ALMA Memo No. 335

Phase Drift Measurements of YIG-Tuned Oscillator Sources for the ALMA LO

Dorsey L. Thacker, Eric W. Bryerton, Richard Bradley, and Kamaljeet Saini
NRAO, Charlottesville, VA 22903
29 June 2001

In this memo, we present measured phase drift data of a YIG-tuned oscillator (YTO) multiplied to mmwave frequencies, a prototype of the baseline ALMA LO source. The test setup and analysis algorithms are described. We also discuss phase drift added by mmwave multipliers and power amplifiers, with the conclusion that their contribution is insignificant. Various components and system concepts are analyzed with respect to their contribution to the phase drift of the array. As phase drift at the level required for ALMA is unique and not well understood, much of the discussion will be of a tutorial nature.

1. Introduction

This memo describes phase drift measurements of a prototype version of the YTO-based driver for the ALMA LO. We report on two prototype driver chains, one at 76 GHz and the other at 108 GHz, consisting of a YIG Tuned Oscillator (YTO) followed by multipliers and power amplifiers. This is a companion memo to [1] E. Bryerton, D. Thacker, K. Saini, and R. Bradley, "Noise Measurements of YIG-tuned oscillator sources for the ALMA LO," ALMA Memo # 311. Please see [2] and [3] for further discussion on phase noise.

Section 2 begins by giving some useful definitions for the characterization of phase drift and discusses the specifications needed for ALMA. The phase drift measurement setup is described and measurements of phase drift are discussed. Section 3 is a general discussion of various components and system concepts with respect to their phase drifts.

2. Phase Drift

It is helpful to separate the concept of phase drift, which is long time-scale (>1 s) fluctuations caused by effects such as thermal drift, from phase noise, which is short time scale (<1 s) zero-mean fluctuations. Phase drift measurements are described in this memo and are given as degrees of phase at a particular frequency or in terms of microns (free space). In addition it is important to have a clear understanding of the time interval and temperature range over which the drift is measured as well as knowing what statistical measure is used to generate the number that is used to characterize the drift.

In order to assess the final effect of the drift measured at some intermediate stage of the system, it is important to distinguish between the two cases of multiplication and mixing. When a signal is multiplied by N , the phase drift of the output is N times the drift of the input. When two frequencies are combined in a mixer the phase drift of the desired output, either sum or difference frequency, is assumed to be sum of the drifts on the input signals in degree measure since we have no a priori knowledge of the sign of the drifts.

2.1 Specifications for Phase Drift

The goal presented in [4] for the phase drift associated with the electronics on time scales greater than one second is 6.9 fs, which translates to 2.1 μm or 2.4 degrees of phase variation at 950 GHz. A less stringent goal of 10 degrees rms at 950 GHz (or 8.8 μm) is given by [5], which we adopt as a tentative specification. A convenient timescale to characterize drifts is 10 minutes, which is long enough to be of practical use for astronomy yet short enough to permit useful engineering measurements. **The goal of 2.4 degrees at 950 GHz is interpreted here to represent the standard deviation of the difference between the true phase of the electronics and an estimate made by linearly interpolating between two calibrations taken 10 minutes apart. The standard deviation of this difference should be less than 2.4 degrees at 950 GHz. This interpretation of phase drift is used throughout this paper.** While the choice of 10 minutes is somewhat arbitrary, preliminary data that will be discussed in section 2.3 reveals that the drifts in LO sources do not change rapidly with the choice of this measurement time scale. At this level of stability the need for phase calibration will be determined by the instability in the atmosphere. (See [6], [7], [8] and [13])

There are several components in the LO system that have appreciable drift on a 10 minute time scale. Each antenna has about a dozen separate quasi independent components with the drifts assumed to add in an rms sense. The drift rates associated with these components are essentially independent in that even when excited by a common change in temperature some units will have a positive temperature coefficient and some a negative temperature coefficient as well as different time constants. The total system phase drift will be less than the sum of the individual drifts, and assuming complete independence of the drifts, equal to the rms sum of the drifts. A working specification for each of approximately ten components is taken to be a factor of three times less than the total allowed phase drift, therefore each component should only cause on the order of 0.8 micron drift (standard deviation of difference between true and estimated over a 10 minute interval). Table 1 shows these drifts allocated equally between these components. It is recognized that the actual drift between antennas may be slightly better than this because each antenna will have identical LO chains and there will be some drifts that will be common to each antenna pair.

We begin counting phase drifts at the moment the common signals diverge in the main control room and continue to the digital sampler. Drift and jitter in the sampler clock is outside the scope of this memo.

Assigning the phase drifts equally among the items in table 1without considering the complexity and cost to implement these drift specifications is only a starting point and adjustments will need to be made. In section 3 we will discuss two items that are likely to be more difficult than first appreciated, the fringe rate synthesizer (aka the direct digital synthesizer, DDS) which provides the offset frequency for the LO Source lock loop and the “delta L * delta F” component of the round trip corrector.

Table 1 Phase Drift Allocation for LO Electronics
Assuming equal contributions to the phase drift for each component at 950 GHz (2.3 fs, 0.7 µm or 0.8 degrees at 950 GHz).

Component	**Drift (degrees)**	**Drift (microns, (fSec) at 950 GHz**
Optical power splitter	0.1 @ 119 GHz	0.7 micron, 2.3 fSec
Round trip corrector and fiber random deltaL*deltaF	 0.1 @ 119 GHz 0.1 @ 119 GHz	0.7 micron, 2.3 fSec
Jumper Cable to Photomixer (uncorrected fiber)	0.1 @ 119 GHz	0.7 micron, 2.3 fSec
Photomixer	0.1 @ 119 GHz	0.7 micron, 2.3 fSec
Lock Loop	0.1 @ 119 GHz	0.7 micron, 2.3 fSec
Fringe Rate Synthesizer (DDS)	0.1 @ 31.25 MHz	0.7 micron, 2.3 fSec
125-MHz Source	0.1 @ 125 MHz	0.7 micron, 2.3 fSec
Cold Multiplier #1	0.2 @ 237 GHz	0.7 micron, 2.3 fSec
Cold Multiplier #2	.39 @ 475 GHz	0.7 micron, 2.3 fSec
Cold Multiplier #3	0.8 @ 950 GHz	0.7 micron, 2.3 fSec
IF amplifiers	0.8 @ 12 GHz	0.7 micron, 2.3 fSec
Second LO	0.8 @ 8 GHz	0.7 micron, 2.3 fSec

