On Calibration of ALMA's Solar Observations

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

January 4, 2007

Abstract

1 Introduction

Solar observations with ALMA will be challenging for a number of reasons: solar observations are wide field and will require adding short spacing information as well as mosaicing. However, solar observations are very dynamic, with structure changing on timescales of seconds. Solar observers are also interested in wide field polarization measurements. As high frequency solar observations have been rare, the solar astronomers don't know exactly what they want or need, so we have to be ready to give them all observing bands and in principle, all baselines. Solar observing will require novel observing modes, such as scanning the sun with several single dishes, looking for flaring regions, and then training the entire ALMA array on the flaring region once it has been found.

In order to observe the sun without frying the receivers, the panels on the dish have been scalopped to diffuse the solar radiation. However, the sun is still too hot for the ALMA electronics, and we plan to equip each antenna with a 13 - 16 dB attenuator that can be inserted into the optics path above each feed in about 2 s.

1.1 The Attenuator Conceptually

Observing with the solar attenuators is quite a bit like observing through a nearly opaque atmosphere: the signal from the astronomical source is greatly diminished, a delay error is introduced above each antenna, and a noise source almost equal to the ambient temperature is added to the system. Of course, the big advantage we have over dealing with a nearly opaque atmosphere is that the attenuators are just sitting on the antenna rather than moving across the sky.

We can think of the attenuator's transmission as the amplitude of a complex gain factor, and the delay as a phasor on this complex gain. Since each attenuator will be a bit different, we need to be concerned with the attenuator's complex gain varying from antenna to antenna and possibly with time or elevation. We also need to be concerned with the distribution of the complex gain across the attenuator; if not uniform, this will be equivalent to introducing illumination or surface errors to the aperture as the subreflector does. Depending upon the dielectric constant and design of the filter, it may affect the focus position of the secondary reflector. It could also affect the antenna pointing.

This document lays out a simple strategy for calibrating these filter parameters and how to observe the sun.

1.2 Quiescent Spectrum of 3C273

Our main strategy will be to observe a bright calibrator through the solar filter, so a good estimate of how bright such a calibrator may be is essential to understanding if such a calibration will work.

Back when we had a library at NRAO/Tucson, I went through the literature and looked at millimeter and sub-millimeter wavelength continuum flux measurements taken with single dish telescopes on the bright quasar 3C273. These guys were all interested in bright quasars' high frequency variability, so they presented time series of flux measurements over a year or longer. I noted the lowest measured flux value at each frequency in building an estimate of 3C273's quiescent spectrum. I then interpolated and extrapolated to fill out an estimate of 3C273's flux at each of the 10 ALMA observing bands, which is shown in Figure 1. 3C273 may never have this flux profile, as small and large flares are always occurring, starting at the high frequencies where the emission becomes optically thin far down the jet, and then moving to lower frequencies. However, the quiescent spectrum is very useful for our purposes, because it represents the minimum flux that we are more or less certain of finding on 3C273. Usually we will find more flux.

Looking at the quiescent spectrum, we see at low frequencies the spectral index is -0.45. Above a spectral break around 200 GHz, the spectrum of the quiescnet emission steepens to -1.1, which is a problem for us in the submillimeter. 3C273 does have violent flares in the submillimeter, and the flux could be several times higher than the quiescent spectrum, and I'm sure ALMA staff astronomers will make good use of 3C273 when it is flaring.

1.3 A Noise Calculation

A useful calculation to perform is the SNR which we might expect for gain solutions while observing 3C273 through the solar filter over 10 s. If we can see that source through the filter with high SNR, we will be able to characterize the various parameters of the solar filter astronomically. While we should be able to estimate most of these paramaters based from lab measurements of from calculations based on lab measurements, an astronomical check is always a good thing to do, as is an astronomical check on variability.

The exact atmospheric conditions we choose to use in our calculation are not very important because the noise will be dominated by the solar filter. We use the first quartile of opacity conditions and an airmass of 1.3 for the SNR calculations. We consider the ambient temperature of the filter to be 273 K. If the filter is 15 dB, then 0.03 of the radiation will pass through, and the filter will present a load of 273 K * 0.97 to the receiver - or 265 K.

