ALMA Memo 402
Illumination Taper Misalignment and Its Calibration

M.A. Holdaway
National Radio Astronomy Observatory
949 N. Cherry Ave.
Tucson, AZ 85721-0655
email: mholdawa@nrao.edu

Dec 12, 2001

Abstract

The alignment of the tapered illumination with the antenna's primary dish will not be perfect, and for ALMA has been specified as having an rms error of 0.1 dish radii at all frequencies, and the resulting image errors are therefore independent of frequency. If uncorrected, simulations indicate the effects of this illumination offset will dominate both pointing and surface errors for ALMA's wide field imaging at frequencies up to about 500 GHz.

The main effect of a shift in the tapered illumination on the dish is a phase gradient across the far field voltage pattern, which is given by the Fourier transform of the illumination (neglecting surface errors). A secondary effect will be a distortion in both the voltage pattern's amplitude and phase, caused by the asymmetric illumination pattern on the dish. The phase gradient in the voltage pattern is identical in effect to an error in the baseline. However, the geometry of the baseline and the geometry of the voltage pattern phase gradient are different and will have different time dependences. The phase gradient can be effectively treated by changing the (u, v) coordinates of each visibility to reflect the weighted antenna center accounting for the dish illumination offset. About 85% of the deviation of the offset far field voltage pattern can be corrected by removing the phase gradient; the residual deviation is due to the asymmetric illumination, and is dominated by a bipolar pattern. This indicates that the majority of the effect of the offset illumination can be removed by adjusting the (u, v) coordinates prior to imaging.

The calibration parameters required to perform the (u, v) correction (ie, the phase gradient in the voltage pattern) should be constant with time for each antenna/feed. Thermal noise and atmospheric phase errors will not hamper the measurement of the phase gradient parameters at low frequencies, but will be of concern in determining these parameters at high frequencies. However, both types of error should average down if the calibration observations are performed correctly.

Detailed numerical imaging simulations indicate that after correction of the (u, v) coordinates, the image quality is restored to the level expected when illuminations offsets are not present. In fact, often the image quality of the corrected images was better than when illuminations offsets are not present. This paradoxical result may be explained by the fact that feed leg and subreflector blockage were also included in the simulations. Without the illumination offsets, all antennas' voltage patterns are affected in the same way by the feed legs, resulting in low level asymmetric side lobes that are not reflected in the symmetric beam model. However, the affects of the random asymmetric illuminations of the different

antennas tends to smooth out these low level side lobes, making the problem less systematic and less damaging.

1 Introduction

An offset in the feed from its desired location at the Cassegrain focus will result in a pointing offset. This pointing offset can, of course, be calibrated. If the feed is not offset from the desired position, but instead the feed is not centrally illuminating the primary, there will not be a pointing error, but the far field voltage pattern will have a phase gradient across it. A really gross mispointing of the feed will substantially decrease the sensitivity, as the feed will illuminate less of the dish and more of the ground. For an offset of 0.1 radii and our 10dB Gaussian taper, the integrated illumination will decrease by about 2% over the no offset case. The noise will increase by more than 2%, as the hot ground is being more illuminated. While this is a modest loss of sensitivity, it seems that it is relatively easily corrected. If the rms offset is 0.1 radii, there will be some antennas with much worse offsets and integrated illuminations being decreased by about 6%. It will certainly be beneficial to realign the feeds to eliminate such outliers.

However, imaging, rather than sensitivity, is the primary concern of this memo, so we refocus on the properties of the far field voltage pattern. To first order, the illumination offset and the phase gradient in the voltage pattern will not affect single dishe imaging, as the primary beam is formed from the voltage pattern times its own complex conjugate, which removes the phase gradient. The effect of illumination offset has never been observed on the voltage patterns of existing interferometers; it vanishes for on-axis observing, so will not effect the dynamic range of most VLA or VLBI observations. It is probably present in wide field imaging with present day interferometers, but wide field imaging is often dynamic range limited by other problems, such as errors in the total power data or ineffective deconvolution algorithms.

When the effect of "illumination offset" was first pondered last year, the antenna group was charged with placing a requirement on the accuracy with which the feeds' orientation must be set. The reasoning was to equate the phase errors which result from the illumination offset to the phase errors which result from pointing errors at 230 GHz. So, the specification that the peak of the illumination would be within 0.1 dish radii (rms) of the primary's center was adopted.

However, there were two problems with this specification: most glaring is that pointing errors (as a fraction of the beam) are proportional to frequency, while the illumination offset specification is independent of frequency. This results in the situation where mosaics at frequencies lower than 230 GHz (ie, frequencies at which nonthermal objects are brightest and noise is lowest, leading to a potential dynamic range of 10^6) will be limited by the illumination offsets. Additionally, the illumination offsets may not be as randomized as antenna pointing errors; a given antenna is stuck with its illumination offset for the entire observation, while several aspects of pointing errors are expected to be much more random, decreasing the negative effects of pointing errors, and possibly leaving observations at frequencies higher than 230 GHz being dominated by the illumination offsets.

The surface error simulation campaign which was launched to justify the budgeting of the ACA (ALMA Compact Array) included illumination offsets almost as an afterthought. It was simple enough to include: why not, I thought, disregarding my marching orders. However, when I looked at the results of the full-blown simulations including pointing and surface errors alongside the illumination offsets, the image quality was essentially flat with frequency, not very good at low frequencies and degrading slightly only at the highest frequencies (see Figure 1). If image quality was limited by pointing errors, we would expect image quality to go like $1/\nu$. If image quality was limited by surface errors, we would expect an even steeper relationship, approaching $1/\nu^2$ (Cornwell, Holdaway, and Uson, 1994). However, the flat nature of the simulated image quality with frequency indicated that the illumination offset, with a specification that did not depend upon frequency, was limiting the images up to about 500 GHz (see Figure 2).

While the severe impact of the illumination offset on mosaicing image quality was a surprise, the simple method of correcting the offset was a relief.

2 Simulating Voltage Patterns

Consider this thought experiment: a very large dish is underilluminated with a Gaussian taper such that the illumination actually falls to zero before the edge of the dish. The far field voltage pattern is given by the Fourier transform of the voltage illumination of the dish. Now, consider our shift of the dish illumination. An offset in the aperture illumination will result in a sky voltage pattern with the exact same amplitude, but a phase gradient across it, in the direction of the shift and in amplitude proportional to the shift.

In our simulations, we use the more realistic illumination of a Gaussian taper down 10dB in power at the edge. An example of the voltage illumination, complete with an illumination offset, central blockage due to the secondary of 0.7 m, and blockage due to feed legs of 0.08 m, is shown in Figure 3. The dominant effect of the illumination offset will be a phase gradient in the far field voltage pattern. However, because this illumination does not fall to zero by the time it gets to the edge of the primary reflector (and also due to the asymmetric feed leg and secondary blockage), other second order effects will be seen in the voltage pattern. For large enough offsets, significant amplitude effects will be seen in addition to departures from a simple gradient in the phase. A cut through the phase and amplitude of a far field voltage pattern is shown in Figure 4. We see here that the residual phase after the gradient is removed is fairly small. If the gradient is not removed, phase errors of tens of degrees can be made on a source at the half-power point, potentially limiting image quality.

3 Overview of Corrections for the Illumination Offsets

3.1 A Complicated and Slow Correction Algorithm

Consider the situation we are in: we have voltage patterns which are different in detail from one antenna to another. These differing voltage patterns are essentially antenna dependent, position dependent complex gains which are messing up our data and limiting our ability to reconstruct it with a single primary beam model.

