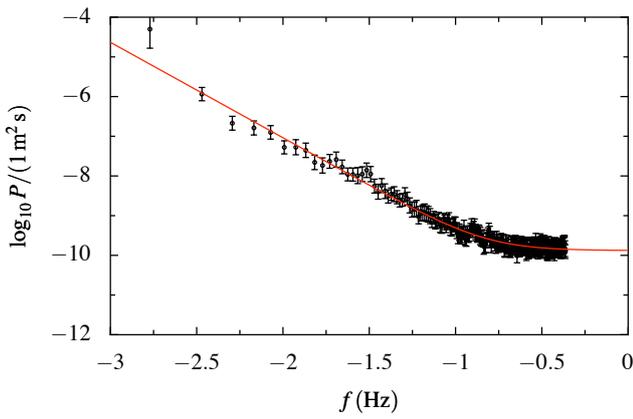
Further work in this area should include an investigation into the effects of smoothing and the baseline dependent effectiveness of WVR application. A radiometer continuously observing zenith would be very helpful in providing consistent monitoring of the atmospheric conditions, for the search for an outer scale and to seek further evidence of the effect of beam-averaging on the WVR data.

ACKNOWLEDGEMENTS

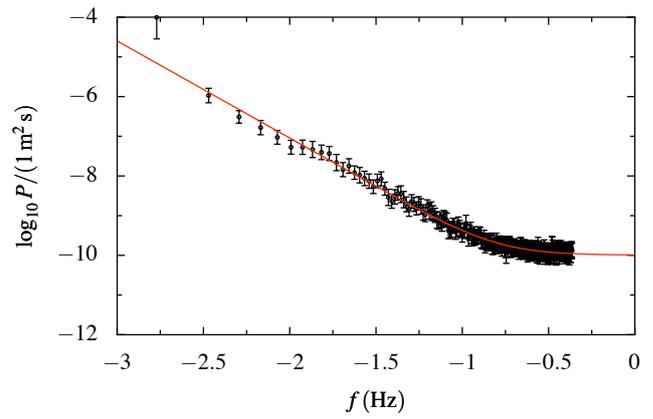
This work was supported by the European Commission's Sixth Framework Programme as part of the wider 'Enhancement of Early ALMA Science' project, managed by the European Southern Observatory (ESO).

REFERENCES

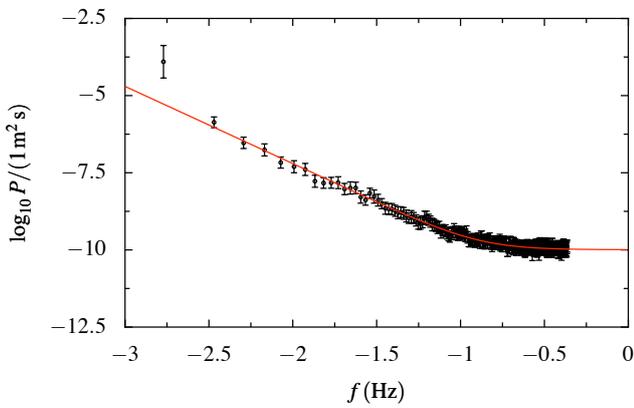
- Carilli C. L., Holdaway M. A., 1999, *Radio Science*, 34, 817. [arXiv:arXiv:astro-ph/9904248](https://arxiv.org/abs/astro-ph/9904248)
- Evans N., Richer J. S., Sakamoto S., Wilson C., Mardones D., Radford S., Cull S., Lucas R., 2003, Site properties and stringency. ALMA Memo Series 471, The ALMA Project. Available from: <http://www.alma.nrao.edu/memos/html-memos/abstracts/abs471.html>
- Holdaway M. A., 1995, Velocity of winds aloft from site test interferometer data. ALMA Memo Series 130, The ALMA Project
- Jaynes E. T., 2003, *Probability Theory : The Logic of Science*. Cambridge University Press
- Jeffreys H., 1961, *Theory of Probability*. Oxford University Press
- Kolmogorov A., 1941, *Akademiia Nauk SSSR Doklady*, 30, 301
- Lay O. P., 1997, *A & A Supplement series*, 122, 535
- Nikolic B., 2009, ArXiv e-prints. [arXiv:0912.2317](https://arxiv.org/abs/0912.2317)
- Nikolic B., 2011, Design and Implementation of the WVRGAL program. ALMA Memo Series 593, The ALMA Project, in preparation
- Nikolic B., Hills R. E., Richer J. S., 2007, Limits on phase correction performance due to differences between astronomical and water-vapour radiometer beams. ALMA Memo Series 573, The ALMA Project
- Radford S. J. E., Reiland G., Shillue B., 1996, *PASP*, 108, 441
- Sivia D. S., Skilling J., 2006, *Data Analysis—A Bayesian Tutorial*, 2nd edn. Oxford Science Publications
- Tatarskii V. I., 1971, *The effects of the turbulent atmosphere on wave propagation*. Jerusalem: Israel Program for Scientific Translations, 1971
- Taylor G. I., 1938, *Proceedings of The Royal Society of London. Series A, Mathematical and Physical Sciences (1934-1990)*, 164, 476