With the solar filter in, the SNR is decreased by two factors: the signal is greatly decreased by the filter, and the noise will increase greatly as an ambient temperature load (ie, the filter) is inserted into the optics path.

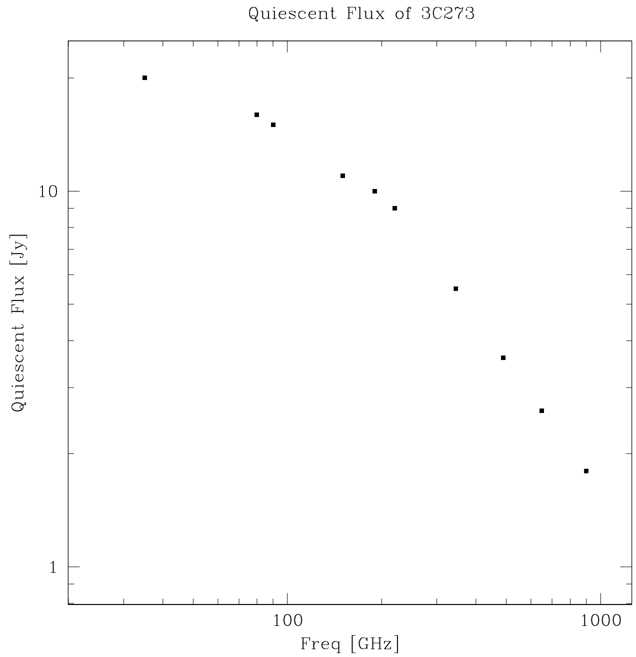

Figure 1: Quiescent flux of 3C273.

The minimum SNR we expect on 3C273 as a function of observing frequency is shown in Table 1. At any given time and observing frequency, 3C273 will usually be brighter than the queiscent level, by a factor typically between 1 and 4, which will increase the SNR. While 3C273 is only in the sky half the time, there will usually be a few other similarly bright quasars hanging about, and other quasars will flare and outshine 3C273. However, it is our hope that "the brightest quasar in the sky" is never actually required during real astronomical observations, but only for test observations for calibrating "constant" parameters of the solar filters. In this case, we can just get test time scheduled for when 3C273 or another bright or flaring source is up and at high elevation.

The SNR of the gain solution is equal to the SNR of a single visibility multiplied by $\sqrt{N_{ant} - 2}$. For these calculations, we have used $N_{ant} = 50$, though there is reason to use a larger number that reflects the total collecting area of the ACA, which could help us get slightly better gain solutions.

While 10 s is a canonical calibration time scale, we can always achieve better SNR by integrating longer, barring systematic errors which we have not considered. (When integrating longer, we may need to perform some sort of atmospheric switching on short timescales and averaging up. The relatively slow 2 s insert/removal time may be the limiting factor for such switching observations.)

1.4 Planets, Anyone?

We prefer a compact source such as 3C273 for any interferometric measurements, and we need to perform some calibrations interferometrically. However, there could also be a significant role

Freq	Band	Noise WITHOUT Filter [Jy]	Noise WITH Filter [Jy]	Flux of 3C273 [Jy]	Flux of 3C273 /32	SNR
35	1	0.0003	0.0027	20	0.625	232
80	2	0.0004	0.0028	16	0.500	179
110	3	0.0005	0.0029	15	0.469	162
150	4	0.0006	0.0030	11	0.344	115
190	5	0.0010	0.0033	10	0.313	94.9
220	6	0.0009	0.0033	9	0.281	85.2
345	7	0.0018	0.0042	5.5	0.172	41.0
490	8	0.0111	0.0135	3.6	0.113	8.37
650	9	0.0110	0.0134	2.6	0.081	6.05
900	10	0.0215	0.0239	1.9	0.059	2.47

Table 1: Question: is 3C273 bright enough to calibrate the solar attenuator complex gain? The "Flux" in this table is the quiescent level of 3C273 between flares, and "Flux/32" is the flux after 15 dB of attenuation. The noise level represents the noise on each gain solution after a 10 s integration with the additional load of a warm attenuator adding 265 K to the system temperature. The integration time can be increased to improve SNR without significant contamination by the atmosphere by switching the attenuator in and out and using the unattenuated observations to correct for the atmosphere. If we desire 1% errors, we need a gain SNR of 100. We can easily do that up to 345 GHz, but it becomes increasingly difficult to achieve this accuracy in the submillimeter.

played by planets. Due to their extended nature, lanets are optimally observed in total power, but if they are very bright, there could be enough flux on short interferometer baslines to be useful. In the submillimeter, the planets begin filling a larger fraction of the beam, which is bad in that we want the source to be small - but its good in that the planets become brighter in Jy.

