

ALMA MEMO 582

Simulating Atmospheric Phase Errors, Phase Correction and the Impact on ALMA Science

B. Nikolic¹, J. S. Richer¹, R. E. Hills²

¹Mullard Radio Astronomy Observatory, Cavendish Laboratory, Cambridge CB3 0HE, UK

²Joint ALMA Office, Santiago, Chile

10 October 2008

ABSTRACT

We present a framework for modelling atmospheric phase errors and their correction by both the fast-switching and water vapour radiometric techniques. Notable features are simulating three dimensional turbulent volumes instead of flat phase screens, considering three representative configurations of ALMA and parametrisation in terms of phase fluctuations on a 300 m baseline, allowing referencing to the site-testing interferometer data. We use this framework to simulate relative point source sensitivity and effective resolution for a range of atmospheric conditions, with and without phase correction. We also consider the effect of phase fluctuation on short ‘snapshot’ observations, where the variance and correlation of atmospheric phase fluctuations between the antennas becomes important.

1 INTRODUCTION

Atmospheric phase fluctuations are an important factor limiting the performance of (sub-)millimetre aperture synthesis arrays. They arise primarily due to fluctuation of the water vapour content in the troposphere along the line of sight of each of the telescopes in the array. The phase fluctuations may also be, in smaller part, due to fluctuations of temperature of the dry air (‘dry fluctuations’, e.g., Stirling et al. 2005).

As summarised by Evans et al. (2003), conditions at the site of ALMA are such that the atmospheric phase fluctuation will significantly limit observations that are possible unless they are effectively corrected. The ALMA project plans carry out this phase correction by a combination of:

- (i) Fast-switching to near-by point sources and calibrating the phase
- (ii) Inferring the residual phase fluctuations from measurements with 183 GHz water vapour radiometers mounted on each of the telescopes in the main array.

Atmospheric phase fluctuations in the context of ALMA have already been studied by a number of authors (e.g., Holdaway 1992; Holdaway & Owen 1995; Carilli & Holdaway 1999; Asaki et al. 2005). In this memo we extend on earlier work by:

- Making full imaging simulations
- Considering three ALMA configurations: the most compact configuration, an intermediate configuration, and the most extended configuration
- Generating a three dimensional turbulent volume to generate phase screens
- Investigating the effect of phase fluctuations on absolute flux measurement and astrometry, as well as sensitivity and resolution.
- Considering phase correction using fast-switching and water-vapour radiometers.

In this note we restrict ourselves to considering only the effect of atmospheric phase fluctuations – we do not consider the implications of thermal-like noise in the receivers at all, nor do we consider potential electronic phase errors.

The goal of this study and the framework for simulations we have developed is to understand the impact of the phase fluctuations on the science that ALMA will do. Besides the impact on sensitivity and maximum resolution, we also consider the effect of the phase fluctuations on astrometry and flux-calibration. We include analysis of simulations without phase correction, simulations with fast-switching phase calibration and simulations with water-vapour radiometer based phase correction.

We do not consider here observations which can be self-calibrated. The reason for this is that although self-calibration will be possible with ALMA for some *specific* observations, at sub-millimetre frequencies the primary beam of ALMA telescopes is small (around 10 arcseconds) and surface brightness of most science targets is relatively low and hence in general self-calibration will not be practical.

2 METHOD FOR SIMULATIONS

The simulations presented in this note may be conceptually split into the following independent parts:

- (i) Simulating the tracks (i.e., the points sampled) by ALMA in the *uv*-plane
- (ii) Simulating the visibilities that would be measured on these tracks if there was no atmospheric or instrumental effects
- (iii) Simulating the corruption introduced by the atmosphere
- (iv) Simulating the inference of the antenna-based phase errors from calibrator observations
- (v) Correcting for atmospheric effects using the calculated antenna based errors

- (vi) Making the dirty image
- (vii) Analysing the dirty image to determine qualities of the observation

2.1 Simulation of tracks and visibilities

We simulated the tracks that ALMA will produce using the `sm` tool in the CASA package (i.e., the `sm.observe` call). For all of the simulations we assumed that the target was at a declination of around -40 degrees and that it was observed around its transit. The total time of observation was around one hour, with integration of times of either 1 s (for simulations of fast-switching) or 10 s (for non-switched simulations). Three different configurations were used (they are shown in Figure 4), spanning the full range of baseline length distributions that the main array of ALMA will have. The position of antennas in these configurations were obtained from R. Reid's web site: <http://www.cv.nrao.edu/~rreid/telesims/almaconfigs.tgz>.

As we are interested in quantifying the basic qualities of observations, we only simulated single unresolved sources. The visibility of such a source is the same on all baselines of the array and so predicting the measured visibilities is a trivial step in this case. We use the `sm.predict` call to do this. In the present simulations we did not consider the effect of thermal noise in the visibility measurements or any instrumental phase effects.

After the generation of the visibilities, they are exported to the standard uv -FITS file format for the next stages of processing.

2.2 Atmospheric phase fluctuations

The basis for our simulation of the effect of the atmosphere on the ALMA observations are large three-dimensional statistical realisation of Kolmogorov fields. These were previously used by Nikolic et al. (2007) to quantify some of the inherit limitations of WVR measurements of atmospheric phase. Here we use these volumes to directly simulate the corruption of phases of the measured visibilities.

The three dimensional volumes have dimensions of $4097 \times 513 \times 513$ volume elements and an intrinsic structure function, defined as:

$$\left\langle [q(\mathbf{r}') - q(\mathbf{r}' + \mathbf{r})]^2 \right\rangle = D_q(r) = 6.88 \left(\frac{r}{r_0} \right)^\xi \quad (1)$$

with an exponent of $\xi = 3/2$. From these volumes we use a number of subsets of layers to simulate different possible atmospheric conditions, i.e., different thickness of the atmospheric turbulent layer. The varying thickness naturally reproduce the observed steepening of the structure function at baselines comparable to layer thickness. The actual method for generation of the turbulent volume is the same as in Nikolic et al. (2007) and will be described in a separate note.

In these simulations we assume that the extra atmospheric path of for radiation received by each antenna is simply proportional to the line integral of the turbulent field q . We denote this quantity as $q'(x, y; \theta, \zeta)$ where x and y are the position of the antenna relative the turbulent volume and θ and ζ are the angles describing the direction of the line integral. In contrast to the approach by Nikolic et al. (2007), we do not consider the finite horizontal extent of the astronomical beam, but simply use the single value of q' at x, y to calculate the phase error at each antenna.

