

A Measurement of X-Band Front-End Phase Dispersion for Delta-Differenced One-Way Range (DDOR) Experiments

Stephen T. Lowe*

ABSTRACT. — This article presents the first direct measurement of the station-differenced nonlinear phase behavior of the front-end electronics used in delta-differenced one-way range (DDOR) measurements. The dispersion's stability and spectral components are measured, and preliminary results indicate that two of the largest DDOR error sources may be approximately halved with simple experiment and calibration modifications. Spacecraft–spacecraft relative navigation can also be improved with similar calibrations. The measured phase deviations induce a systematic delay error in DDOR measurements due to different spectral footprints between spacecraft and quasar sources. The observed phase behavior of the front-end electronics is consistent with indirect estimates of phase variations based on DDOR repeatability tests. The dispersion's spectral behavior is shown to be low-frequency peaked, making self-calibration schemes feasible. The stability of these effects was measured at DSS-26 and DSS-43 on timescales of up to 1.5 hr and for antenna pointing differences of up to 20 deg, and found to be better than this experiment's resolution of ~ 0.05 deg over 62.5 kHz.

I. Introduction

Spacecraft navigation beyond Earth's orbit may employ a variety of techniques, but for the highest accuracy, both ranging and angular measurements are typically required for their orthogonality. At the Jet Propulsion Laboratory (JPL), angular measurements of a spacecraft's location are made using a technique called delta-differenced one-way range, or delta-DOR (DDOR), as shown in Figure 1. Here, the difference in arrival times of spacecraft radio signals at two distant Deep Space Network (DSN) sites is accurately measured, along with similar measurements of a nearby quasar, whose position is known in the International Celestial Reference System (ICRS) quasar frame. In this way, the spacecraft's location is measured angularly relative to the ICRS frame.

The spacecraft group-delay measurements, indicated in Figure 1, involve differencing the measured phase of two or more discrete tones transmitted by the spacecraft. The measured

* Tracking Systems and Applications Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. © 2011 California Institute of Technology. Government sponsorship acknowledged.

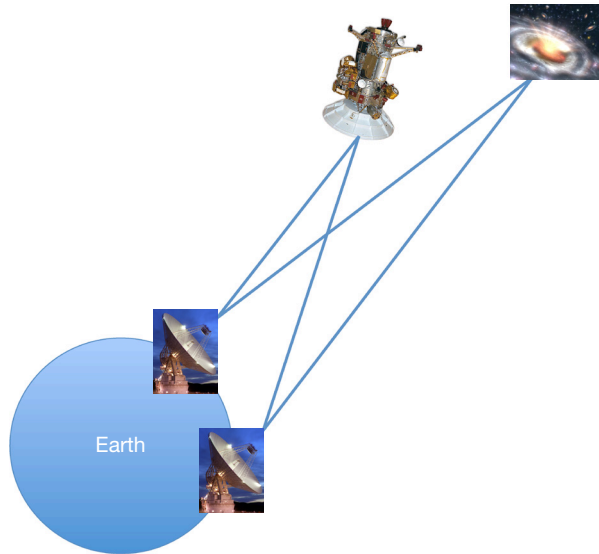


Figure 1. The DDOR measurement involves observing the differences in arrival times between two DSN stations of a spacecraft and a quasar.

phase difference between a pair of tones, divided by their frequency separation, is the primary group-delay observable. This is in contrast to the several-MHz bandpass used to measure the quasar group delays. Quasar observations are necessarily wideband in order to increase the signal-to-noise ratio (SNR) of their quasi white-noise transmission to an acceptable level in terms of phase and delay precision. The result is a group delay derived from the average phase slope across the several-MHz bandpass. If a nonlinear phase response as a function of frequency exists in the front-end electronics, the two different methods used to process spacecraft and quasar data can induce a systematic phase error, which then induces a systematic delay error. The systematic phase error for one channel is illustrated in Figure 2, where the phase at the spacecraft tone frequency differs whether you look at that frequency with very narrowband data, as is done with the spacecraft tone, or whether the phase is obtained from a linear fit across the whole channel, as is done with the quasar wideband noise.

The systematic phase errors shown in Figure 2 induce group-delay errors, as shown in Figure 3. The full hardware bandpass, usually at 8.4 GHz (X-band) or 32 GHz (Ka-band), and the associated phase ripple, are shown with two recorded channels. The difference between the narrowband tone phases and the wideband average linear phases are shown for both channels. The phase difference divided by the frequency difference is the group-delay observable, shown as the slope of the dashed lines, and differs between the spacecraft and quasar. In this exaggerated example, one can see how the difference in measurement bandwidth induces systematic group-delay errors in the presence of bandpass phase deviations. One can also see that the quasar delays across each channel may differ from each other and the overall interchannel group delay. Another mission scenario is to navigate one spacecraft relative to another, for example, a spacecraft in cruise to Mars relative to a Mars orbiter. In this case, differences in the spacecraft transmit frequencies, which are typical, would induce a similar effect.

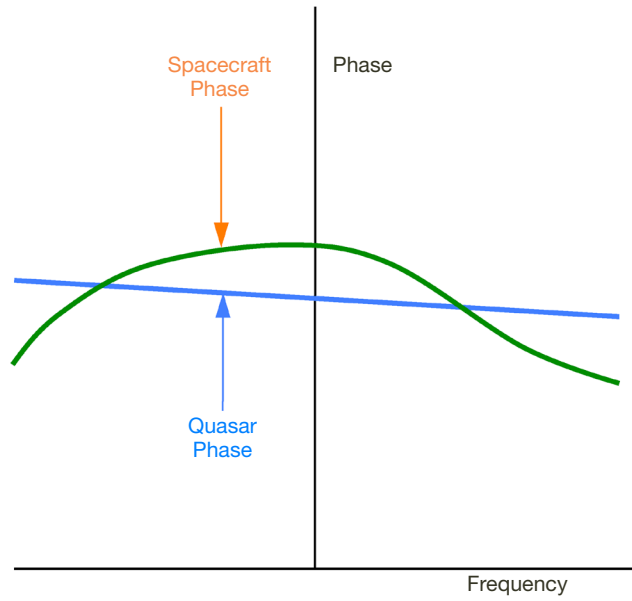


Figure 2. The systematic phase error induced by nonlinearities in the front-end electronics in a single channel. The green curve shows the phase dispersion as a function of frequency, while the blue line shows the average linear behavior. At the spacecraft tone frequency (arrows), the measured phase differs between the narrow-band spacecraft measurement (orange) and the wideband quasar measurement (blue) resulting in a systematic phase error.