2.2 Phase Lock Loop and Measurement Setup

The HP E5500 Phase Noise Measurement System was configured for drift measurements. The phase noise test set as described in ALMA Memo 311 was modified by replacing the 103.65 MHz oscillator with a 140 MHz oscillator of the same general type but phase locked to the system master 10 MHz standard. The setup used to take measurements is shown in Figure 1. A 600-MHz reference signal, provided by the test set, is used both to down-convert the W-Band output to the 140 MHz which is processed by the E5500 and as the mm-wave reference for the driver PLL which also has a 140 MHz IF. Since this signal is common to both the test down-converter and the reference to the W-Band lock loop, its drift cancels in the phase detector. This architecture parallels that of the photonic reference which also has a common reference oscillator. Note that the 140 MHz from HP8663A used with the E5500 and the 140 MHz crystal reference oscillator used for the driver PLL are separate; therefore, their drifts combine in the phase detector instead of canceling. The voltage output of the phase detector in the E5500 test set

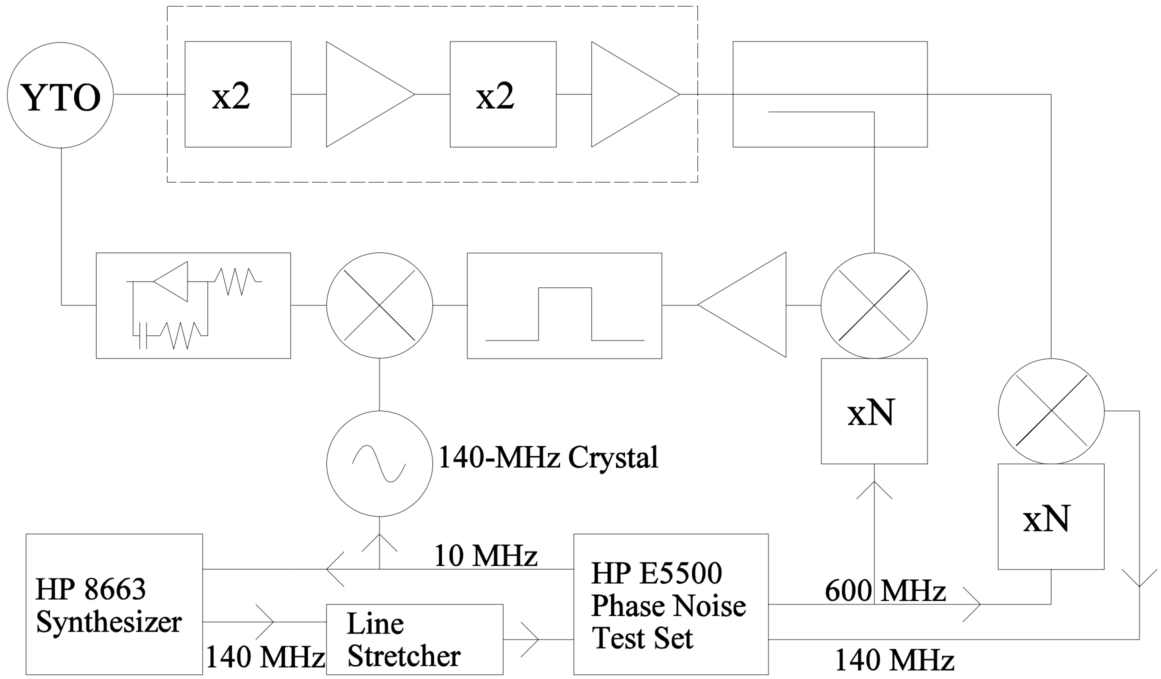

Figure 1 Block Diagram of Phase Drift Test Set

was set to zero by adjusting a calibrated mechanical line stretcher in the 140 MHz path. The stretcher was calibrated with a section of precision airline of known length. This output of the phase noise test set was processed in two different ways: 1) using its internal data processing of the E5500 (figures 4a and 4b) and 2) with the phase difference sampled at 10 second intervals, recorded to disk and post processed in a spread sheet (figures 2,3, 5, 6, 7, and 8.) This system is fully coherent with all oscillators locked to the internal crystal standard of the HP E5500 test set.

2.3 Phase Drift Measurements

The HP E5500 works well measuring down to 10 milliHertz (100 Seconds) as can be seen from the sample plot shown in figures 4a and 4b, but for longer drift measurements we sample the output of the HP E5500 with an A/D running under LabView and process the data in a spreadsheet.

Figure 2 shows the raw data from a long (40 hour) run of a complete system at 76 GHz measured by the E5500 using the common reference generator from the E5500. The 76 GHz source was a YTO locked at 19 GHz followed by a doubler, power amplifier, another doubler and a final power amplifier. Other than the fact that the multipliers and power amplifiers were outside the lock loop, such a system is very close to the architecture of the photonic reference scheme and represents typical phase drifts of the ALMA LO source. Shown in red is the air temperature measured just above the base plate upon which were mounted the components of the LO source. The temperature varied over a 13 degree Celsius range and is a much larger temperature variation than we expect in the receiver cabin. The rapid increase that begins after two and a half hours into the experiment is caused by the building air conditioners shutting off at 6 pm Friday evening.