The angles θ and ζ are in principle defined by the azimuth

and elevation of the antennas as they track the simulated source. However, in order to simplify the calculations we assume that the science target source is in a direction directly perpendicular to the phase screen, i.e., $\theta = \zeta = 0$ for the science source. Therefore we can not expect to reproduce change of properties of phase fluctuations as the source rises or sets. As the simulations presented here are one hour or shorter in length, and as the source reaches close to zenith, this approximation should not be too poor.

We do, however, in general want to take into account the different lines of sight through the atmosphere of the telescopes when they are observing the calibrator instead of the science target source. We can do that by assuming that the calibrator is offset by a constant angle and direction; typically we assume $\theta = 1.5$ degrees offset. In order to be evaluate the impact of the divergence of the lines of sight to science and calibrator sources, we also need to assume a height for the turbulent volume, h , which we typically assume to be at 800 unless the thickness of the turbulent layer is larger than 1.6 km, in which case we assume the bottom of the turbulent layer to be at ground level. In these simulations we therefore only need to compute two 'flattened' phase-screens, corresponding to $q'(x, y; 0, 0)$ and $q'(x, y; 1.5, 0)$. In order to quantify the significance of this angular offset of the calibrator, in the results sections we run simulations with the calibrator offset by 1.5 degrees and the calibrator at same position as the science source.

The important final parameter is the vertical thickness of the turbulent layer, which we denote w . This quantity determines the balance of phase fluctuations that are on short spatial and temporal scales compared to magnitude fluctuations on longer time scales. Since the value of this parameter is not clear from observations so far, and probably varies with time of day, we always carry out our simulations for a range of its values.

Once the flattened phase screen has been constructed, it is easy to compute the effect it will have on the observed visibilities. To this we extract the positions of the antennas in the array from the antenna table and translate these to positions on the local tangent plane. Some of the details of this calculation are presented in Appendix A. We then assume that the phase screen is uniformly translated across the tangent plane, in the East-West direction, at a wind speed of 12 m s^{-1} , and compute the path delay to each antenna. We then phase rotate each visibility according to the difference of the path delays to the antennas forming that visibility.

These procedures result in an uv -FITS file with corrupted visibilities.

2.3 Phase Calibration

In a part of this report we analyse some simulations without any attempt to correct the effect of atmospheric phase fluctuations. This is typically not going to be relevant for ALMA but provides a useful baseline against which to measure the efficiency of calibration schemes. In this case there is no phase calibration step and processing proceeds directly to imaging.

The phase calibration step is however applied to the fast-switched simulations. For these simulations we assumed a calibration cycle of a total of 15 s, consisting of 13 s observation of the science target and a 2 s observation of the calibrator. As in the rest of the simulations shown here we assume there is no thermal noise in either observation. We also neglect the slewing time between the calibrator and source.

From each observation of the phase calibrator we calculate the antenna-based phase error. The simulated calibrators are completely unresolved so we can use the complete visibility set (that is




Figure 1. Illustration of the effect of atmospheric phase fluctuations and correction with fast switching. Top plot: antenna-based phase errors (antenna numbers are shown along vertical axis) as a function of time (integration number, horizontal axis) for the medium configuration and 200 m thick turbulent layer. Bottom plot: antenna-based phase errors after correction using the fast-switching calibration, assuming no thermal noise or systematic error. The 15 s calibration cycle is evident from the periodic bands of columns with zero phase errors in the corrected plot (i.e., the bottom panel).

all baselines) to do this. The computation of antenna based errors is based on a direct least-squares approach implemented in a python-based processing stage. It consists of solving for the antenna based errors, θ_i , given the observed visibilities ϕ_j , with a constraint that the mean phase error is zero $\langle \theta \rangle = 0$. The computed antenna based phase errors are stored in separate FITS file for inspection and application for calibration of the simulated visibilities data.

When the phase errors are comparable to a full turn, the algorithm begins to fail to find the true delays to each antenna. Otherwise it finds the same delays that have been applied as described in Section 2.2. The main reason for implementing the phase solving stage is to be able to later investigate how it will inter-operate with the WVR observations and to investigate the effect of systematic errors and thermal noise.

Once the antenna-based errors have been calculated from calibrator observations, they are linearly interpolated in time and the simulated visibilities phase rotated to correct for phase errors as far as possible. The uncorrected and corrected antenna-based phase errors for the medium configuration are illustrated in Figure 1. The phase errors in the compact array are shown in Figure B1 in Appendix B.

2.4 Imaging

The requirements for the imaging stage are extremely simple, and consist only of making dirty images of from the corrupted (and possibly calibrated) *uv*-FITS data. This was conveniently done by calling the `Obit` package directly from python.

The data were uniformly weighted when computing the dirty images and the pixel size was selected automatically by the software. The only variable parameter was the time range of the data to be included in the dirty image. Either we included all of the data, representing longer observations, or very short sections to simulate snapshot observations.

The output of this stage is a FITS-based dirty image of the science source.

The effect of the simulated phase fluctuations on the synthesised beams is illustrated in Figures 2 and 3. The first figure shows the beam as function of the magnitude of phase fluctuation while the second figure illustrates how the snapshot synthesised beams vary with time due to variance of the turbulent volume.

2.5 Analysis

We are interested in the following results from these simulations:




Figure 2. Illustration of the degradation of the synthesised beam, when no phase correction is applied, as the magnitude of phase fluctuations increases (they increase left to right, top to bottom). The beams have been shown for the medium configuration, 30 minute long observation and thickness of turbulent volume of 200 m. The intrinsic strength of the source is 2 Jy and the axes are labelled in radians. The magnitude of fluctuations are in the ratio 3 (top-left) : 10 (top-right) : 50 (bottom left) : 100 (bottom-right).

- (i) Point source sensitivity relative to sensitivity of observations un-corrupted by the atmosphere
- (ii) Resolution relative to un-corrupted observations
- (iii) Flux calibration accuracy
- (iv) Astrometric accuracy

We estimate the relative point source sensitivity, S , as simply the magnitude of the maximum pixel value in the dirty beam map divided by the intensity of the simulated science target.

We estimate the resolution and astrometry by fitting a single Gaussian to the dirty image. We assume that the resolution is given by the major radius of the best-fitting Gaussian. We compute the astrometric accuracy from the offsets of the Gaussians as:

$$\sqrt{\langle \Delta P^2 \rangle} = \sqrt{\langle \delta x_i^2 + \delta y_i^2 \rangle} \quad (2)$$

where $(\delta x_i, \delta y_i)$ is the position of the best-fitting Gaussian to the dirty map made from i -th snapshot observation.