In principle, if the voltage pattern for each antenna is known accurately, given a model for the brightness distribution, we can accurately calculate the effects of the voltage patterns on the visibilities which correspond to that brightness distribution. We could create an iterative algorithm which thereby removed the effects of the deviant voltage patterns on the data and imaged the data cleanly. However, this algorithm is not only very complicated, but very slow, since it needs to perform a Fourier transform for each *baseline* (as in simulating data

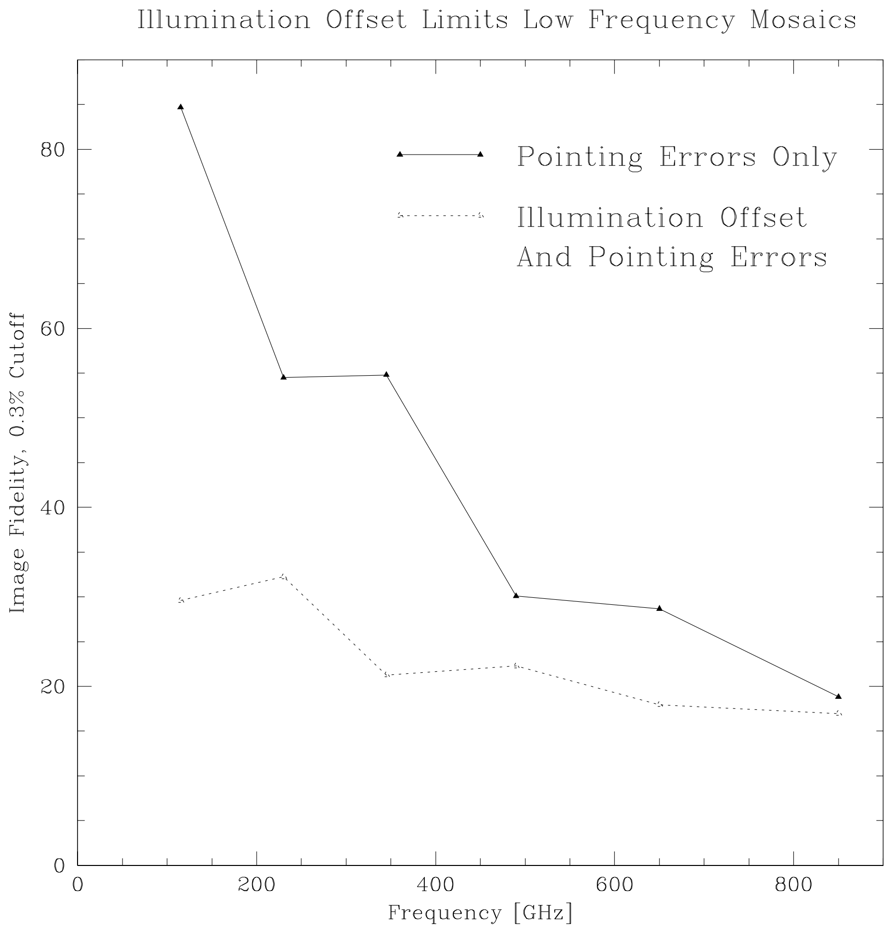

Figure 1: Simulation results for a homogeneous array with only pointing errors and then with both pointing errors and the 10% illumination offset. For high SNR observations, the errors below 500 GHz will be dominated by the illumination offset if it is not corrected.

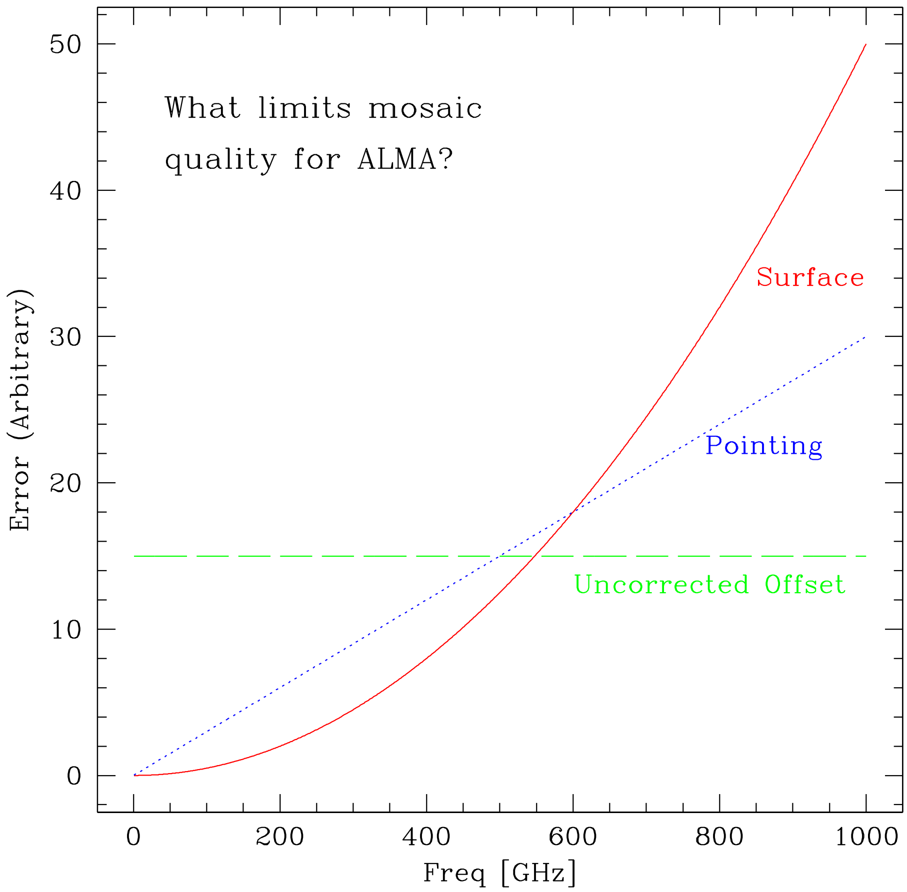

Figure 2: A schematic representation of how the imaging errors due to illumination offset, pointing errors, and surface errors vary with frequency for mosaiced observations. (Imaging errors will be inversely related to dynamic range or image fidelity.) If left uncorrected, the illumination offsets could limit the quality of wide field imaging up to about 500 GHz.