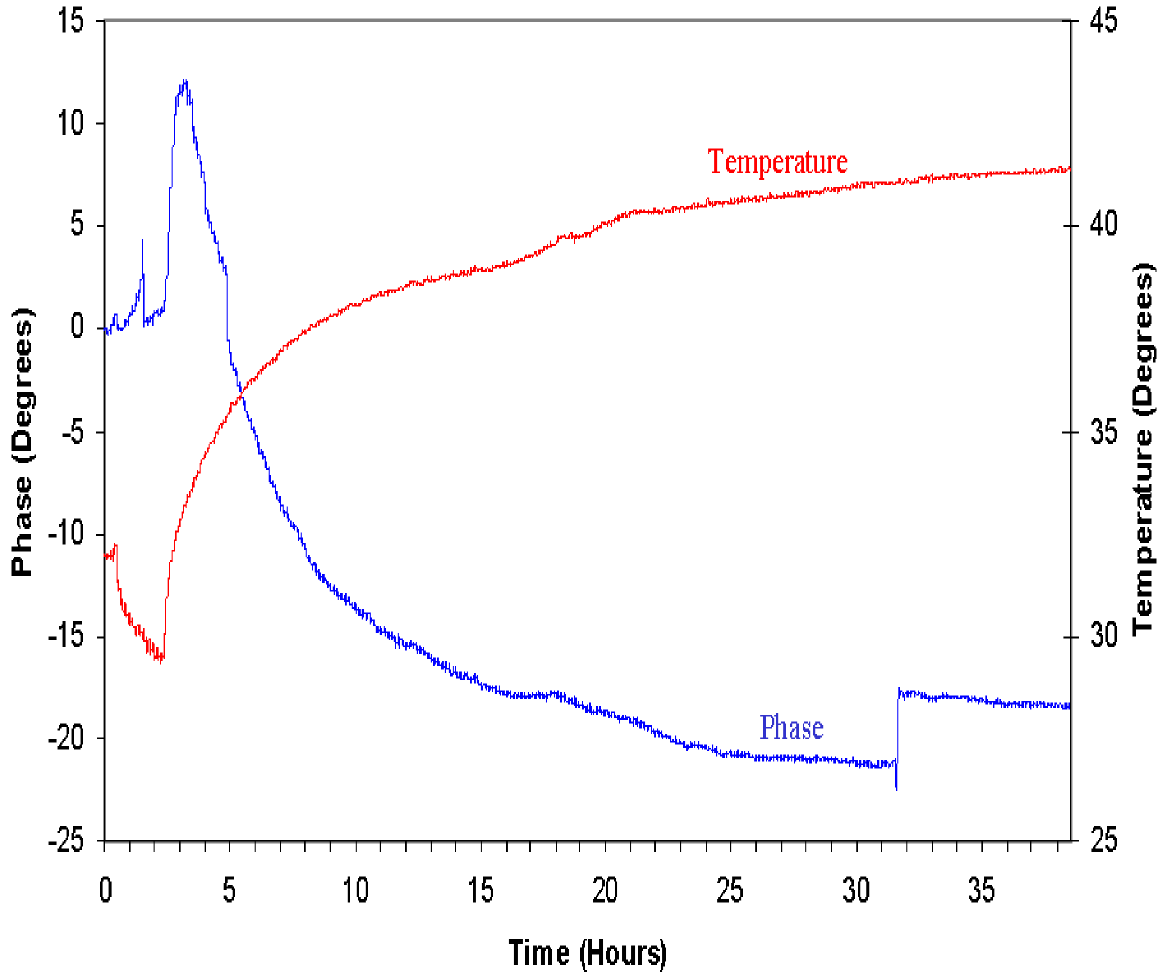

Figure 2 Raw Data for 76 GHz Source Locked at 19 GHz

Figure 3 shows the standard deviation from an estimate of the phase made by taking a linear fit between points 10 minutes apart. Examination of this plot shows long periods of time where the phase error is less than 0.2 degrees of phase at 76 GHz. The goal for individual component (see table 1) is only 0.1 degree at 119 GHz and here we have the entire system including two multipliers and two power amps that are outside the lock loop with less than 0.2 degrees. This is an extremely encouraging result. There are three unexplained jumps in phase at 1.5 hours, 5 hours and 31 hours. Other data sets have had very frequent jumps in phase of this magnitude which were traced to a cracked capacitor on the lock board. However, in this case these events are likely to be external interference or perhaps spontaneous relaxing of mechanical connectors or coaxial cables. These jumps are approximately 0.001 inch path length changes.

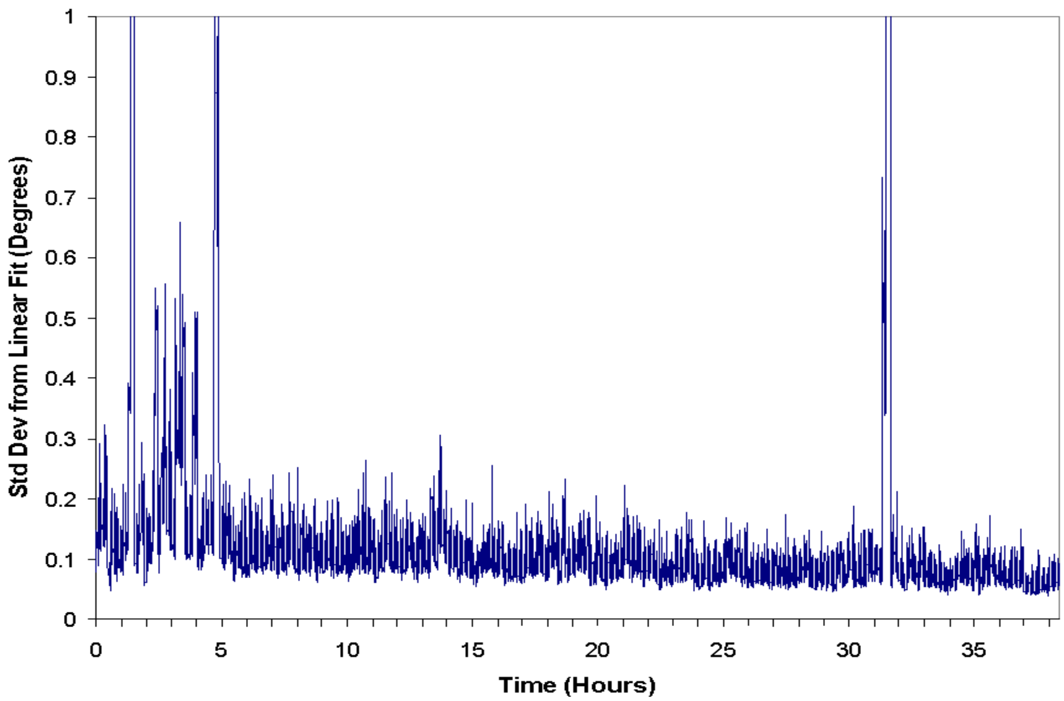

Figure 3 Standard Deviation from 10 min fits

In addition to the 76 GHz driver data shown in figures 2 and 3 we have also tested a driver at 108.14 GHz. Figures 4a and 4b show the performance of this driver as measured by the E5500 test set down to a frequency of 2 milliHz (500 seconds or 8.3 minutes.) This driver was configured with the phase lock loop closed at 108 GHz just before the final amplifier. The "birdie" at 1.5 MHz in figure 4a is a reminder of how sensitive to EMI these loops are. Figure 4b shows the Allan variance of the 108.14 GHz source. The points marked with X's were taken by the E5500 measurement system and processed by the E5500 software from the data set of figure 4a. Points that are circles were added from the spread sheet data shown in figure 6. The deviation from a linear estimate of phase as shown in figure 7 is a more appropriate statistical measure for the drift performance of an interferometer than the more familiar Allan variance and we show the Allan variance for comparison purposes and for the physical insight that it can give to the noise processes involved.