Finally the flux calibration accuracy is computed as the standard deviation of point source sensitivities of a series of snapshot relative to the point source sensitivity of a long observation.

3 RESULTS

All of the simulations were carried out for three configurations of ALMA: the most compact configuration (configuration number 1), a medium configuration (configuration number 15) and the most extended configuration (configuration number 28). The positions of the antennas in these configurations are shown in Figure 4 and the distribution of baseline lengths is shown in Figure 5.

We also considered a range of atmospheric conditions in the simulations, parametrised by two variables: the magnitude of phase fluctuation on a baseline of length 300 m (ϕ_{rms}) and the thickness of the turbulent layer, w . The reason for quantifying the magnitude of fluctuations by ϕ_{rms} is that the site-testing interferometer at Cha-




Figure 3. A sequence of dirty snapshot beams (each calculated from a single one second integration) separated by 200 s in time. There is no phase correction applied, and as in all other results, no thermal noise. Other parameters are as described in Figure 2.

inator has measured the distribution of phase fluctuations on this baseline over a number of years. The results presented below are displayed so that the parameter ϕ_{rms} runs along the abscissa of the plots and the different thicknesses of the turbulent layer are represented by different line-styles. There are three plots in each figure, representing the three configurations.

3.1 Observations without phase calibration/correction

We first consider simulations without any phase calibration applied. In actual operation ALMA will operate with a number of calibration steps always in place but these results are nevertheless of interest in order to establish the ‘baseline’ effect of atmospheric phase fluctuation if no attempt was made to correct them, and so be able to quantify the improvement made by various calibration procedures. They also provide information on effect that may be present when observing the phase calibrators.

Since we are not simulating any instrumental effects there can be no systematic errors in the phases, but the atmospheric phase

fluctuations will degrade the quality of the observations in various ways. The most basic quantities to consider are the relative sensitivity to point sources and the resolution: these are shown in Figures 6 and 7. We are considering relative sensitivity only since calculating absolute sensitivity would have required assumptions about the system temperature and the overall gain of the array, which are functions of frequency, final performance of receivers and the atmospheric transparency, and so would complicate our calculations unnecessarily.

Qualitatively, the plots shown in Figures 6 and 7 show the expected behaviour: sensitivity and resolution are close to their ideal values for a wide range of fluctuation magnitudes and then they decline rapidly close to a critical magnitude of fluctuations. The easiest to understand are the middle plots, i.e., the plots for the medium configuration. This is because we parametrised the results in terms of the magnitude of fluctuation on a 300 m baseline, and since the distribution of baseline lengths in the medium configuration peaks at about 250-500 m (Figure 5). It is clear that for the medium con-




Figure 4. ALMA configuration used in the simulations; from top to bottom: the compact, medium and extended configurations.

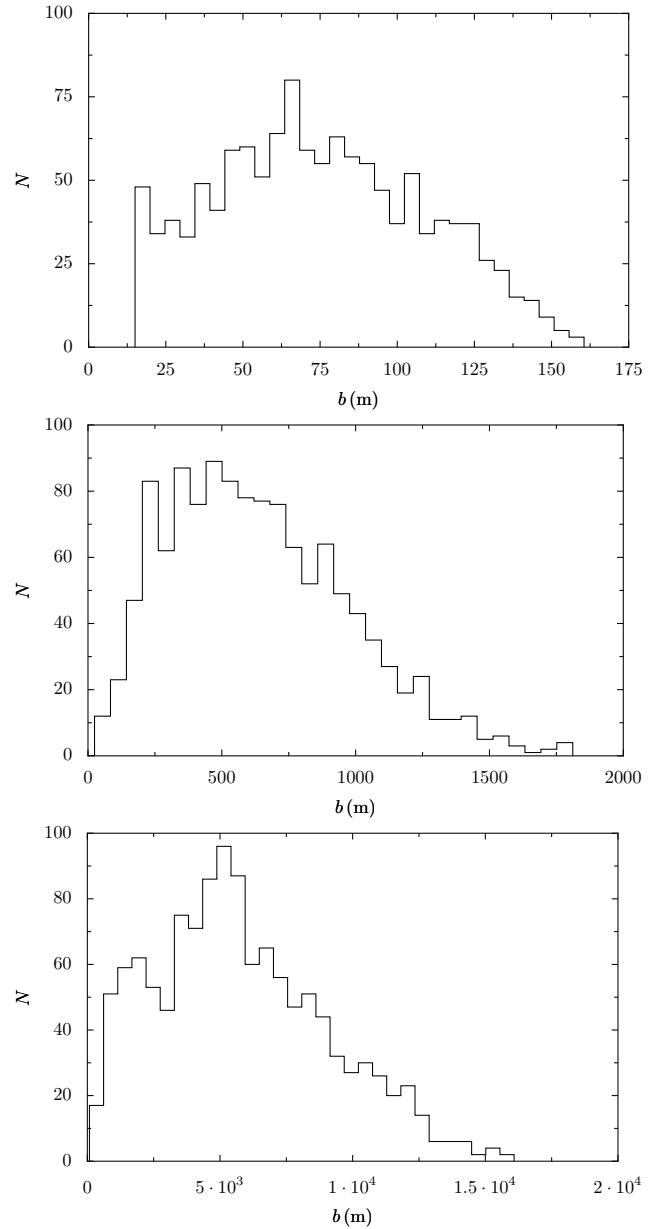


Figure 5. Distribution of the lengths of baselines, measured along the tangent plane, for the compact (top plot), medium (middle plot) and extended (bottom plot) configurations in our simulations (Figure 4).

figuration, the critical magnitude of fluctuations is around 0.5 radians (about 30 degrees), as expected from theory.

The behaviour of sensitivity plots for the compact and extended configurations (top and bottom plots in Figure 6) is driven by their distribution of baselines relative the 300 m baseline from which we calculate the fluctuation magnitude. The thickness of the turbulent layer also becomes important since it determines how fluctuations at 300 m scale to shorter or longer baselines. For example, the lengths of baselines in the extended configuration are about 10 times greater than in medium configuration but otherwise similar (middle and bottom plots of Figure 5). The solid line in the extended configuration sensitivity plot shows the results for the thinnest layer; for this layer we expect the root-mean-square of fluctuations to grow as $1/3$ power of the baseline length, hence to be approximately two times greater for median baseline of extended




Figure 6. Sensitivity to a point source (with no phase correction) as a function of phase fluctuations (in radians) toward the zenith on a 300 m baseline. The plot at the top of the figure is for the compact configuration, the middle for the medium configuration and the bottom plot is for the extended configuration. In the top two plots, the four lines correspond to turbulent layer thicknesses of 25 m (full line), 200 m (dashed), 800 m (dotted) and 3 km (dash-dot-dash). In the bottom plot the four lines correspond to thicknesses of 50 m (full line), 400 m (dashed), 1600 m (dotted) and 5 km (dash-dot-dash).

configuration compared to medium configuration. This is indeed what the plot shows. For thicker layers, fluctuations grow faster with baseline length and so the falloff of sensitivity occurs sooner.