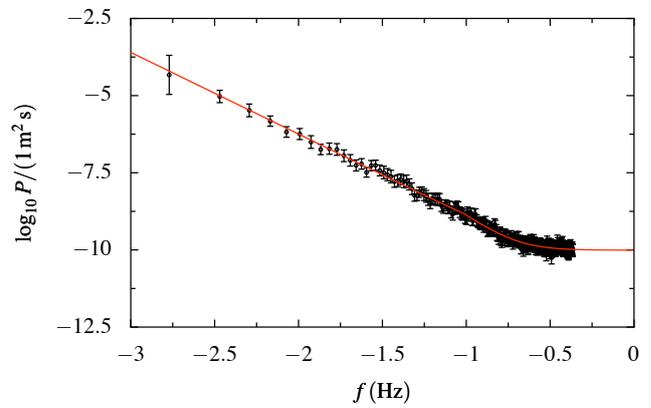
(a) A002 X181196 X7ff Unbroken sloping component. Slope of -2.4.



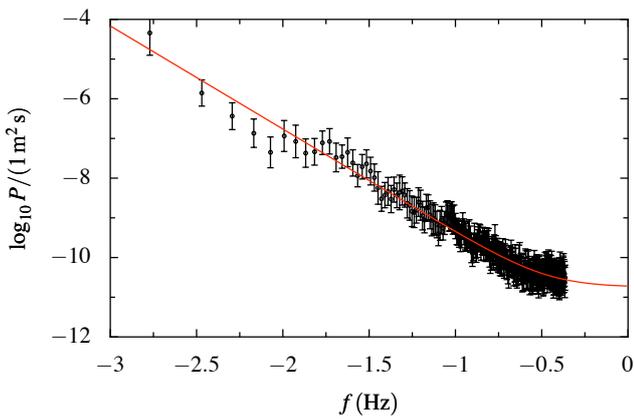
(b) A002 X181196 X834 Unbroken sloping component. Slope of -2.5.



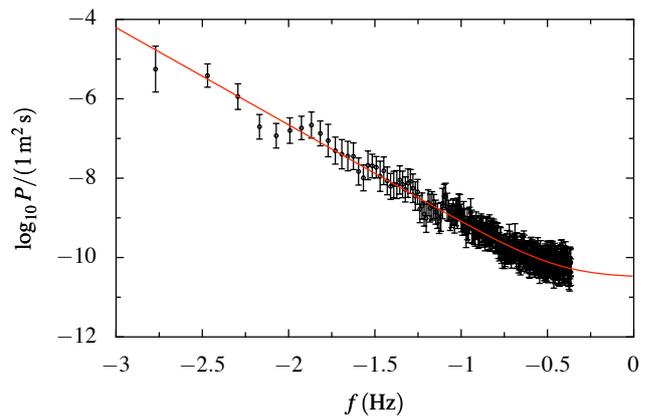
(c) A002 X181196 X869. Unbroken sloping component. Slope of -2.5.