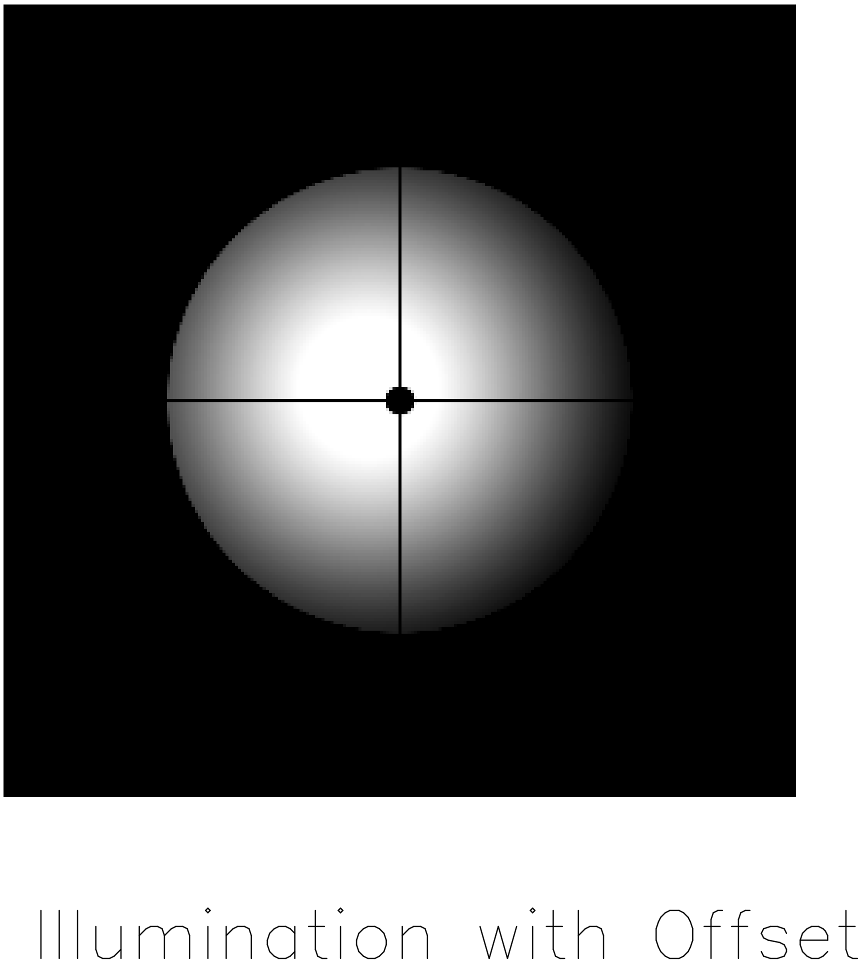

Figure 3: Example of the offset illumination used to simulate the voltage patterns used here.

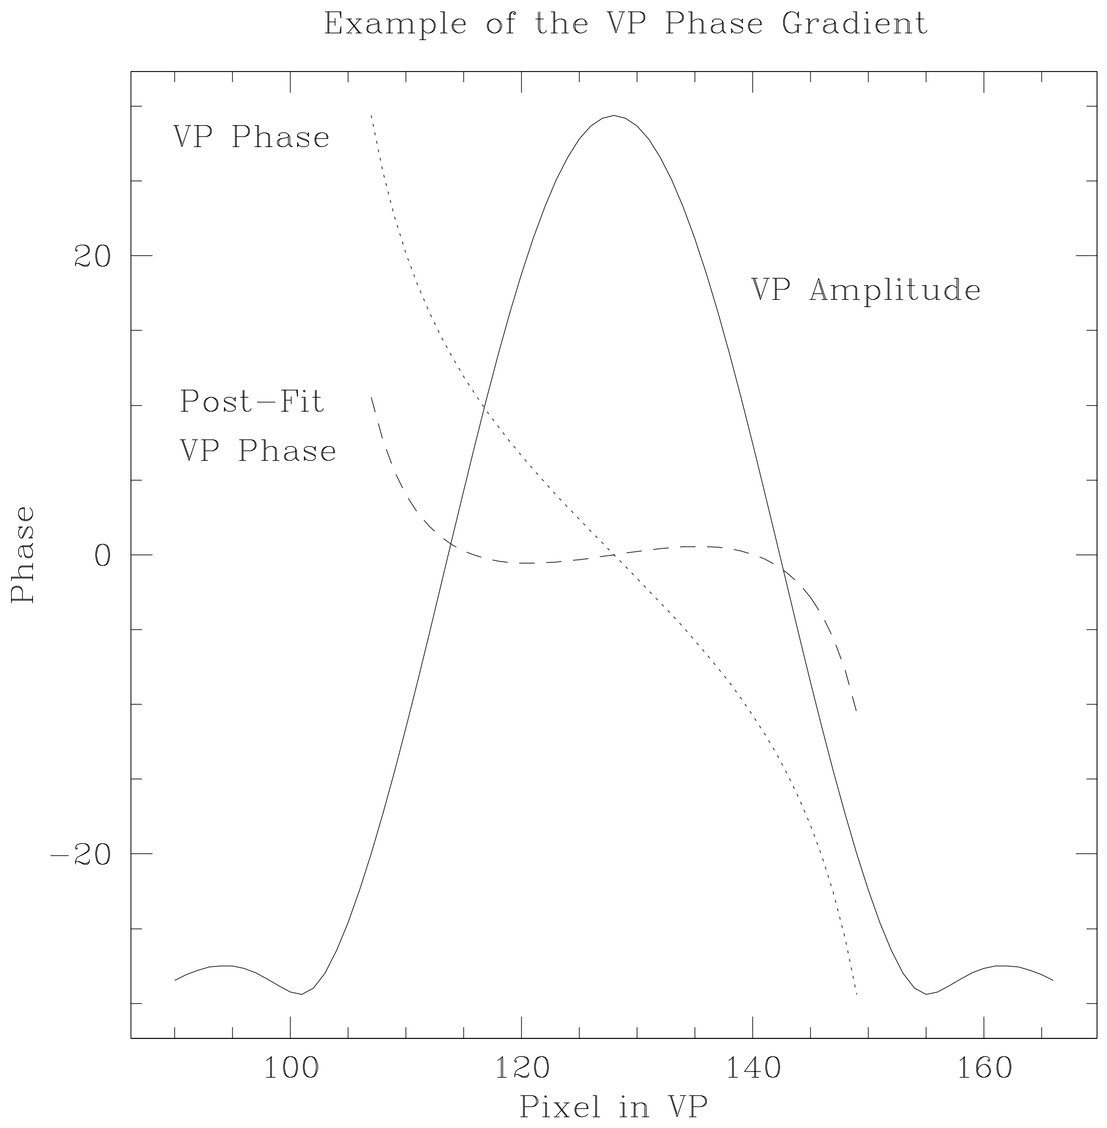

Figure 4: Example of the phase gradient in the voltage pattern, and the residual phase after the phase gradient has been removed, for one of the VP's simulated in the 10% offset case.

with different voltage patterns for each antenna). There will probably be some situations where this sort of algorithm is required, such as imaging in the presence of pointing errors (Holdaway, 1993). There are likely modifications to the basic algorithm, such as only dealing with the brightest sources in the field with this specialized antenna-dependent machinery, and performing the usual FFT on all the rest of the field, However, this is still a scary algorithm, to be avoided if possible.

3.2 A Simple and Fast Correction Algorithm

The heart of a simple, fast, and highly effective correction to the illumination offset lies in the recognition that the offset of illumination is really an offset of the physical baseline.

A shift in the dish *location* would correspond identically to a shift in the baseline. A shift in the baseline would have three consequences:

- a change in w, or the geometrical delay, resulting in phase errors if not corrected.
- decorrelation across the band due to the phase errors which change with frequency.
- the incorrect (u, v) would be attached to each visibility, equivalent to an image plane phase error equal to
$$e^{2\pi i \Delta\mathbf{u}\cdot\mathbf{x}}, \tag{1}$$
where $\Delta\mathbf{u}$ is the vector error in the (u, v) coordinate, and $\mathbf{x}$ is the vector distance from the pointing center.

In our situation, the dish itself is not shifted, but the illumination of the dish is shifted. The first two effects, due to a change in w, should not be observed for an illumination offset, but the error in the (u, v) coordinate will be observed.

Consider a very large antenna with no blockage and an illumination taper which drops to zero at the edge of the dish. In such a case, the full phase gradient described by Equation 1 will be seen in the voltage pattern, and the magnitude of the change in the (u, v) coordinate will be the same as the illumination offset on the dish; however, the geometry will be different, but more on that later. A simple simulation of a very large dish with zero illumination at the dish edge produces a voltage pattern with the expected phase gradient corresponding exactly to the illumination offset. This basically insures that our software is working as expected.

