Noise Pedestal Effects in Turnaround Ranging Channels on Conscan Performance

Joseph Griffin*

ABSTRACT. — Many deep space missions retransmit received uplink signals back to Earth in order to enable round-trip light time measurements. These measurements are used to calculate the spacecraft's distance to Earth, and the retransmission channel is referred to as a turnaround ranging channel. During downlink tracks, Deep Space Network antennas are repetitively scanned in a conical pattern to cause small changes in residual carrier signal-to-noise ratio. The antenna pointing error can be inferred from these changes. Shortly after launch, turnaround ranging channel uplink noise can dominate noise power in the downlink signal, which results in noise power that shifts to match residual carrier power changes during a conical tracking scan. When turnaround noise dominates downlink noise, the carrier signal-to-noise ratio does not change over a conical scan, and the pointing error estimate is ruined. The dominant turnaround noise appears on a power spectrum as a pedestal raised above the downlink noise floor, so the phenomenon is named the "pedestal effect". The pedestal effect is studied here in detail, and an algorithm for mitigating the pedestal effect is proposed.

I. Introduction

The Deep Space Network (DSN) estimates the distance between spacecraft and ground-side antenna by measuring the time a radio signal takes to make a round trip between them. This estimation is important for many applications, such as planning maneuvers, estimating locations of heavenly bodies, and updating orbital trajectory models. To perform this measurement, the ground tracking station transmits an uplink carrier onto which a ranging signal is phase modulated. The spacecraft transponder demodulates the uplink carrier and phase modulates the downlink carrier with the received ranging signal. Some missions build transponders to rebuild the ranging signal and

*Communications Ground Systems 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.

remove noise from the uplink in the process, but the spacecraft of interest to this study simply modulate the downlink carrier with a filtered version of the ranging signal as it is received in the uplink channel. A ground tracking station receives and demodulates the downlink carrier, and the round-trip light time (RTLT) of the ranging signal is calculated to estimate the distance between spacecraft and ground tracking station. Retransmitting a received uplink signal, instead of a reconstructed one, is referred to as turnaround ranging. Reconstructing uplink signals is referred to as regenerative ranging.

In order to maintain antenna pointing precision toward deep space vehicles, the DSN measures pointing error with a technique called conscan. The boresight of the antenna is repeatedly scanned in a conical pattern to cause small changes in received signal power. The antenna pointing error can be inferred from the relative signal power profile over the conical sweep. By default, conscan is completed in 2-minute intervals, and residual carrier power is measured at 1-second intervals during conscan. The presence of automatic gain control (AGC) modules in the signal processing chain typically ruins relative power measurements, so pointing correction calculations are typically done with the residual carrier signal-to-noise ratio (SNR). Typically, the noise floor at DSN ground tracking stations holds a steady power level, a power level that receives all the same scaling as the residual carrier power. Using the residual carrier SNR allows the noise floor to normalize AGC effects out of the residual carrier power profile. Thus antenna pointing error can be determined correctly even with AGCs active.

Turnaround ranging and conscan are technologies with significant flight legacies, and both are used regularly by the DSN. During early mission phase, turnaround ranging signals are generally picked up at the ground tracking station with very high SNR. With turnaround ranging channels, the downlink signal can be so strong that the noise floor from the uplink ranging signal dominates the noise floor in the downlink received signal. The uplink noise floor, filtered to match the spacecraft's downlink bandwidth, appears in a power spectral plot as a raised pedestal above the ground tracking station noise floor. This noise pedestal in the downlink power spectrum is the source of the name "pedestal effect".

Early in a mission, when the spacecraft is close enough to earth to produce a pedestal effect, noise power estimates at the ground tracking station can be dominated by the turnaround noise pedestal. As the received signal power changes during conscan, the noise power over the pedestal bandwidth rises and falls to match. When the ground tracking station estimates residual carrier SNR, it is normalizing carrier power against a noise floor that changes, not the fixed reference point usually available later in the mission. The residual carrier SNR appears not to change over the course of the scan, and the antenna pointing correction fails. This paper studies the failure of conscan correction due to the pedestal effect, and proposes a technique for mitigating it.

The pedestal effect is a result of two complex technologies interacting with each other in a deep space communication setting. Both turnaround ranging and conscan will be

discussed to build background for pedestal effect analysis. The pedestal effect will be explored with some reliance on graphical evidence generated from a simulation model, and then the same model and further discussion will be used to justify an algorithm for mitigating the pedestal effect. Discussion on turnaround ranging will draw heavily from [1], which covers turnaround ranging channels in a mathematically robust analysis. Due to the complexity of conscan pointing error estimation, that discussion will be largely limited to quality of input data.

II. Ranging Uplink Signal Description

Ranging measurements are performed in the DSN using one of two techniques. The older technique is based on sequentially transmitting tones of decreasing frequency, which can then be correlated against a local model to find a phase offset and determine RTLT. Sequential ranging, as this is known, will be less common among future missions, and pedestal effect analysis with this ranging form is particularly difficult. As described in [1], the uplink power spectrum changes significantly between tones, often causing wild gain swings in the AGC on the spacecraft's turnaround ranging channel. Characterizing these sudden changes is beyond the scope of this analysis, so discussion is limited to the newer ranging technique: pseudonoise ranging. Pseudonoise ranging is performed by phase-modulating onto the carrier a ranging signal that is constructed, in part, from pseudorandom noise codes. This technique has the advantage that the uplink power distribution remains fixed over the entire ranging measurement, which means the pointing error estimate only fails during early mission phase due to the pedestal effect. Limiting scope to pseudonoise turnaround ranging channels that exhibit the pedestal effect therefore makes for a focused and manageable analysis. For a similar reason, this analysis also assumes that the command signal is turned off.

Similarly to the sequential ranging discussion in [1], a fixed bit sequence is pulse-shaped into a squarewave, optionally filtered to limit bandwidth, then phase-modulated onto an uplink carrier, with some residual carrier power. Using the same mathematical model from Chapter 3 of [2], we start with a residual carrier phase-modulated uplink signal as received at the spacecraft.

$$x\left(t\right) = s\left(t\right) + n_u\left(t\right) = \sqrt{2P}\sin\left[\omega_{\mathrm{u}}\mathrm{t} + \theta\left(\mathrm{t}\right)\right] + \mathrm{n_u}\left(\mathrm{t}\right) \tag{1}$$

In (1), P is the received uplink signal power, ω_u is the carrier frequency in radians, $\theta\left(t\right)$ is the signal phase-modulated onto the uplink carrier, and $n_u\left(t\right)$ is noise. Specifically,

$$n_u\left(t\right) = \sqrt{2}[n_c\left(t\right)\cos\left(\omega_{\mathrm{u}}\mathrm{t}\right) + \mathrm{n_s}\left(\mathrm{t}\right)\sin\left(\omega_{\mathrm{u}}\mathrm{t}\right)] \tag{2}$$

is a narrowband process that appears white across the uplink frequency band. Both $n_c\left(t\right)$ and $n_s\left(t\right)$ are independent, stationary, white Gaussian noise processes.