The plots in Figure 7, showing the calculated resolution, largely show similar structure to the sensitivity plots. Typically, if fluctuations are sufficient to cause a loss of two in sensitivity they also cause a very large drop in resolution.

The second set of investigations for uncalibrated simulations

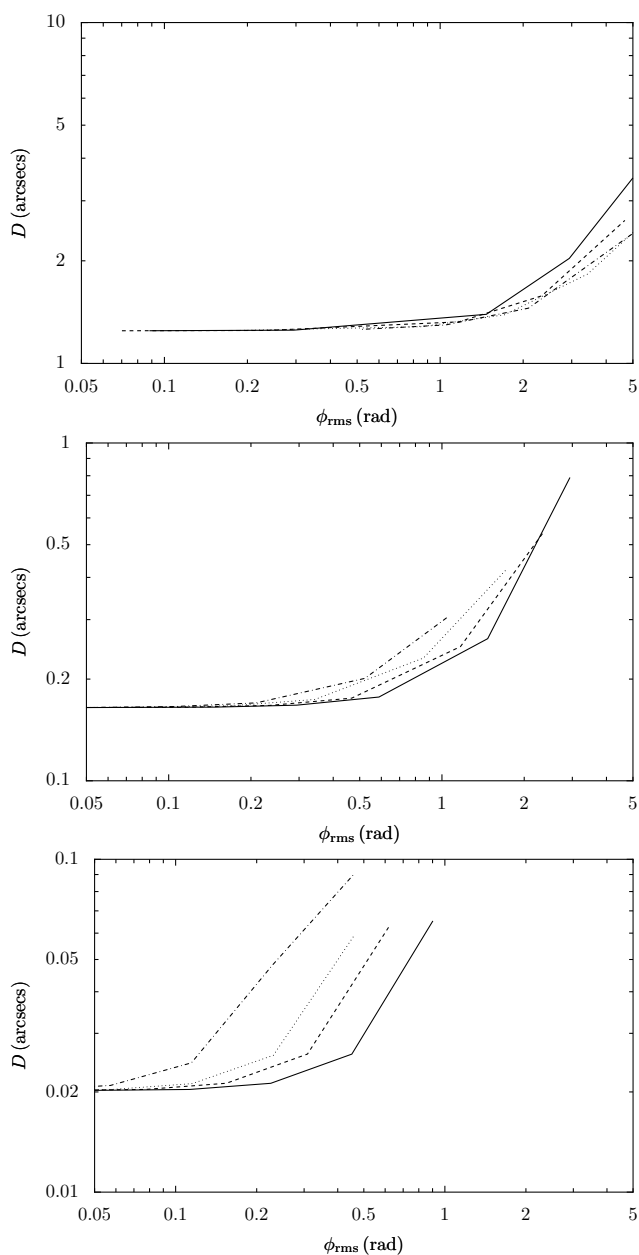


Figure 7. Major radius of the best-fitting Gaussian to the dirty point source maps (made without phase correction) as a function of phase fluctuations on a 300 m baseline. We have assumed a wavelength of 1 mm. Line styles have the same meaning as in Figure 6. Top to bottom: compact, medium and extended configurations.

consisted of analysing the statistical properties of snapshot observations. These were calculated by imaging each integration of an one-hour observation individually and analysing sensitivity and astrometric accuracy as described in Section 2.5.

The variation of sensitivity for the same ranges of conditions as discussed previously is shown in Figures 8. It can be seen that for a wide range of conditions this statistical variation in sensitivity is approximately a power law function of the magnitude of phase fluctuations. Comparison with Figure 6 shows that this variation is significant (given ALMA's aim for 1 per-cent flux calibration accuracy) in that even when the long-time average of sensitivity is only five or less per-cent smaller than ideal sensitiv-

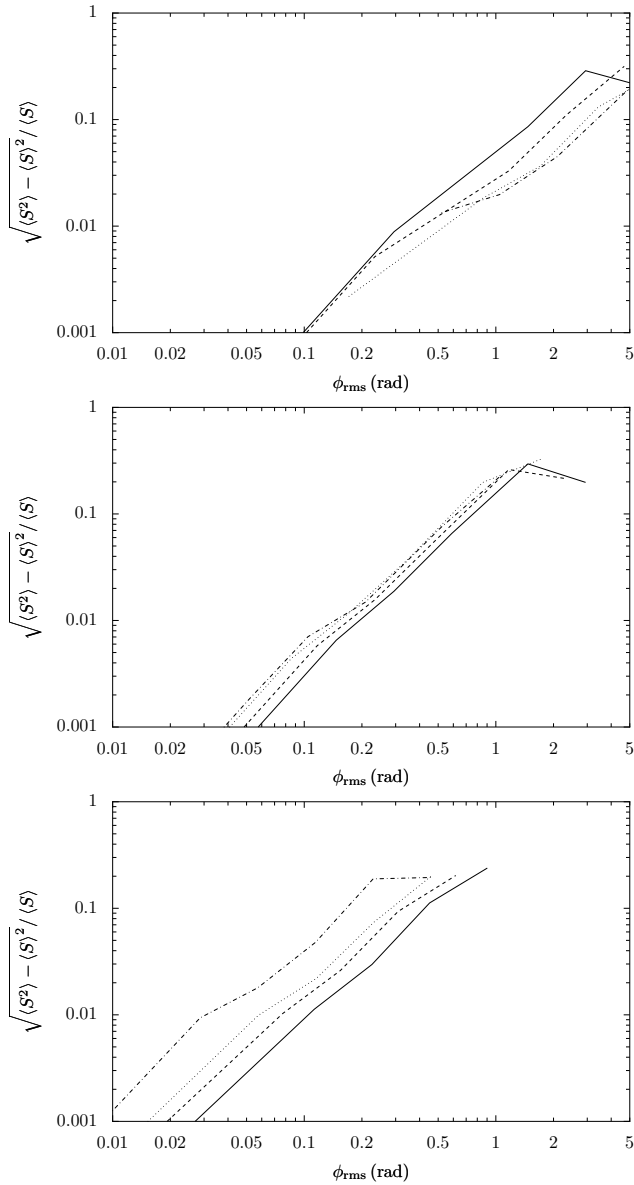


Figure 8. Standard deviation of the fractional sensitivity to point sources (without phase correction) as a function of the magnitude of phase fluctuations on a 300 m baseline. Line styles have the same meaning as in Figure 6.

ity, the accuracy with which this sensitivity is known in any one integration is already one-percent or worse.