The final amplifier was outside of the PLL and mounted on a heat sink whose temperature we could vary on time scales of a few minutes. The amplifier was isolated from the rest of the driver and measurement system by two stainless steel waveguide sections. The phase data versus time for the power amplifier as the amplifier temperature was cycled over a 10 degree range is shown in figure 5. The measurements indicate a temperature coefficient for this amplifier of approximately 0.4 degrees of phase per degree centigrade. Note that this is less than a half meter long section of typical uncompensated fiber. A typical uncompensated fiber will have 7*10e-6/C temperature coefficient [9]. All of the mm-wave power amplifiers used in the ALMA LO will be of the same general topology except for the 100-120 GHz amplifier which will be InP rather than GaAs. A similar measurement will be performed with the InP type when it is available, but it is not expected to differ significantly.

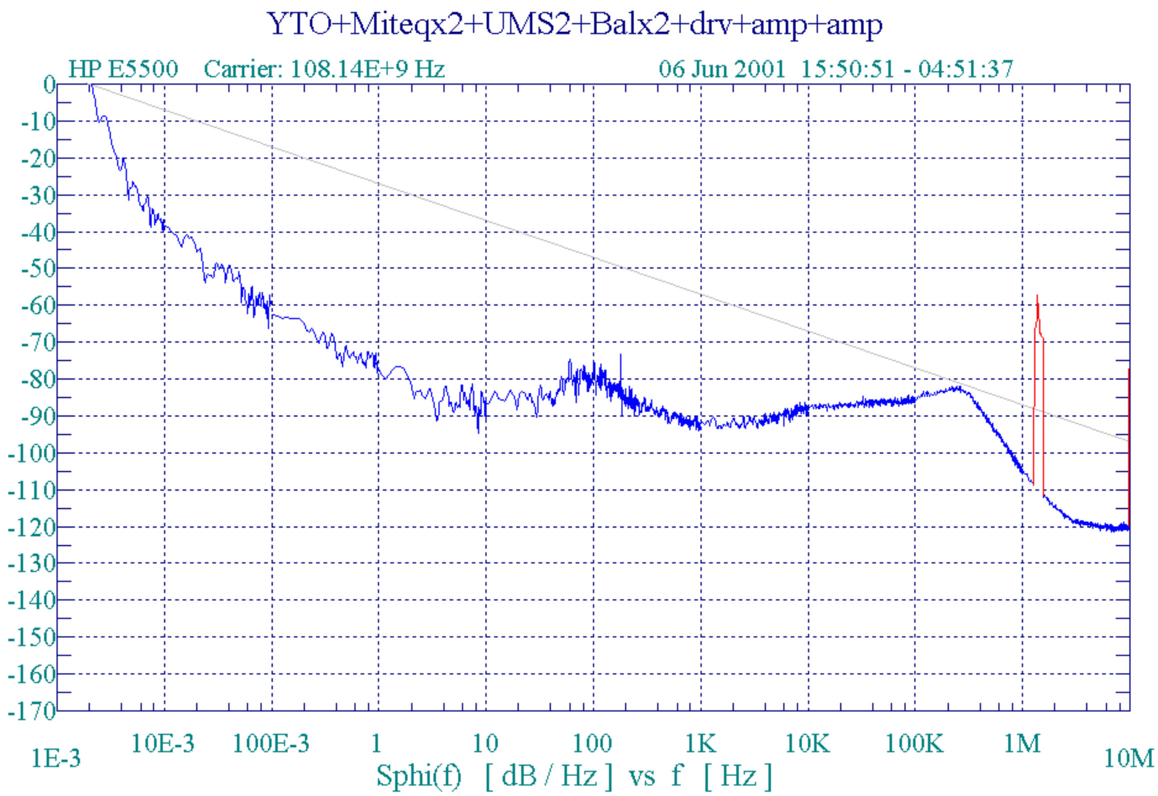

Figure 4a Phase Noise to 1 milliHz of System Phase Locked at 108.14 GHz with the final power amplifier outside the PLL

