NISTIR 8073

Simulated Sinewave Testing of Data Acquisition Systems using Sine Fitting and Discrete Fourier Transform Methods
Part 1: Frequency Offset, Random, Quantization, and Jitter Noise

Jon Geist
M. Yaqub Afridi

This publication is available free of charge from:
http://dx.doi.org/10.6028/NIST.IR.8073

National Institute of Standards and Technology
U.S. Department of Commerce

NISTIR 8073

Simulated Sinewave Testing of Data Acquisition Systems using Sine Fitting and Discrete Fourier Transform Methods Part 1: Frequency Offset, Random, Quantization, and Jitter Noise

Jon Geist
M. Yaqub Afridi
Semiconductor and Dimensional Metrology Division
Physical Measurement Laboratory

This publication is available free of charge from:
http://dx.doi.org/10.6028/NIST.IR.8073

July 2015

U.S. Department of Commerce
Penny Pritzker, Secretary

National Institute of Standards and Technology
Willie May, Under Secretary of Commerce for Standards and Technology and Director

Simulated Sinewave Testing of Data Acquisition Systems using Sine Fitting and Discrete Fourier Transform Methods Part 1: Frequency Offset, Random, Quantization, and Jitter Noise

ABSTRACT

This report studies the effect of frequency offset, quantization error, random additive noise, and random phase jitter on the results of sine fitting and performing Discreet Fourier Transforms (DFT) of measurements of sinewaves with Data Acquisition Systems (DAS). Quantization error is the only effect of non-linearity that is considered.

The majority of the simulations are carried out for full-scale signals on a 16-bit DAS, but the simulations do not simulate the effect of clipping on portions of noisy signals that require more than 16 bits. Instead, the signals are rounded in a way that provides the minimum number of bits required to represent the integer part of the signal. The results of this approach are easier to interpret and simple corrections for this shortcoming are given. Only quantization error is studied as a function of available bits, ranging from 8- to 64- bits in increments of 8 bits.

Among other results, it is shown that probability plots of the residuals of fits of sinewaves to simulated distorted sinewaves can distinguish among random additive noise, random jitter, and frequency offset error when the effect of quantization noise is small compared to the effects of these other sources of distortion.

Acknowledgements:

This work was partially supported by the U. S. Army WIAMan Warrior Injury Assessment Manikin Project

Simulated Sinewave Testing of Data Acquisition Systems using Sine Fitting and Discrete Fourier Transform Methods Part 1: Frequency Offset, Random, Quantization, and Jitter Noise

July 17, 2015

Introduction

The purpose of this report is to describe the results of simulations that compare frequency-domain and time-domain sinewave-testing metrics for embedded data acquisition systems (DAS) that will be used in anthropomorphic test devices (ATD) subjected to pyroshock. But these results are applicable in a much wider range of applications. The metrics are inspired by those described in Institute of Electrical and Electronics Engineers Standard 1057 (IEEE 1057) for Digitizing Waveform Recorders[1, 2], but have been modified slightly for harmonization with standard statistical concepts. The simulation results give insights into the interactions of different sources of error in DAS measurements of sinewaves and other signals. Part 1, which is given here, is restricted to frequency error, quantization error, random additive noise, and random phase jitter noise. Quantization is the only source of non-linearity that will be treated here, and are the only distortion investigate for available bits ranging from 8 to 64 bits. Non-linearities will be treated in more detail in a future Part.

For the purposes of the simulations described here, consider an ideal sinewave described in terms of peak-amplitude analog-to-digital-converter (ADC) count A, frequency f in Hz, and phase angle ϕ in radians relative to zero phase at time $t = 0$

$$v(t) = A\cos(2\pi f t + \phi). \tag{1}$$

Also consider a non-ideal implementation of that ideal sinewave,

$$s(t) = A\cos(2\pi f_1 + \phi) + \sigma_N\, r_N(t), \tag{2}$$

where t is time, f_1 is the frequency of the non-ideal sinewave, $r_N(t)$ is a random number drawn from a standard normal distribution (zero mean, unit standard deviation), and σ_N is the standard deviation of the noise voltage level.

Assume that the non-ideal sinewave is sampled by a DAS at a nominal sampling frequency f_s, at times

$$t_n = n\Delta t + \sigma_J\, r_J(n), \tag{3}$$

where $n = 1, \ldots, M'$, M' is the number of samples, $\Delta t = 1/f_s$ is the average time between two adjacent samples, $r_J(n)$ is a random number drawn from a standard normal distribution, and σ_J is the standard deviation of the time-base jitter.

In this case, the discrete non-ideal sinewave is given by

$$x[n] = Q_b(s(t_n)) = Q_b\big(A\cos(2\pi\, f_1\, t_n + \phi) + \sigma_N\, r_N(t_n)\big), \tag{4}$$

for $n = 1, \ldots, M'$, where b is the number of bits available to represent the peak-to-peak ADC count $2A$, and Q_b is the function that quantizes the signal into the appropriate bit pattern.