The above model describes a general phase-modulated signal measured at the spacecraft antenna, leaving room for many different communication standards to be substituted

in for $\theta(t)$. In pseudonoise ranging channels, the DSN modulates a fixed bit sequence with a pulse shape $p(t)$ onto the carrier with some modulation index.

$$\theta(t) = \theta_m m(t) \tag{3}$$

$$m(t) = \sum_{n=-\infty}^{+\infty} a_n p(t - nT) \tag{4}$$

$$a_n = \pm 1 \tag{5}$$

In (4), the parameter T can be selected by a mission from a list of DSN-supported values. Elements of the sequence a_n are referred to as "chips", and the rate of transmission of these chips is called the chip frequency, f_{chip}.

$$f_{chip} = \frac{1}{T} \tag{6}$$

DSN-supported chip rates span down to 32 kHz and up to 24 MHz.

For simplicity of analysis, this paper considers a square pulse shape.

$$p(t) = \begin{cases} 0, & t < 0, \\ 1, & 0 \le t < T, \\ 0, & t \ge T. \end{cases}$$

The power distribution between the carrier and data modulation is controlled by the modulation index θ_m, with $0 < \theta_m < \frac{\pi}{2}$.

A. Residual Carrier Power

Under the square pulse shape model, the received signal $x(t)$ can be decomposed as

$$x(t) = \sqrt{2P_c}\sin(\omega_{\mathrm{u}}\mathrm{t}) + \mathrm{m}(\mathrm{t})\sqrt{2\mathrm{P_d}}\cos(\omega_{\mathrm{u}}\mathrm{t}) + \mathrm{n_u}(\mathrm{t}) \tag{7}$$

which highlights power distribution between the residual carrier P_c and data modulation P_d, according to

$$P_c = P\cos^2(\theta_{\mathrm{m}}) \tag{8}$$

$$P_d = P\sin^2(\theta_{\mathrm{m}}) \tag{9}$$

as set by the modulation index θ_m.

B. Modulation Power Distribution

Understanding the power distribution across the ranging signal helps the reader interpret power spectra of signals presented later in the analysis, especially at the downlink receiver. A thorough description of the ranging uplink spectrum can be found in [3], but for this paper the reader need only recognize the presence or absence of a ranging signal in a power spectrum plot.

Because the ranging sequence a_n is periodic, the ranging signal spectrum consists of discrete spectral lines. Additionally, ranging measurement precision tends to improve with the number of bit transitions in a_n. Therefore a_n is designed to approximate an alternating sequence $\{+1, -1, +1, -1, \dots\}$, with a few of the bits being inverted to allow for ambiguity resolution in a measurement against a periodic signal.

DSN ranging supports a few different sequences a_n, all of which similarly apply rare bit inversions to a long alternating sequence. As a result, $m(t)$ approximates a squarewave tone, and the majority of the power in a ranging signal appears at the frequency of this tone and its harmonics. This tone's frequency, which is half the chip rate, is called the range clock frequency, a name derived from the equivalent signal in a sequential ranging context.

$$f_{rc} = \frac{f_{chip}}{2} = \frac{1}{2T} \tag{10}$$

This range clock tone in the ranging signal can be observed in the power distribution of a_n, shown in Figure 1.

While most of the ranging signal power is concentrated at the range clock frequency, the uncommon, irregular bit inversions throughout each period of a_n spread the remaining power across the remaining spectral lines. After pulse-shaping according to $p(t)$, the spectrum is scaled according to a sinc pattern with nulls at multiples of f_{chip}, which negate the range clock harmonics at the even multiples of f_{rc}. Although the spectral spread of a_n is not uniform, the center and side lobes of the ranging signal's pulse shape are easy to visually distinguish in a plot in logarithmic scale, as in Figure 2. Additionally, the harmonics of the range clock are visible as spikes centered on the sidelobes of the pulse shape $p(t)$.

Once the ground tracking station has assembled the ranging signal, it phase modulates the signal onto the uplink carrier. The power spectrum shifts to center on the carrier frequency, and the residual carrier power appears as a new spike among the range clock harmonics. Phase modulation shifts the uplink spectrum away from the simple ranging signal in Figure 2 to match that in Figure 3. After phase modulation, the signal is transmitted to the spacecraft, which measures the transmission in the presence of noise process $n_u(t)$ to form $x(t)$.

III. Turnaround Downlink Signal Description

After demodulating the uplink carrier, a spacecraft will perform some signal processing to maintain the downlink signal quality. An AGC is used to minimize uplink noise effects on the downlink phase modulation. Next, the spacecraft will apply a low-pass filter to limit downlink transmission bandwidth before modulating onto the downlink carrier. Ranging channel filters vary between spacecraft, so the discussion here will not characterize the filter in detail. The filter sets a ceiling on the ranging channel

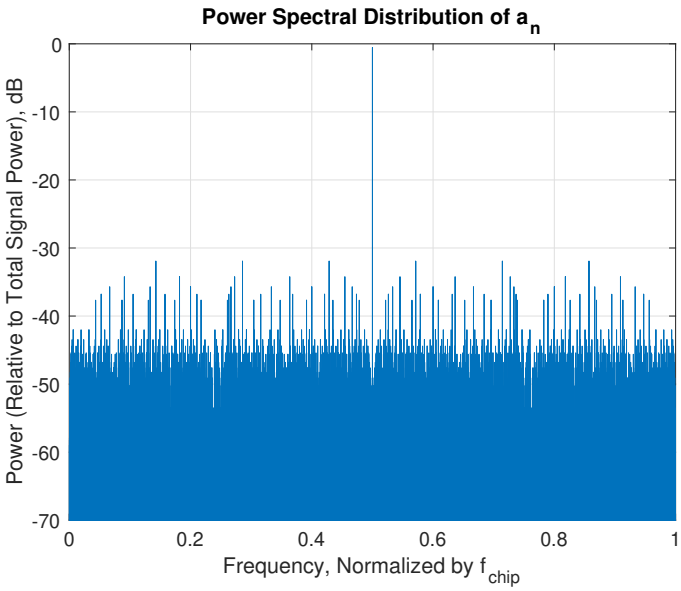

Figure 1. Power distribution of a_n, without pulse-shaping.