In the submillimeter, where we need an alternative to "the brightest quasar" (which just isn't bright enough to be seen well through the solar filters), all of the planets except Neptune and Uranus are too big in angular size (well, Mars is about 1/3 of a beam between 345 and 490 GHz, depending upon how far it is from the earth). Neptune will be about 1/3 of a beam or less up to 900 GHz, and Uranus is 1/3 of a beam or less at 490 and 345 GHz. And these planets will be about 100-200 Jy. If we observe in total power, we must pay a penalty in not having 50 antennas to perform a gain solution with (loss of 7 in SNR), as well as a switching penalty of about $\sqrt{2}$ in SNR. The factor of about 100 greater flux in the planet as compared to 3C273 means we come out way ahead observing the planet in single dish - by abut a factor of 10 over 3C273 observed interferometrically at 900 GHz.

In principle, planets observed in total power mode could be used to calibrate pointing, focus, the amplitude of the solar filters' gains, and the primary beam shape with the attenuators in. On the other hand, total power continuum observations will be virtually impossible owing to the large 1/f fluctuations and the lack of nutators on most of the antennas, so total power observations of planets are probably not really a viable option.

On the other hand, we could try to observe the planets interferometrically. We will have to model the visibilities using a disk model for the planet. The shortest baselines will have something like 5% of the planet's flux - even so, the planet likely comes out way ahead of the brightest quasars for calibration in bands 9 and 10, and can be used for all interferometric calibrations, though the algorithms may need to be modified to deal with the visibility model for the planet.

2 Various Calibrations

2.1 Pointing

The solar filter may introduce a pointing offset (for example, if it did not have a uniform thickness, but had a slight gradient in thickness). The solar filter should probably have a specification set on the uniformity of thickness so as to not have any pointing issues.

However, if there are pointing deviations caused by the filter, we can calibrate that effect. We will be able to perform a pointing cal on the brightest sources in the sky with the filter in, and then do a pointing cal with the filter out, and thereby solve for the offset. However, this offset must be constant - ie, we hope to perform the pointing offset calibration once during test time, and then use the pointing offset to adjust for observations when using the solar filter.

Software requirement: if the solar filter results in a pointing offset, the telescope control software must be able to take the pointing offset in and out as the solar filter is inserted and taken out.

High frequency warning: we won't have enough SNR observing 3C273 in the submillimeter to measure any pointing offset through the filter. One way arround this is to use an offset measured for lower frequencies plus a calculation which accounts for any geometrical differences between the low and high frequencies. Or we could use Neptune or Uranus, in total power or interferometrically.

2.2 Focus

The solar filter will affect the focus. Again, we expect the effect on the focus to be constant in time. The strategy for correcting the focus is the same as for pointing: during test time we will perform focus observations at each band with and without the solar filter, and the focus offset will be applied whenever we are observing with the filter in.

Software requirement: the telescope control software must be able to take the focus offset in and out as the solar filter is inserted and taken out.

High frequency warning: we won't have enough SNR to perform this calibration in the sub-millimeter. Presumably it is a dely-like effect and the filter is non-dispersive.

2.3 Amplitude Calibration

The solar filter basically introduces an extra complex multiplicative gain into the calibration equation. It is anticipated that the amplitude gain (ie, the solar filter's 13-16 dB reduction in amplitude) will be more or less constant, though provisions can be made for calibrating the gain if we need to.

Either the telescope control software or the offline reduction software must be able to insert the complex gains for the solar filters when and only when the solar filters themselves are inserted.

Software requirement: Since there is the possibility that the amplitude or phase of the filters could change with time, it might be a better idea to insert the gains in the offline system where we can modify the gains with calibration observations if required.