We can look at things more analytically if the only effect upon the voltage pattern is a phase gradient. A single visibility $V_{1,2}(\mathbf{u})$ is related to the sky brightness distribution $I(\mathbf{x})$ and the far field voltage patterns $A_1(\mathbf{x})$ and $A_2(\mathbf{x})$ as

$$V_{1,2}(\mathbf{u}) = \int A_1(\mathbf{x}) A_2^*(\mathbf{x}) I(\mathbf{x}) e^{-2\pi i \mathbf{u}\cdot\mathbf{x}} d\mathbf{x}. \tag{2}$$

Consider that the voltage patterns for antennas 1 and 2 are identical in amplitude and just have a gradient for the phase:

$$A_1(\mathbf{x}) = A_0 e^{2\pi i \mathbf{g_1}\cdot\mathbf{x}}. \tag{3}$$

Then the calculated visibility is given by

$$V_{1,2}(\mathbf{u}) = \int A_0^2(\mathbf{x}) I(\mathbf{x}) e^{-2\pi i (\mathbf{u} - (\mathbf{g_1} - \mathbf{g_2}))\cdot\mathbf{x}} d\mathbf{x}, \tag{4}$$

where $A_0^2(\mathbf{x})$ is now the primary beam. This shows that the baseline vector $\mathbf{u}$ needs to be adjusted by the voltage pattern's phase gradients as $\mathbf{g_1} - \mathbf{g_2}$.

Now, in the case of the ALMA dishes, we have been assuming a Gaussian taper of 10 dB in power at the edge of the dish. Additionally, we have four feed legs each 0.08 m across, and a central blockage of 0.7 m (see Figure 3). The net effect of all this is to result in a weighted mean position of the dish surface which is not as far removed from the dish center as the offset center of illumination. Additionally, this asymmetrically illuminated aperture will result in more complicated phase errors and some amplitude errors in the voltage pattern (see Figure 4). These facts combine to give two implications:

- the phase gradient in the voltage pattern will be less than that expected by the illumination offset
- in addition to the simple phase gradient, there will be more complicated deviations in the voltage pattern characterized by both amplitude deviations and phase deviations which depart from a simple gradient.

Again, using aperture illumination simulations to calculate the voltage pattern, we have quantified these effects. First, we found that for a typical case, the magnitude of the phase gradient for the 10 dB taper ALMA antennas was only 60% of the phase gradient implied by the 10% radius (ie, 0.6 m) offset in the illumination. Second, we found that after fitting and removing the phase gradient from the voltage pattern, the deviation from a "perfect" voltage pattern unhampered by illumination offsets was reduced to only 18% of the pre-fit value. Naively, this indicates that the majority of the effects of the illumination offset on imaging can be removed simply by measuring the phase gradients in each antenna's voltage pattern, converting it into a deviation from the antenna position which varies with parallactic angle, and correcting the (u, v) coordinates accordingly.

4 Calculating and Applying the Baseline Correction

The baseline correction due to the illumination offset should be constant with time. If it is not, we have yet another short term calibratable effect to worry about, but it appears that this effect is fairly easy to calculate. The pointing error is measured by performing a "five point" procedure (or one of its more modern equivalents, such as moving the dish in a circle centered on the pointing source) and comparing the amplitudes at the supposed half-power positions with the central pointing. Similarly, the phase gradient in the voltage pattern can be calculated by performing a five point interferometrically and solving for the phases at the five positions and fitting a plane to the phases. Formally, only three positions are needed, as we are just solving for the two coordinates of an offset, but five positions will help reduce systematic errors.

4.1 Noise and Atmospheric Phase Errors

Thermal noise and atmospheric phase fluctuations will both limit the accuracy of the phases which we measure in a five point procedure. But before we calculate limits on the errors due to noise and atmosphere, we need to understand how accurately we need to determine the phase gradient. After removing the phase gradient, the simulated voltage patterns have a residual

freq GHz	$\sigma_V/\sqrt{N-2}$ in 0.25s [Jy]	typical cal Flux [Jy]	typical phase error in 0.25s [deg]
115	0.0046	10.0	0.02
230	0.0052	5.95	0.06
345	0.012	4.39	0.14
415	0.016	3.82	0.24
650	0.043	2.73	0.90
850	0.120	2.23	3.0

Table 1: Phase errors expected due to thermal noise in 0.25 s. Phase errors at the half power point will be a factor of 2 larger.

phase of 0.5 deg rms for our 10dB taper case. The phase at the half power point in the gradient direction is about 10 deg. Hence, measurements with an accuracy of about 0.5 deg, or a bit better, are required.

From a thermal noise perspective, this is not a problem. Table 1 shows the expected thermal noise in the gain (ie, $\sigma_V/\sqrt{N-2}$, where σ_V is the noise in a single visibility in 0.25 second and N is the number of antennas, taken to be 60), the expected flux of the calibrator (assuming a spectrum of $\nu^{-0.75}$), and the equivalent phase error given the thermal noise. A calibrator flux of 10 Jy has been chosen at 115 GHz, assuming we are using the brightest quasar in the sky. (There is no reason to use a source particularly close to a target source since the calibration should be applicable to all observations for a long time.) One 0.25 s of integration will be sufficient to adequately determine the voltage pattern's phase at frequencies up to about 400 GHz. Higher frequencies will require more integration to achieve the desired accuracy in the phase gradient measurement.

Atmospheric phase errors will be a much tougher problem. Typical fast switching will have residual phase errors of the order 10-20 degrees, which is much too large. The "five-point" procedure will have smaller residual atmospheric phase errors, as there won't be much dead time between using the central pointing as a phase calibrator and measuring the phase on one of the off positions. The fastest way to proceed would be to do some sort of on-the-fly interferometric observing, linearly crossing over the bright source in something like a figure "8".

Assuming this calibration procedure is performed during the best tenth percentile of the phase stability conditions, and that 0.5 s elapses between the observation of one of the half power points and the calibration source straight on, we can use the interferometer site testing data to estimate the phase errors (see Table 2). The atmospheric phase errors are higher than acceptable for all frequencies. Faster on-the-fly slewing will improve the atmospheric phase errors, but if the time on the central pointing keeps on decreasing, we won't detect the point source with sufficient accuracy to calibrate the phases (Table 1). An optimum strategy would probably equalize the thermal and atmospheric noises, resulting in faster scans at low frequencies and slower scans at higher frequencies (this counterintuitive result comes about because the noise-frequency dependence is steeper than the atmospheric phase-frequency dependence due to the falling calibrator spectrum). In the event that the combined thermal and atmospheric phase errors are larger than the desired 0.5 deg level, multiple scans will be

freq GHz	Atmospheric Phase error [deg, over 0.5s]
115	0.80
230	1.61
345	2.41
415	2.90
650	4.54
850	5.93

Table 2: Atmospheric phase errors resulting from an "on-the-fly" observing strategy with 0.5 s between the observation at the half-power point and the beam center. At the beam center, the voltage pattern should have no phase errors, and so can be used as a phase calibrator. Phase dispersion in the submillimeter has not been factored in.

required to calculate the voltage pattern phases with sufficient accuracy. The atmospheric phases should average down, especially if we reverse the direction of the scans to eliminate any systematic component in the residual phase errors.

The analysis above assumes that the physical baselines between antennas have been very well determined. If they have not, there will be some confusion between phase gradients caused by errors in the physical baselines and phase gradients resulting from the illumination offset. Such an ambiguity can be removed by observing over a wide range of parallactic angles, as the projected physical baseline depends upon hour angle and declination, but not parallactic angle, and the effective illumination baseline depends upon parallactic angle.