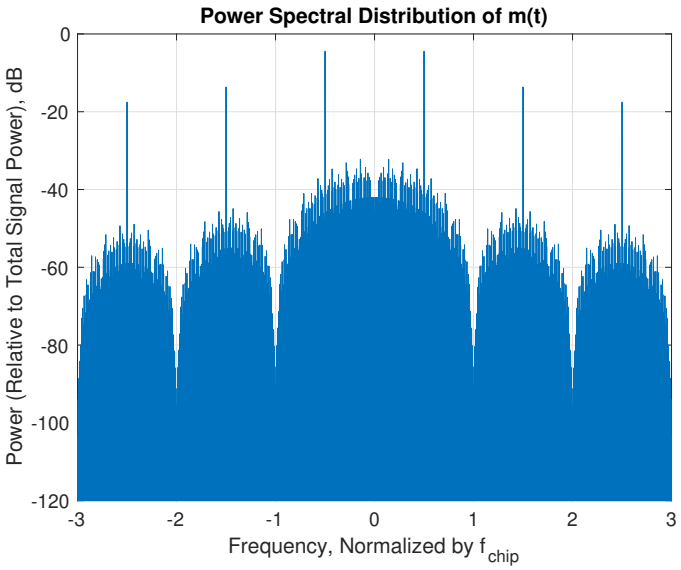

Figure 2. Power distribution of $m(t)$, after pulse-shaping.

downlink bandwidth. The uplink noise process $n_u(t)$ is modulated onto the downlink carrier as well, and the downlink manifestation of the uplink noise that appeared in the turnaround ranging band is renamed $n_{ta}(t)$. Similarly, the composite uplink received signal, after filtering and shifting to the downlink carrier, is renamed from $x(t)$ to $y(t)$.

As with the uplink, turnaround noise power is uniform across the ranging channel by assumption and helps to visually indicate the bounds of the downlink band in a power spectrum. Figure 4 shows the power spectrum of an example turnaround ranging signal $y(t)$ that contains noise $n_{ta}(t)$ originally collected on the uplink as $n_u(t)$.

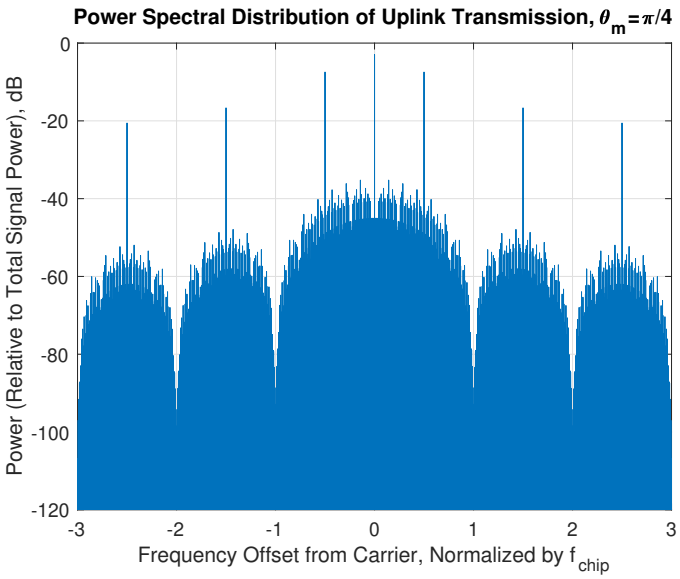

Figure 3. Power distribution of uplink transmission, after phase modulation.

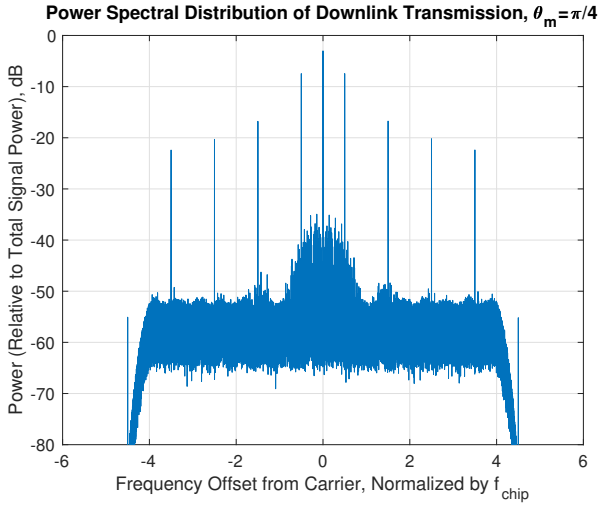

Figure 4. Example power distribution of downlink transmission. Noise floor $n_{ta}(t)$, spanning $8f_{chip}$ and centered on the carrier, is displayed very low in this example to keep spread-spectrum components of $m(t)$ visible.

Once the filtered signal has been modulated onto the downlink carrier, the spacecraft can retransmit back to Earth. Let $n_d(t)$ denote the downlink noise process measured at the antenna. Additionally, let P_{nt} denote the total power of the turnaround noise $n_{ta}(t)$, and P_{nd} be the same for $n_d(t)$. For the majority of the mission, most spacecraft tracked by the DSN are sufficiently far away that the turnaround noise power P_{nt} is much weaker than the downlink noise power P_{nd}. Indeed, even the spectral spreading of the ranging signal $m(t)$ often falls below the noise floor, in which case the ranging processor must rely on long integration intervals in order to reliably calculate the rang-

ing estimate. However, missions that have been recently launched and are still moving away from Earth tend to exhibit much higher downlink signal power, and in some cases P_{nt} even exceeds P_{nd}. In these cases, the downlink received signal's power spectrum displays the composite turnaround signal with a secondary, higher noise floor that spans the ranging channel. The elevated noise floor is called a noise pedestal, and Figure 5 shows a power spectrum with an example of a noise pedestal. As the spacecraft continues away from Earth, the downlink signal grows weaker. Both the noise pedestal and ranging signal sink toward the noise floor, until the downlink noise floor matches and eventually exceeds the noise pedestal power, as in Figure 6.

Noise pedestal effects on turnaround ranging are minimal, because the measurement assumes very little about the composite noise process $n_{ta}(t) + n_d(t)$ that accompanies the ranging signal on the downlink. The conscan pointing correction, however, performs a normalization that relies on an assumption violated by noise pedestal behavior.

