Detailed Simulations of Atmospheric Emission Fluctuations' Effects on Continuum Total Power Observations with ALMA

M.A. Holdaway
National Radio Astronomy Observatory

DRAFT: August 18, 2003

Abstract

Atmospheric water vapor emits strongly at millimeter wavelengths. Fast temporal and spatial fluctuations in the atmospheric water vapor results in problems in removing the atmospheric emission, which is done by very quick position switching or beam switching. Most phase fluctuations seen with the 11.2 GHz site testing interferometer are due to fluctuations in water vapor, so we can use the statistics of the phase fluctuation measurements to infer how well the cancellation of variable water vapor emission will be for a given observational strategy and a given atmosphere.

We present a suite of software written in AIPS++ created to simulate this effect, along with the first simulations from the suite. If care is taken in the dynamic scheduling of total power continuum observations, OFT observations will have noise levels only marginally higher than the thermal noise.

Attempts to simulate gain fluctuations with 1/f noise have not produced understandable results, as simulated signals corrupted by 1/f noise still average down reasonably well.

1 The Model

The basic model we use for the atmospheric simulations is a frozen 2-D screen of inhomogenously distributed water vapor flowing over a telescope which is dancing on and off source. The details of the distribution of water vapor in the screen are chosen such that the screen reproduces a common phase structure function observed with the 11.2 GHz site testing interferometer. To acocunt for better of worse conditions, the fluctuations can be reduced or accentuated, but we keep the root structure function power law exponent α the same (0.58). The height of this screen is chosen to be 500 m, consistent with what we think we know about the turbulent layer above the site. We also assume the screen moves at a velocity of 14 m/s. While we have taken care to get the details correct, the 2-D assumption is probably the weakest aspect of the the simulations we report on here. The level of sophistication of other parts of the simulations is a bit awkward compared to the shortcomings of the 2-D screen.

The near field pattern of the antenna is considered to be just a circular column, and the intersection with the atmosphere and the beam are considered to be a circle. The near field pattern of an astronomical ON and an astronomical OFF will show significant overlap, and that is taken into account in our simulations. Elevation effects are in large part "calculated" rather than "simulated" because a full 3-D atmosphere and a more realistic near field beam pattern would be required to simulate them. In order for the granularity of the digital atmosphere to result in minimal problems, we use a cell size of 0.2 m (ie, 60 cells across the 12 m beam column). Our atmosphere model is 600 by 2048 pixels.

2 The Software

Rather than merely a detail, I take advantage of the organization of the software to present many important aspects of the simulations.

The software has been written in AIPS++ as glish scripts, and can be downloaded from `http://www.tuc.nrao.edu/ mholdawa` and run on any AIPS++ installation. AIPS++ is freely available and can be obtained from `http://aips2.nrao.edu/docs/aips++.html`.

1. **almatau.g** encapsulates the approximate behavior of the opacity spectrum above the Chajnantor site into a glish tool. Using the most recent version of J. Pardo's ATM program (current as of 2003-June: see http://www.aoc.nrao.edu/ bbutler/work/alma/calibration/ATM), I calculated wet and dry opacity terms for the Chajnantor site with 1 GHz resolution. The almatau.g script defines a tool which uses the wet and dry term spectra from 1 to 1000 GHz to perform various calculations. Since the site testing data archives opacity at 225 GHz, some special functions are defined on the 225 GHz opacity, scaling it to other frequencies based on the model. The wet and dry terms used in scaling from 225 GHz were calculated with 20 MHz resolution over the correct 500 MHz upper and lower sidebands used by the 225 GHz radiometer to better take into account some dry lines in the USB.