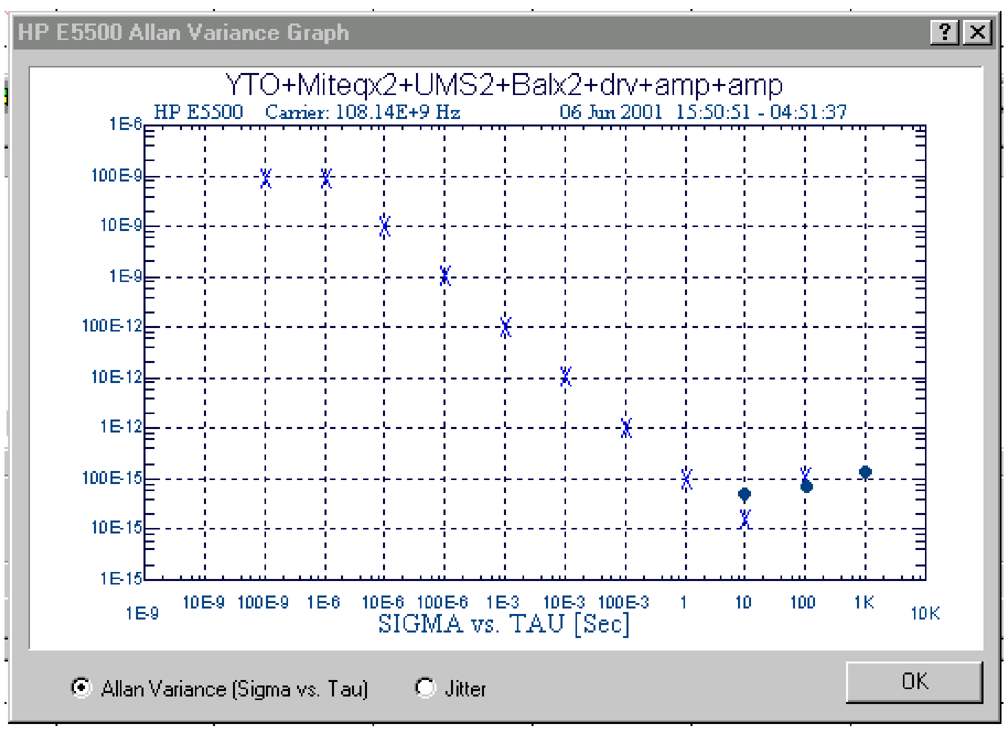

Figure 4b Allan Variance of data in figure 4a

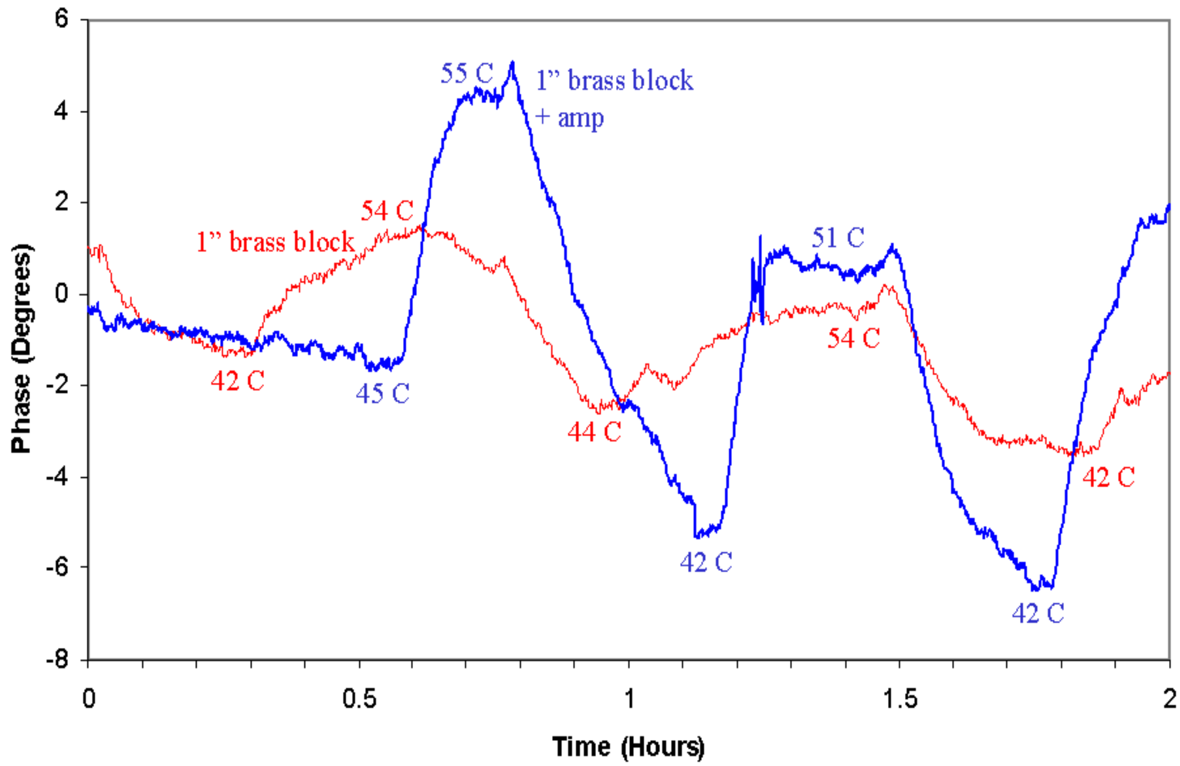

Figure 5 Phase of Amplifier #4 for different temperatures ==> 0.4 degree phase/degree C at 108 GHz

Figures 6 and 7 are data taken with the LabView phase drift system and are similar to the previous figures 2 and 3. In addition to the phase drift, in figure 6 we plot the correction voltage to the YTO. In the PLL for this oscillator we have placed a nonlinear element to reduce the DC gain of the loop for correction voltages that are not near zero volts in order to simplify lock acquisition. We clearly see the phase drift increases for correction voltages more negative than 1.5 volts where this nonlinear element has reduced the DC loop gain. This increase in phase drift is an artifact of this particular test setup and is not a part of the final design in which lock acquisition will be handled by firmware. The step changes in correction voltage at 30 minutes and 10 hours were inserted by a the coarse tuning algorithm in the LabView system in order to keep the correction voltage centered.

Early in this project we chose 10 minutes as our baseline time for characterizing drifts. In figure 8 we plot the deviation from a linear model for 10 minutes, 30 minutes and 1 hour. For a substantial part of the 20 hour data set, acceptable performance is obtained with 60 minutes fit, degrading only about a factor of two. For times where there were large excursions of temperature in the time frame of an hour the performance degraded significantly; see for example, the period of time between three and five hours of figure 8. *Indeed, phase drifts of the electronics on times scales from 5 minutes to an hour are within a factor of two of the 10 minute data for periods when the room temperature does not significantly change during the measurement.*