4.2 Application Algorithm

The software for applying the (u, v) correction to the data is pretty straightforward. We build up a calibration table from the fits in the image plane to the phase gradients for each antenna's voltage pattern. These phase gradients then scale to give the (u, v) offset which should be used (see Equation 4). Finally, the (u, v) offsets need to be rotated by the parallactic angle and added to (or subtracted from, for the second antenna in a baseline) each baseline's (u, v) coordinate. The different rotations of the projected baseline and the incremental (u, v) offset are demonstrated in Figure 5

To test the accuracy of the algorithm, we simulated voltage patterns whose defects could be essentially completely described by the phase gradient: we made a 24 m dish with a Gaussian taper down to 100 dB at the edge, so we were not concerned at all by edge effects. Then we simulated data for an off-center point source as observed through the surface error simulation machinery (ie, different voltage patterns for each antenna). We then adjusted the (u, v) coordinates as instructed by the voltage patterns and resimulated the off-center point source with software which simply took the DFT without concern for the various complex voltage patterns. The phases of the visibilities calculated with the voltage patterns and the shifted (u, v) coordinates agreed to less than 0.01 degrees. This level of disagreement is probably due to small errors in the HGEOM procedure which rotates the voltage patterns, or the pixellation in the aperture illumination and voltage pattern construction process. Anyway, the procedure works very well for the case in which it should work perfectly. With this algorithm in hand, we are

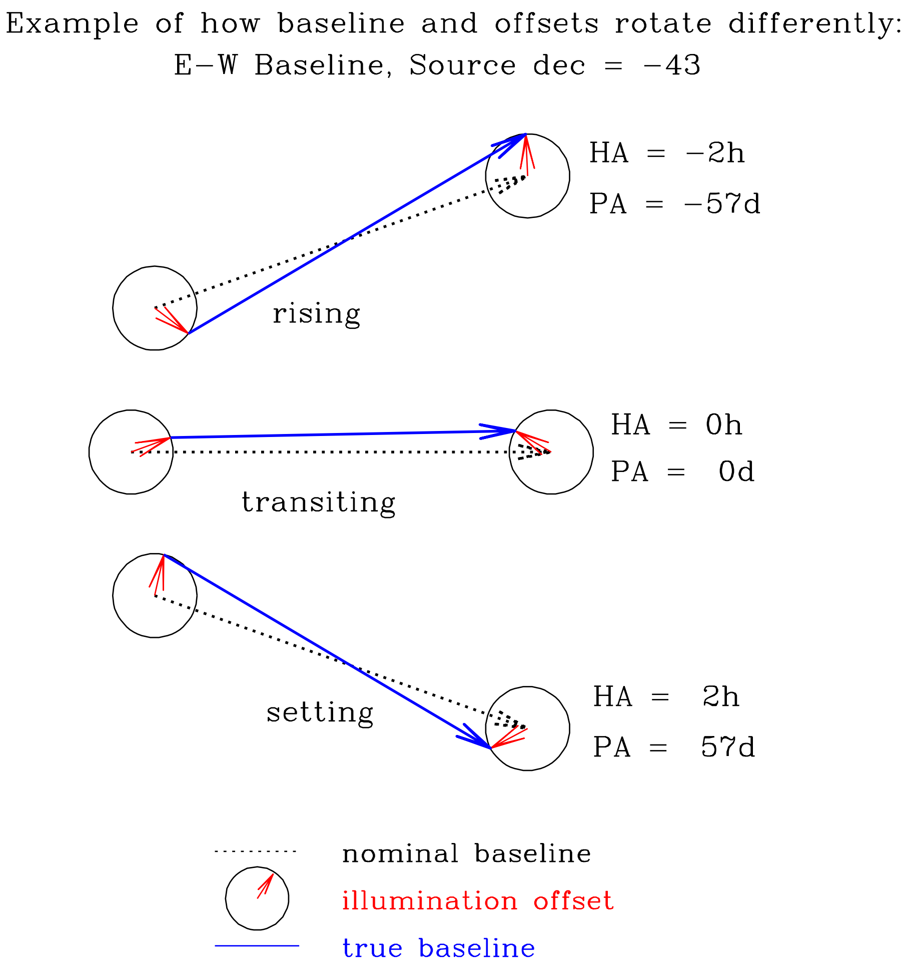

Figure 5: An example of how an E-W baseline rotates against the sky, and how the incremental baseline correction rotates with the parallactic angle.

ready to move on to correct more realistic illumination tapers (ie, 10 dB at the dish edge), and also simulated data which is affected by pointing errors, surface errors, and the illumination offset.

5 Imaging Simulations

To gauge the level at which the (u, v) correction scheme will work, we have performed numerical imaging simulations of the effect of illumination offsets and their correction on ALMA mosaicing. As the offset specification is independent of frequency, we are free to simulate at any frequency.

5.1 Simulation Details

The details of the imaging simulations include:

- simulations were performed in SDE.
- homogeneous array mosaics were performed. Total power was measured by 4 antennas spending four times as much time on the source as the interferometer. The total power extended over a larger region than the interferometric observations covered. No ACA (ALMA Compact Array) data was added to the regular ALMA interferometric and total power data.
- A 3 x 3 pointing mosaic of the CLUSTER model image from the ACA simulation campaign was used for this study. The image was scaled such that a 3 x 3 pointing mosaic would cover the model source at 115 GHz.
- no thermal noise, phase noise, pointing errors, or surface errors were included in the simulated visibilities.
- illumination offsets ranging from 0% to 20% of the dish radius, rms, were simulated.
- for each value of the illumination offset, a separate voltage pattern was calculated for each antenna. The voltage patterns were generated for a 12 m circular aperture, four feed legs with blockage of 8 cm wide, blockage from a 75 cm secondary, no panel gaps, a 10dB Gaussian illumination taper at the dish edge, and Gaussian offset errors, centering the tapered illumination off from the primary's center (see Figure 3). The model apertures, with zero phase errors, were Fourier transformed to yield the far field voltage pattern.
- Once the far field voltage patterns for each antenna were simulated for a given rms illumination offset, they were used to calculate visibility data on a baseline by baseline basis. For a given baseline at a given time, the voltage patterns were regridded onto the model brightness distribution's grid, rotated by the parallactic angle and shifted to the pointing position; the second regridded voltage pattern is conjugated and they are multiplied by each other to form the primary beam. The primary beam is multiplied by the model brightness distribution, and the result is Fourier transformed onto the single (u, v) coordinate appropriate to that baseline at that time. (An FFT is not efficient, as only one baseline has this exact primary beam with this orientation, so a DFT is used.)

- The (u, v) corrections were derived from the individual voltage patterns by fitting a plane to the phase of the voltage pattern, considering all pixels above the half power point. No errors (such as thermal noise or atmospheric phase errors) were considered in the gradient fitting process, indicating that very good measurements have been performed here. This should be the case if the illumination offsets are indeed constant with time. The phase gradients were converted into antenna-based (u, v) offsets which were then rotated by the parallactic angle and incremented pairwise to the nominal (u, v) coordinates for the corrected simulations. These (u, v) increments were generally about half as large as the simulated illumination offsets when a 10dB taper was used, but essentially equal to the illumination offsets when a 100dB taper was used (ie, the effect on the voltage pattern was entirely a phase shift, and the illumination pattern, dropping to zero before the edge of the dish, was symmetric about the illumination offset.

- Because of a problem relating to phase rotation which is still not understood, we could not use the standard SDE simulation tools, which reference all visibility phases and (u, v) coordinates to the center of the model image. Again using the 100dB taper case as a debugging tool, we found that the correction worked perfectly (ie, visibility phases were correctly reproduced down to the 0.01 degree level) for the central pointing (where there were no phase or (u, v) rotations as the simulated interferometer phase center was the same as the reference value of the image). However, none of our attempted phase rotations could get the correction to work for the other pointings. We were able to make a work around. Since everything worked when we were looking at the central pointing (ie, aligned with the center or tangent point of the image), we used the SDE simulation software to simulate a series of single pointing observations, with the single pointing shifting to the different mosaic pointings. We accomplished this by using HGEOM to regrid the model image to be tangent to the celestial sphere at the location of each of the 9 mosaic pointings, controlling the details of the pointings and transformations with a shell script which ran HGEOM and the simulation program. After each pointing is simulated, all the data are concatenated and the subsequent (u, v) offset correction works perfectly for the 100dB taper case, and very well (like 85%) for the 10dB taper case.