The variance in the position of the source measured from each integration is shown in Figure 9. Similarly to the variance in sensitivity this shown remarkably power-law-like behaviour across a wide range of conditions.

4 FAST-SWITCHING PHASE CALIBRATION

We made a number of simulations of fast-switching observations. All of them have a calibration cycle of 15 s, but differ in the way antenna based phase errors are derived from the visibilities, and differ in the angular distance between the direction of the science target and the calibrator. We primarily show the point source sensitivities.

In Figure 10, the antenna-based phases are derived in the straightforward described in Section 2.3. In Figure 11, the phases

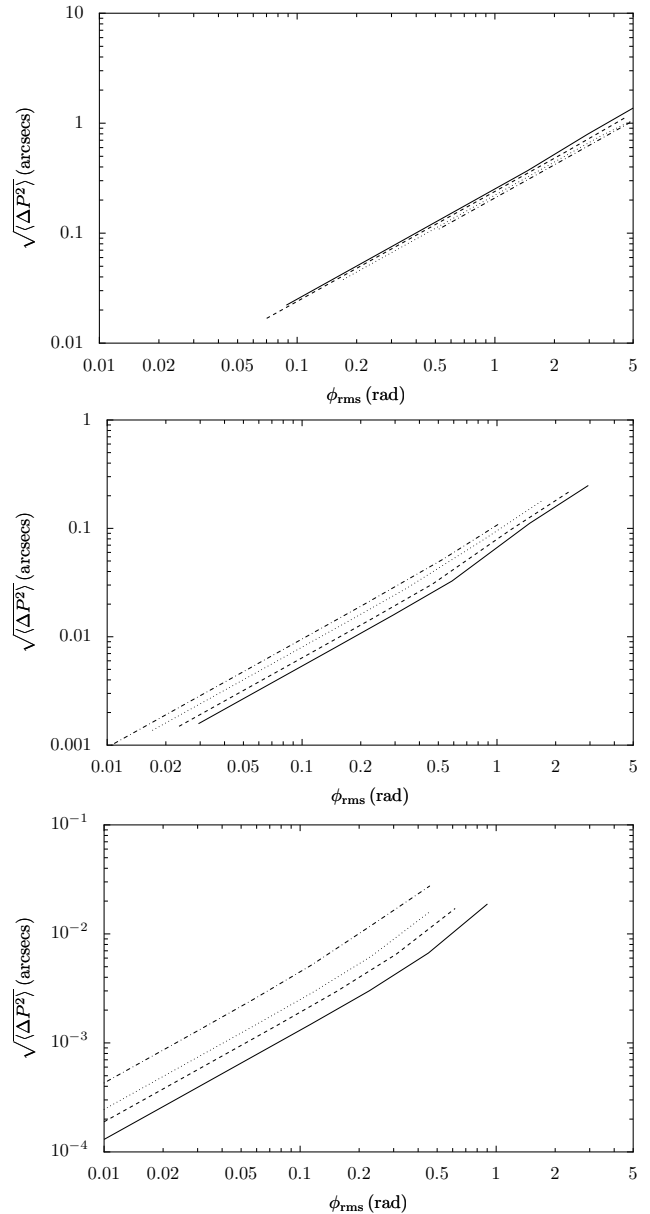


Figure 9. Root-mean-square deflection of the centre of the dirty beam (without phase correction) from the expected position. We have assumed a wavelength of 1 mm for these simulations, as in Figure 7. Line styles have the same meaning as in Figure 6.

are derived not from the observation being simulated, but from a different simulated visibility set that has much smaller fluctuation magnitude. The latter case corresponds to observing the calibrator at a lower frequency than the observations of the target source (a capability that has been specifically designed into ALMA), and so ensures that the phase wraps can not corrupt the derived antenna based phases. In Figure 12, the phase calibration is done in the same way, but we have considered a calibrator which is 1.5 degrees away from the science target. This introduces a limit on the accuracy with which the phase correction can be done, even if the calibration cycle were arbitrarily short.

It can be seen in Figure 10 that sensitivity falls-off quickly as magnitude of fluctuations increases, i.e., at about 0.2 to 0.5 radians. The reason for this is precisely the corruption of the inferred




Figure 10. Relative point source sensitivity after fast-switching phase calibration with a 15 s time cycle. For these results it is assumed that there is no offset to the phase calibrator, i.e., the best-case scenario. The frequency used for phase calibration is the same as the science frequency, and so the antenna-based phase solutions can be seen to suffer from instabilities and phase wraps, leading to rapid fall-off in performance with increasing ϕ_{rms} . Line styles as Figure 6.

antenna based phase errors once the fluctuations become significant enough that phase wraps may occur between the different antennas in the array.

In contrast in Figure 11, sensitivity only begins to fall-off when fluctuation magnitude on a 300 m baseline is about 1.5 radians. Furthermore, it can be seen that the critical fluctuation magnitude is the remarkably similar for the compact, medium and extended configuration, i.e., it has become almost independent of typical length of baselines in the configuration. This result is derivable theoretically and has been shown experimentally (Holdaway 1992;