(d) A002 X1849a5 X191. Broken sloping component. Slope of -3.6 for the majority of the PSD the slope is actually -2.6.

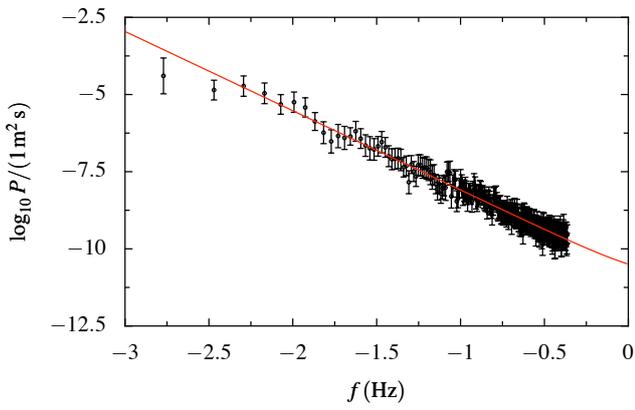


(e) A002 X1cd467 X45. Unbroken sloping component. Slope of -2.6.

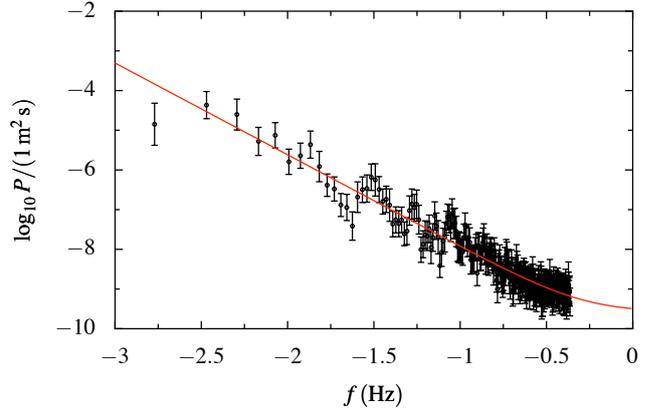


(f) A002 X1cd62f X4f. Unbroken sloping component. Slope of -2.4.

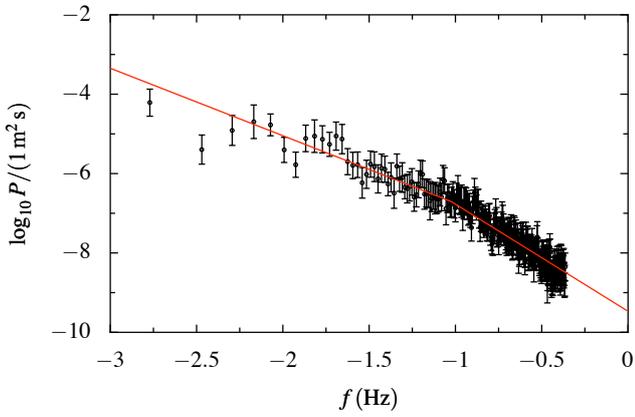
Figure 9. PSDs and fits for the first 6 runs presented in this paper.



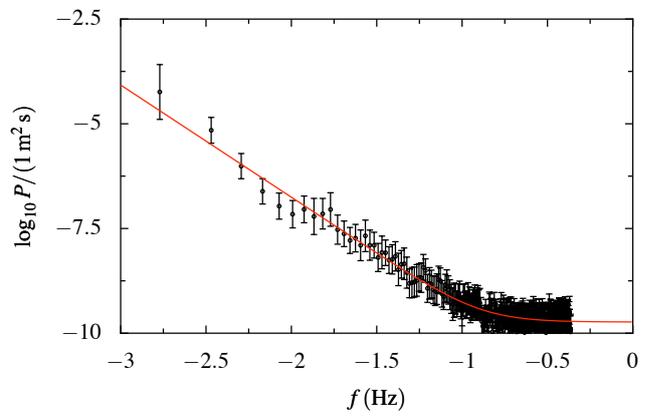
(a) A002 X1cd62f X253 Unbroken sloping component. Slope of -2.6.