2.4 Phase Calibration

The solar filters give us two problems for phase calibration. First, they introduce their own delay which must be calibrated, and second, they interfere with the regular ways of performing phase calibration - WVR and fast switching.

The solar filters will render any fast switching phase calibrator too weak to detect. WVR cannot be used while the solar filter is in, so we have to switch the filters in and out as quickly as possible. We have two options: we can switch the filter in for solar observing and out again for WVR observations and interpolate the water vapor signals - or we can switch the filter in for solar observing and out again for fast switching phase observations. As WVR observations would require some sort of fast switching anyway (ie, once every 1-5 minutes), and since the fast switching time scale is comparable to the time required to move the filter in or out, it makes sense to ignore WVR all together.

There will be a significant delay (hopefully a non-dispersive delay) which the filter will add when it is moved into the beam resulting in phase errors and phase gradients across the band, so we will need to calibrate each filter's delay. We will measure the filter's delay in the lab, but we will also check for filter phase (or more accurately, delay) fluctuations by observing the brightest quasar in the sky, both with and without the filter in. The atmosphere will change during the time it takes the filter to move in or out, and we will have to do this observation over and over for 30 minutes or so, and average to get the delay. Furthermore, in order to extract the delay, rather than just the phase, we will want to observe this at different frequencies.

Hopefully the delay will be independent of frequency and will also be constant with time, but these assumptions will be checked observationally. If the delay is constant, we can just lookup the delays for each solar filter. If not, then we will have to measure the delay during the solar observations.

Software requirement: Either the telescope control software or the offline reduction software must be able to add the solar filter's delay while the solar filter is inserted, and must take it out when we take the solar filter out. Furthermore, it must be easy to update the table holding the solar filter's delay.

2.5 Bandpass Effects????

It is possible that we could get reflections within the solar filter, in which case we'll have standing waves and such. Standing waves would result in frequency-dependent gains, or a non-trival bandpass curve associated with each solar filter. A first thing to do would be to measure the bandpass shape of the solar filter in the lab.

2.6 Effects on the Primary Beam

The solar filters will not be perfectly homogeneous and may have an effect on the overall primary beam during solar observations. As many solar observations will need to be mosaiced, we will need to have a good idea of the primary beam shape.

Astronomical measurement of the beam should work on 3C273 at frequencies of 300 GHz or lower, but at higher frequencies we may be in trouble with SNR.

We may be able to perform a simple total power measurement - observing the limb of the moon or the sun with the filter in, with different sky offsets. Starting with a model for the primary beam that has been derived from measurements without the filter, we can predict what total power we should see from the lunar or solar limb as a function of sky offset. If these measurements agree well with the prediction, this indicates that the solar filters do not affect the primary beam very much. If the solar filter affected the primary beam in a symmetric way, such observations could even help provide a new model for the primary beam.

2.7 Effects on Polarization and the Polarization Beam

Typically for determining the on-axis polarization propoerties, we need SNR of 100:1 to 1000:1 or greater. As this is difficult to achieve astronomically through the filter, we may have the best luck measuring the polarization properties of the filter in the lab. We should have sufficient SNR to measure the on axis polarization properties on 3C273 at frequencies below 300 GHz. Between 300 and 500 GHz, we may struggle for enough SNR for on-axis polarization calibrations, but band 9 and 10 polarization observations through the solar filters might have to make do with lab measurements.

For the effects of the solar filters on wide-field polarization measurements of the sun, we are really running out SNR - we will have the SNR problems we see for the on-axis polrization, but then we multiply down by the primary beam, so SNR will be down by a factor of 2 - 10. Again, lab measurements may be our best bet, and should reveal if the solar filter's properties are uniform enough to not affect the polarization beam.

3 Calibration Timelines

3.1 A Best Case Example Timeline for Solar Observations

Lets assume that none of the parameters associated with the solar filters change significantly from one test session to the next. In this case, the parameters which we measured at the last test session are simply applied to the solar observations. At the test session we measure the focus offset, the pointing offset, the gain amplitude and the delay of the filter.

Before observations: We need to pick a good fast switching calibrator that is near the sun. It will be about 0.1 Jy and will be located a few degrees from the sun, We will detect the phase on the fast switching calibrator with sufficient SNR in about 1 s.