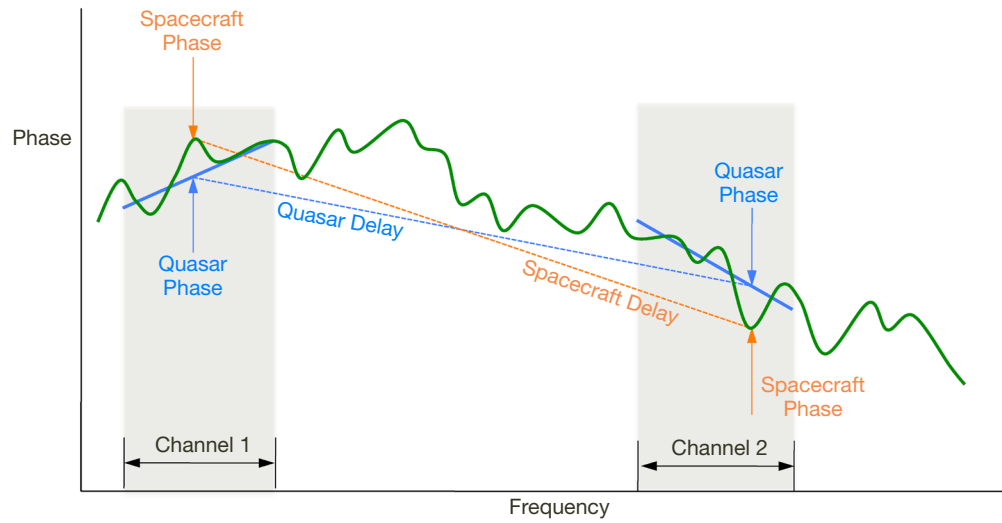


Figure 3. An exaggerated example of how phase deviations across the full bandpass (green) can induce systematic delay errors when comparing the spacecraft group delay (slope of orange dotted line) with the quasar group delay (slope of blue dotted line).

The large number of DDOR measurements performed over the years has provided indirect evidence that the average size of the systematic phase error in a channel is about 0.2 deg.¹ This estimate comes from repeatability tests, and induces a delay error that represents one of the larger DDOR error sources [1]. This small number has made direct measurement of these dispersive effects difficult in the past.

II. Data Set and Initial Processing

In order to measure the magnitude of phase nonlinearities across the bandpass, very high signal-to-noise ratio (SNR) quasar observations, obtained from a recent DDOR measurement, were used to measure the cross-correlation phase with high phase precision and frequency resolution. The data set used for this analysis consists of four quasar OJ 287 observations taken on DOY 108, 2010, using stations DSS-26 and DSS-43, which are 34-m- and 70-m-diameter antennas, respectively. Each of the four observations recorded four 4-MHz channels with 2-bit complex (I/Q) sampling. The channel local-oscillator frequencies were approximately 8406.384 MHz, 8387.277 MHz, 8368.171 MHz, and 8444.597 MHz, respectively, for channels one through four. The first two observations were 8 min in duration while the other two were 6 min, for a total of 28 min of data, spanning about 1.5 hr.

The data were processed with standard very long baseline interferometry (VLBI) correlation techniques using SoftC, JPL's software VLBI correlator,² with the exception that 1024 lags, or sample offsets, were computed rather than the usual four. As in usual VLBI processing, the cross-correlation sums as a function of lag were Fourier-transformed to the frequency domain by SoftC. These frequency-domain complex quantities are referred to as "bins," and roughly divide the channel into equal-width subchannels. In this case, the bins are 4 MHz/1024 in width, or about 3.9 kHz wide. The phase of each bin represents a sinc-function weighed average phase centered over that bin's subchannel.

One SoftC processing error that requires attention is referred to as finite-lag phase errors. Bin phases show a small systematic offset that alternates in sign with bin number, and is the result of having only a finite number of lags: the more lags that are computed, the smaller this phase error becomes. Monte Carlo simulations show this error to be about 0.01 deg for 1024-lag processing. In addition, the odd and even bins show the same average error, with opposite sign, to better than a nanodegree.

Figures 4 (a) and 4(b) show channel 1's bin amplitude and phase from raw SoftC output, as a function of frequency over the 4-MHz baseband channel, for a single 1-s integration. Note: the cross-correlation amplitude drops to near zero at the bandpass edges due to hardware filtering, and the phase scatter increases there as expected. The phase slope across the channel is the residual delay with respect to the SoftC's delay model. SoftC's model includes an arbitrary clock parameter, which could have been adjusted to better remove this slope; however, this phase slope will be more precisely removed in the processing noted below.

¹ J. Border, personal communication.

² S. Lowe, *SoftC: A Software VLBI Correlator* (internal document), Jet Propulsion Laboratory, Pasadena, California, 2003.

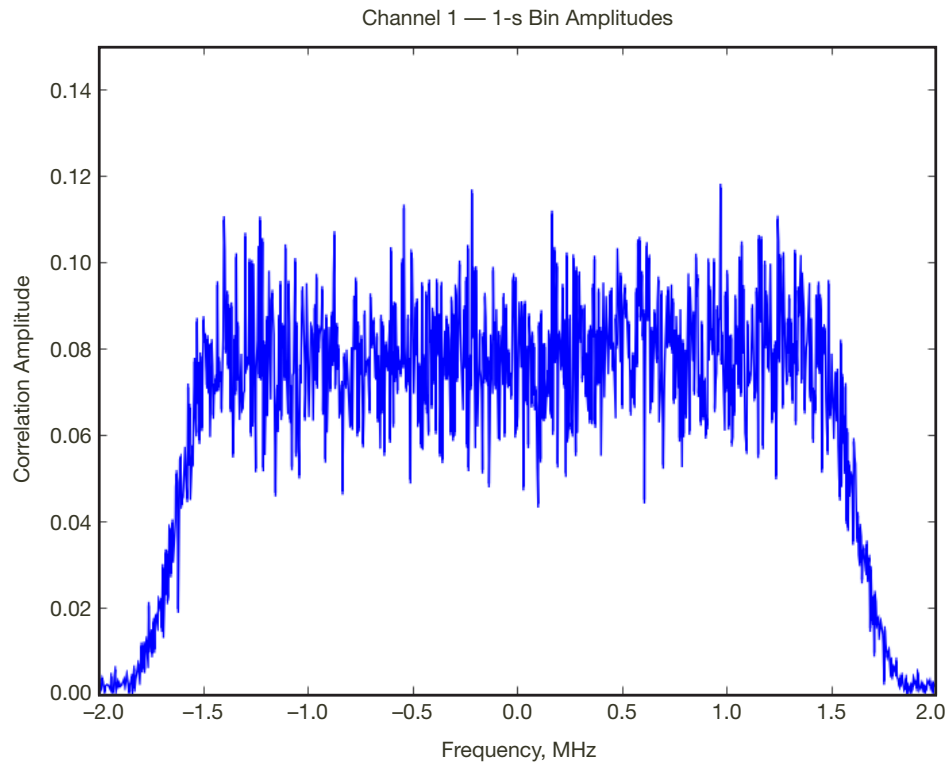


Figure 4(a). The 1-s bin amplitudes from the first second of channel 1's data as a function of baseband frequency. Hardware bandpass filters attenuate the signal at the bandpass edges.