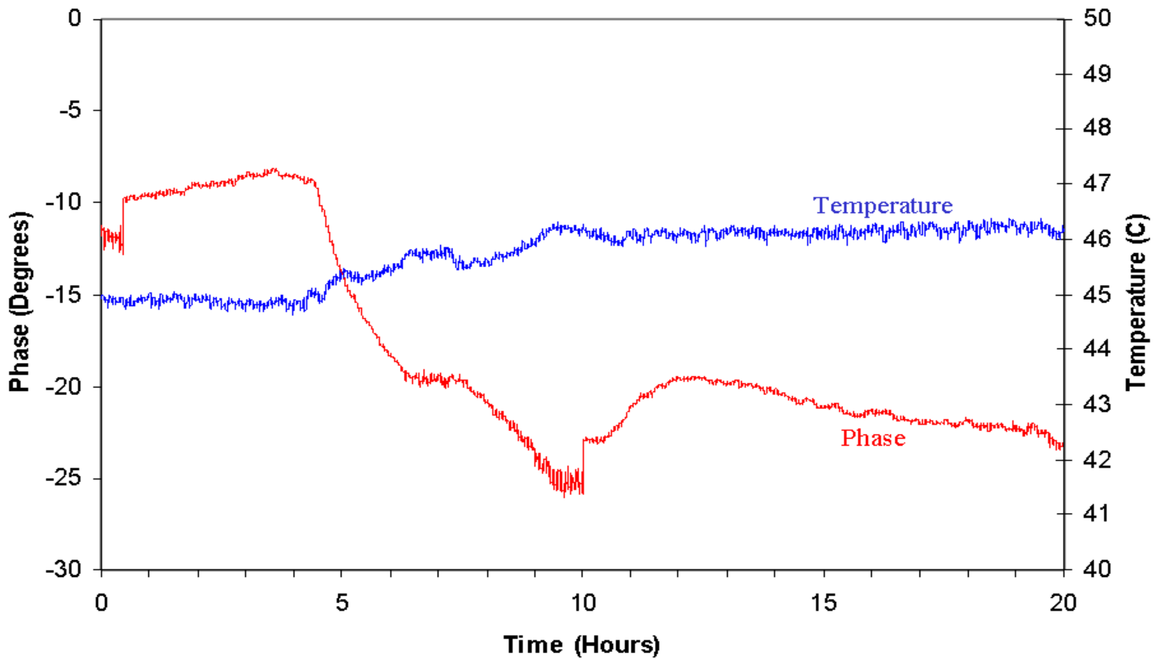

Figure 6 Raw phase data for 108 GHz Driver and Physical Temperature vs time in Hours

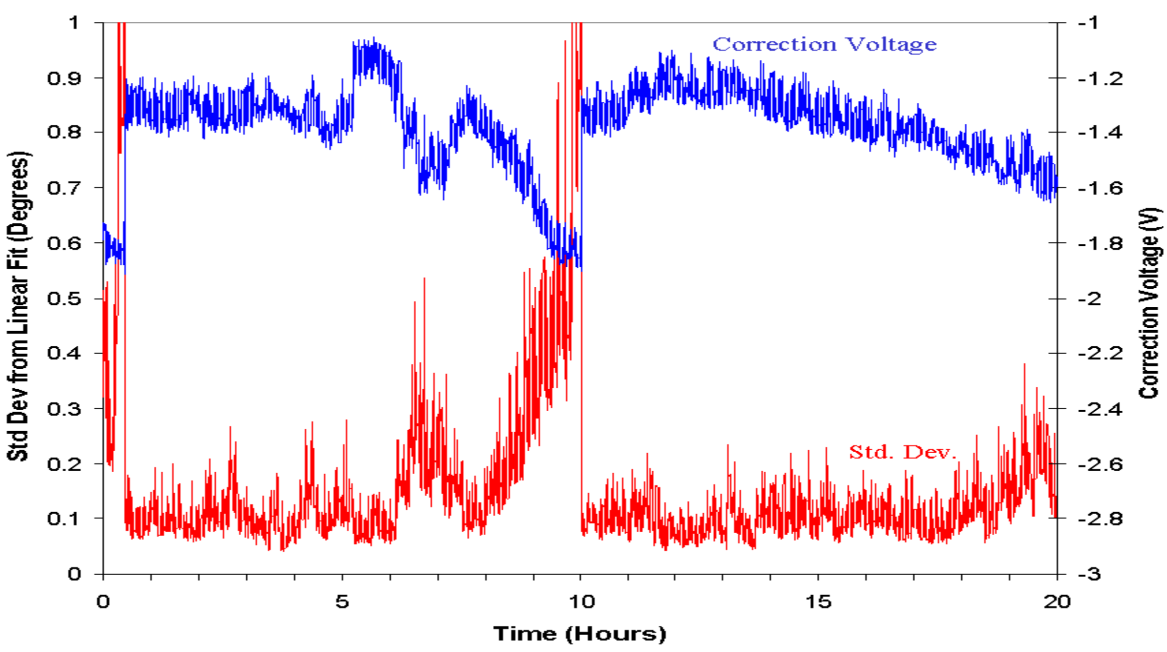

Figure 7 Standard Deviation from linear fit for 10 minutes and Correction Voltage

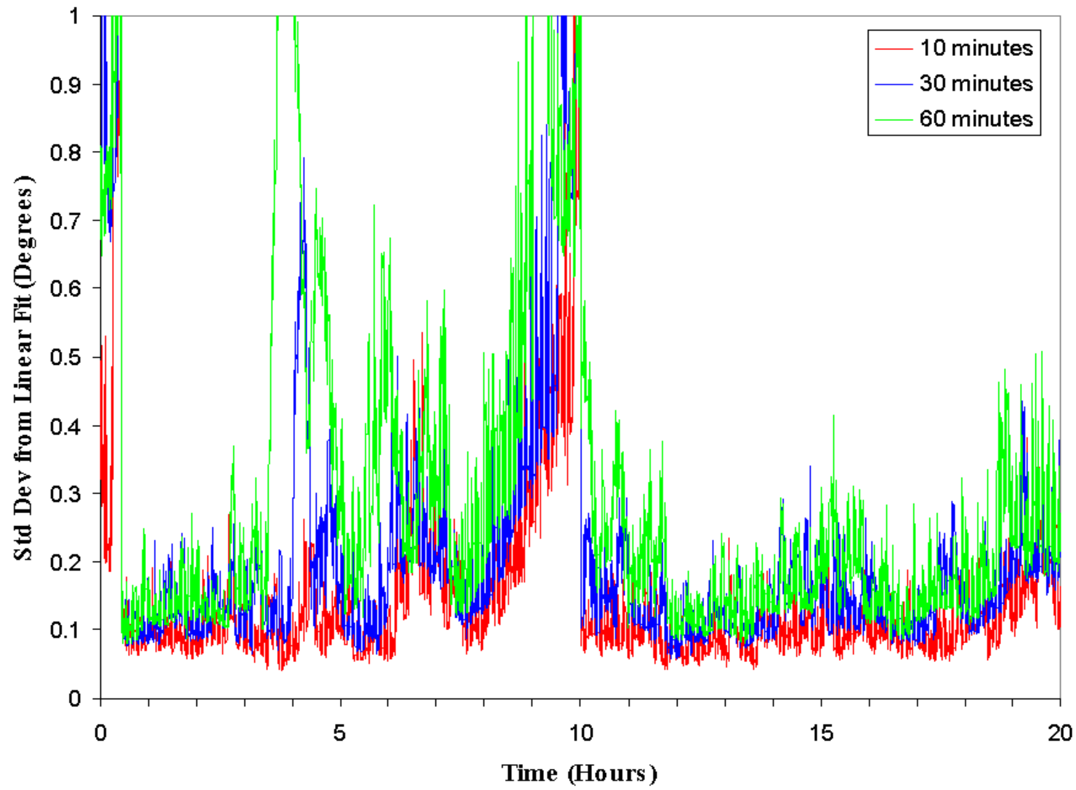

Figure 8 Standard deviation from linear fit for 10 minutes, 30 minutes and 1 hour