2. **almasensitivity.g** encapsulates the sensitivity of the ALMA array or individucal ALMA telescopes and permits very fine control by the user. The main fine control which we need for the current work is the input of the PWV, which is converted into an opacity at the observing frequency by **almatau.g**. Currently, the noise figures generated by **almasensitivity.g** are somewhat lower than those generated for continuum observations by the ESO web-based sensitivity tool (typically 10-40%). Some of this difference is due to differences in assumptions about the PWV, but some of the difference is real. This memo, which seeks to illustrate the magnitude of the residual sky brightness fluctuations, finds that the noise is generally larger than the residual brightness fluctuations. Hence, if a revised version pf **almasensitivity.g** results in somewhat larger thermal noise levels, this will still be true.

   A couple of sensitivity-related details that we should mention: first, we choose an integration time which is equal to Nyquist sampling at the observing frequency as the beam slides over the source, so the noise in each of our simulated integrations will be the same as the noise level in an OTF map made from our scans. (Our simulations software can't actually make any maps.) Second, standard calculations for the noise increase the system temperature by a factor of $e^{\tau airmass}$. The system temperature doesn't actually increase by this factor, but as this is the gain factor we must multiply by to remove the effects of absorption by the atmosphere, our noise level also increases by this amount. One of the things we are doing in these simulations is observing the atmosphere, so we use noise calculations without the factor of $e^{\tau airmass}$, add that noise to the observation of the atmosphere, and then scale the observed power by $e^{\tau airmass}$.

3. **otfpath.g** encapsulates a realistic path for the OTF antenna motion. We are constrained by maximum slew velocity and maximum accelerations. The antenna motion profiles were derived from top sectret Vertex documents which provided simulations of the Vertex prototype antenna motions for fast switching. Analysis of these fast switching profiles indicated, in addition to the velocity and acceleration limits, a maximum JERK of 100 deg/s^3 (remember that JERK is the time derivative of acceleration; the JERK is the main instigator in exciting vibrations in the antenna). The otfpath tool is given maximum slew velocity and a source size, and it returns detailed profiles for the antenna's constant velocity slew over the source and a quick turn-around which does not exceed the maximum jerk, acceleration, or velocity.

   The oftpath tool also indicates which parts of the profile are designated as ON-source and which are OFF-source so that the OFF-source parts can be used for sky subtraction. We can actually fit an n-th order polynomial to the two OFF parts of the scan surrounding the ON part. When the noise is very low (ie, the signal from the sky fluctuations is very large), the ON-source sky residuals are effectively dimenished by increasing n up to about 7 or 8. However, with more realistic noise, fitting such a large order polynomial results in increasing the ON-source sky residuals, as we are just chasing the details of the thermal noise. **In actual observations, the structure function of the OFF-source fluctuations could be investigated to infer if they are dominated by thermal noise or by sky fluctuations to determine the optimal polynomial fit order. This would make a simple research project that could be carried out with this software package in the future.** For the simulations reported on in this work, we have just fit a linear term to the OFF-source fluctuations. In noise-dominated cases, this does not negatively affect the ON-source residuals. In some cases, a higher order fit would improve the ON-source residuals.

   Internally, the otfpath works on a 0.001 s time grid, but the motion profiles can also be provided at other time resolutions (such as the integration time required for Nyquist sampling at a particular observing frequency and slew velocity).

   In our simulations, we make 15 OTF scans, going back and forth, and average the ON-sky residual to give a representative figure for that simulation. A few of those will just happen to have very large residuals and will be dominated by the atmosphere. Significant improvements could be made by identifying the scans which have anomalously large residuals and performing a higher-n sky fit, assigning a representative weight for that scan which reflects the higher fluctuation level, or just throwing that scan out.

4. **phasemonitor.g** encapsulates the behavior of the site testing interferometer. It can "observe" a phase screen and fit a phase structure function to the screen. This is used not only to verify that the model phase screen is a realistic one, but also to determine how to scale the fluctuations in magnitude to account for better or worse phase conditions.

5. **oneoverfnoise.g** provides the simulation machinery the service of creating a gain time series which has 1/f noise. The varying gains, when applied to the simulated observations,

should limit our ability to correct the ON-source atmospheric emission using the OFF-source atmospheric emission.