Next, define the relative offset of the sinewave frequency f_1 from f as

$$\epsilon_F = \frac{f_1 - f}{f}, \tag{5}$$

which is a useful dimensionless parameter to express f_1 in eq. 4 in terms of f in eq. 1. With this definition, the parameter ϵ_F and the relative standard deviations σ_N/A and $\sigma_J/\Delta t$ are physically meaningful dimensionless parameters that are convenient for expressing the variation of IEEE 1057 noise and distortion metrics with frequency offset error, random noise, and jitter.

For the purposes of this report, the quantization function Q_b was implemented by rounding the argument $s(t_n)$ of Q_b in eq. 4 to the nearest integer. The error introduced by this overly-simplified model of signal-clipping, which is easily corrected by signal scaling as shown later in this report, make it much easier to understand the relations among the errors caused by frequency error, additive noise, phase jitter, and number of bits.

It is important to understand that the DFT is extremely sensitive to the frequency offset error ϵ_F. DFTs were calculated from sinewaves defined by eq. 4 and 3 for sinewaves having the same parameters except that $\epsilon_F = 0$ in Fig. 1 and $\epsilon_F = 10^{-6}$ in Fig. 2.

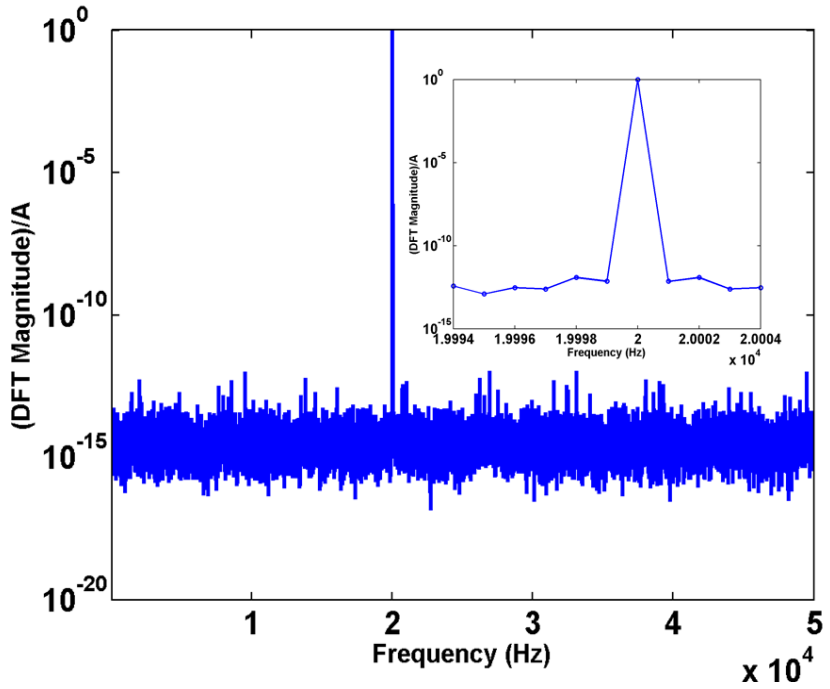

Figure 1: The DFT of a sine wave that was calculated with the MATLAB `fft` function. The sinewave and sampling frequencies were, respectively, $f_1 = 20,000.0$ and $f_s = 100,000.0$ Hz to within the precision of MATLAB double-precision numbers. Only the first half of the DFT, which was plotted at 50,000 frequencies from 0 to 49,999 Hz, is shown. The inset shows the shape of the peak at higher frequency resolution.

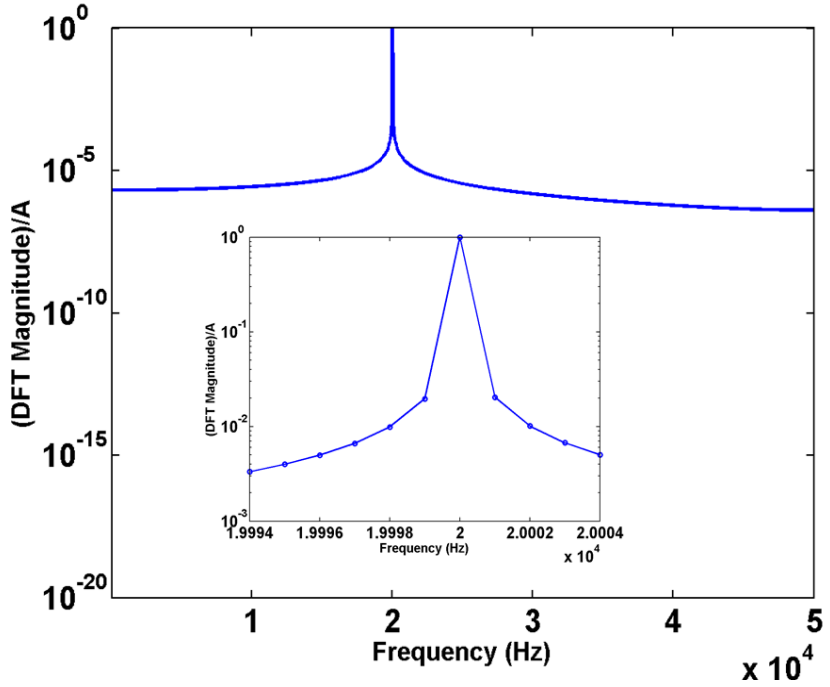

Figure 2: The DFT of a sinewave having the same parameters as the sinewave whose DFT is shown in Figure 1 except that the frequency of the sinewave was 1 part in 10^6 greater than that of the sinewave used for that figure. Specifically, the parameters $\sigma_J/\Delta t, \sigma_N/A$, and ϵ_F were all set to zero for both figures, and only ϵ_F was changed to from zero to 10^{-6}, which increased the frequency f_1 from 20,000 Hz to 20,000.02 Hz for this figure. The inset shows the shape of the peak at higher frequency resolution.