3. Potential Problem Areas

3.1 Fringe Rate Synthesizer

Even though the output frequency of the Fringe Rate Synthesizer is relatively low, in the range of 20 to 40 MHz (nominal value 31.25 MHz), the proposed use of direct digital synthesis (DDS) with digital logic delays that are a function of temperature for the Fringe Rate Synthesizer is likely to be a problem. The specification for the DDS in terms of digital logic delay times is that the propagation delay from the master timing signal for the DDS to its output stage must be constant to 7.1 picoSeconds. The data sheet for a typical DDS chip (Analog Devices 9852) is silent with respect to long term change in this propagation time but it does specify short term jitter as 12 picoSeconds rms. Expect the long term drift to be significantly worse than 12 picoSeconds. Note the 12 picoSeconds short term jitter is acceptable in that it is a small part of the phase noise error budget for this device. In this example, the phase noise specification is easily met while the phase drift is likely to be more difficult. Obviously, the phase vs temperature of the DDS chip that is chosen for the fringe rate synthesizer needs to be measured at the earliest possible stage in the prototype design.

The preferred frequency for the LO Source offset is 156.25 MHz which would be generated by mixing the 31.25 MHz signal with the ultra stable 125 MHz reference frequency. Since the drift in the fringe rate synthesizer adds directly to the W-Band output, the fringe rate synthesizer and the 125 MHz reference must be stable to 0.08 degrees of phase at their output frequencies over the temperature range variation in the receiver cabin over the assumed 10 minutes between calibrations. In the laboratory (Refer to Table 1), we have verified that a HP8663A synthesizer (designed around 1982) will meet these specifications at a frequency of 140 MHz and we use this synthesizer as part of our test set to measure phase drift. However, our measurements indicate that the HP8664A synthesizer (designed around 1990) is approximately 30 times worse than this specification *presumably because the newer synthesizer uses modern digital techniques which are not as phase stable as the old fashioned analog techniques*.

3.2 Round Trip Corrector

The second item is the delta L*delta F of the Round Trip Corrector. There are two basic schemes to implement the round trip corrector. The first is to shift (with an acoustic optic Bragg cell) the frequency of the master laser, send it back on the fiber, and servo the length of the fiber to essentially hold the number of wavelengths of the master laser in the round trip to be constant ref [10, 11]. As Bill Shillue has pointed out [10], in order to maintain the actual length of the fiber to one half micron in a 25 km long fiber the wavelength (frequency) of the master laser needs to be much more stable than the existing laser which gives a phase drift of a couple hundred microns per minute in a 25 km fiber [10]. The laser must be stable to a couple of parts in 10^{11} for a difference in fiber lengths of 25 km. {df/df0 = .5e-6/25e3 = 2e-11} This implies that the master laser needs to be locked to an atomic frequency standard if the difference in fiber path lengths is allowed be large. Indeed the frequency drift spec on an unlocked laser places a constraint on the matching of lengths of the fiber runs for the array. A good laser [12] has a drift spec of about 10 MHz per hour for a 194 THz (1550 nanometer) laser for a df/df0 = 5e-8 which implies that matching the cable lengths to a part in 1e4 should be more than adequate.

The phase of the W-band photonic reference signal due to the dispersion of the signal is about 1.2 picoSecond for a 25 km cable. If this is stable to a part in 1e4 or if the cables are matched to this precision then the changes due to dispersion will not be significant.

The obvious engineering solution of locking the master laser to an atomic frequency standard is sufficiently difficult (expensive) that the alternative of a software solution should be investigated. The solution of installing 25 km of buried cable to connect antennas that are only a few meters apart also has a cost impact. A software correction could be implemented by first determining the difference in fiber lengths to the part in 1e4 accuracy either from direct engineering measurements of possibly from astronomical observations and then applying a software correction to the phase based on solving for the frequency of the master laser. This information coupled with the amount of correction added by the fiber length servo could be used also to estimate the dispersion error. While these methods are highly speculative, they represent an opportunity to save a little money and are not necessary for the array to meet specifications since the straight forward approach of matching the fiber lengths to a part in 1e4 is a viable alternative.

The other basic scheme for implementing a round trip corrector is to use a microwave reference frequency that was derived from the hydrogen maser to servo the round trip signal so that the number of wavelengths of this frequency in the round path is held constant. This frequency could be one of the reference frequencies transmitted by AM modulation of the master laser or in the photonic reference scheme it could be the difference frequency between the master laser and the slave laser. For the case of using the 2 GHz reference signal, a half micron is 1.2 millidegrees which is a very precise and difficult loop. We agree with John Payne (John Payne, private communication) that this scheme in not really practical. In the photonic reference design the difference frequency of the two lasers is a W-Band frequency where 0.5 micron is 0.06 degrees of phase and is the same for all antennas in a subarray. While difficult, this is not impossible and has the advantage of servoing the phase of the signal in which you are most interested. This scheme automatically eliminates any dispersion effects in the fiber. Servoing the fiber length to the W-band reference frequency has some undesirable operational consequences. It is not clear how you could do any of the operations that require a change in the W-band reference frequency and still maintain phase coherence. This impacts between band phase calibrations (the so called fast band switching calibration) and frequency switching. Using the difference frequency between the master and slave laser for a round trip corrector would be difficult for the photonic reference B or the direct photonic case because of the fact that this frequency will be different for each antenna due to the fringe rotation that is added to this frequency.

We feel that line length corrector based on holding the number of wavelengths of the master laser constant in a round trip path with equal length buried fiber to each antenna, with the best thermally conpensated fiber for the above ground jumpers, with sufficient optical isolation to eliminate reflections and with the best stability master laser that money can buy should meet the drift requirements.