6. **sdatmosim.g** is the main simulation engine for single dish observing through the atmosphere. You give it the model atmosphere and the details of the observation (atmospheric conditions, source size, the details of the switching method, elevation angle, azimuth angle for switching, etc). It is really complicated. It supports OTF scans and three sorts of beam switching

7. **The 1996-2001 Site Testing Database** We have created a database from 6 years of site testing data which tabulates the phase structure function amplitude and exponent as well as the opacity. The data are reported on a one hour grid. The natural time scale of most of our data is 10 minutes, so we could have as many as 6 samples per hour. All data that fall into a given one hour bin are sorted and the median value is taken. If a given hour has no data, but the preceeding and following hour do have data (such as can exist when the tipper performs a 60 minute stability run), we interpolate for the missing hour.

8. **explore2.g** is used to derive the statistics of the Site Testing Database which are relevant for each observing band.

3 Our Simulation Strategy

The problem of understanding total power observations and the residual (ie, uncancelled) emission fluctuations due to inhomogeneously distributed water vapor is really very complicated. The phase space this problem inhabits contains at least 10 dimensions. In the atmospheric conditions, we have τ, σ_{phi} and α (parameterization of the phase structure function), and atmospheric temperature. Astronomically, we have the observing frequency ν and the source size θ_s (ok, we are a bit silly to do so, but consider a square astronomical source...). And observationally, we have the source elevation angle and azimuth angle with respect to the wind direction, the switching method (ie, on-the-fly (OTF) or a variety of beam switching methods) and the details of the switching speed (either the OTF maximum slew velocity or the beam switching frequency). Obviously we need to make some assumptions and integrate out some of these dimensions.

1. **We assume T_{sky} = 250 K.** Residual sky brightness fluctuations may vary by on order 10% due to this assumption.

2. **Observing frequency** ν is discretized as illustrated in Table 1 and will endure as a (nearly) free dimension in our treatment of this problem. Hence, we ignore complications due to increased opacity at the window edges.

3. We sssume the **phase structure function exponent** α **is 0.58**. This is the median value in the 1996-2001 Site Testing Database, though values range from about 0.33 (smaller values are likely due to the very fast ionospheric fluctuations) to 0.83. There is a slight correlation between α and the amplitude of the structure function (ie, when the phase stability is very good on 300 m baselines, α tends to be smaller, indicating that

the phase errors don't decrease so quickly as you go to shorter baselines), but we ignore this.

4. **The opacity τ and the phase fluctuations** are not particularly correlated. In fact, if we select the best 10% of *tau* conditions and define a_τ such that

$$\langle \tau \rangle = a_\tau \langle \sigma_{phi} \rangle$$

for the τ-selected conditions, and then select the best 10% of σ_{phi} conditions and define a_ϕ such that

$$\langle \tau \rangle = a_\phi \langle \sigma_{phi} \rangle$$

for the σ_ϕ-selected conditions, we find a_ϕ / a_τ to be about 20.

We want to minimize the overall noise in an observation. For total power observations, we need low τ to minimize the thermal noise, and low σ_{phi} to minimize the residual atmospheric emission. Hence, when it is time to select the dynamically scheduled observations, we need to consider both opacity and phase fluctuations. To determine the optimal way of picking an observing frequency, we need to already have solved the problem at hand. In the absence of this solution, I have taken an *ad hoc* approach. We assume that each of the 10 observing bands will be observed 9% of the time. The worst 10% of the time is deemed worthless. We select the best 9% of the observing time for band 10 by finding those conditions with the smallest values of $\tau_\nu + a_\nu \sigma_\phi$, and subsequently pick the next-best conditions for each remaining frequency band. The values of a_ν are listed in Table 1.

This way of simulating dynamic scheduling and band-demand on ALMA is not perfect, but if we can produce good observations which are not limited by the atmosphere with this non-optimal method, a better method will not do any worse. This effectively collapses three dimensions of our problem phase space, opacity, phase stability, and observing frequency, into a single dimension, frequency. For each observing frequency band, we consider median conditions for the best half and the worst half of the conditions (ie, the 24% and the 75% conditions in each band).