IV. Conscan Pointing Correction

Conscan measurements are used by the DSN to provide real-time pointing error feedback. This feedback is used for increasing pointing precision, which in turn improves uplink and downlink signal quality. To build pointing error estimates, the antenna boresight is swept through a conical pattern, and downlink signal power variations over the period of the sweep are used to estimate the pointing error of the cone's center. Figure 7 shows how direction vectors for the conscan pointing operations and the spacecraft might project onto a unit sphere, with some exaggeration for visual clarity. DSN antennas perform conscan over a cone pattern much narrower than that depicted in the figure.

For this paper, we name angles in the analysis relative to spacecraft's true direction in the sky and the minimum-elevation point in the conscan pattern. Let θ_{scan} denote the angle measured from the centerline of the cone pattern to any point on the cone's edge. Let ϕ_{scan} denote the angle measured from the minimum-elevation point in the conscan pattern, in the plane orthogonal to the cone's centerline. Over the course of a conscan interval, θ_{scan} will remain constant and ϕ_{scan} will vary from 0 to 2π. Additionally, let $\phi_{s/c}$ denote the value of ϕ_{scan} at which the antenna points the closest to the spacecraft. Finally, $\theta_{s/c}$ shall refer to the angle between the spacecraft and the cone's centerline. As conscan is meant to provide feedback relative to predicted spacecraft trajectories, it is reasonable to assume that $\theta_{s/c}$ and $\phi_{s/c}$ change little over the scanning interval. Figure 8 indicates how these definitions fit into the diagram in Figure 7.

From the above definitions, we can build an approximation for the most meaningful angle in the analysis: the angle between the spacecraft and the antenna's current pointing position.

$$\theta_{err} \approx \theta_{scan}\sqrt{1 + \frac{\theta_{s/c}^2}{\theta_{scan}^2} - 2\frac{\theta_{s/c}}{\theta_{scan}}\cos\left(\phi_{\mathrm{scan}} - \phi_{\mathrm{s/c}}\right)} \tag{11}$$

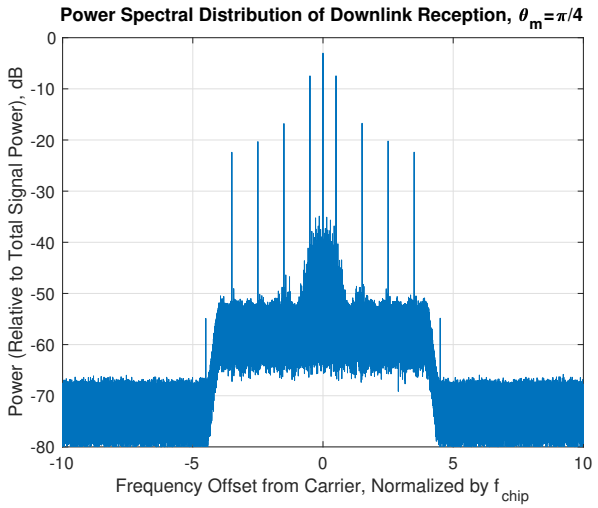

Figure 5. Power distribution of downlink reception during early mission phase. Note that in-band noise floor $n_{ta}(t)$, within $4f_{chip}$ of the carrier, is much stronger than the out-of-band noise floor $n_d(t)$, more than $5f_{chip}$ from the carrier.

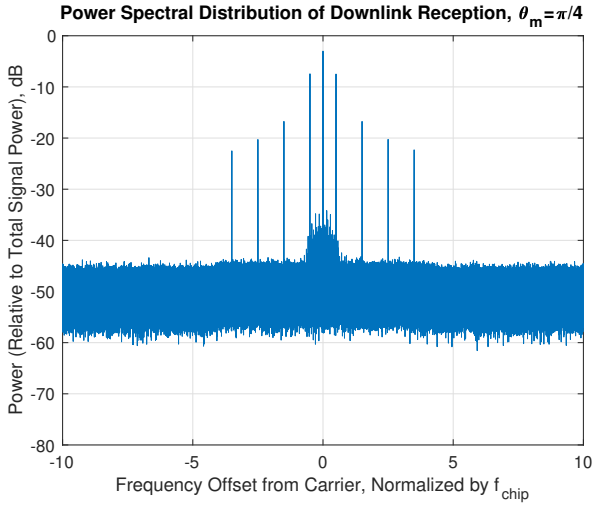

Figure 6. Power distribution of downlink reception during late mission phase. In-band noise floor $n_{ta}(t)$, within $4f_{chip}$ of the carrier, barely stands out from the out-of-band noise floor $n_d(t)$, more than $5f_{chip}$ from the carrier, in contrast with early mission phase.

Approximating θ_{err} this way holds well for small values of θ_{scan} and $\theta_{s/c}$, when the portion of the unit sphere spanned by these angles is nearly planar. A derivation of this approximation is outlined in the appendix.

Calculations for the pointing correction assume larger pointing errors result in strictly lower signal power, and the angular width of the cone is set to keep the spacecraft on the primary gain lobe over the entire sweep. The pointing error estimate is derived from relative variations in signal power. Using relative power measurements automatically normalizes out differences across carrier frequency, antenna size, and hardware quality.

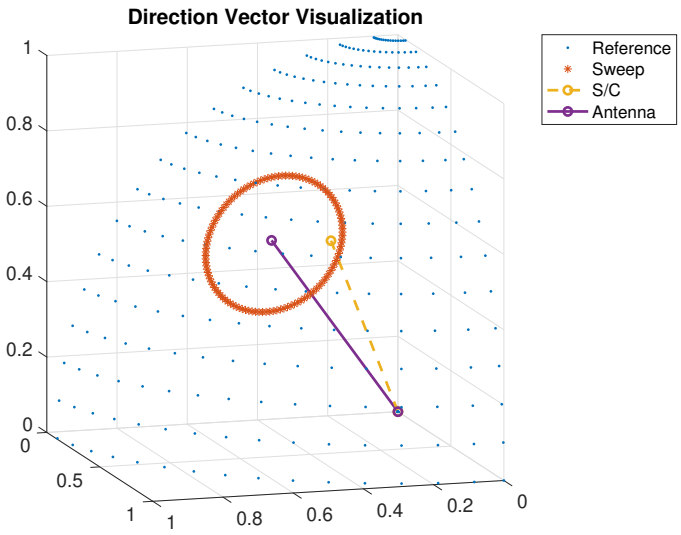

Figure 7. An example plot of antenna pointing direction vectors involved with conscan. The solid line is the center of the conical pattern, the dashed line is the true direction of the spacecraft, and the stars are the conical pattern the antenna sweeps through. Dots are plotted across the unit sphere for visual reference. For visual clarity, the sweep size and pointing error are greatly exaggerated from typical values in the DSN.