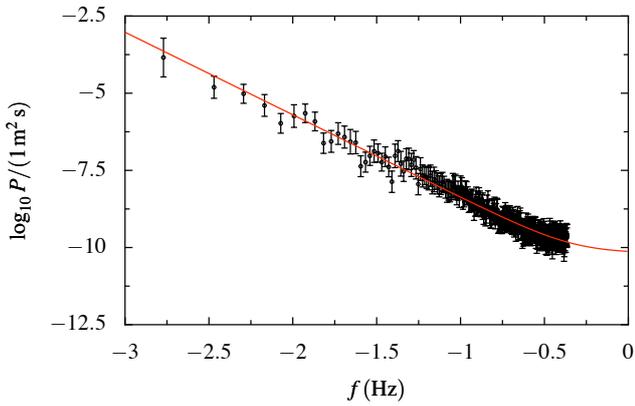
(b) A002 X1ce2d9 X1e. Unbroken sloping component. Slope of -2.3.



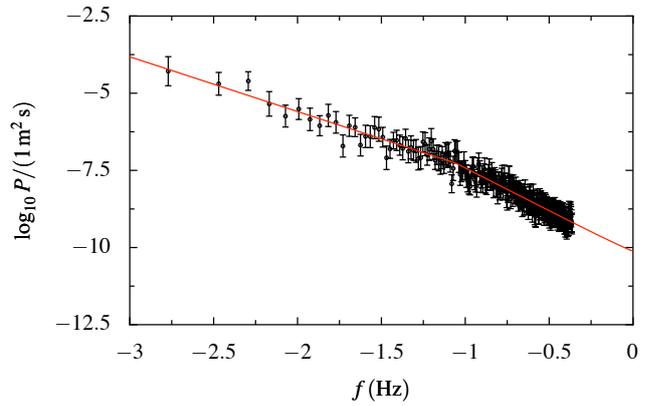
(c) A002 X1ce560 Xc6. Broken sloping component. Slope of -2.7.



(d) A002 X1cf8d4 X11. Unbroken sloping component. Slope of -2.7.

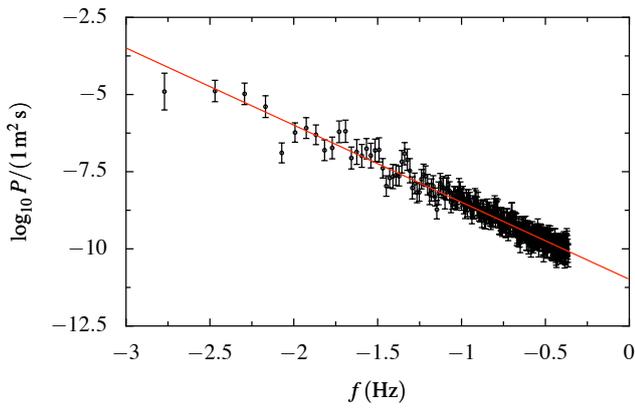


(e) A002 X1cf9f0 X13d. Unbroken sloping component. Slope of -2.6.

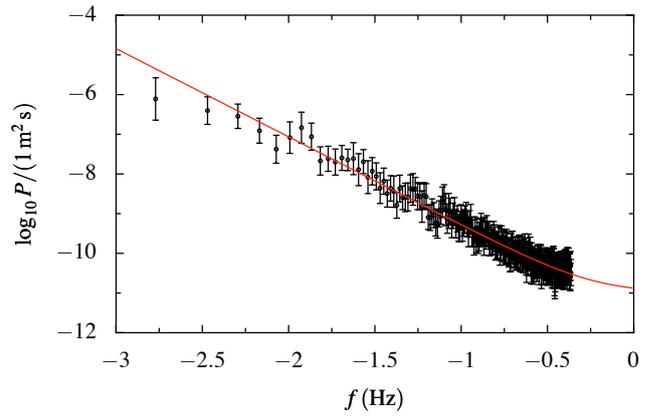


(f) A002 X1d1cbf X3b. Broken sloping component. Slope of -2.8.