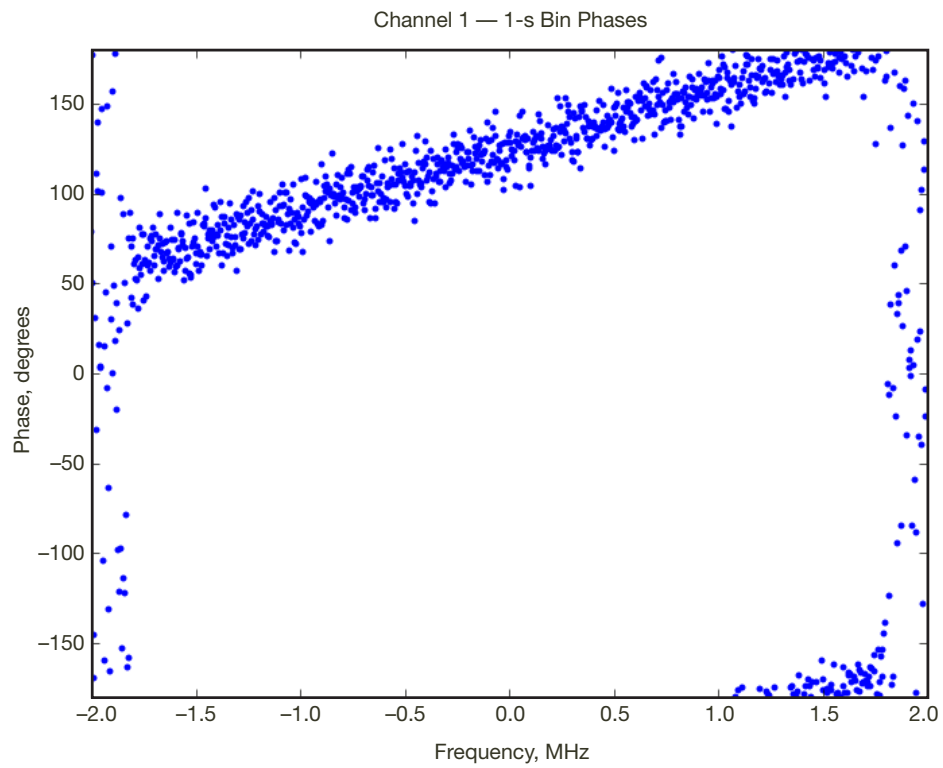


Figure 4(b). The 1-s bin phases from the first second of channel 1's data as a function of baseband frequency. The upward slope indicates a delay residual relative to the correlator model. The large phase scatter at the bandpass edges is due to the signal SNR dropping to near zero there.

III. Data Analysis

Because the goal of this analysis is to observe phase nonlinearities as a function of frequency, the phases within each 1-s integration were fit to a line, ignoring the low-amplitude phases near the channel bandpass edges. The resulting linear trends were removed from each second of data, and independently for each of the four channels. This was done to remove any time-variable delays, such as atmospheric or signal-propagation effects, that could smear a bin's time-averaged phase. The bin phases in each channel were then averaged over the entire 28-min data set. Figure 5 shows channel 1's averaged bin phase, where the 64 low-amplitude bins on each side of the bandpass were removed. The error bars on the points show the expected phase error based on white noise statistics derived from the bin SNR. These errors are also consistent with the scatter of the phases that enter into each time-averaged phase, indicating there is no obvious time-dependent systematic phase behavior observed beyond the expected white system noise. Note that since these points are cross-correlation phases, they represent the station-differenced bandpass nonlinearities.

In order to increase the phase measurement precision, though at the expense of frequency resolution, each group of four adjacent bins was averaged. Figure 6 shows the 4-bin averaged bin phases for all four channels, where each channel is offset by 1 deg in the plot. Figure 7 shows a similar plot for 16-bin averaged phases. Simply running SoftC with 64 lags would produce a similar plot, but would increase the finite-lag processing effects from

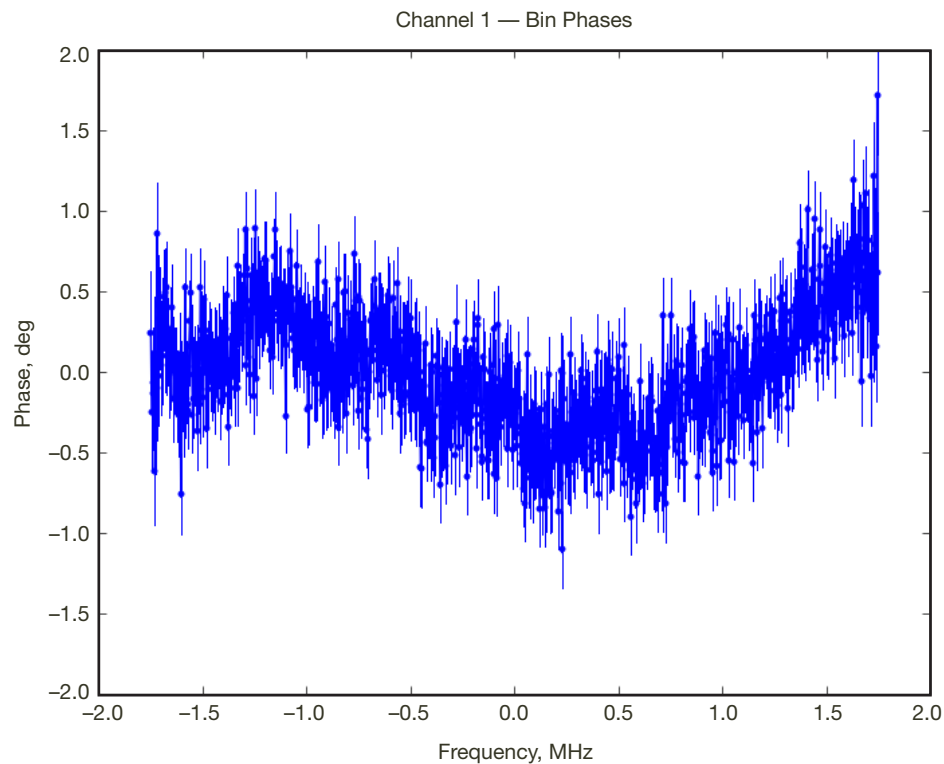


Figure 5. The time-averaged bin phases using all 28 min of data for channel 1 as a function of baseband frequency. Nonlinear phase behavior is evident at about the ~1-deg level for large structures and at ~0.5-deg level for ~0.5-MHz structures.