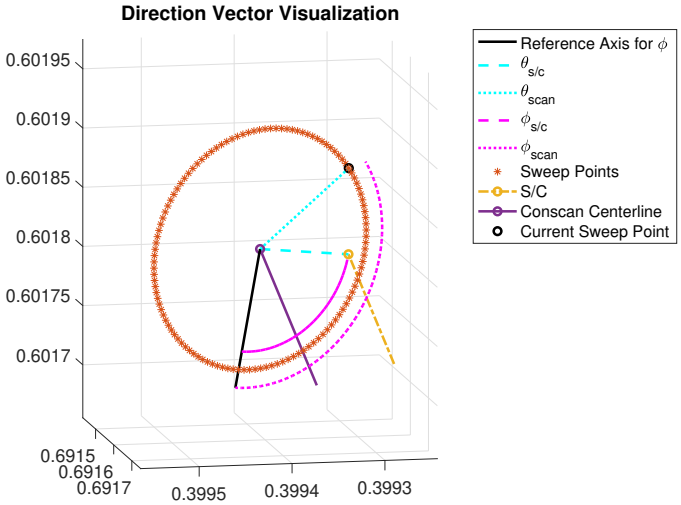

Figure 8. An example plot of antenna pointing direction vectors, with angles diagrammed for definition clarity. This figure depicts a conscan size θ_{scan} more appropriate for the DSN, so angles measured about the center of the unit sphere appear nearly flat.

Because the main lobe width, as measured in degrees, varies between antennas and carrier frequencies, the pattern is designed to induce a specific gain reduction between the cone's center and edge. By default, this gain difference is a factor of 0.1 decibels, though missions can configure the scan to be wider or narrower. Similarly, scanning through the pattern faster allows for more frequent updates, but signal power measurements are "blurred" together by the smooth, continuous motion of the antenna during

the sweep. By default, each scan through the conical pattern is configured to take 120 seconds. Power measurements for the pointing correction estimator are performed over 1-second intervals.

Gain patterns vary across carrier frequencies and antennas, so the conical patterns are designed according to a Gaussian approximation to the center of the pattern's main lobe.

$$G\left(\theta_{err}\right) \approx 10\log_{10}\left(\exp\left\{\frac{-2.773\theta_{\mathrm{err}}^2}{\mathrm{W}^2}\right\}\right), \mathrm{dB} \tag{12}$$

Here, $G\left(\theta_{err}\right)$ is the antenna pointing loss in decibels, and W is the half-power beamwidth of the antenna at the spacecraft's downlink carrier frequency. The use of relative power measurements in the correction estimator means the approximation only needs to fit the shape of the main lobe, not the amplitude. Therefore the only parameter in the approximation is the half-power beamwidth W. This approximation holds well for $\theta_{err} \ll W$.

The default radius of the cone, in radians from the centerline, is straightforward to calculate.

$$\theta_{scan} = W\sqrt{\frac{0.01\ln\left(10\right)}{2.773}} \tag{13}$$

Because each iteration around the cone occurs at a fixed rate, it is simple to rewrite ϕ_{scan} as a function of time.

$$\phi_{scan}\left(t\right) = 2\pi f_{scan}t \tag{14}$$

In (14), f_{scan} defaults to $\frac{1}{120}$ Hz. Building a composite function of time from (14), (11), and (12), it is straightforward to design an approximation for antenna pointing loss as a function of time over the course of the conscan interval. A great deal of cancellation and simplification occurs in this composite function.

$$G\left(t\right) = \frac{10\left(-2.773\right)}{\ln\left(10\right)\mathrm{W}^2}\theta_{scan}^2\left(1 + \frac{\theta_{s/c}^2}{\theta_{scan}^2} - 2\frac{\theta_{s/c}}{\theta_{scan}}\cos\left(\phi_{\mathrm{scan}} - \phi_{\mathrm{s/c}}\right)\right) \tag{15}$$

$$= \frac{10\left(-2.773\right)}{\ln\left(10\right)\mathrm{W}^2}\left(W^2\frac{0.01\ln\left(10\right)}{2.773}\right)\left(1 + \frac{\theta_{s/c}^2}{\theta_{scan}^2} - 2\frac{\theta_{s/c}}{\theta_{scan}}\cos\left(\phi_{\mathrm{scan}} - \phi_{\mathrm{s/c}}\right)\right) \tag{16}$$

$$= 0.1\left(2\frac{\theta_{s/c}}{\theta_{scan}}\cos\left(2\pi f_{scan}t - \phi_{s/c}\right) - \left(1 + \frac{\theta_{\mathrm{s/c}}^2}{\theta_{\mathrm{scan}}^2}\right)\right), \mathrm{dB} \tag{17}$$

Therefore the composite function $G\left(t\right)$ is sinusoidal (as measured in decibels) when (12) is accurate, specifically when $\theta_{err} \ll W$. By design, when $\theta_{s/c} = 0$, $G\left(t\right)$ is fixed at a constant -0.1dB. When $\theta_{s/c} = \theta_{scan}$, $G\left(t\right) = 0$ precisely when $\phi_{scan} = \phi_{s/c}$. The sinusoidal profile of $G\left(t\right)$ is depicted in Figure 9.

Sinusoidal behavior in downlink power is useful for inferring spacecraft correct direction. A phase estimate can be formed from the entire sequence of 120 measurements. When $\theta_{scan} \ll W$ and the error/gain relationship is linear, the amplitude of $G\left(t\right)$ indicates $\theta_{s/c}$. These relationships are the driving reason to keep $\theta_{scan} \ll W$, which in turn drives the small default gain difference of 0.1 decibels. The linear relationship between pointing error and gain variation is depicted in Figure 10.

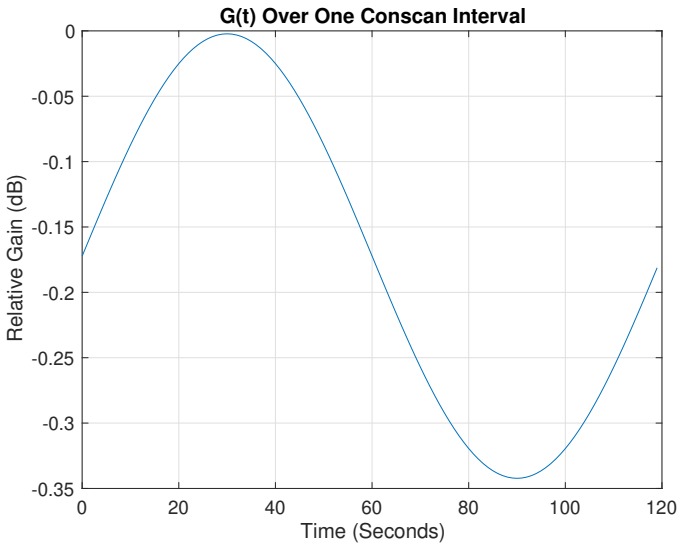

Figure 9. An example plot of gain values $G(t)$ applied to a downlink signal $y(t)$ during conscan.