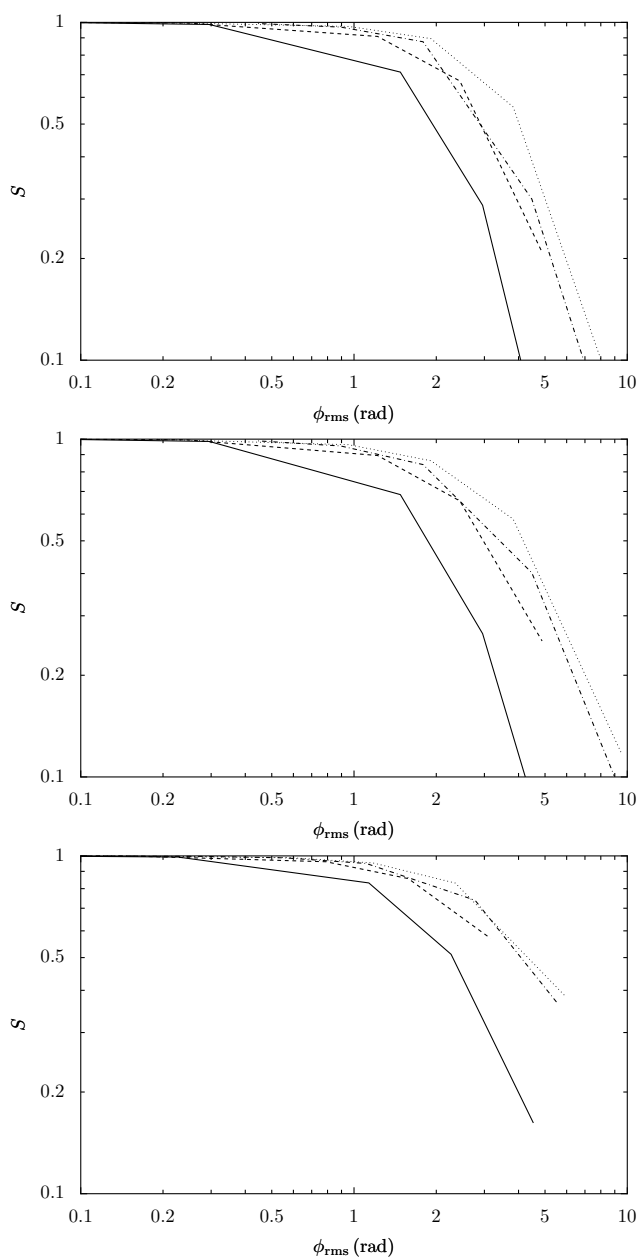


Figure 11. As Figure 10, but the phase calibration was derived from low frequency observations, preventing problems in retrieval of antenna base errors.

Holdaway & Owen 1995; Carilli & Holdaway 1999) and is one of the main motivations for the fast-switching technique.

To reiterate, the difference between Figures 10 and 11 is whether the phase solutions are derived from data with same, or smaller, phase fluctuations. Since we are not considering other sources of errors, or thermal noise, the phase solutions used for Figure 11 are *exact* and the results shown there represent the idealised best possible performance of the fast-switching technique. The solutions used in Figure 10 are exact until phase wraps become possible between antennas; the exact way the phase-wraps affect the solutions depend on the precise algorithm used and should be investigated further.

For the observations calibrated from lower fluctuation visibilities (i.e., corresponding Figure 11) we again calculated the width of

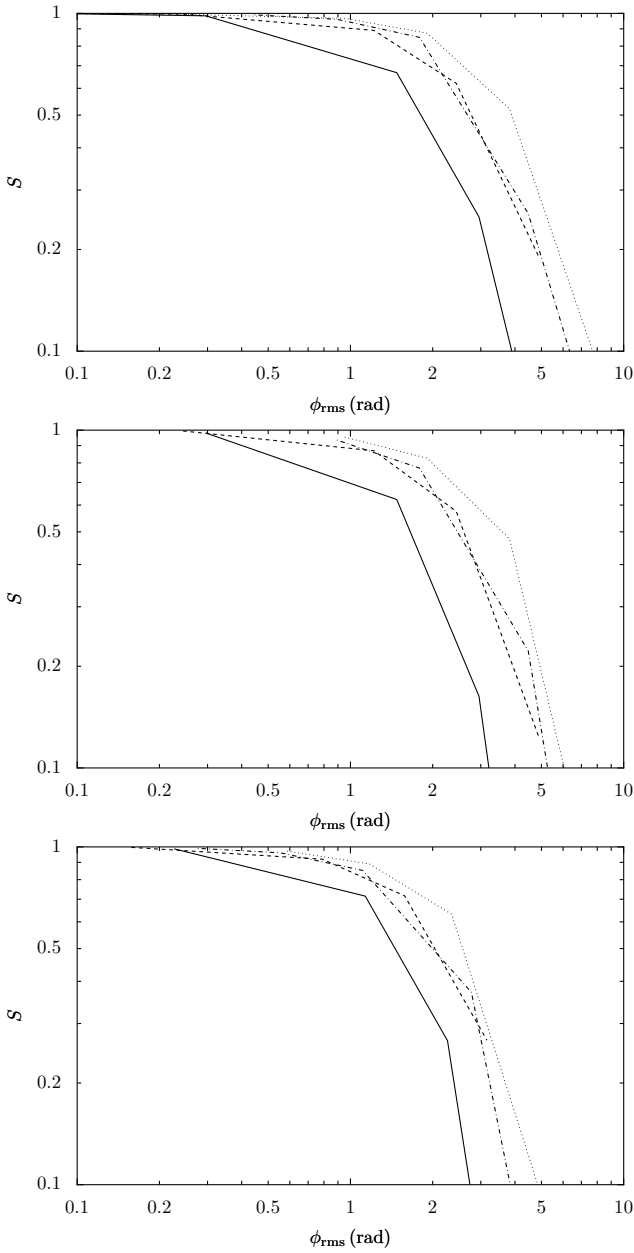


Figure 12. Relative point source sensitivity after fast switching with a 15 s cycle (i.e., as Figure 11), but with the phase calibrator offset by 1.5 degrees from the science target.

the best-fitting Gaussian, that is, the approximate resolution. This is shown in the plots in Figure 13. Again, qualitatively different behaviour is seen compared to uncalibrated observation (Figure 7).

For the medium and extended configurations, it can be seen that the resolution is independent of magnitude of phase fluctuations. For the compact array however, we can see that the resolution begins to fall off at around 2 rad fluctuation magnitude. The reason for this can be understood from the relative magnitudes of the effective calibration length, $l_C = vt_C/2$, (here v is the wind speed and t_C is the calibration cycle time), and the distribution of baseline lengths. For the medium and extended configurations, baselines are almost all longer than l_C , and so the fluctuation on all of them to first order depends on l_C only and so is roughly the same. In the compact configuration however, some baselines are shorter than l_C




Figure 13. Major radius of the best-fitting Gaussian to the dirty beams, i.e., as in Figure 7, but for the fast-switched observations at 15 s cycle time and calibrated at low frequencies (i.e., corresponding to sensitivity plots shown in Figure 11).

and hence the fluctuations increase with baseline length; as a result, longer baselines lose coherence more quickly than shorter baseline leading to loss of resolution.