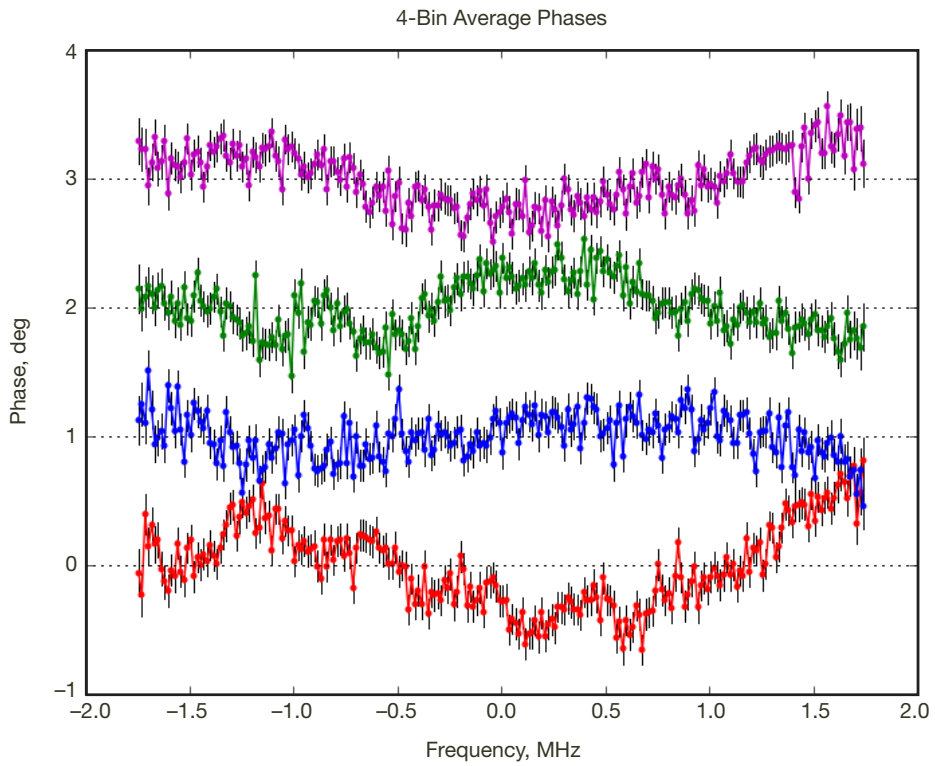


Figure 6. The 4-bin averaged phase as a function of baseband frequency for all four channels (each channel is offset 1 deg above the previous channel). These data average over the full 28 min of quasar data. Red is channel 1, blue is channel 2, green is channel 3, and purple is channel 4.

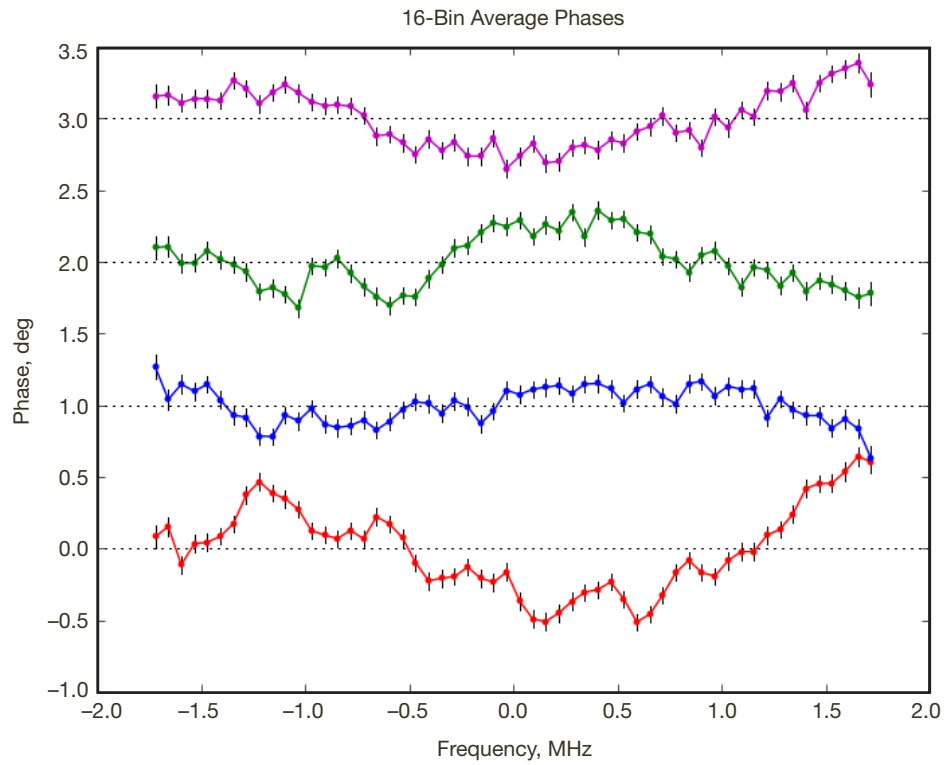


Figure 7. The 16-bin averaged phase as a function of baseband frequency for all four channels (each channel is offset 1 deg above the previous channel).

~0.01 deg to ~0.1 deg. Averaging with an even number of bins, as was done here, cancels these effects to first order, reducing this error to well below the 0.01-deg level. Bin averaging also results in a flatter phase weighting across the bin and better separates the averaged bins, compared to bins with overlapping sinc-function weights, as in the 64-lag case. The error bars in Figure 7, which indicate the white noise component of the phase behavior, are much smaller than systematic phase structure seen. The RMS scatter of these points about the removed linear trend depends on the channel, but is in the range of 0.15 to 0.31 deg (see Figure 13), consistent with the assumed value of 0.2 deg from repeatability tests, as noted above.

To assess the level of software processing effects in these plots, SoftC’s Monte Carlo was used to generate a similar data set, but with a somewhat higher SNR for greater phase resolution. These Monte Carlo scans were processed identically to the real DDOR quasar data and resulted in the phase plot shown in Figure 8. These phases show no nonlinear behavior across the bandpass. This indicates that SoftC is not contributing to systematic nonlinearities at the level shown, and that the phase behavior in Figure 7 is very likely hardware related.