Clearly a 1 part in 10^{-6} time-base error in measurements of the response of ATDs to pyroshock is negligible. So from the point of view of that application, the sinewaves whose DFTs are shown in Figs. 1 and 2 are equally good, which is also the conclusion that would be obtained from four-parameter sine-fitting in the time domain as described in Annex A of IEEE 1057[3]. However, because the DFT does not treat them equally, the DFT of Fig. 2 cannot be used to characterize DAS performance for this application. Windowed and zero-padded DFTs, as described in Section 4.5.2 of IEEE 1057[6], are the solution when ϵ_F is large enough to interfere with other pertinent information in the DFT.

Noise and Distortion Metrics

IEEE 1057 defines the noise and distortion metric (NAD) of a digitized sine wave in both the time[4] and the frequency[7] domains. For the simulations described here, the theoretical value of NAD in the time domain is approximated by

$$\mathrm{NAD} = \sqrt{\frac{1}{M'-1}\sum_{n=1}^{M'}\left[x[n]-v[n]\right]^2} = \sigma_{x,v} \tag{6}$$

where $x[n]$ is the non-ideal digitized sinewave defined in eq. 4, and $v[n]$ is given by

$$v[n] = A\cos(2\pi\, f\, n\Delta t + \phi), \tag{7}$$

which is the ideally sampled implementation of eq. 1, and $\sigma_{x,v}$ is the estimated standard deviation of the distribution from which the values of r_N in eq. 4 are randomly drawn. The prefactor $1/(M'-1)$ could be replaced by $1/M'$ to be consistent with IEEE 1057[4], but establishing the identity with the standard deviation is preferable.

With this definition, in the case where the contribution of all other sources of distortion are negligible compared to that of random noise,

$$\mathrm{NAD} = \lim_{M\to\infty} \quad \overline{\frac{1}{M'-1}\sum_{n=1}^{M}\left[\sigma_N\, r_N(t_n)\right]^2} = \sigma_N \tag{8}$$

because the random variable $r_N(t_n)$ in eq. 4 is drawn from a normally-distributed random distribution whose standard deviation is 1.

In the time domain: Define NAD_x as[1]

$$\mathrm{NAD}_x = \min\left(\quad \overline{\frac{1}{M'-3}\sum_{n=1}^{M} x[n]-x'[n]\big)^2} \right) = \sigma_{x,x}$$

$$\approx \min\left(\quad \overline{\frac{1}{M'}\sum_{n=1}^{M} x[n]-x'[n]\big)^2} \right) = \mathrm{NAD}_{1057} \tag{9}$$

where $x[n]$ are the M' values of the sine wave defined in eq. 4, and NAD_{1057} is the definition for NAD that is used in IEEE 1057.

The quantity $\sigma_{x,x}$ in eq. 9 is the standard deviation of the distribution of the residuals of a least squares fit of

$$x'[n] = Q_b \quad A_1\cos(2\pi\, f_1\, t_n) + B_1\sin(2\pi\, f_1\, t_n) + C\big), \tag{10}$$

to the sampled sine wave $x[n]$ defined in eqs. 4. The quantities A_1, B_1, and C are the free parameters that are varied to minimize NAD_x. The constant term C is needed even though no constant term occurs in eq. 4 because the mean of a finite sample of the normally distributed random numbers $r_N(t_n)$ is not identically zero even though the mean of the distribution from which the $r_N(t_n)$ are drawn is identically zero.

The prefactor $1/(M'-3)$ in the first line of eq. 9 is appropriate because three degrees of freedom have been used in fitting. Therefore, the theoretical relation of NAD_x to NAD and NAD_{1057} is given by

$$\mathrm{NAD}_x = \sqrt{\frac{M'-1}{M'-3}}\mathrm{NAD} = \sqrt{\frac{M'}{M'-3}}\mathrm{NAD}_{x1057}. \tag{11}$$

The quantity $b' \gg b$ in eq. 10 is the number of bits available for representing $x'[n]$ in the computer program being used to carry out the least-squares fitting procedure. Since b in eq. 4 is accurately known for a modern DAS, it is not adjusted to improve the fit, nor is this likely to prove practical for determining the effective number of bits (ENOB) available from the DAS.

The results of the fitting procedure give time-domain estimates, denoted here by

$$A_x = \sqrt{A_1^2 + B_1^2} \tag{12}$$

from eq. 10 and NAD_x from eq. 9. Later in this report, these will be compared with the theoretical values of A and NAD from eqs. 7 and 8 for various values of ϵ_F, σ_N/A, $\sigma_J/\Delta t$, and b.