5. For the time being, **we are only investigating OTF observations**. This is just to get this project airborne. Earlier unpublished work (Holdaway, 1998) indicated that OTF observations are superior to beam switched observations for all source sizes except for the most compact. While we intend to investigate beam switched observations in the future, clearly OTF observations cover a wider swath of phase space than beam switching observations.

6. **Switching rate (or v_{slew} for OTF)**: For an OTF observation, we will need to select a reasonable maximum v_{slew}. However, for given observation and atmosphere, we can calculate the optimal v_{slew}. Very slow v_{slew} will result in observations which are limited by residual sky emission. Very fast v_{slew} will result in more time being spent turning around off the source, or a low duty cycle and hence a higher thermal noise. An example of this effect is seen in Figure 1, which shows the noise level as a function of v_{slew} for a case with no atmospheric fluctuations (open boxes) and a case with representative atmospheric fluctuations (filled boxes). The atmosphere jumps up very quickly at low v_{slew}, but the duty cycle effect is fairly gradual. Hence, if we err on the side of too high

v_{slew}, our observations are compromised less than if we err on the side of too low v_{slew}. Another detail to consider: at the optimal v_{slew}, the observations' noise level will have a non-negligible contribution from the atmospheric fluctuations. Higher v_{slew} may have noise which is totally dominated by thermal noise (ie, no sign of the atmosphere), but will have a higher noise level. This detail may be misleading, as we have not yet studied how the residual atmospheric errors average down spatially or with time.

7. **Elevation angle**: there are a number of competing factors concerning elevation. Generally speaking, millimeter observations degrade rapidly as the elevation angle decreases. The phase fluctuations increase as the square root of the airmass; furthermore, at low elevation angles, the angular offset between an ON and an OFF beam translate into a larger separation through the atmosphere at the elevation of the water vapor differences. Both these effects increase the magnitude of the residual sky fluctuations in the noise, and can be partially compensated for by switching more quickly.

   Of course the opacity increases proportionally to airmass. Thermal noise takes a double hit from the opacity, as the opacity adds thermal sky noise and the noise is scaled up to account for the signal lost due to opacity.

   The interesting thing is that things again degrade at the higher elevations. In order to cover an angular distance on the sky of θ_s, an antenna with an AZ-EL mount must move $\theta_s/cos(el)$. Hence, as we approach the zenith, our antennas need to slew faster and faster just to keep up with what they were easily doing at lower elevation angles, resulting in a lower duty cycle and higher noise.

   While the elevation effects are important and interesting, we do not include a full treatment of them here. We do show a single slice along the elevation dimension in Figure 2 which illustrates the competing forces we have mentioned. At least for the other parameters used in this work (300 GHz, 1mm PWV, 1 degree of phase at 300m and 11.2 GHz), the elevation angle of 60 degrees is optimal, and we use that elevation angle for all other simulations reported on in this document.

4 Simulation Results

4.1 Optimal Slew Velocity

The slew velocity for the OTF observing mode is one of the dimensions in our simulations which we are able to ignore by picking the velocity which minimizes the ON-source residuals. Figures 1 and 2 illustrate how we select the optimal value of the slew velocity.

Hey, we are still missing the numbers for how well we do for all observing conditions, split up among all 10 bands. I initially ran the numbers for very large source sizes – 0.2 degrees; but now I am looking at smaller sources. As you can see with the size-plot, things work pretty well for small or large source sizes, though.