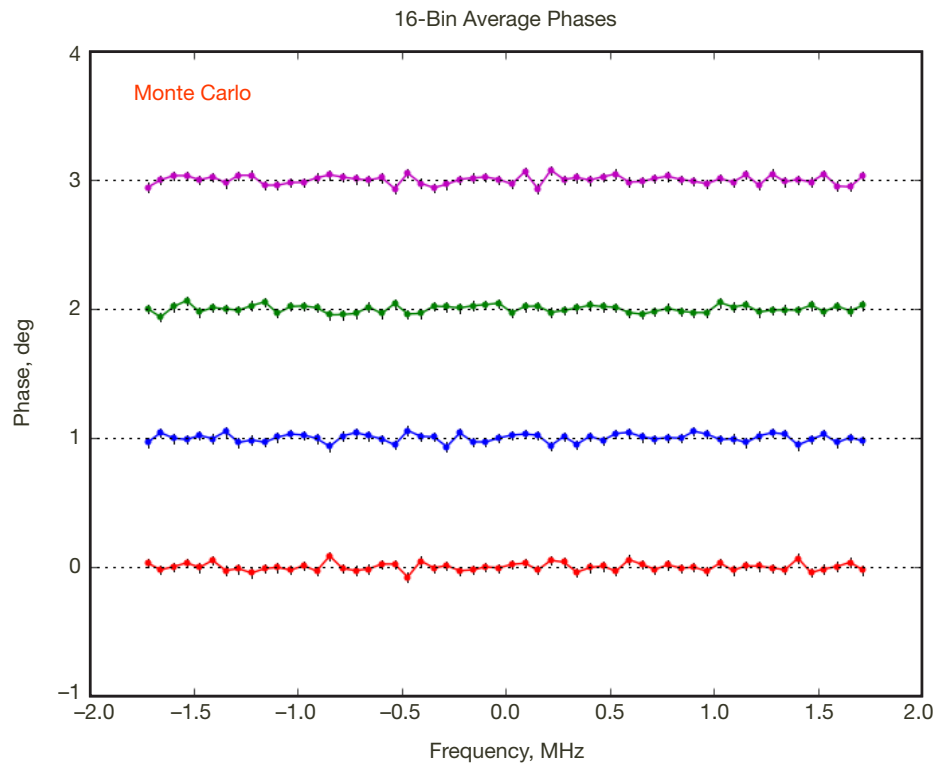


Figure 8. Monte Carlo simulated data. The 16-bin averaged phase as a function of baseband frequency for all four channels, processed identically to the real data (each channel is offset 1 deg above the previous channel). The phase RMS scatter is 0.03 deg, consistent with the Monte Carlo SNR.

IV. Dispersion's Spectral Properties

The spectral components of the phase behavior shown in Figures 5 to 7 are important for assessing calibration possibilities, as will be discussed below. Looking at Figure 6, for example, one sees that the lowest-frequency components of the figure have the largest phase fluctuations, indicating a low-frequency peaked systematic error.³ The bin phases shown in Figure 5 were Fourier transformed, as were the other three channels. Only the very low-frequency components were found to be significant, and any higher-frequency nonlinearities in phase are below the sensitivity of these data. Figure 9 shows the spectrum of all four channels superimposed over each other for comparison. In Figure 5, one can observe power in channel 1 at ~ 1 cycle per 4 MHz, which corresponds to the peak seen in Figure 9 at $1 \text{ cycle} / 4.0 \times 10^6 \text{ cycles/second} = 0.25 \mu\text{s}$. The phase behavior is explicitly low-frequency peaked (the DC point is expected to be zero since a linear trend was explicitly removed) and essentially all observed power is below $2 \mu\text{s}$. This implies that relatively simple functions with a small number of parameters can fit the systematic phase behavior.

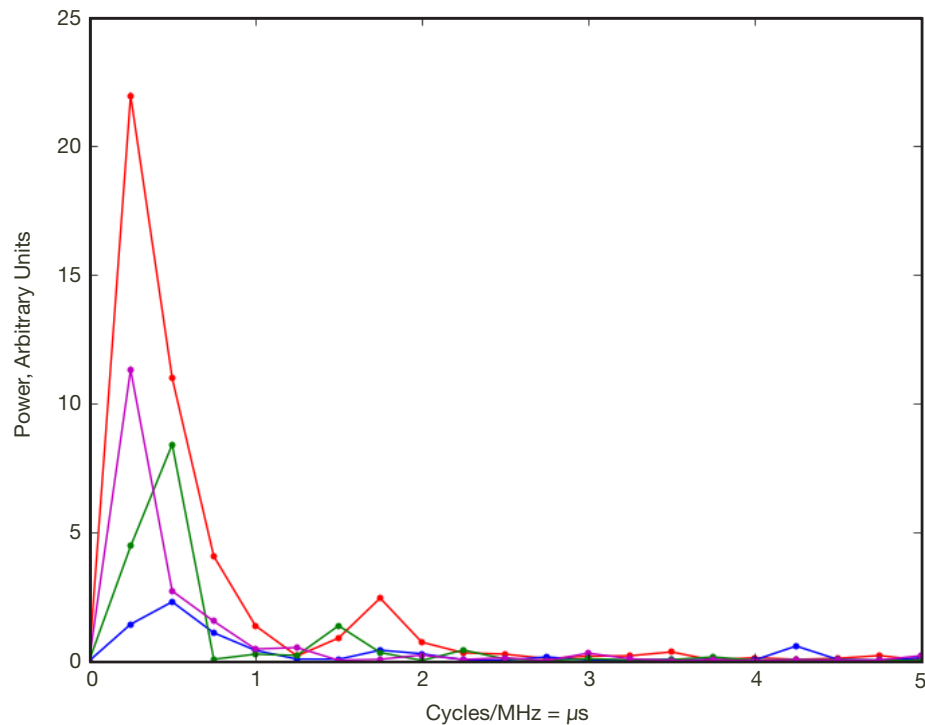


Figure 9. The spectrum of the dispersive phase behavior for all four channels. Except for the very low-frequency end of the spectrum, the measurement noise floor dominates. Above about $2 \mu\text{s}$, any systematic phase errors are below the detection threshold of these data. The bump just below $2 \mu\text{s}$ in channel 1 (red) corresponds to the $\sim 0.5\text{-MHz}$ oscillation evident in that channel in Figure 7.

³ The terms “frequency component” or “low-frequency” in this case refer to the relative rate of change of phase across the bandpass and borrow from the common time-domain terminology, but these quantities are not true frequencies, but rather times. The same comment applies to the discussion of the dispersion's spectral components.