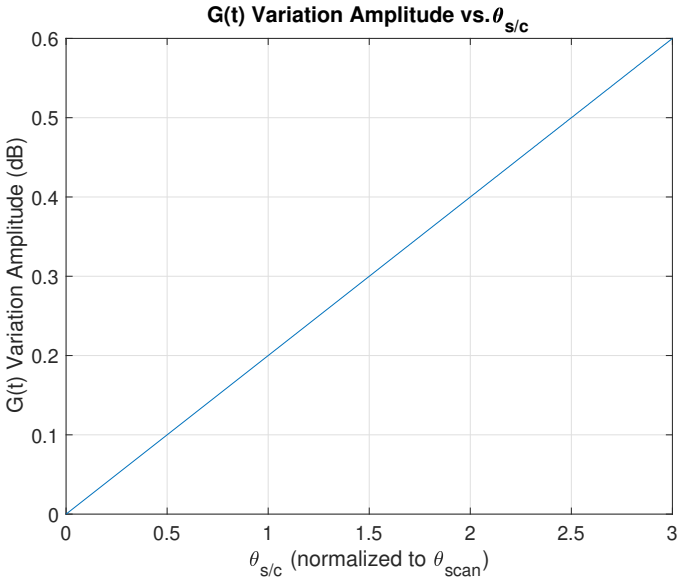

Figure 10. A plot of the relationship between gain variation during conscan and the true pointing error of the cone's centerline from the spacecraft. Note that zero pointing error produces no gain variation.

Thanks to the linear relationship depicted in Figure 10, an estimate of $\theta_{s/c}$ can be calculated from a good measurement of $G(t)$. Similarly, based on the sinusoidal behavior depicted in Figure 9, an estimate of the phase of $G(t)$ allows $\phi_{s/c}$ to be calculated. These calculations are more involved than the analysis for this paper requires, so discussion is limited to whether estimates of $G(t)$ conform to the sinusoidal profile expected.

When a spacecraft's turnaround ranging signal is active and the DSN receiving antenna is performing conscan iterations, the groundside antenna's output signal is a composite of turnaround ranging signal $y(t)$ and downlink noise process $n_d(t)$. For the purpose of

discussing the gain effects directly on the signal, define $G_l(t)$ to denote the conversion of $G(t)$ from decibels to linear scaling units. Because $G(t) \leq 0$, $G_l(t)$ should be between the unitless values of 1 and 0, calculated appropriately for scaling voltage values, not power values.

$$G_l(t) = 10^{\frac{1}{20}G(t)} \tag{18}$$

Only $y(t)$ is scaled according to $G_l(t)$, so the formula for the DSN antenna received signal $z(t)$ is easy to build.

$$z(t) = G_l(t)\, y(t) + n_d(t) \tag{19}$$

Continuing to utilize only relative signal behavior over the conscan interval, the propagation attenuation and absolute antenna gain are ignored. Over the course of conscan, it is assumed that the power of $y(t)$ and $n_d(t)$ remain nearly constant. This constant-power assumption has been verified historically, so the only power variations should be produced by $G_l(t)$, reflecting variations in $G(t)$.

A. Automatic Gain Controller Effects

Power variations are expected to be produced exclusively from variation in $G(t)$, so the $G(t)$ estimate is built up from signal power measurements, which would be taken anyway and used for other applications. To generate these signal power measurements, the processor at the ground tracking station measures residual carrier power. Residual carrier power measurements require the use of a phase-locked loop, which is implemented digitally. To prevent clipping and minimize quantization effects in the digital processing, the loop is preceded by both analog and digital AGCs. As the transmitted downlink signal power varies over the conscan period, the AGCs tune the gain values they apply to $z(t)$.

Let $A(t)$ denote the net gain applied by all the AGCs in the signal processing chain over the conscan interval, and let $w(t)$ be the signal processed by the phase-locked loop.

$$w(t) = A(t)\, G_l(t)\, y(t) + A(t)\, n_d(t) \tag{20}$$

Denote the carrier power in $y(t)$ as P_{ct}, assumed to be constant. A carrier power measurement of $w(t)$ would produce $A^2(t)\, G_l^2(t)\, P_{ct}$. The extra signal power variation $A^2(t)$ ruins the $G_l(t)$ estimate.

While carrier power at the AGC input varies over the conscan period, the power in the downlink noise floor $n_d(t)$ remains unchanged. Assume for the remainder of this Section that $G_l^2(t)\, P_{nt} \ll P_{nd}$. Because it tends to maintain a constant power level at the AGC input, the downlink noise floor power can normalize out the AGC effects on the carrier power during conscan. A measurement of noise power in $w(t)$ would produce $A^2(t)\, P_{nd}$. Instead of estimating pointing error from absolute carrier power, pointing error can be estimated from the square root of a carrier SNR measurement on $w(t)$.

$$\sqrt{\frac{A^2(t)\, G_l^2(t)\, P_{ct}}{A^2(t)\, P_{nd}}} = \sqrt{\frac{G_l^2(t)\, P_{ct}}{P_{nd}}} = G_l(t) \sqrt{\frac{P_{ct}}{P_{nd}}} = 10^{\frac{1}{20}G(t)} \sqrt{\frac{P_{ct}}{P_{nd}}} \tag{21}$$

Normalizing AGC effects out of the $G(t)$ measurement with SNR calculation is convenient because it requires no synchronization. The carrier power and noise power can be simultaneously measured from the same signal, without the need to coordinate timing of the two measurements. The importance (and difficulty) of measurement coordination is described in Section VI.

V. Pedestal Effects on Conscan

Estimating pointing error from carrier SNR is an effective way to normalize out AGC effects, but it relies on the assumption that the noise power in $w(t)$ is fixed over the conscan interval. When a turnaround ranging channel produces a noise pedestal early in the mission, this assumption is violated. Because the noise pedestal is a part of the spacecraft's downlink transmission, the pedestal power in $w(t)$ rises and falls with the carrier power during conscan. If the $P_{nt} \geq P_{nd}$, the carrier SNR appears doesn't scale linearly with $G_l(t)$. Noise measurements in $w(t)$ vary as $A^2(t)\left(P_{nd} + G_l^2(t) P_{nt}\right)$, and the SNR measurement fails to match the behavior of $G(t)$.

$$\sqrt{\frac{A^2(t) G_l^2(t) P_{ct}}{A^2(t)\left(P_{nd} + G_l^2(t) P_{nt}\right)}} \neq G_l(t) \sqrt{\frac{P_{ct}}{P_{nd}}} \tag{22}$$

If $P_{nt} G_l^2(t) \gg P_{nd}$, the $G_l^2(t)$ factor appears to cancel out of the expression in (22), and the carrier SNR measurement remains nearly constant over the conscan measurement.