Band	ν [GHz]	fraction	a_ν	Good		Bad	
				PWV	σ_ϕ	PWV	σ_ϕ
0	–	0.20	0.004	2.337	10.69	1.847	12.685
1	43	0.08	0.004	7.438	5.190	3.012	8.295
2	80	0.08	0.013	3.874	4.120	2.802	5.730
3	90	0.08	0.020	2.500	3.270	2.360	4.225
4	145	0.08	0.025	1.731	2.120	1.847	3.150
5	190	0.08	0.030	0.683	5.090	1.055	4.315
6	230	0.08	0.030	1.102	1.970	1.102	2.610
7	345	0.08	0.2	0.822	1.715	0.892	2.055
8	500	0.08	2.0	0.613	1.440	0.706	1.580
9	680	0.08	2.0	0.496	1.010	0.566	1.235
10	880	0.08	2.0	0.310	0.690	0.403	0.880

Table 1: This table shows the frequencies we chose for our total power simulations, and the fraction of time each band was observed (band 0 represents the 20% of the time which was discarded as unusable). The parameter a_ν is used to form an objective function $J = \tau_\nu + a_\nu \sigma_\phi$, where τ_ν is the opacity at the observing frequency and σ_ϕ is the 11.2 GHz site testing interferometer phase (this is a heterogeneous measure, using one parameter at the observing frequency and the other parameter at the 11.2 GHz reference frequency, but "oh well!"). This objective function determines how observing conditions are divied up among the 10 bands, with the 9% of conditions with lowest J allocated to band 10, the next lowest 9% of conditions allocated to band 9, and so on.