Verify that the phase conditions are well-matched to the observing frequency.

Begin observations

- Slew to a bright calibrator
- Focus observation on nearby bright quasar (0.1 - 1 Jy, 10 deg away) with the filter out to get the native focus. A predetermined focus offset will be added to this native focus whenever we observe with the filter in. This will take less than a minute.
- Pointing observation on nearby quasar (0.1 - 1 Jy, 10 deg away) with filter out - for filter-in observations, we apply a predetermined pointing offset to this pointing solution. This observation takes less than a minute, and will be done on the focus source.
- Amplitude Calibration - perform amplitude calibration with the solar filter out using whatever technique or cal source is most accurate. This could take a few minutes, depending upon the source and frequency.
- Ongoing: Opacity will be monitored to determine if opacity fluctuations need to be considered in the amplitude calibration.
- While solar observations are generally continuum in nature, the ALMA correlator is spectral line in nature, and we'll need to do some sort of bandpass calibration to weight the spectral channels appropriately. We hope to perform the bandpass calibration without the solar filter, and apply the filter bandass as another frequency-dependent gain when it is in. We probably want a very bright source for bandpass calibration.
- Perform Several Fast Switching Cycles
  - Slew to the Fast Switching Phase Calibrator and observe the calibrator with the filter out. Detect phase with sufficient SNR. This will take of order 1 s. WVR data may be taken while off the sun, but it may or may not be useful.
  - While slewing back to the sun, we insert the solar filter, and the online system knows that all filter offsets (ie, focus, pointing, and if we do it online, amp and delay) need to be applied. Slewing will take about 2 seconds, and can be matched to the time scale of the solar filter insertion.
  - Observe the sun with the solar filters in for about 20-30 s. The length of time will be determined by the atmospheric coherence. The sun will be bright enough to use for

self-calibration, but the source structure of the sun may often be too complex to use self-calibration. If experience indicates that self-calibration can be used routinely, the time on sun can be increased, perhaps to 30-60 s or longer. This is good, as solar observers would like to get as much continuous time on the sun as possible.

- When the target phase of the fast switching cycle is done, remove the solar filter as we slew to the fast switching calibrator source to finish this fast switching cycle or to start a new one.
- Note: there is nothing stopping us from using WVR data while we observe the fast switching calibrator (ie, while the solar filter is out of the optical path) - though there appears to be no advantage to doing so - the idea behind the WVR is that it will track the short time-scale atmospheric fluctuations while on the traget source, but we will not be cable to use the WVR while the solar filter is in or while observing the sun.

We have left out the instrumental or cross-band calibration phase which is required if we observe the calibrator at a lower frequency than target solar observations.

An obvious drawback of the fast switching strategy is that the solar observer has only short continuous time series of the sun, interupted by 5 or more seconds of calibration ever 20-60 seconds.

3.2 A Mild Worst Case Example Timeline for Solar Observations

In the worst case, some or all of the offsets (ie, focus offset, pointing offset, gain amplitude, and delay) are not constant over many months or years so that earlier service or test observations of these quantities are not suficient to meet the scientific demands of the observations. A "mild worst case" would be where the varying parameters need be observed just once during the observations. In a "worst worst case", the varying parameters would need to be calibrated repeatedly during the solar observations. This would probably make solar observing impossible, so we should put specs on the time scales of the variability of any focus offset, pointing offset, amplitude gain, or delay associated with the solar filter. This, in turn, can be translated into specs on the precision of reinsertion of the solar filter.

bf Before observations: We need to pick a good fast switching calibrator that is near the sun. It will be about 0.1 Jy and will be located a few degrees from the sun, though it may need to be further away - can we observe a few degrees from the sun with the solar filter out? We will detect the phase on the fast switching calibrator with sufficient SNR in about 1 s.

Verify that the phase conditions are well-matched to the observing frequency.

Begin observations:

- Slew to a very bright calibrator
- Focus observation on the brightest quasar in the sky (1-10 Jy, 100+ deg away) first with the filter out, and then with the filter in (this is why we need the brightest quasar) to measure the focus offset. If the focus varies with elevation or angle to the sun, we will need to focus again on a bright (0.1- 1 Jy) quasar close to the sun with the filter out, and apply the focus offset we just measured. It is assumed that this will be sufficient for the entire observations.