V. Stability

The phase nonlinearities' temporal stability over a time scale of about 1.5 hr can be examined by comparing the data from the first and last observations. These observations were processed identically to that leading to Figure 7, and are shown in Figures 10 and 11. These plots appear quite similar, with most differences appearing within error bars. To make this more quantitative, these plots were differenced and the chi-squared per degree of freedom computed for each channel. Figure 12 shows the channel phase differences between observations. No obvious phase structure is seen, and the behavior appears consistent with white noise. The chi-squared per degree of freedom for each the four channels is 0.96, 0.63, 1.05, and 0.80, respectively, where white noise has an average value of 1. Thus, there is no evidence of temporal changes in the phase systematics over the 1.5 hr between the two scans being differenced.

The stability of the phase deviations with antenna movement is important for optimizing calibration schemes. The stability with antenna pointing is likely good for few-degree movements, as these effects are already in the quasar-to-spacecraft antenna slewing typical of most DDOR observations. The temporal stability test above is also implicitly a test of stability with respect to antenna pointing, as the antenna pointing changed by about 20 deg in the sky between the first and last observations. In terms of azimuth and elevation, both stations moved 7 to 8 deg in elevation, while in azimuth they moved about 12 and 22 deg, respectively.

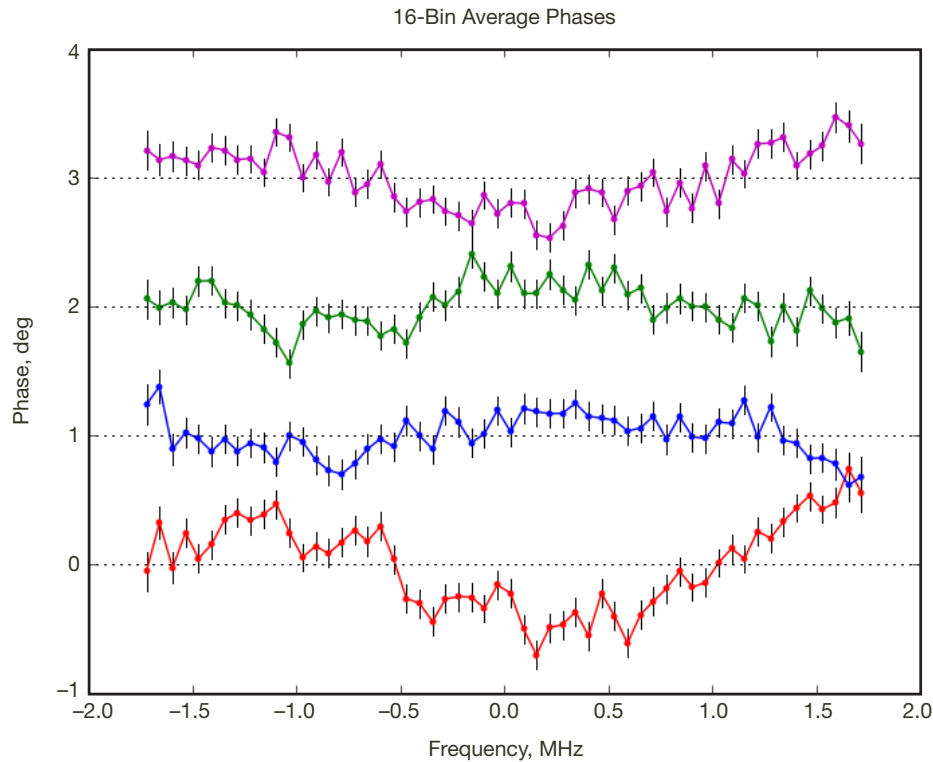


Figure 10. The first 8-min observation processed as in Figure 7. The mid point of this observation was at DOY 108 07:03:00.

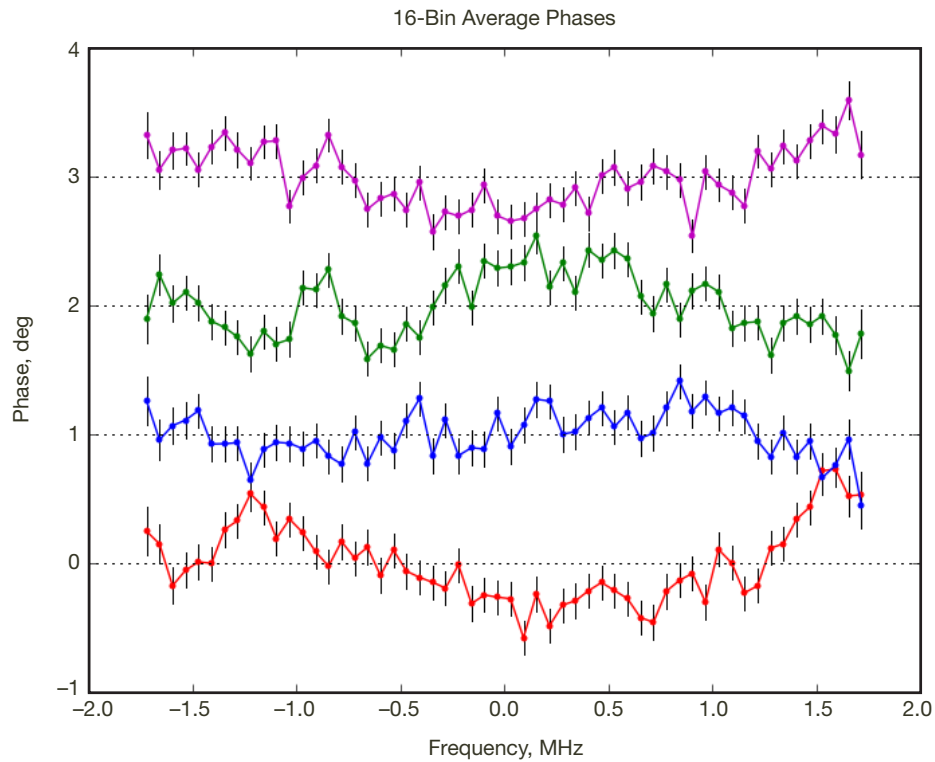


Figure 11. The last 6-min observation processed as in Figure 7. The mid point of this observation was at DOY 108 08:30:00, or 1:27:00 later than the first observation.