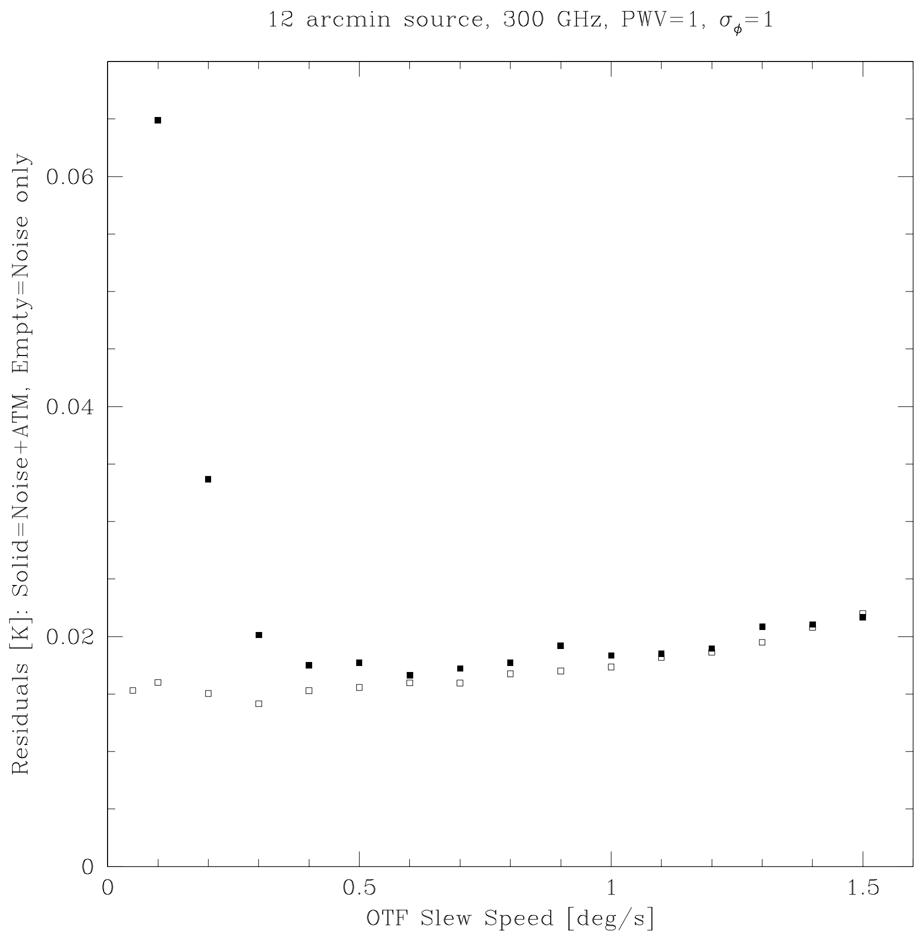

Figure 1: On-source residual noise levels for a single scan through a 12 arcminute source observed at 300 GHz with 1mm PWV and 1 deg rms phase noise on the 11.2 GHz interferometer. In order to compare the different slew velocities, we normalized the residuals to 1 second scans by multipltying the noise per Nyquist integration time by $\sqrt{t_{scan}}$. Empty boxes represent thermal noise only, solid boxes have atmospheric fluctuations included. For very slow slewing velocities, the residuals are dominated by the atmosphere, while thermal noise dominates for very fast slewing.

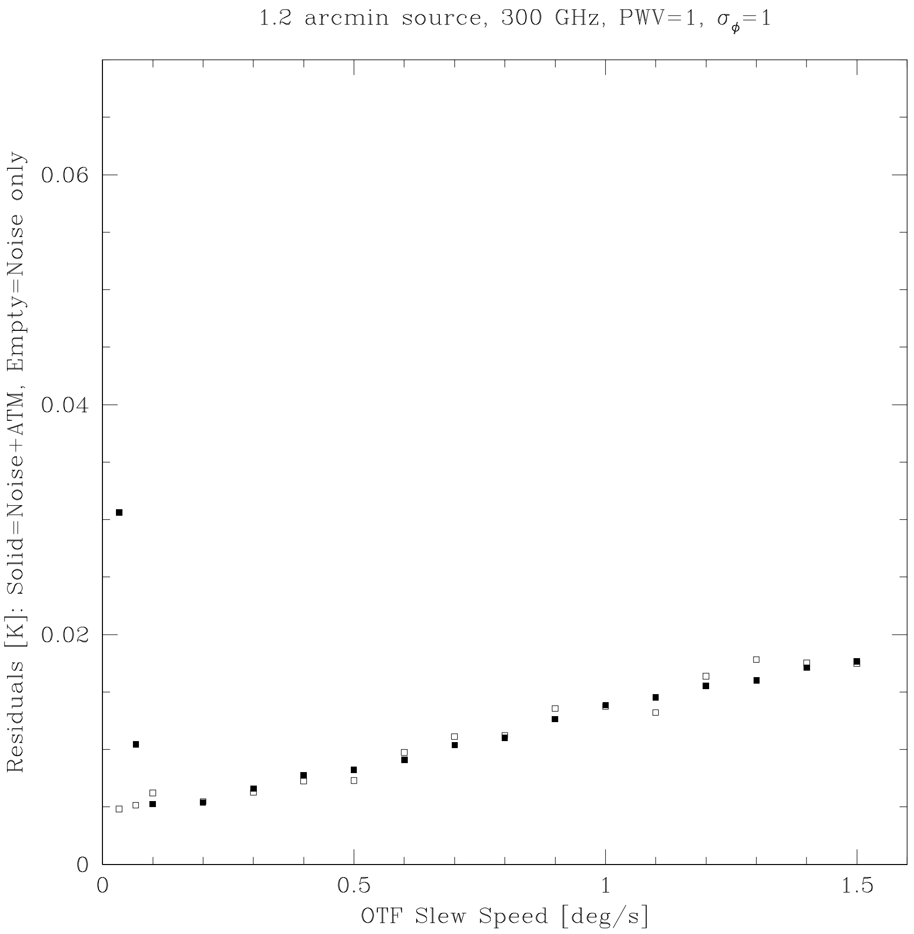

Figure 2: On-source residual noise levels for a single scan through a 1.2 arcminute source (ie, 1/10 the size of the previous plot), observed at 300 GHz with 1mm PWV and 1 deg rms phase noise on the 11.2 GHz interferometer. See the previous plot for the meaning of the solid and empty boxes. This source is much smaller, so (a) the optimal slew velocity is much smaller, (b) as high velocities run off the source much further (as a fraction of the source size), the thermal noise-dominated high v side of the curve increases more steeply, (c) as the source is 10 times smaller than the previous case, in a second (our chosen normalization) we spend approximately 10 times more integration per Nyquist sampling interval, so our optimal noise level should be $\sqrt{10}$ times lower, which is true.