[1] The subscript x will be used to identify quantities that are calculated in the time domain.

In the frequency domain: Start by adopting the MATLAB definition and normalization of the DFT, which is given by

$$Y[k] = \sum_{n=1}^{M} x[n] \exp\left(\frac{-j2\pi[n-1][k-1]}{M}\right), \tag{13}$$

where $k = 1, \ldots, M$, where M is an even positive integer. With this normalization, the peak amplitude $A_Y = |Y[k_A]|$ is $M|Y[k_A]|/2$.

To proceed, set $Y[1] = Y[k_A] = Y[M - k_A] = Y[M] = 0$, and define

$$\mathrm{NAD}_Y = \sqrt{\left[\frac{1}{M[M-6]} \sum_{k=1}^{M} \left|Y[k]\right|^2\right]} \approx \mathrm{NAD}_{Y1057}. \tag{14}$$

Note that NAD_{Y1057} differs from NAD_Y in that it has $M - 3$ instead of $M - 6$ in its prefactor. The use of $M - 6$ is justified on statistical grounds because six degrees of freedom have been removed from the sum in NAD_Y, namely the first and last points of the DFT (which in general are non-zero for the same reason that C was needed in $x'[n]$ in eq. 10), the fundamentals at k'_A and $M - k'_A$, and two degrees of freedom for the two identical sums from $k = 1, \ldots, N = M/2$ and $k = N + 1, \ldots, M$ in eq. 14.

It is often inconvenient for A_Y to depend upon the number data points M in the time series. In this case, it is convenient to use only the first half of the DFT in eq. 13 and express as

$$\begin{aligned}\mathrm{NAD}_Y &= \sqrt{\frac{1}{M[M-6]} \sum_{k=1}^{M} \left|Y[k]\right|^2} = \sqrt{\frac{2}{M[M-6]} \sum_{k=1}^{M/2} \left|Y[k]\right|^2} = \sqrt{\frac{2M^2}{4M[M-6]} \sum_{k=1}^{N} \left|\frac{2Y[k]}{M}\right|^2} \\ &= \sqrt{\frac{M}{2[M-6]} \sum_{k=1}^{N} \left|\frac{2Y[k]}{M}\right|^2} = \sqrt{\frac{N}{2[N-3]} \sum_{k=1}^{N} \left|X[k]\right|^2} = \mathrm{NAD}_X,\end{aligned} \tag{15}$$

where

$$X[k] = \frac{2Y[k]}{M} = \frac{1}{N} \sum_{n=1}^{M} x[n] \exp\left(\frac{-j2\pi[n-1][k-1]}{M}\right) \tag{16}$$

for $k = 1, \ldots, N$ and

$$A_X = |X[k_A]| = \sqrt{X_{re}[k_A]^2 + X_{im}[k_A]^2} \tag{17}$$

as desired, where the re and im subscripts denote the real and imaginary parts of the DFT. This is a particularly convenient normalization when $M = f_s/U$ is an even integer, f/U is an integer, where U is the frequency unit being used (Hz, kHz, etc.). These two constraints on f, f_s, and M are necessary to avoid DFT artifacts that can obscure important information contained in the DFT even when $\epsilon_F = 0$.

NAD calculated from the DFT noise floor: The inset in Fig. 3 shows the effect (blue curve in inset) of noise with $\sigma_N = 1$ on a 1 ms segment of the digitized sinewave of eq. 4 (black curve in inset) at frequency $f_1 = 20$ kHz with $f_s = 100$ kHz, $\sigma_J = \epsilon_F = 0$, and $b = 16$. The main figure shows the effect of the same noise on the scaled DFT of the same sinewave segment (blue curve in main figure). Notice how the clean sinewave in the time domain (black curve in inset) is obscured by the noise when the SNR is 1, in striking contrast to how the amplitude of the same sine wave rises out of the noise floor of the DFT (blue curve in main figure). This illustrates the important fact that the noise floor of the DFT is not NAD_X, but is approximately $\mathrm{NAD}_X/\sqrt{M}$, which can result in major underestimates of NAD if forgotten. The green curve is the noise floor of the scaled DFT increased by the factor required to make it comparable to NAD_X.