5 WATER VAPOUR RADIOMETRIC PHASE CORRECTION/CALIBRATION

In addition to phase calibration through fast switching, as described in the previous two sections, ALMA will also use the dedicated water vapour radiometers (Hills et al. 2001) mounted on each of the 12 m-diameter antennas in the array to correct the phase fluctuations due to the troposphere. There are two effects which the data from the radiometers can correct:

- (i) Phase fluctuations during observation of the science target
- (ii) The error in transferring a phase solution from a phase calibrator due to the difference in the lines of sight to calibrator vs science target

When the angular distance between the science target and the calibration source is not too large, these two effects look essentially the same to the radiometers. That is, the slew to the calibrator and back to the science targets looks just like a change in atmospheric properties along the line of sight. Therefore, the resulting path change should be corrected by the radiometer and no special simulation is required beyond what has been described in the previous sections and a model for the residual errors after correction by the radiometers.

Since the WVRs will not give a perfectly accurate measure of path delay due to the atmosphere, there will be residual phase fluctuation in the corrected data. The specification for the ALMA WVRs allows for two sources of inaccuracies¹:

- (i) ‘thermal error’ of $10\mu\text{m}[1 + c/(1\text{mm})]$ RMS per antenna, where c is the column of precipitable water vapour
- (ii) ‘proportional error’ of 2% of the actual fluctuation on any baseline.

The ‘thermal error’ is expected to be completely uncorrelated between antennas and independent of baseline length. Its effect on relative sensitivity is easy to model: the sensitivity decreases as $\exp(-\phi_{\text{rms}}^2/2)$. Because the thermal errors are independent between the radiometers they do not cause a significant loss of resolution, or astrometric or flux-calibration inaccuracies.

The effect of the proportional error is also easy to understand: the residual error after phase correction can be modelled simply as the scaled down uncorrected fluctuation. This scaling down relationship should apply equally to the phase errors while tracking a source and errors due to phase transfer, i.e., both of the cases (i) and (ii) above. As a result, residual proportional errors can be simulated by simply multiplying the magnitude of fluctuations on a 300 m baseline by an adjustment factor equal to the inverse of the proportional error.

On these assumptions, the relative point-source sensitivity is simulated in Figure 14, now including in the simulated the phase corrections from the WVRs.

The first panel of this of Figure 14 repeats the results from in Figure 6, showing the fall off in point-source sensitivity when no phase calibration or correction is made. The second panel shows relative point-source sensitivity after fast-switching phase calibration. In contrast to Figure 11, the fast-switching calibration cycle in Figure 14 is set to 200 s (rather than 15 s), since this approximately the expected cycle time when the WVRs are in use. As 200 s is longer than the time for a phase screen to blow across the array in this configuration, the phase calibration makes no improvement to

the sensitivity, as can be seen from this figure (see also, for example, discussion by Holdaway 2004).

The third panel of Figure 14 shows the point-source sensitivity with the WVR phase correction applied, on the assumption that the radiometers have proportional errors only. An analysis including only proportional would be in line with the analyses of the previous sections, where we have been ignoring thermal-like sources of errors (but we note here that according to the specifications, the performance of the radiometers will almost always be dominated by this thermal noise). In this scenario, the sensitivity performance remains good when the phase fluctuations are as large as 20 rad on a 300 m baseline (when observing at $\lambda = 350\mu\text{m}$, this corresponds to 1.1 mm path RMS on the 300 m baseline).

The final panel of Figure 14 includes both the proportional error and thermal noise corresponding to a phase fluctuation of 0.25 rad. This amount of thermal noise corresponds to the specified performance of the WVRs when there is negligibly little water and the observing wavelength is $\lambda = 350\mu\text{m}$, or when the water column is 1 mm and the observing wavelength is $\lambda = 850\mu\text{m}$. As discussed above we model this thermal error as independent Gaussian random phase errors on each baseline; the resulting loss of coherence can be seen by noticing that the calculated relative point sources no longer reach unity even in the limit of very small phase fluctuations. Therefore when the atmospheric phase stability is extremely good, a small gain in efficiency could be realised by smoothing the WVR data to a few seconds timescale to further reduce the decorrelation arising from WVR thermal noise (Holdaway 2006).

The specification of 2% proportional error for the WVR phase correction excludes the effect of dry fluctuations, which may introduce a larger error. Also, any inaccuracies in atmospheric modelling and the divergence between astronomical and WVR beams (discussed by Nikolic et al. 2007), could make this specification difficult to meet. Therefore, in Figure 15, we extend the results shown in last panel of Figure 14, showing the relative sensitivity for proportional errors higher than the specification.

Stability of WVRs with a change in the airmass can also be important. In the simple approximation that we used above, a change of air-mass introduces zero path change and so we also assume no error from the WVRs. This would be consistent with a very strict interpretation of the proportional error budget.

In reality however, small differences between the radiometer devices will inevitably produce an error in the path correction that is applied as the antennas change position.

It can be shown that for calibrator sources that are reasonably close, this should not be a very big effect. The rate of change of airmass with angle from zenith is:

$$\frac{d}{d\theta} \sec(\theta) = \tan(\theta) \sec(\theta), \quad (3)$$

which at 30 degrees from zenith evaluates to $2/3$. For example, a two degree change in elevation will produce an air-mass change of about 2.3%. For 1 mm of water this corresponds very roughly to $130\mu\text{m}$ of path. If the 2% proportional accuracy holds for absolute changes of path as well as relative, then this will result in a maximum of $3\mu\text{m}$ path error, which is not significant compared to thermal errors even in the best conditions.

6 DISCUSSION AND SUMMARY

In Section 2 we described the framework we use to simulate atmospheric phase fluctuations and their correction. The framework is

¹ These specifications are only strictly applicable to phase fluctuations due to water vapour, i.e., the ‘wet’ phase-fluctuations.




Figure 14. Illustration of the effect of WVR correction in the medium configuration (see middle plot in Figure 4). The top-left panel shows the relative point source sensitivity with no phase calibration (it is the same plot as the middle plot of Figure 6). The top-right panel shows the efficiency after fast-switching phase calibration with a switching cycle time of 200 s. The bottom-left panel shows the relative point source sensitivity after further correction with the WVRs, assuming only that there is a 2% proportional error. The bottom-right panel also shows the efficiency after correction with the WVR but now assuming 2% proportional error and a thermal error corresponding to 0.25 rad rms (this leads to a small but observable loss of efficiency at even small phase fluctuations, i.e., on the left of panel).




Figure 15. As the bottom right plot of Figure 14, but showing performance with 5% (left) and 10% (right) proportional errors in the performance of the WVRs.

based on a number of tools (*casa*, *Obit*, and our own code), and so can take time to assemble and setup. However, the end result is highly efficient and ergonomic for use, allowing investigation of wide range of atmospheric conditions and phase correction techniques. The source code for the majority of the part of the framework that has been designed by us is being made available under the GPL licence – see <http://www.mrao.cam.ac.uk/~bn204/alma/memo-turb.html>.