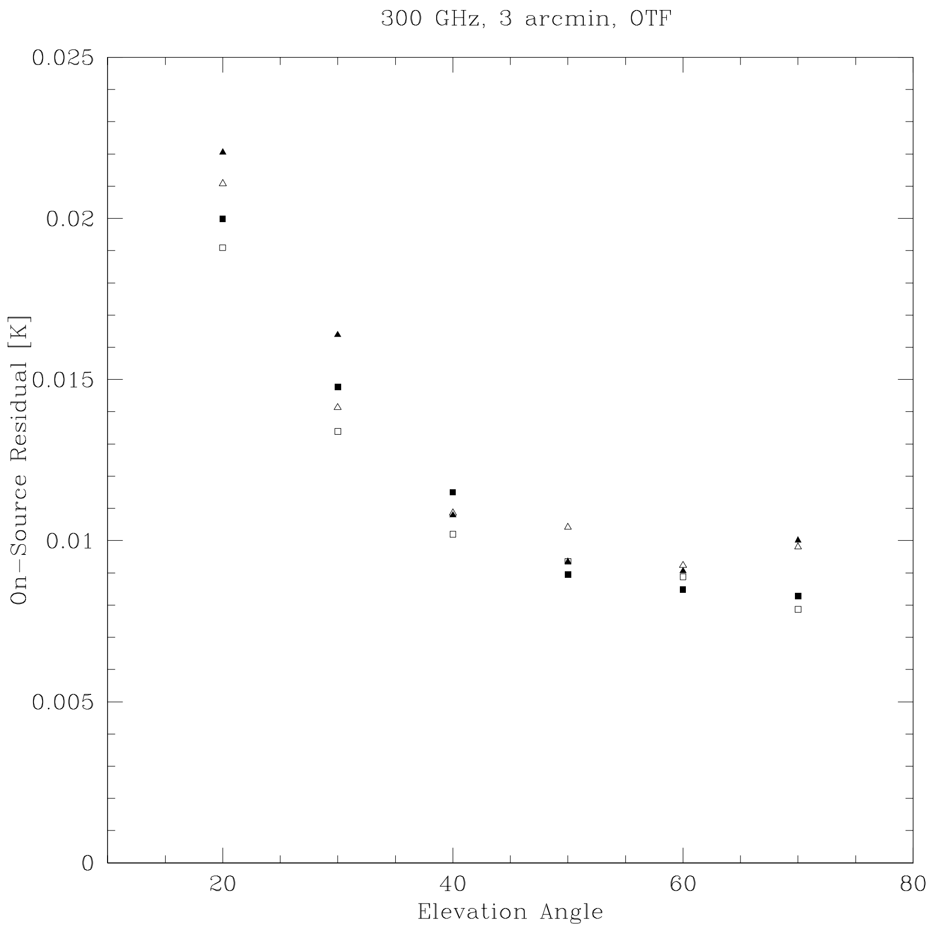

Figure 3: Residual on-source noise in a single OTF scan, normalized to 1 s, as a function of elevation for a 3 arcmin source observed at 300 GHz with 1mm PWV and σ_ϕ of 1.5 deg (boxes) or 3 deg (triangles). Open symbols have no atmosphere added, filled symbols have the atmosphere added. According to our ad hoc manner of assigning atmospheric conditions to the various observing frequencies, 1.5 deg phase errors (as measured with the 11.2 GHz iste testing interferometer) would be good conditions for 300 GHz observations, and 3 deg phase errors would be worse than most conditions which would be assigned to 300 GHz. As expected, the 3 deg (triangle) cases indicate somewhat higher noise than the 1.5 deg cases. The atmosphere-free cases (open boxes and open triangles) do not produce exactly equal residual levels because (a) the thermal noise process is random and (b) these values have been calculated with the optimal v_{slew}, and as the 3 deg case requires a larger v_{slew}, the thermal noise will be systematically higher. Yes, things degrade at low elevation, but in all cases, the noise is dominated by thermal noise and not the atmosphere.