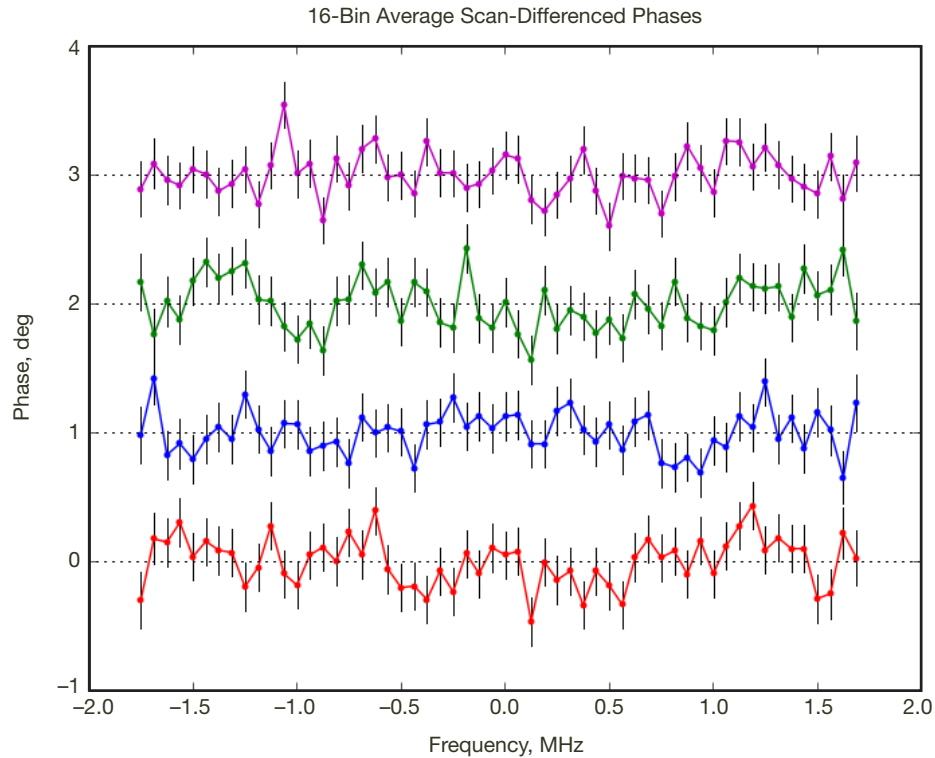


Figure 12. Bin phase differences between the first and last observations, which were 1.5 hr apart. The chi-squared per degree of freedom for each channel is consistent with 1, indicating no measurable change in the phase nonlinearities over the 1.5 hr is observed. The antennas at both stations moved about 20 deg over this interval.

VI. Calibration Possibilities

Although the preliminary stability tests presented above derive from only a single experiment, the level of stability observed implies that particularly simple calibration schemes may be feasible. With the 1.5-hr temporal stability, and 20-deg pointing stability measured above, adding a 30-min high-SNR quasar observation immediately before and/or after a DDOR experiment would allow measuring the front-end phase systematics in the same way as presented here, and thus allow their removal. The DDOR quasar observation itself may also allow a “self-calibration” at some level if its SNR is great enough to remove the lowest-frequency components of the dispersion. This idea can be tested with archived DDOR data.

In order to assess the benefit of self-calibration schemes, or the amount of data needed on an additional high-SNR quasar, the phase RMS as a function of polynomial degree that is removed from the data was measured. The data in Figure 7 had successively higher order polynomials removed and the phase RMS computed. Figure 13 shows the residual phase RMS of Figure 7’s data as a function of the degree of the removed polynomial. After removing a fourth-order polynomial, the phase RMS is below about 0.1 deg in all channels. The phase systematics in channels 1 and 3 have a noticeable oscillation in frequency, as seen in Figure 7, and this appears in their spectrum in Figure 9 as small bumps between 1.5 and 2 μ s. These effects are removed by polynomial fits of degree 18–20, and result in phase RMSs near 0.05 deg, the noise floor for 16-bin averages. This study indicates that fitting and removing even a fourth-order polynomial from the quasar front-end phase nonlinearities can halve the systematic effects from \sim 0.2 deg to \sim 0.1 deg. Higher SNR-observations, such as used in this study, can remove an additional factor of two, for an RMS of \sim 0.05 deg. An assessment of this improvement in relation to the other errors in the DDOR error budget, and the additional resources required to perform this calibration, is needed to determine the optimal level of calibration.

The systematic error resulting from the observed phase dispersion across the bandpass is closely related to the quasar system noise error, and these are the two largest error sources in the DDOR error budget, as shown in Figure 14. These errors are closely related because for a given quasar, antenna aperture, and system temperature, the only way to reduce the quasar noise error in a fixed amount of time is to widen the bandpass and collect more samples. However, this will have the effect of averaging the phase over larger systematic features, given the low-frequency-peaked nature of the dispersion. Lowering the bandpass width to lessen the impact of these systematic phase errors increases the quasar’s white system noise. It has been determined empirically⁴ that channel widths of about 4 MHz best balance these competing effects. This study indicates there is the possibility of recording wider quasar bandwidths to reduce the quasar white noise contribution, while simultaneously using simple fits of the phase dispersion to reduce the systematic phase effects. This could significantly reduce the two largest DDOR error sources, shown in Figure 14 (orange border), with relatively simple experimental modifications.

⁴J. Border, op. cit.