- Pointing observation on the brightest quasar (1-10 Jy, 100+ deg away) with filter out and then with the filter in, to solve for the pointing offset due to the filter. This calibration source will be the same one we did the focus observations on. We will need to find a bright source close to the sun (0,1 - 1 Jy) for pointing calibration with the filters out, and apply the newly found pointing offsets.

- Amplitude Calibration - perform the standard amplitude calibration, plus observing the brightest quasar in the sky, determine the amplitude gain of the solar filter by observing the bright quasar with and without the solar filter.

- While solving for the amplitude gain of the solar filter, we can also solve for the delay of the filter on that same brightest quasar.

- Solar Bandpass - yeah, yeah, yeah.

- Ongoing: Opacity will be monitored to determine if opacity fluctuations need to be considered in the amplitude calibration.

- For the fast switching cycle, apply the current focus, pointing, amplitude, phase, and bandpass offsets due to the solar filter as required.

- Perform several fast switching cycles

  - Slew to the Fast Switching Phase Calibrator and observe the calibrator with the filter out. Detect phase with sufficient SNR. This will take of order 1 s. WVR data may be taken while off the sun, but it may or may not be useful.
  - While slewing back to the sun, we insert the solar filter, and the online system knows that all filter offsets (ie, focus, pointing, and if we do it online, amp and delay) need to be applied. Slewing will take about 2 seconds, and can be matched to the time scale of the solar filter insertion.
  - Observe the sun with the solar filters in for about 20-30 s, or longer if possible.
  - When the target phase of the fast switching cycle is done, remove the solar filter as we slew to the fast switching calibrator source to finish this fast switching cycle or to start a new one.

3.3 A Note on Solar Self-calibration

One can imagine clever algrothms that take into consideration that the residual phase errors will be small at the beginning and at the end of each fast switching target-source phase on the sun. Hence, we will have a high quality image to use as a model for the first round of self-calibration for integrations near the beginning and near the end. However, at some point in the middle of the target-source phase of the fast switching cycle, the residual phase errors may be rather large, and it is imaginable that the previous integration time's model image could be a better starting model than the image made from the current time's data, even considering solar variations. In this way, one can imagine bootstrapping a model image across the integrations to the middle of the switching cycle.

4 Imaging

4.1 A Note on Total Power Observations

As the sun is much larger than the primary beam and we will need total power for successful imaging of many types of solar observations. This provides us with some logistical problems. First of all, we will need to calibrate the total power and ACA antennas as well as the rest of the 12m ALMA antennas, which probably requires that we cross-correlate the ACA and total power dishes with the rest of ALMA to get sufficient SNR. Second, as solar emission is highly variable, we will need to observe in total power at the same time as we observe the sun interferometrically. Furthermore, the fact that the sun is changing while we are observing it will lead to errors in the deconvolved image. We can minimize these by zipping over the region of interest as quickly as possible - ie, attempting to cover the entire region of interest in less time than it takes for the solar emission to change appreciably. On-the-Fly mosaicing will speed up the imaging process significantly.

Another consideration which solar variability leads us to is that the total power and ACA should observe each field at the same time as the rest of ALMA does. We will still have some variability issues from one field to another, but at least the short spacing data for each field will be consistent with the interferometric data for that field. The ACA's larger primary beam permits the ACA to sample the sky with a coarser grid of pointings, but the tighter grid should be used to minimize the effects of variability.

For compact array observations, we will have excellent (u,v) coverage in a single snapshot, so we will be able to make high quality mosaic observations of the sun in this manner. However, for larger configurations, we will run into two issues: first, we may not get sufficient (u,v) coverage in a single snapshot, and second, the time scale for variations in the solar structure at this higher resolution will be smaller. So, things will be good for short baseline observations but not so good for long baseline observations.