  (When I first joined NRAO and the MMA project, I was told that Robert Braun had attempted to perform mosaic simulations by using HGEOM to shift the model image and then use a single pointing simulation and concatenating the single pointing visibility sets. However, this approach was limited by errors in the regridding to about 200:1 dynamic range, and SDE was written in part to address the need for high quality simulations for the MMA. I checked my HGEOM+single pointing simulation loop with no illumination offsets and no other errors against the standard SDE mosaic simulation code, and found that there were no substantial differences, so imaging limitations seen in the simulations can be attributed to the illumination offset or the deconvolution method, but not to the simulation method. The fact that these simulations oversampled the synthesized beam with almost 6 pixels across probably greatly reduced the level of the regridding errors in HGEOM, so we did not see the problem reported in the past.)

- The simulated data were imaged with the MEM-based mosaic program (Cornwell, 1988). An effective primary beam for mosaic reconstruction was formed by averaging all of the baseline primary beams (Holdaway, 1992), calculated from the individual voltage patterns:

$$PB(\mathbf{x}) = \frac{1}{n(n-1)/2} \sum_{i>j} ABS(VP_i(\mathbf{x})VP_j^*(\mathbf{x})). \quad (5)$$

Furthermore, the effective primary beam was azimuthally averaged to give a single radial profile for the beam. A small reconstruction error was made here due to the asymmetric feed leg side lobes and second order effects due to the illumination offset's asymmetric illumination.

- The success of the image reconstructions was gauged with the image plane fidelity (Cornwell, Holdaway, and Uson, 1994; J. Pety, F. Gueth and S. Guilloteau, 2001), and the Fourier plane fidelity (J. Pety, F. Gueth and S. Guilloteau, 2001).

5.2 Simulation Results

The image plane fidelities (broken up by pixel brightness) and Fourier plane fidelities (broken up by range in the (u, v) plane) are plotted as a function of the illumination offset as a fraction of the radius for mosaic images made from uncorrected and corrected data in Figures 6 and 7.

The Fourier plane and image plane fidelities bear some investigation, for they are rich and strange to say the least.

In the Fourier plane, the image fidelity of the uncorrected images is roughly inversely proportional to the fractional offset (see below). After applying the correction to data affected by the illumination offset, the resulting images are essentially of the same quality as the case with no illumination offset. In other words, to first order, the correction scheme works perfectly.

However, if we look closely at the simulation results (often a dangerous thing), we see that sometimes the corrected images are actually better then the zero offset (ie, error free) case. Around offsets of 5% and on baselines less than 37 m, we see a modest increase in image quality after the correction has been applied. How can this be, that the fix improves the images to be better than if there were no errors? Then, as expected, the quality of the fix gradually decays for the larger offsets. (Remember, the fact that the illumination taper doesn't fall to zero at the dish edge results in an asymmetric illumination about the illumination center, which then results in voltage pattern errors which are mostly, but not entirely described by a phase gradient. Those non-gradient voltage pattern errors increase with larger offsets, so it is expected that the method would begin to break down for larger offsets.)

In the image plane, the results appear more puzzling. The uncorrected image plane fidelities decrease with illumination offset as expected (except around an 0.05 fractional offset clipping for pixels brighter than the 10% level, where a bump in the image fidelity is observed; that bump is not really understood). However, the image fidelity for the corrected images *increases* with the illumination offset. The first paradox is that the corrected images show Fourier fidelities which generally *decline* with offset as expected, while the image plane fidelities for the corrected images *increase* with offset. How can this be? It seems that the Fourier plane and the image plane should mirror each other here. This paradox is resolved by the fact that the Fourier plane fidelity looks at the entire image, while the image plane fidelity is only looking at on-source pixels. Apparently, the decline in the Fourier plane fidelity is dominated by off-source errors. The slight rise in Fourier plane fidelity is probably explained below in the image plane fidelity explanation.

The second paradox in the image plane fidelity results is the increase in the fidelity with illumination offset for the corrected images. The increase is very modest for the faintest pixels, but about 50% for the brightest pixels. What is going on? I think I can explain it, but it is a complicated story.

The voltage patterns were created from simulated aperture illumination patterns, which included the feed legs. The feed legs result in asymmetrical side lobes which differ from the rotationally symmetric primary beam model by a fraction of a percent. These side lobes will be systematic, as all antennas have the same feed leg structure. Hence, for the effects of the feed legs, zero illumination error is a worst case. Now, as we increase the illumination offset, we get the phase gradient and other voltage pattern errors which occur at random angles. These residual effects tend to wash out the systematic effects of the feed legs. The errors they make are less systematic than the feed leg errors. Hence, the deviation between the mean primary beam used in reconstruction and each baseline's primary beam will be less than in the zero offset case. This is a very small fractional error, and we can estimate its magnitude by the increase in the image plane fidelity for the brightest pixels, namely on the order of 0.2%. This is the same order of magnitude of the change in the primary beam caused by the feed legs.

Typically, the dynamic range of an image (not reported in these simulations) is much higher than the on-source image fidelity. For an imaging process such as mosaicing, the on-source errors are due in large part to things like errors in the primary beam model or pointing errors, where different data on a given pixel disagree and they average out to a pixel value which is in error. The error is not localized to that pixel, however, but gets scattered across the entire image like the side lobes of the synthesized beam. As the synthesized beam side lobe level is typically a few percent, the dynamic range, limited by off-source errors, is typically one to two orders of magnitude higher than image fidelity, which is limited by on-source errors.

The brighter pixels will have errors which are dominated by the localized on-pixel error, in this case due to the deconvolution errors and the incorrectly modeled primary beam sidelobes due to the feed legs in about equal parts. The fainter pixels will still have the low level of error caused by the unmodeled primary beam sidelobes, but will also have larger errors due to deconvolution errors and also errors scattered from the brighter pixels. Hence, we don't see such a large increase in image fidelity with illumination offset for the weaker pixels of the corrected images.

5.3 Oh Really? Prove It!

We limit our discussion here to the corrected images, which are the ones that show the paradoxical behavior of improving with illumination offset.

There are two more sets of simulations we could do to basically prove this point of view, but the simulations take a very long time to complete, and I am quite tired of doing them. Furthermore, the main points of this memo were to quantify the effects of an illumination offset on imaging and to show that there is an algorithm which will fix the problem, and those objectives have been met.

One type of simulation which could clarify the increasing image quality of the corrected images is to use a 2-D primary beam model for the mosaic reconstruction. Using a 2-D beam model which included the details of the feed leg side lobes should result in somewhat improved images. Specifically, for low illumination offsets, the bright pixels (ie, cutoff of 3% and 10%) would show much higher fidelities, and would gradually *decrease* with increasing offset, as the average details of the feed leg side lobes (and the average primary beam) became more and more smoothed out by the random offset errors and the detailed primary beam model became less applicable. Meanwhile, at low pixel values (1% and 0.3%), the errors are already dominated by other effects such as deconvolution errors. As a modest improvement is seen at these low brightness pixels as the offset increases and a 1-D primary beam model is used, only a modest decrease in image fidelity should be observed with increasing offset if a 2-D primary beam model were used.