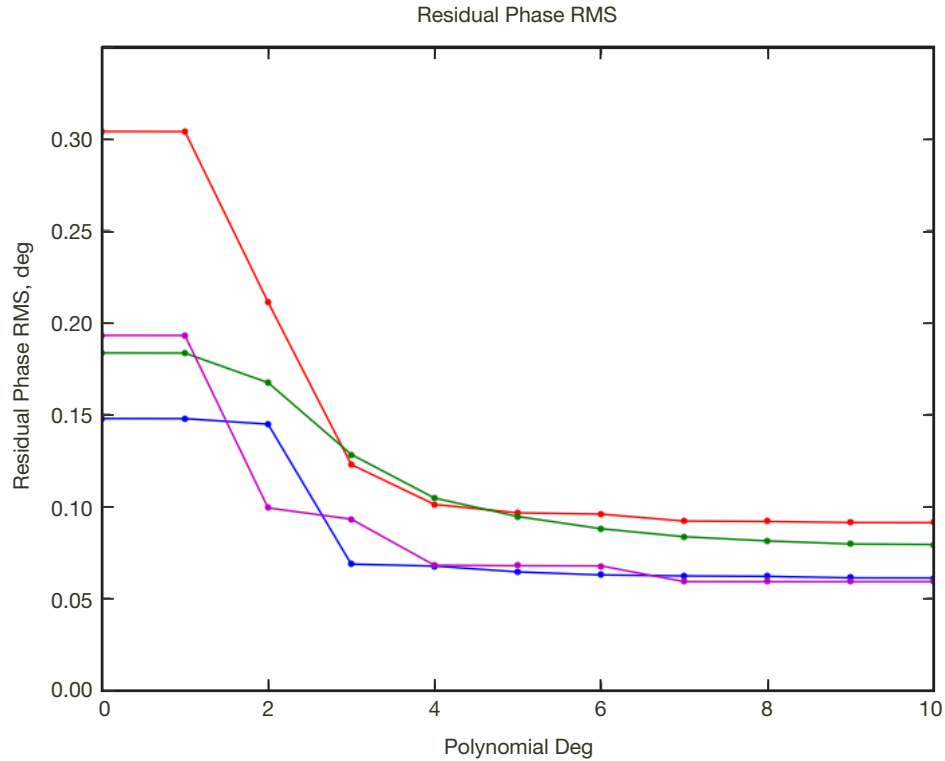


Figure 13. Residual phase RMS as a function of removed polynomial degree. A polynomial of degree 4 removes most of the lowest-frequency dispersive phase behavior. The noise floor is at about 0.05 deg in this 16-bin average case. Channels 1 and 3 (red and green) are somewhat higher, but drop to near 0.05 deg at about degree 18–20 (not shown). This drop corresponds to the removal of the higher-frequency components shown in Figure 9 between 1.5 and 2 μ s as additional small bumps. The data in Figure 7 has a linear trend removed, thus there is no change in the RMS scatter removing a 0th or first-order polynomial.

VII. Follow-On Work

The next steps include experiments to better understand the origin of Figure 7’s phase dispersion, its temporal stability, and stability with antenna movement. Experiments to measure the stability on day and week timescales, and over full-sky antenna pointing, seem essential. Self-calibration schemes where only the typical DDOR data are used, but processed to also remove the largest low-frequency dispersive components, should be studied with archived and current DDOR data. Experiments where all channels observe the same frequency, overlap in frequency by 1/2 channel width, or use 2- or 8-MHz bandpasses, may help characterize the dispersive effects and their sources.

Calibrating the front-end dispersion more accurately than this study may not be required in the near term, given the level of other errors in the DDOR error budget ([1] and Figure 14). On the other hand, a factor-of-two improvement over this study could be realized by repeating this experiment with 2 hr of data rather than the 28 min presented here. If even more accuracy is required, or if higher-frequency components of the dispersion require study, other signal sources will likely be required. Ideas for this include tower-mounted transmitters near each antenna, an unmanned aerial vehicle (UAV)–mounted source flying

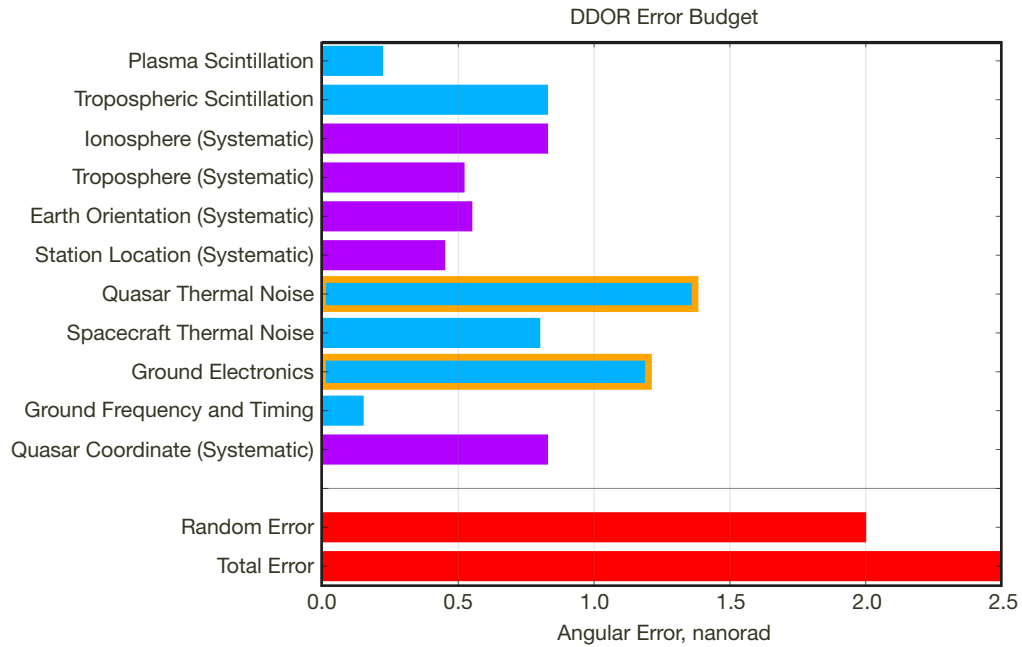


Figure 14. The DDOR error budget as published in [1]. Error sources indicated with blue bars are assumed to have Gaussian statistics, while purple bars indicate systematic error sources. This article has focused on dispersion in the hardware front-end electronics, indicated by “Ground Electronics.” Simple modifications to DDOR experiments may significantly reduce the two largest error sources (orange border), as described in the text.

over each antenna, radar reflections from a Earth-orbiting satellite visible to both antenna, an Earth orbiter transmitting spread-spectrum telemetry, among others. A problem with some of these ideas is finding a signal with sufficient stability and in the protected portion of the spectrum used by the DSN.

VIII. Conclusions

Figures 6 and 7 represent the first measurement of the station-differenced nonlinear phase behavior of the front-end electronics used for DDOR measurements. The phase scatter is consistent with indirect estimates based on DDOR repeatability tests. Care has been taken to remove finite-lag effects from SoftC to well below the 0.01-deg level. The spectral behavior of the nonlinear phase deviations is shown to be low-frequency peaked, as indicated in Figure 9. The temporal stability of these effects was measured on time scales up to 1.5 hr, and for pointing changes of up to 20 deg, and found to be stable to better than this experiment’s resolution of ~0.05 deg over 62.5 kHz. These preliminary stability measurements indicate that using larger quasar bandwidths with a simple calibration scheme may significantly reduce the two largest DDOR error sources.

Acknowledgments

Thanks to Jim Border for providing the quasar data used in this study and discussions on DDOR processing, Larry Young for a careful review of this article, and Charles Naudet and Gabor Lanyi for their discussions on this topic. Thanks to Chris Jacobs for discussions on SoftC's finite-lag phase corrections and their relation to this work. This work was supported by the Radiometric Tracking Work Area, under the Interplanetary Network Directorate DSMS Technology Program.

References

- [1] J. Border, G. Lanyi, and D. Shin, "Radiometric Tracking for Deep Space Navigation," presentation no. 08-152, 31st Annual AAS Guidance and Control Conference, Breckenridge, Colorado, February 2008.