The modelling of the phase fluctuations is, as in previous

work, based on the frozen-screen model with Kolmogorov statistics. The improvement in this work is that a three (rather than two) dimensional Kolmogorov field is simulated. This means there the transition from ‘thick-screen’ to ‘thin-screen’ regime is naturally reproduced and that phase change due to a change in the of direction of observation (as when switching to a calibration source) is more accurately modelled.

Section 3 presents some results of simulations without phase correction and with fast-switching phase correction. All of the sim-

ulations parametrise the atmospheric conditions in terms of phase fluctuation on a 300 m baseline, allowing easy referencing to results from the site-test interferometer campaigns.

The results of the simulations without phase correction are relevant because when the fast-switching cycle is longer than the time it takes for the frozen screen to cross the array, phase calibration will not decrease the atmospheric phase fluctuations and we expect similar results to the uncalibrated observations. (This is effect is illustrated in Figures 14 and B1 which are covered later in the discussion.) Therefore these simulations should be a useful guide to when fast-switching and WVR phase corrections should be employed.

Besides calculating the two well appreciated effects on point source sensitivity and the effective resolution of the array (Figures 6 and 7), we also consider single integration snapshot observations and calculate the errors in the inferred position and the source strength. These results, displayed in Figures 8 and 9, show that even when phase fluctuations are small enough that they have a small effect on overall efficiency, they can cause significant fluctuations in apparent strength and position of a source in snapshot observations.

Figure 1 illustrates the principle of fast-switching calibration and shows some of the important effects, which are quantified in Figures 10–13. The main result is that, as expected, fast-switching greatly extends the range of atmospheric conditions that are usable in medium and extended configurations (Figure 11). However, in our simulations, this is only true when phase transfer from lower frequencies can be accurately done. Without phase transfer, path fluctuations can be large enough that deriving good antenna-based phase solutions becomes difficult. In our simulation this badly affects the extended configurations (Figure 10) although we expect that this may be improved by a more sophisticated phase solving algorithm.

One potentially important source of error with fast-switching phase correction is the error due to the different lines of sight to science and calibrator sources. This effect can be seen in Figure 1: the phase error discourteously jumps after the calibration scan, due to the different lines of sight. The magnitude of this error for calibrator 1.5 degrees is quantified in Figure 12.

Finally, in Section 5 we simulate observations with WVR phase correction, assuming the performance of the WVRs is close to their top-level specification. It can be seen in Figure 14 that if the WVRs work to these specifications, they will largely remove atmospheric phase fluctuations as a scheduling constraint on ALMA observations.

REFERENCES

- Asaki A., Saito M., Kawabe R., Morita K., Tamura Y., Vila-Vilaro B., 2005, Simulation series of a phase calibration scheme with water vapor radiometers for the atacama compact array. ALMA Memo Series 535, The ALMA Project
- Carilli C. L., Holdaway M. A., 1999, Tropospheric phase calibration in millimeter interferometry. ALMA Memo Series 262, NRAO
- 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
- Hills R. E., Gibson H., Richer J. S., Smith H., Belitsky V., Booth R., Urbain D., 2001, Design and development of 183ghz water vapour radiometers. ALMA Memo Series 352, The ALMA Project
- Holdaway M. A., 1992, Possible phase calibration schemes for the mma. ALMA Memo Series 82, The ALMA Project
- , 2004, Does the aca need phase compensation? ALMA Memo Series 491, The ALMA Project
- , 2006, Phase calibration steps. Specification document, The ALMA Project, <https://wikio.nrao.edu/pub/ALMA/CalExamples/PhaseCalStepByStep.pdf>
- Holdaway M. A., Owen F. N., 1995, A test of fast switching phase calibration with the vla at 22 ghz. ALMA Memo Series 126, The ALMA Project
- 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
- Stirling A., Richer J. S., Hills R. E., Lock A., 2005, Turbulence simulations of dry and wet phase fluctuations at chajnantor. part i: The daytime convective boundary layer. ALMA Memo Series 517, The ALMA Project

ACKNOWLEDGEMENTS

We would like to thank W. D. Cotton for making the Obit package publicly available and for assistance in making the best use of it.

APPENDIX A: GEOCENTRIC TO TANGENT PLANE COORDINATES

Coordinates of antennas in FITS IDI are stored as array centre position x, y, z and the antenna offsets dx, dy, dz . The coordinate system is Cartesian with the positive z axis running toward North Pole, and the x running going toward the intersection of the Greenwich Meridian and the Equator. If we assume that the atmosphere is translating parallel to local tangent plane, then we need convert to local tangent plane coordinates. Define cylindrical and spherical coordinates:

$$\phi = \arctan(y/x), \quad (\text{A1})$$

$$\rho = \sqrt{(x^2 + y^2)}, \quad (\text{A2})$$

$$\theta = \arctan(z/\rho), \quad (\text{A3})$$

$$r = \sqrt{(x^2 + y^2 + z^2)}. \quad (\text{A4})$$

Using the relation:

$$\frac{d \arctan(x)}{dx} = \frac{1}{1+x^2}, \quad (\text{A5})$$

their relevant differentials are:

$$d\rho = \frac{x \cdot dx + y \cdot dy}{\rho}, \quad (\text{A6})$$

$$\rho d\phi = \frac{x \cdot dy - y \cdot dx}{\rho}, \quad (\text{A7})$$

$$rd\theta = \frac{\rho \cdot dz - z \cdot d\rho}{r}. \quad (\text{A8})$$

Clearly, the plane to use is the plane that is tangent to the surface of the Earth at the array centre, so that antenna coordinates in this frame of reference are:

$$x' = \rho d\phi \quad (\text{A9})$$

$$y' = rd\theta. \quad (\text{A10})$$

The x' axis is parallel to the local East-West direction and y' to the local North-South direction. Increasing x' corresponds to the East direction. The coordinate system rotates with the Earth, hence it is a suitable reference for a phase screen due to the atmosphere.

APPENDIX B: ANTENNA-BASED PHASE ERRORS FOR COMPACT ARRAY

In Figure B1, we show the simulated phase errors with and without the fast-switching phase calibration for the most compact configuration of ALMA and for a wind speed of 12 m s^{-1} .




Figure B1. Like Figure 1, but for the compact array. The high filling factor of the compact array gives rise to strong correlation between antenna-based phase errors. On the other hand, the short crossing-time for the screen across the array means that fast-switching makes little improvement to the phase errors (see for example the discussion by Holdaway 2004).