The other type of simulation that could prove our view of the corrected images' fidelities is to simply perform the simulations using an aperture illumination which does not include feed legs. This is simple, as it requires no additional software, just time. My assertion is that without feed legs and their associated asymmetric side lobes which are not included in the 1-D primary beam model, the image fidelity of the corrected images would degrade slightly with illumination offset as the non-gradient phase and amplitude deviations in the voltage pattern became larger and larger. Maybe this will get done some week when I am really bored.

5.4 Analysis of Errors in the Fourier Plane

We have investigated the Fourier plane fidelity for the uncorrected images in detail to gain a quantitative understanding of how the illumination offset affects the reconstructed image. First, we have converted the fidelity into its reciprocal, or the fractional error in the Fourier plane. Next, we hypothesize that there is a deconvolution error which is present in all simulations, but masked by the illumination errors. The 0.0 illumination offset case will show just the deconvolution errors. So, subtracting this deconvolution error in quadrature from the errors of the other simulations, we are left with a reduced error, something which may represent the imaging error, viewed in the Fourier plane, caused by the illumination offsets. Fitting a line to the reduced error verses the offset on a log-log graph indicates a power law relationship between the offset and the imaging error (see Figure 8). This exercise indicates that the imaging error is very close to directly proportional to the fractional illumination offset, or that image quality is inversely related to the fractional illumination offset.

6 Conclusions

- The illumination offset results in errors which show up at all spatial frequencies except those measured by total power observations. The errors are linear with the offset magnitude, over the range of offsets studied (0 to 20% of the dish radius).
- If left uncorrected, the illumination offset of 0.6 m (10% radius) rms will limit the quality of ALMA mosaics up to a frequency of about 500 GHz. This means it is more damaging than pointing errors or surface errors up to frequencies of 500 GHz.
- The fact that the illumination offset effect has never been seen on any existing interferometer underscores the subtlety of the effects we are concerned with in trying to justify the need for the ACA (ALMA Compact Array). This is significant in that it is a purely interferometric effect, not affecting total power at all. (Often, it is argued that pointing errors adversely affecting total power data will limit the quality of homogeneous array observations.)

- The phase gradients in the voltage pattern which result from the illumination offset correspond to antenna-based offsets in the (u, v) plane which rotate like the parallactic angle, and a correction of the (u, v) coordinates is algorithmically simple, computationally fast, and results in excellent image reconstruction.

- In fact, when illumination offsets are simulated and the correction is applied, the image reconstruction is in many aspects superior to simulated reconstructions with no error. This is apparently due to the residual amplitude and phase errors in the voltage pattern helping to randomize the systematic side lobes caused by the feed legs.

- The phase gradient parameters for each antenna which are required for the correction can be simply measured by studying the phase of a "five-point" pointing observation performed interferometrically. Thermal noise and atmospheric phase noise will complicate the measurement of the phase gradient at higher frequencies (ie, above 500 GHz), but these sources of measurement error will average down for many repeated measurements.

Acknowledgements

Peter Napier helped initiate this work, and Darrel Emerson had a key insight into the correction algorithm.

References

Cornwell, 1988, "Radio-interferometric imaging of very large objects", A&A 143, 77.

Cornwell, Holdaway, and Uson, 1993, "Radio-interferometric imaging of very large objects: implications for array design", A&A 271, 697-713.

Holdaway, MMA Memo 74, "Surface Accuracy Requirements for Mosaicing at Millimeter Wavelengths", 1992.

Holdaway, MMA Memo 95, "Imaging with Known Pointing Errors", 1993.

J. Pety, F. Gueth and S. Guilloteau "Impact of ACA on the Wide-Field Imaging Capabilities of ALMA", ALMA Memo 398, 2001.

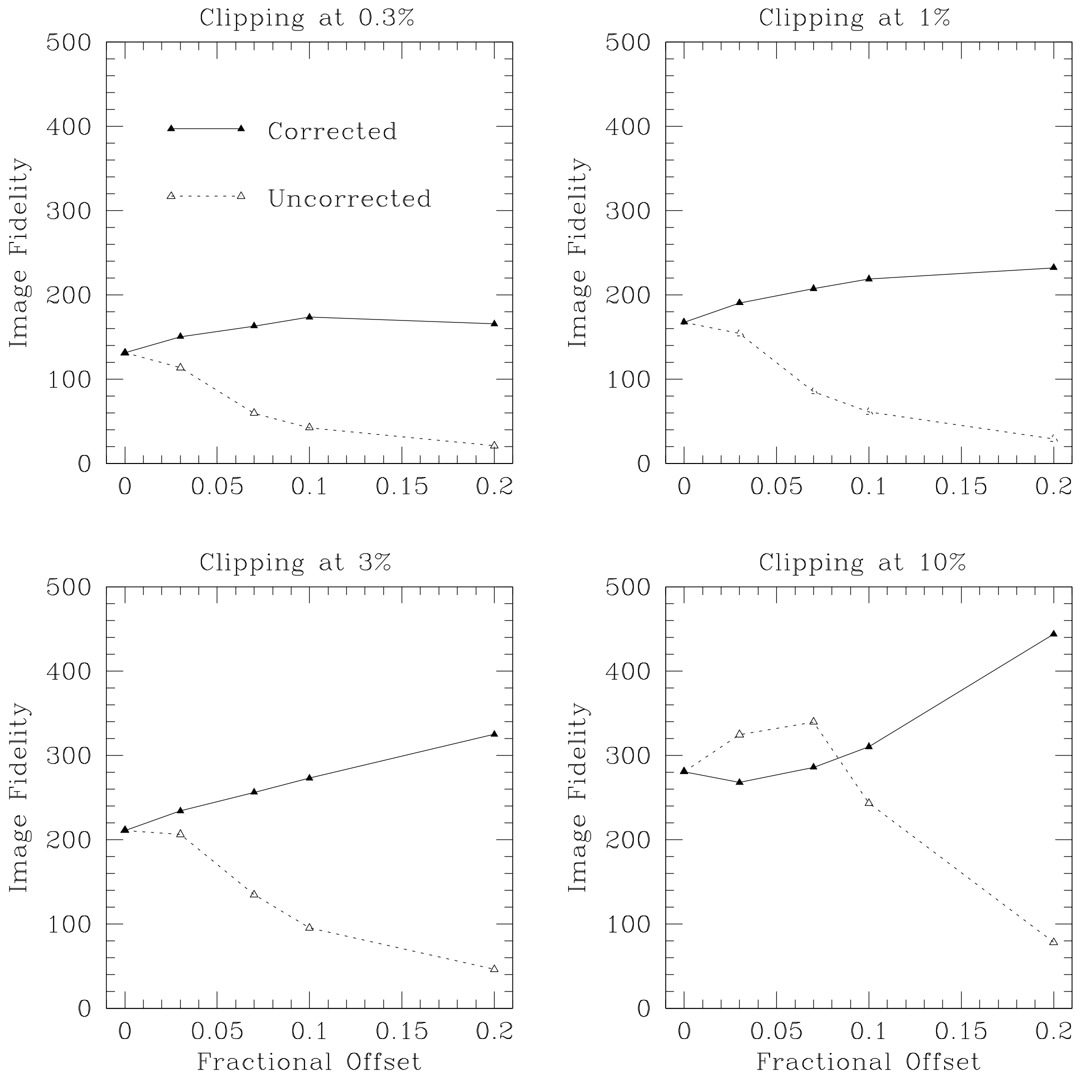

Figure 6: Image plane fidelity as a function of illumination offset as a fraction of the dish radius, for images made from corrected and uncorrected data.