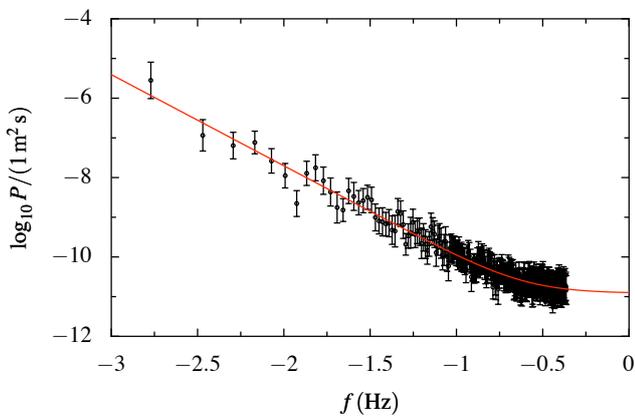
Figure 10. PSDs and fits for the 7th to 12th runs presented in this paper.



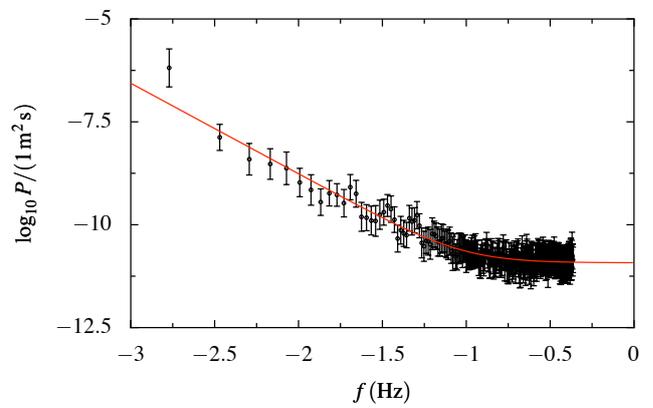
(a) A002 X1d214c X143. Unbroken sloping component. Slope of -2.5.



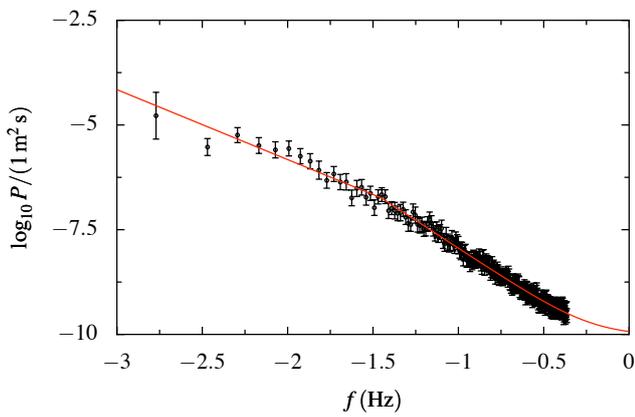
(b) A002 X1d214c X2f3. Unbroken sloping component. Slope of -2.2.



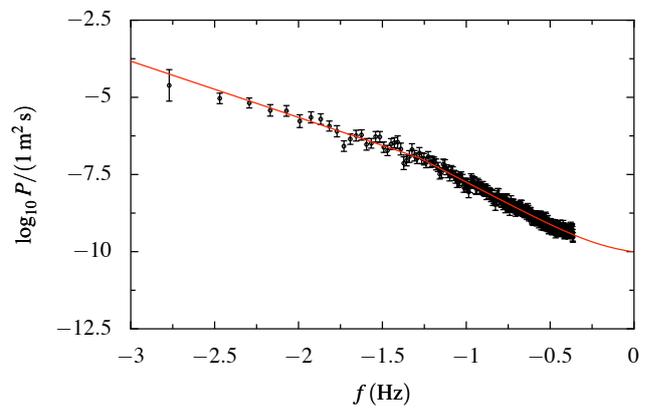
(c) A002 X1d4d38 Xb. Unbroken sloping component. Slope of -2.3.