4.2 Notes on Variability and Deconvolution

This could become an active area of research if millimeter wavelength solar imaging becomes important. One can imagine imaging algorithms that image two channels at once: an average or queiscent solar image plus an image which takes all of the fluctuations. If we had zero sidelobes, we would make the average solar image automatically, so the effects of the variability only degrade the image through the spreading of emission through sidelobes. So, if the variable emission is somehow modelled, we can remove it from the visibilities and get that average or static solar image, which permits us to integrate for a long time to get good (u,v) coverage. Or, a poor-man-solution would be to reweight the visibilities to reduce the sidelobe level and just image all the visibilities without worrying about variability. And last, if variability is what the astronomer is most interested in, they may want to form that quiescent or average solar image, subtract that from the visibilities, and then image the residual visibilities into a time-cube to catch the fluctuations, some of which will show up negative in this method unless a re-baselining of some sort is performed, trading some of the quiescent emission for variable emission.

5 What Might Help the SNR Problems at High Frequencies?

We will have marginal or insufficient SNR for calibrating the higher frequency observations as seen through the solar filter. When SNR becomes an issue, we have some choices:

- **Pray** - there is a significant chance that everything will work out nicely and all of the effects of the solar filter will be constant and we will seldom need to observe calibration sources through the solar filter. If this is so, some of the characteristics of the solar filters as measured in the lab could be used during observations years later. However, the higher the observing frequency, the higher the chance that the delay fluctuations will be large enough that we need measure them.

- **Wait** - if the filter parameters are stable, we only need to solve for them once, and we can just integrate for a long time (ie, during one or more test time sessions) to get the required SNR.

- **Ignore** - solar scientists say they will be interested in observing the sun at all frequencies. However, at the higher frequencies, it seems that the resolution of ALMA may exceed the intrinsic angular size scale of structure on the sun. Certainly at higher frequencies we will have trouble mapping a very large region on the sun due to the decreased primary beam size. There may be a critial frequency above which solar astronomers are not so interested.

- **Hedge** - the situation could be improved a bit by having a filter that is more like 13 dB than 16 dB - or even a filter which becomes more transparent at the higher frequencies.

- **Skew** - we can get much higher SNR on the solutions by observing with some antennas with the filters in and some with the filters out. I think the optimal thing to do is to have half the antennas with filters in and solve for the solar filter's gain, then take the solar filters out on those antennas while inserting it for the other antennas and solveing for the solar filters' gains again on the second half of the array. This will improve the SNR by something like a factor of 3 as compared to solving for the solar filter gains all at once, which will just push up the frequency at which we run out of SNR. This strategy would not be so good for fast switching as it will take enough extra time that significant atmosphere blows over during this calibration, but we don't need to solve for the filter gains very often and we would not want to use this during fast switching - rather we'd want to use it during test or service observations to check on the focus or pointing to make sure that the solar filter offsets have not varied since the last time they were measured.

- **Cool it** - the SNR of observations with the filter in are decreased due to loss of signal as well as increase of noise temperature from the ambient temperature load. If the filter were cooled, it would still work well for reducing the solar flux, but would not decrease the SNR so much. However, the place we really need to increase the SNR is at the highest frequencies, where the system temperature is already much higher than the ambient temperature and the improvement due to cooling the filter will be minimal.

6 Laboratory Measurements

We will want to measure the amplitude gain and the delay of the solar filters at all frequencies of use.

It will be simple to measure the amplitude gain of the filters, but measuring the delay will require some sort of interferometry - comparing the phase of a signal which did not go through the filter with the phase of raditaion which did go through the filter. Presumably someone can design this measurement.

From the measured delay of the filters, we should be able to calculate the focus offset we will have to apply when we observe with the filters in.

We will want to measure some of the polarization properties of the filter in the lab. We can't measure the polarization leakage beam in the lab, but we can get a sense of the filter's contribution to the on-axis polarization leakage with laboratory measurements.

We can't measure the primary beam in the lab, but we can measure the distribution of the amplitude gain and delay across the filter, which presumably can be input into some optics software for an estimate of the effects of the solar filter on the primary beam.

Similarly, it should be possible to measure the distribution of polarization properties across the filter. These measurements could be used as input to optics modelling software to get an idea of the effect on the polarization beam.

7 Solar Manpower

The coordination of the Laboratory Measurements, the coordination of the software aspects of the solar filters, required test observations, and the testing of the solar filter calibration system are significant tasks. We will probably need to have someone who who spends most of their time on the solar filters for a few years to ensure that solar observations can take place and can be calibrated.