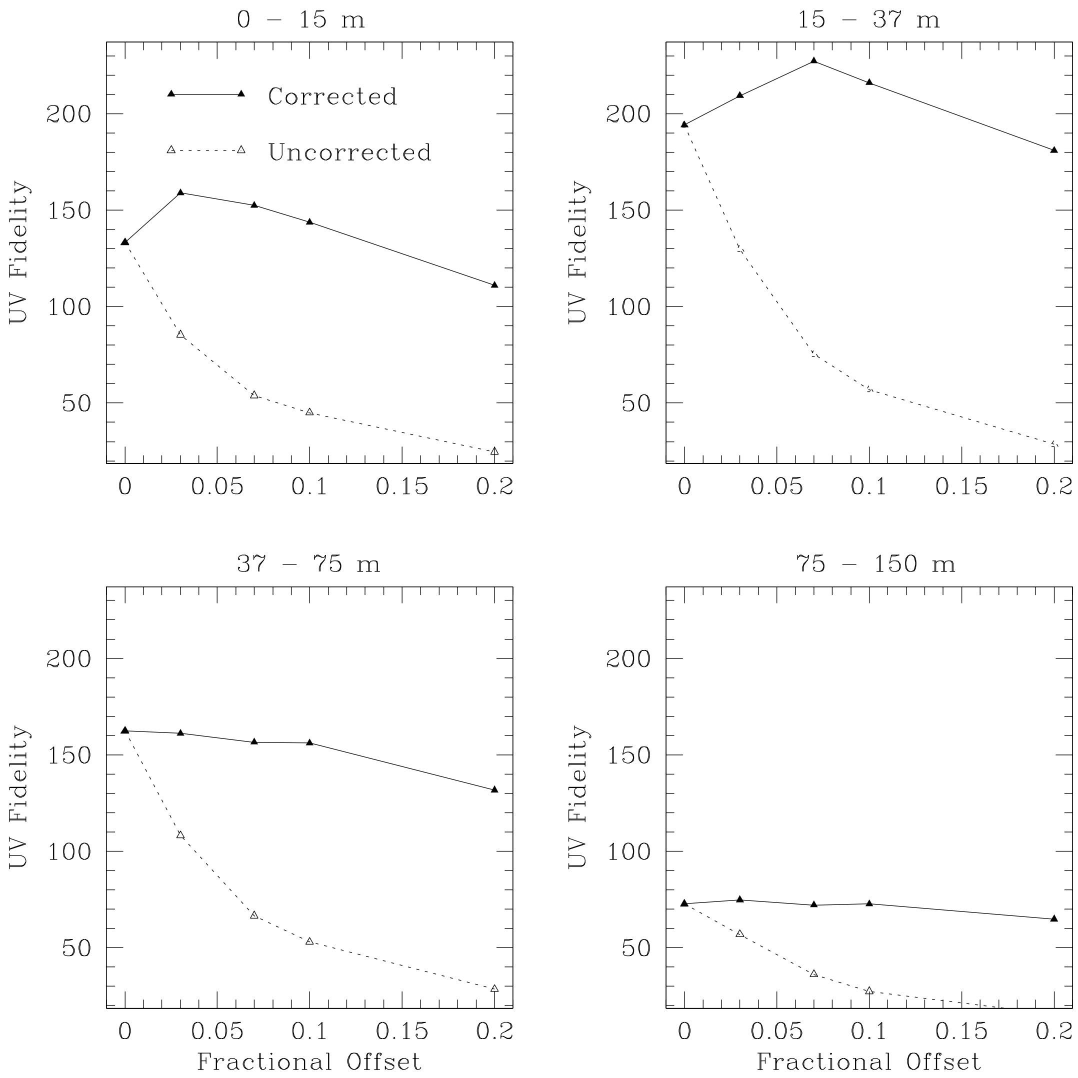

Figure 7: Fourier plane fidelity as a function of illumination offset as a fraction of the dish radius, for images made from corrected and uncorrected data.

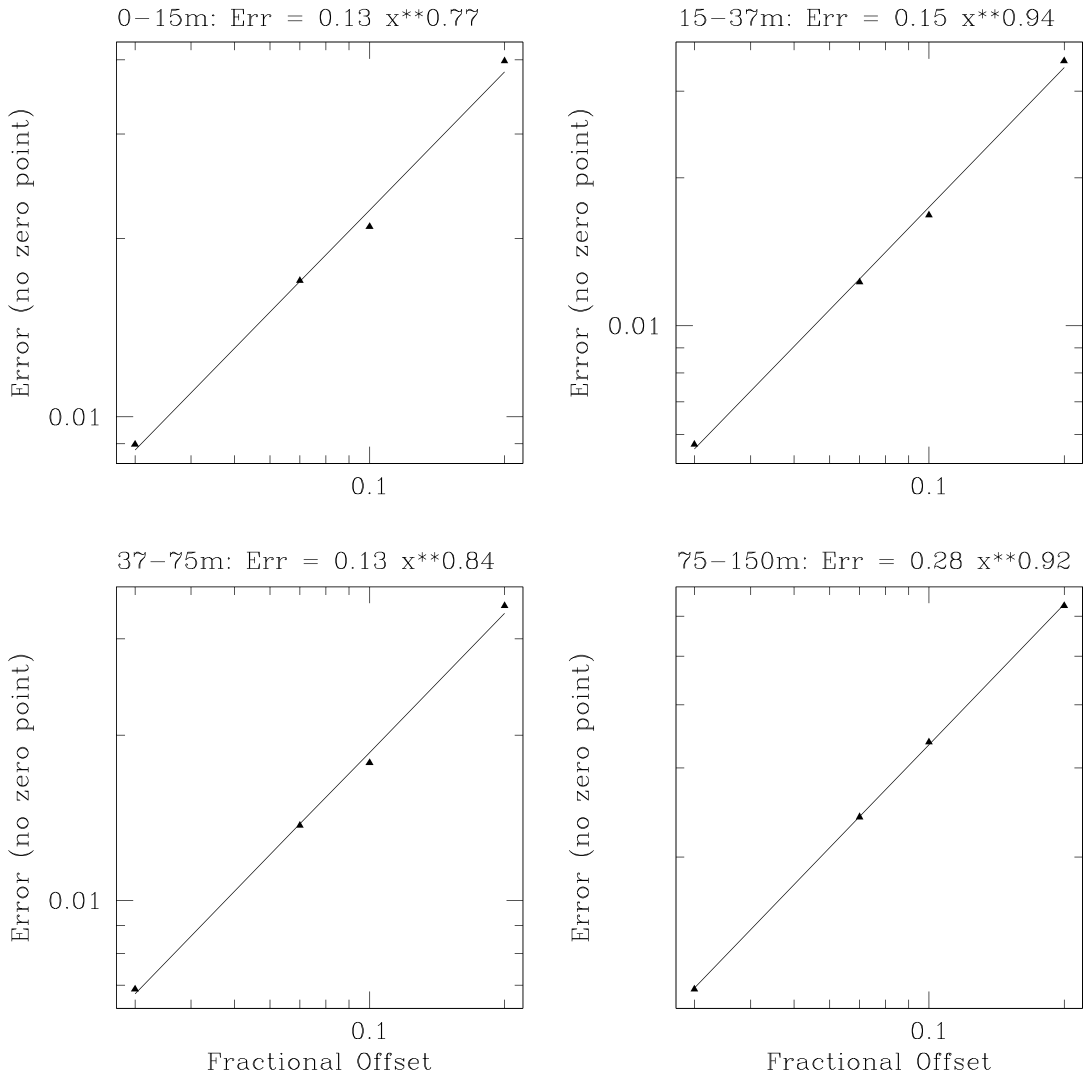

Figure 8: Fourier plane fidelity as a function of illumination offset as a fraction of the dish radius, for images made from corrected and uncorrected data.