In this case the pointing correction estimator, which is designed under the assumption that carrier SNR always changes with the antenna gain applied to the downlink signal, interprets unchanging carrier SNR as an indication that the spacecraft is centered in the cone. In Figure 11, the SNR variation over a conscan interval is compared between early mission phase, when $P_{nt} G_l^2(t) \gg P_{nd}$, and late mission phase, when $P_{nt} G_l^2(t) \ll P_{nd}$. Pedestal effects reduce the SNR variation to the point that it is almost imperceptible in comparison to the late-mission case.

As the mission continues and the received power in the ranging channel decreases, the noise pedestal in the power spectrum sinks toward the antenna noise floor, and the apparent noise power $P_{nd} + G_l^2(t) P_{nt}$ varies less over the conscan measurement. The carrier SNR begins to vary properly as this happens, but the amplitude of this variation is still diminished.

VI. Direct AGC Cancellation

Normalizing carrier power estimates against noise power is a very clever technique for automatically ignoring AGC effects. During early mission phase however, when noise pedestal power can dominate the noise measurement, this clever normalization ruins the pointing estimator input data. In this case, a straightforward solution would be to

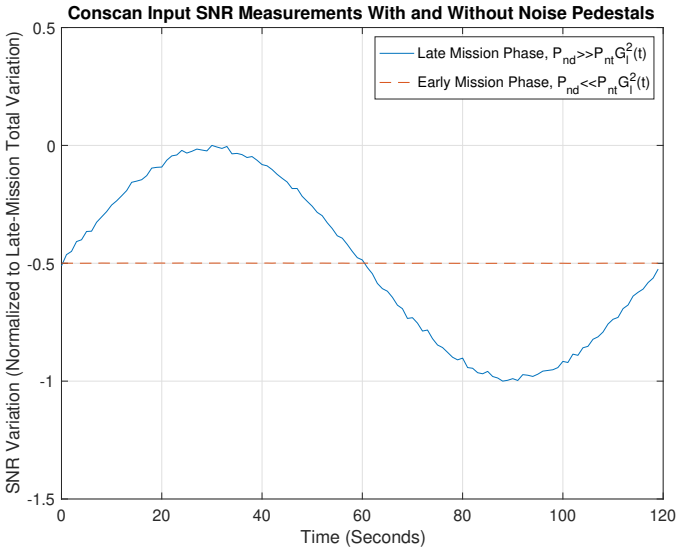

Figure 11. Example plots of the carrier SNR variation (in decibels) measured over one conscan interval, both with $P_{nt}G_l^2(t) \ll P_{nd}$ and $P_{nt}G_l^2(t) \gg P_{nd}$. The two plots are normalized to the total variation amplitude of the SNR in the absence of pedestal effects, when $P_{nt}G_l^2(t) \ll P_{nd}$.

record the AGC gain applied during each carrier power measurement, and then manually divide the recorded gains out of the power measurements. This procedure would remove AGC effects on conscan inputs, but the phase-locked loop and ranging signal processing would still benefit from proper input power.

The success of the direct AGC cancellation approach can be verified in Figure 12. The sinusoid at frequency f_{scan} is extremely clear, and the curve is smoother even than the SNR-normalized measurements in Figure 11 for late-mission conscan. The profile is smoother because direct AGC cancellation ignores statistical variation in noise floor power measurements, unlike SNR normalization.

With direct AGC cancellation, it is ideal to keep the AGC updates synchronized with the power measurements. The interval over which power measurements are taken would act as a filter, generally smearing the effects of AGC gain transitions across pairs of points in the conscan input data. The effect of sloppy synchronization is not studied here; it is instead assumed that AGC gains are perfectly synchronized with power measurements.

In practice, the DSN does not perform direct AGC cancellation during early mission phase because it is difficult to coordinate. The antenna signal $w(t)$ is processed at a very high sample rate, and the analysis and simulation results here have inherently assumed that power measurements are synchronized with gain transitions in AGCs. Synchronizing both analog and digital AGCs with sufficient precision to match the sample rate is difficult. Coordinating power measurements with the timing of gain transitions can also be challenging. Historically, the DSN has used the simpler SNR-normalization technique and simply left the ranging channel turned off during early mission phase.

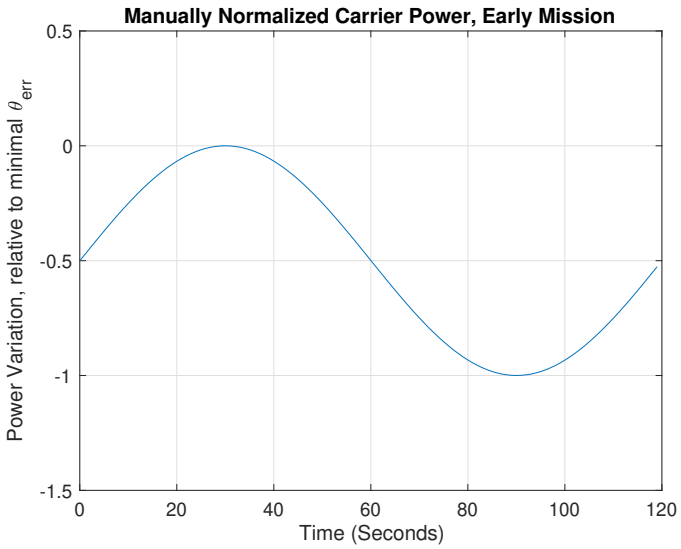

Figure 12. An example plot of power measurements P_{ct} over conscan, with AGC effects cancelled out directly. The sinusoid at f_{scan} is extremely smooth, matching the gain prediction from the antenna model very closely.

VII. Methods

Sample-level models of various DSN subsystems were built to generate the figures in this work and verify the predictions in the discussion. An uplink signal generator model creates a sampled intermediate-frequency (IF) signal to which noise is added to simulate the uplink received signal. The uplink received signal is processed by a model that represents the spacecraft onboard turnaround ranging channel, and more noise is added to the output to form a model of $y\,(t)$.