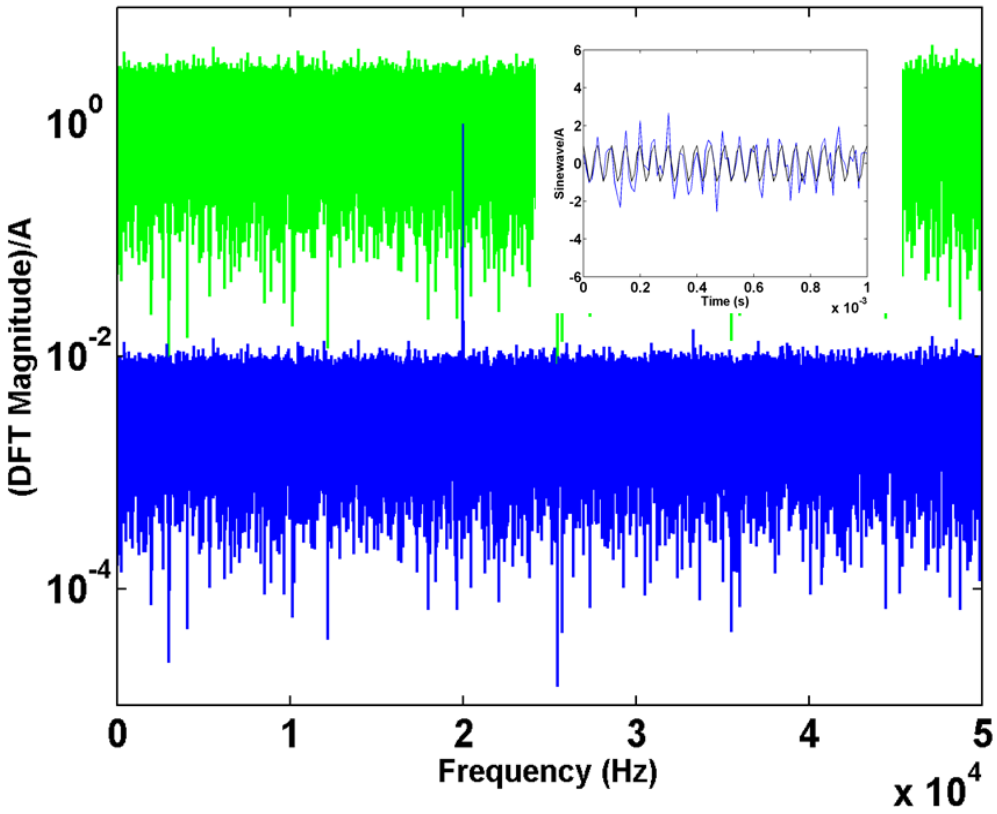

Figure 3: Comparison of the effect of noise giving a SNR equal to 1 on the display of a sinewave in the time domain (inset) and frequency domain (main figure). The green curve is the noise floor of the scaled DFT increased by the factor required to make it comparable to the noise on the sinewave as discussed above. When $x[n] - v[n]$ (eqs. 4 and 7) is drawn from a normally-distributed random distribution, the DFT noisefloor is drawn from a Rayleigh distribution.

SINAD versus Sources of Distortion

This section summarizes the results of setting the additive noise, jitter noise, and frequency offset to zero and the number of available bits to 16 and then independently varying each of the distortion parameters and b to determine the effect on SINAD. In agreement with IEEE standard IEEE 1057[1], define signal-to-noise-and-distortion ratios.

$$SINAD = \frac{A}{NAD}, \tag{18}$$

$$SINAD_x = \frac{A_x}{NAD_x}, \tag{19}$$

$$SINAD_X = \frac{A_X}{NAD_X} \tag{20}$$

for the distorted sinewave in eq. 4, where the theoretical amplitude A is defined in eq. 7, and the estimated amplitudes A_x and A_X are defined in eqs. 12 and 17, respectively. Similarly, the theoretical