(d) A002 X1d4eb3 X16. Unbroken sloping component. Slope of -2.2.



(e) A002 X20c6c7 X4b. Broken sloping component. Slope of -2.6.



(f) A002 X20c6c7 X71. Broken sloping component. Slope of -2.8.

Figure 11. PSDs and fits for the final 6 runs presented in this paper.

Table 4: Dataset recovered parameter summary

A	B	C	D	E	F	G	H	I	J	K
Obs name	Date	Time	PWV (mm)	Elevation (mid point) (degrees)	Wind speed (m/s)	PSD slope	Break for BPL fit (s)	White noise or limit (microns)	wrrgcal noise est. (microns)	Noise / slope transition (m)
A002 X181196 X7f	20-Dec-10	1hr	3.1	69	4.7	-2.4		11.5 +/- 0.6	12.1	33
A002 X181196 X834	20-Dec-10	1hr	2.8	61	4.6	-2.4		10.1 +/- 0.5	11.3	31
A002 X181196 X869	20-Dec-10	1hr	3.3	50	3.7	-2.5		10.1 +/- 0.5	12.5	34
A002 X1849a5 X191	22-Dec-10	1hr	2.5	66	1.8	-3.6*	11.0 +/- 2.4	9.9 +/- 0.5	10.3	9
A002 X1cd467 X45	5-Apr-11	20min	1.3	58	6.1	-2.6		4.1 +/- 1.2	6.2	24
A002 X1cd62f X4f	6-Apr-11	20min	1.8	64	3.7	-2.4		5.3 +/- 3.0	8.1	13
A002 X1cd62f X253	6-Apr-11	20min	3.3	74	6.6	-2.6		20.8 (UL)	13.4	15
A002 X1ce2d9 X1e	7-Apr-11	20min	7.4	51	4.9	-2.3		26.0 (UL)	28.6	15
A002 X1ce560 Xc6	7-Apr-11	20min	7.2	67	3.8	-2.7	10.7 +/- 2.1	79.6 (UL)	18.7	4
A002 X1cf8d4 X11	8-Apr-11	20min	4.0	72	3.7	-2.7		13.6 +/- 1.1	15.2	31
A002 X1cf9f0 X13d	8-Apr-11	20min	4.3	45	3.9	-2.6		29.4 (UL)	16.6	14
A002 X1d1cbf X3b	11-Apr-11	20min	2.9	77	9.4	-2.8	10.6 +/- 4.0	9.2 (UL)	11.8	13
A002 X1d214c X143	12-Apr-11	20min	1.9	70	5.6	-2.5		6.9 (UL)	7.9	12
A002 X1d214c X2f3	12-Apr-11	20min	1.2	66	5.4	-2.2		4.6 (UL)	5.9	15
A002 X1d4d38 Xb	14-Apr-11	20min	0.9	62	4.3	-2.3		3.5 +/- 0.6	4.7	23
A002 X1d4eb3 X16	14-Apr-11	20min	1.0	49	2.9	-2.2		3.5 +/- 0.3	4.8	42
A002 X20c6c7 X4b	5-Jun-11	1hr	2.7	83	10.2	-2.6	38 +/- 16	20.1 (UL)	11.1	18
A002 X20c6c7 X71	5-Jun-11	1hr	2.7	76	10.9	-2.8	20.6 +/- 3.4	20.4 (UL)	11.3	18

Table 5: Key parameters, including fitting results and weather data for the 18 data sets presented in this memo. *Note that data set A002 X1849a5 X191 has a recovered slope of -3.6 in the steep part of its PSD, which breaks to a slope of -2.6 for timescales longer than 11s. Notes on columns: Column D contains the PWV value measured *along the line of sight* at the mid point of the observation. In column I, (UL) denotes 'upper limit', taken from the minimum value of the power spectrum. This is done where the fitting resulted in a one-sided histogram - i.e. the white component was too weak to be detected above the sloping component.