This simulated downlink signal is upsampled by a model for the downlink tracking and telemetry processing, which operates in iterations on data scaled according to the antenna pointing error approximation in (12). Once each second's worth of data is scaled and upsampled, the signal is shifted to baseband and fed through a digital phase-locked loop. After the carrier and noise power are estimated from the baseband signal model for $w\,(t)$, each second's SNR or directly AGC-cancelled carrier power measurement is recorded. As mentioned previously, direct AGC cancellation is much simpler in a simulation than a real-time processor.

Although the true pointing correction estimator is more complicated, a model of the estimator can run with a single array of 120 SNR or power measurements. Accuracy of the estimator is directly linked to the quality of input data, so this study relies on plots of the power or SNR measurements to demonstrate that direct AGC cancellation works as an alternative to SNR-based normalization.

VIII. Future Work

Direct AGC cancellation on carrier power measurements is an effective way to handle pedestal effects, but the synchronization of AGC gain transitions with carrier power measurements can be tricky. SNR-normalization is convenient because it involves data collected at only one point, the phase-locked loop output. Later efforts could reasonably investigate the impact of sloppy synchronization between AGC transitions and carrier power measurements. It could also be valuable to develop a scheme for automatically identifying a suitable transition from direct AGC cancellation to SNR-based normalization.

Direct AGC cancellation is not the only way to mitigate AGC effects in the presence of a strong noise pedestal. The downlink signal processor could alternatively measure noise power with a signal that's filtered to ignore the noise pedestal, perhaps with a high-pass filter. Designing this high-pass filter automatically for various missions may take some effort, but there would be no requirement to record AGC gains, nor synchronize their timing with power measurements.

IX. Conclusion

This paper presents descriptions of DSN ranging signals and conscan as a precursor to discussing noise pedestal effects from the former on the latter. To give a succinct, clear description of noise pedestals, discussion is limited to pseudonoise ranging in turnaround channels, and contrasts are drawn between power spectra with and without noise pedestals. A description of conscan follows, with some discussion about how SNR measurements automatically ignore the effects of AGCs in the absence of noise pedestal effects. With descriptions of both PN turnaround ranging and conscan, noise pedestal effects on conscan are explained. Noise pedestal effects are linked to poor input data for conscan, with counterexamples indicating clean conscan input when pedestal effects are minimal. Finally, a description of direct AGC cancellation is presented, along with data to indicate that it allows conscan to work properly in the presence of a severe noise pedestal.

Acknowledgments

The author would like to thank Jeff Berner for his continued instruction and guidance in developing this work. Additional credit for help in gathering resources goes to Pete Kinman, Zaid Towfic, Robert Navarro, Chau Buu, Scott Bryant, Jeff Buchanan, Bruce Parham, and Barzia Tehrani.

References

[1] P.W. Kinman and J.B. Berner. Two-way ranging during early mission phase. In *2003 IEEE Aerospace Conference Proceedings (Cat. No.03TH8652)*, volume 3, pages 3_1441–3_1455, 2003.

[2] Joseph H. Yuen and Jet Propulsion Laboratory (U.S.). *Deep space telecommunications systems engineering / edited by Joseph H. Yuen.* Plenum Press New York, 1983.

[3] Kinman, P., "Pseudonoise and Regenerative Ranging," Module 214, Rev. B, July 17, 2019, in *DSN Telecommunications Link Design Handbook*, DSN No. 810-005, Jet Propulsion Laboratory, Pasadena, CA. https://deepspace.jpl.nasa.gov/dsndocs/810-005/214/214B.pdf

[4] Slobin, S., "34-m BWG Stations Telecommunications Interfaces," Module 104, Rev. N, June 7, 2021, in *DSN Telecommunications Link Design Handbook*, DSN No. 810-005, Jet Propulsion Laboratory, Pasadena, CA. https://deepspace.jpl.nasa.gov/dsndocs/810-005/104/104N.pdf

APPENDIX

I. Derivation of the θ_{err} Approximation

First assume that all angles θ are sufficiently small that the portion of the unit sphere that conscan spans is nearly planar. Under this assumption, the angle θ_{err} can be derived as the linear distance between the spacecraft location and an arbitrary point on the circle. Next, let the origin of the coordinate system used for this derivation be the intersection of the cone's center axis with the unit circle. The center of the cone is the defined vertex for all angles ϕ, so defining it as the origin allows a derivation in polar coordinates, with ϕ denoting the direction and θ denoting the radius. As a final precursor, note that $\phi_{s/c}$ is a constant over the conscan operation. Subtracting $\phi_{s/c}$ from any appearance of ϕ_{scan} in the final formula for θ_{err} is equivalent to rotating the reference axis so the spacecraft appears at $\theta_{s/c}$ radians along it.

With the spacecraft on the reference axis and the conscan circle centered at the origin, it is straightforward to build Cartesian coordinates for both spacecraft and antenna direction at ϕ_{scan}.

$$\phi_{rel} = \phi_{scan} - \phi_{s/c} \tag{23}$$

$$\left(x_{s/c}, y_{s/c}\right) = \left(\theta_{s/c}, 0\right) \tag{24}$$

$$\left(x_{scan}, y_{scan}\right) = \left(\theta_{scan}\cos\left(\phi_{\mathrm{rel}}\right), \theta_{\mathrm{scan}}\sin\left(\phi_{\mathrm{rel}}\right)\right) \tag{25}$$

Plug these points into the typical Pythagorean formula for distance.

$$\theta_{err}^2 = \left(\theta_{scan}\cos\left(\phi_{\mathrm{rel}}\right) - \theta_{\mathrm{s/c}}\right)^2 + \theta_{scan}^2\sin^2\left(\phi_{\mathrm{rel}}\right) \tag{26}$$

Factor θ_{scan}^2 out of every term and simplify, applying the Pythagorean Theorem to the squared trigonometric functions.

$$\theta_{err}^2 = \theta_{scan}^2\left(1 - 2\frac{\theta_{s/c}}{\theta_{scan}}\cos\left(\phi_{\mathrm{rel}}\right) + \frac{\theta_{\mathrm{s/c}}^2}{\theta_{\mathrm{scan}}^2}\right) \tag{27}$$

Finally, the radical will encompass everything except the factor of θ_{scan}^2.

$$\theta_{err} = \theta_{scan}\sqrt{1 - 2\frac{\theta_{s/c}}{\theta_{scan}}\cos\left(\phi_{\mathrm{rel}}\right) + \frac{\theta_{\mathrm{s/c}}^2}{\theta_{\mathrm{scan}}^2}} \tag{28}$$

Substitute the definition of ϕ_{rel} to complete the proof.