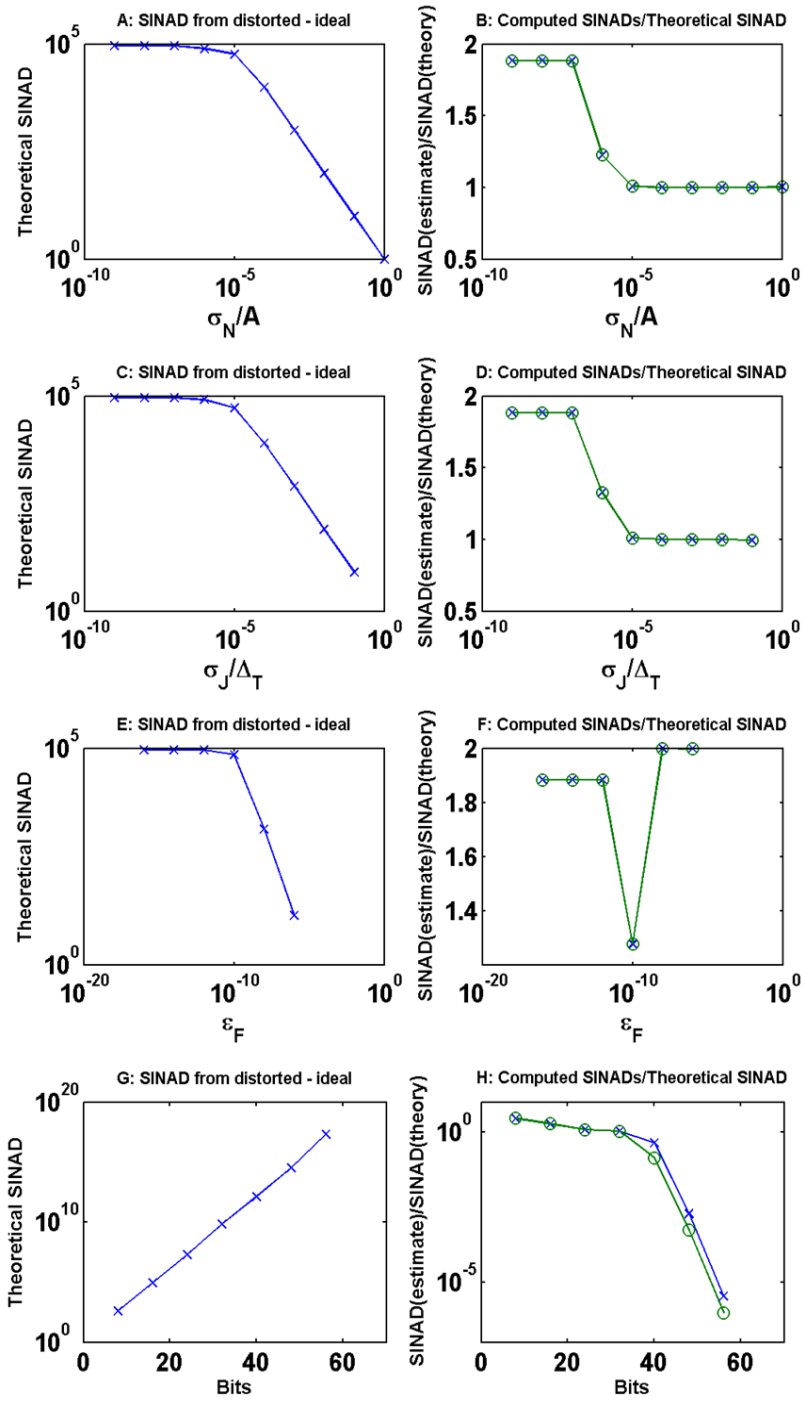

Figure 4: **A:** Theoretical variation of SINAD as a function of the relative, normally-distributed additive random noise parameter σ_N/A. **B:** Ratio of time- and frequency domain estimates of SINAD_x and SINAD_X, respectively, to SINAD as a function of σ_N/A, both for a 16-bit DAS. **C and D:** The same as A and B, but the relative normally-distributed random phase noise $\sigma_J/\Delta t$ is the parameter being varied. **E and F:** The same as A and B, but the relative frequency offset from the target sampling frequency is the parameter being varied. **G and H:** The same as A and B, but the number of bits b is the parameter being varied.

noise-and-distortion metric NAD is defined in eq. 8, and the estimated metrics NAD_x and NAD_X are defined in eqs. 9 and 15, respectively.

The quantities A, A_x/A, A_X/A, $\mathrm{SINAD} = \mathrm{A/NAD}$, $\mathrm{SINAD}_x = \mathrm{A}_x/\mathrm{NAD}_x$, and $\mathrm{SINAD}_X = \mathrm{A}_X/\mathrm{NAD}_X$ were independently calculated for a range of values of the additive noise σ_N, the jitter σ_J, the frequency offset ϵ_F, and the number of available bits b. The results are summarized in Fig. 4 on the previous page. As pointed out previously, the simulations described above do not treat clipping properly. The proper treatment of clipping is described in Appendix **??**.

Frames A and B of Fig. 4 plot the theoretical values of SINAD and the ratios $\mathrm{SINAD}_x/\mathrm{SINAD}$ and $\mathrm{SINAD}_X/\mathrm{SINAD}$, respectively, as a function of σ_N/A from 10^{-9} to 1 in factors of 10 with $b = 16$ and $\sigma_J = \epsilon_F = 0$. Each time the program that calculates the values in Fig. 4 is run with new random values for $r_N(t_n)$ in eq. 4, new values of A_x/A and A_X/A that differ from the old values by much less than σ_N/A are obtained. Both $\mathrm{SINAD}_x/\mathrm{SINAD}$ and $\mathrm{SINAD}_X/\mathrm{SINAD}$ overestimate SINAD by a factor of approximately 1.9 below $\sigma_N = 10^{-7}$ and approximate SINAD quite well above $\sigma_N = 10^{-5}$.

Frames C and D of Fig. 4 plot the same quantities for $\sigma_J/\Delta t$ from 10^{-9} to 0.1 in factors of 10 with $b = 16$ and $\sigma_N/A = \epsilon_F = 0$. These results are qualitatively similar to those in frames A and B, the most notable difference occurring in the transition region around $\sigma_N/A = \sigma_J/\Delta t = 10^{-6}$. No result is shown for $\sigma_J/\Delta_t = 1$ because these results were not reproducible, and in many cases the least squares fitting program did not even converge.

Frames E and F of Fig. 4 plot the same quantities for $\epsilon = 10^{-9}$ to 0.1 in factors of 10 with $b = 16$ and $\sigma_N/A = \sigma_J/\Delta t = 0$. Comparison of these results with those in A and C shows that SINAD is much more sensitive to relative frequency-offset errors than to relative amplitude errors and relative jitter errors. Also, SINAD_x and SINAD_X both overestimate SINAD for $\epsilon_F < 10^{-12}$ by a factor of 1.9 as is the case for random additive noise and random jitter below their transition regions at $\sigma_N/A = \sigma_J/\Delta t = 10^{-6}$. The reason for this is that NAD is dominated by the same quantization error associated with $b = 16$ available bits below all three transition regions.