3.3 Additional considerations regarding phase drift

Phase Detector and OP Amp Drift

A Mini-Circuits phase detector with a typical phase coefficient of 100 milliVolt/radian has a specified temperature coefficient of 3 μV/C. This implies 30 μradian/C (.0017 degrees/C) which is negligible for the W-band source. Note that if this phase detector were used in an 8 GHz lock loop its drift would require temperature stability of 2.3 C to meet spec. This implies a required temperature stability of 23 degrees C for an 80 GHz lock loop. The phase detectors can be improved by specifying a tighter match of the diodes for a small increase in cost. Operational Amplifier drift is also on the order of a few μV/C, with the spec for the LM6171 being 6 μV/C (integrator). Better op amps can be found, particularly for lower bandwidth loops. These temperature stability requirements are not very stringent and can easily be met in the steady state for the op amps or the phase detectors. Even for the case of a 2 GHz loop for round trip correction, with careful design the analog phase detectors and op amps should be adequate In the fast switching mode the phase detectors used in the mm wave lock loop may change their temperature appreciably when the RF power is switched on or changed appreciably when changing bands. (We can always arrange to keep DC power on the op amps so their temperature should not change as much.) We will measure the phase vs time for RF turn on to verify that the present phase detectors are satisfactory.

Phase Change due to changes in reflection coefficients of interconnected components

A sometimes important source of phase drift is the change in reflection coefficients of components with temperature, particularly those components connected by long cables. One can assume that for sufficiently long cable runs, there will be a frequency for which reflections will add in quadrature to the main signal. For a tentative phase budget of 0.025 degrees, this places an upper limit on the sum of return losses of 70 dB for the cables supplying reference signals for the LO sources. If reflections are worse than this, then the stability of the reflections must be considered carefully. The same 70 dB applies to the sum of the return losses for the fiber cable carrying the W-band reference plus an added restriction that the phase change due to dispersion must be stable as well to the level of 0.025 degree of phase. This applies not only to the fiber jumper cable in the central control building and the antenna but also to the optical power splitters.

A similar concern is the leakage of coherent signals into the LO system. For phase stability of 0.025 degrees, a coherent coupled signal must be 70 dB less than the signal in question. In the case of the low level signal at the output of the W-Band mixer which may be as small as 10 nW coherent with a 10 mW signal reference, 130 dB of isolation is required between the two signals. We have seen a large effect in our lab setup when we used flexible braided coaxial cables instead of semi-rigid coax for these signals. System tests should be made with cables types and lengths as close to actual deliverable hardware as possible.

Photo Mixer Phase versus Temperature

The photo mixer is outside of both the LO Source loop and the fiber length correction loop so its phase versus temperature coefficient adds directly to the total phase drift. As stated above, changes in the reflection coefficient due to changes caused by temperature and optical power variations also can influence the phase.

Multipliers (outside the loop)

The mm-wave multipliers are planned to be mounted on the 77 K stage in the Dewar with very good temperature stability. Tests of individual multipliers will be made as these designs are realized. The phase data shown in figure 2 and 3 which have low frequency prototypes of the doublers outside of the loop encourage us to believe that these devices will not be a problem. However, prototypes of each multiplier design will be tested for phase drift as soon as they are available.

4. Conclusions

We have presented measured data of phase drift of various components and analysis of other components that show it should be possible to build an ALMA LO fully compliant with the proposed phase drift specifications using the baseline photonic reference scheme. Two problem areas, the round trip corrector and the fringe rate synthesizer, were identified that may require special effort. It is the opinion of the authors that the phase drift specification as given in the November 2001 version of the project book is as great a technical challenge as the phase noise specification. The phase calibration intervals will likely be determined by atmospheric instability (or the accuracy of measuring and modeling that instability.)

5. Acknowledgments

These measurements would not be possible without the technical support of Dan Boyd and Mike Stogoski. We would also like to thank Larry D'Addario, Bill Shillue, John Payne, Dick Thompson and John Battle for helpful discussions.

6. References

[1] E. Bryerton, D. Thacker, R. Bradley, and K. Saini, "Noise Measurements of YIG-tuned oscillator sources for the ALMA LO," ALMA Memo # 311.

[2] Eric W. Bryerton, Dorsey L. Thacker, Kamaljeet S. Saini, and Richard F. Bradley, "Wideband Low-Phase-Noise High-Power W-Band Signal Sources," *IEEE MTT-S International Microwave Symp. Dig.*, Phoenix Az, May 2001, pp 1817-20

[3] Eric W. Bryerton, William Shillue, Dorsey L. Thacker, Robert Freund, Andrea Vaccari, James Jackson, Robert Long, Kamaljeet S. Saini, and Richard F. Bradley, "Integration of LO Drivers, Photonic Reference, and Central Reference Generator," ALMA Memo #376

[4] R. Bradley, D. Thacker, E. Bryerton, and K. Saini, "Local oscillator: phase-locked source and multiplier system," ALMA Project Book, Chapter 7.2, edited by L. D'Addario, 2001 Feb 05

[5] Robert Brown, "ALMA Science Requirements ," ALMA Project Book, Chapter 2.1, 2000 April 25

[6] Bryan J. Butler, Simon J. E. Radford, Seiichi Sakamoto, and Kotaro Kohno, "Atmospheric Phase Stability at Chajnantor and Pampa la Bola," ALMA Memo 365

[7] M.A. Holdaway, Satoki Matsushita, and M. Saito, " Preliminary Phase Stability Comparison of the Chajnantor and Pampa la Bola sites," MMA Memo #176

[8] Guillermo Delgado, et al., "Phase Cross-Correlation of A 11.2 GHz Interferometer and 183 GHz Water Line Radiometers at Chajnator," ALMA Memo #361

[9] John C. Webber and D. L. Thacker, "Phase Distribution on Fiber Optic Cable," Proceedings of the 21st Annual PTTI Meeting,

[10] B. Shillue "Round Trip Correction," ALMA Test Interferometer Project Book, Chapter 7.3.2, 2000-02-15

[11] J. Payne, B. Shillue, and Andrea Vaccari, "Photonic Techniques for Use on the Atacama Large Millimeter Array," ALMA Memo #267

[12] Bill Shillue, private communication

[13] C.L. Carilli,and M.A. Holdaway, "Tropospheric Phase Calibration in Millimeter Interferometry," ALMA Memo 262