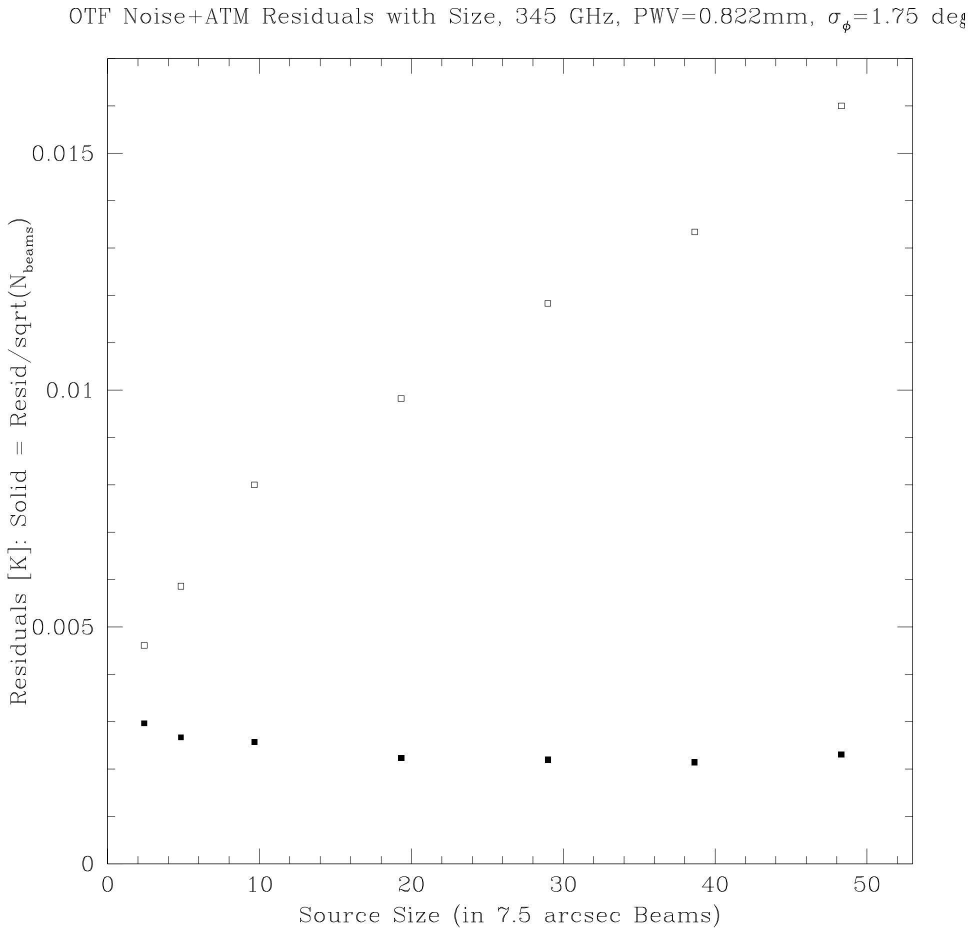

Figure 4: On-source residuals for a single OTF total power scan, as a function of source size. The open symbols have been normalized to 1 s scans (ie, $\sigma\sqrt{t_{scan}}$). However, as larger source sizes are serviced by faster v_{slew}, and hence less time per Nyquist integration, comparison across source size can be better achieved by a further normalization to account for the increase in observed points across the larger source sizes. Hence, the filled symbols indicate $\sigma\sqrt{t_{scan}/n_{beams}}$. The gradual increase in this quantity for the smaller source sizes indicates the increased inefficiency in OTF for very small sources due to the antennas spending more time off source. The phase and opacity conditions are comensurate with the 345 GHz observing frequency.