Frames G and H of Fig. 4 plot the same quantities for $b = 8$ to 64 bits in increments of 8 bits with $\sigma_N/A = \sigma_J/\Delta t = \epsilon_F = 0$. SINAD is exponential as expected, but the behavior of SINAD_x and SINAD_X as a function of the number of available bits in frames G and H is very different from that shown in frames BD, and F because the number of bits is a completely different type of parameter. Increases in this parameter increase SINAD whereas increases in the other parameters decrease SINAD. Also the number of available bits was set at 16 for the other simulations, but was varied in this one. Over the range from $b = 8$ to $b = 24$ $\mathrm{SINAD}_x/\mathrm{SINAD}$ and $\mathrm{SINAD}_X/\mathrm{SINAD}$ decrease from approximately 3 to approximately 1, which is also its value at $b = 32$ bits. Above the transition region $b = 24 - 32$ bits, the relative accuracy of these ratios deteriorates rapidly.

Given a measured value of SINAD_x or SINAD_X and based only on the information provided in Fig. 4, it would not be possible to distinguish among any of the source of distortion considered here. However, a closer look at the measured data from which SINAD_x and SINAD_X are calculated provides much more information as shown below.

SINAD versus Additive Noise

Table 1 in Appendix B compares the $A_x/A, A_X/A, SINAD = A/\mathrm{NAD}$, $\mathrm{SINAD}_x = A_x/\mathrm{NAD}_x$, and $\mathrm{SINAD}_X = A_X/\mathrm{NAD}_X$ for times equal to 1 s, 10 s, and 100 s as a function of the relative additive noise σ_N/A from 10^{-9} to 1 with $b = 16$ and $\sigma_J/\Delta t = \epsilon_F = 0$. This table shows that $A/NAD = 1/\sigma_N$ (1/sigN in the table) to within a small fraction of σ_N when $\sigma_N > 10^{-5}$ and that SINAD_X is an excellent approximation to SINAD_x for all conditions. Also, the excellent agreement between the corresponding quantities at the three different simulated times shows that NAD is independent of time for this type of error, which justifies estimating the mean noise as the time average of NAD and the standard deviation uncertainty in the mean noise as NAD/M. This is just one of the reasons for adopting the constraints on f, f_s, and M for the DFT.

A closer look at the $\sigma_N/A = 1$ results: For this simulation, Fig. 5 presents in A: a probability plot of the set of differences $x[n] - v[n]$ for $n = 1, \ldots, M$ from eqs. 4 and 7, respectively, in B: a probability plot of the set of residuals $x[n] - x'[n]$ for $n = 1, \ldots, M$ from the fit defined in eq. 9, in C: a plot of the noisefloor of the DFT of $x[n]$ defined as the set of $X[k]$ for $k = 1, \ldots, N$ in eq. 15 with $X[k_A]$ and $X[1]$ set to zero, and in D: a probability plot of the noisefloor of the DFT of $x[n]$.

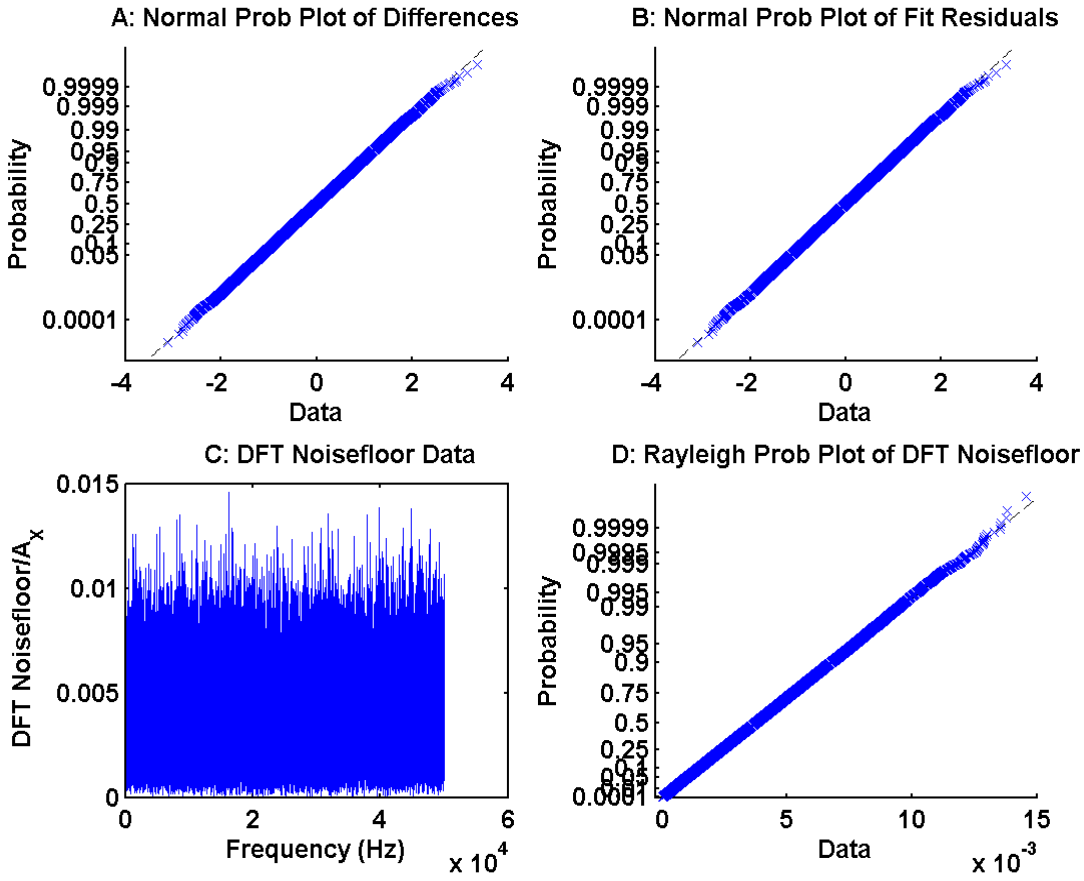

Figure 5: From Top Left moving clockwise: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

The probability plots of the set of differences $x[n]-v[n]$ and the residuals of the fit $x[n]-x'[n]$ are very similar and normally distributed for all practical purposes. This is expected since the data were simulated with normally distributed random noise with a standard deviation of 1, which rendered the quantization noise completely negligible. Similarly, the probability plot of the noisefloor of the DFT follows a Rayleigh distribution very well, as expected for magnitudes calculated from normally distributed data.

A closer look at the $\sigma_N/A = 10^{-9}$ results: For this simulation, Fig. 6 presents the same type of graphs used in Fig. 5 above. These four graph types, which will be used in the remainder of this report

to examine some instructive examples of the simulation results in detail, will be referred to as the four standard graphs.

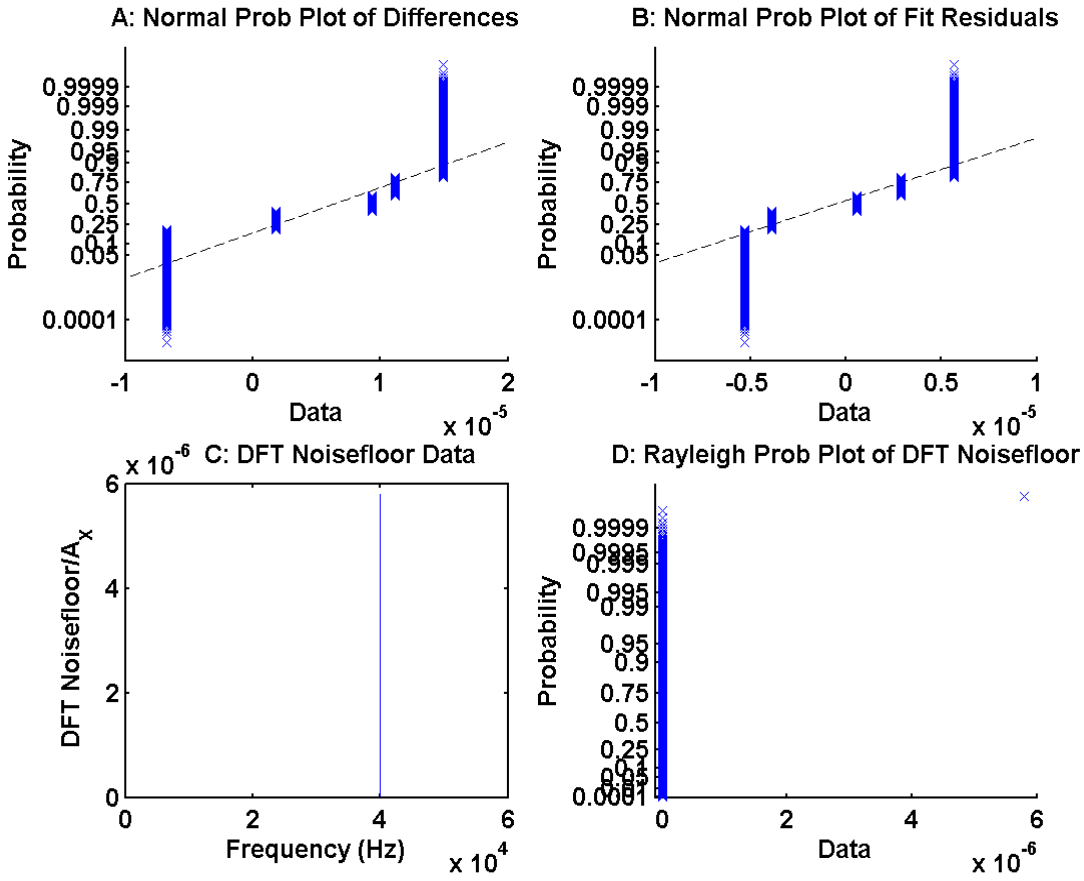

Figure 6: From Top Left moving clockwise: Probability plot of the theoretical difference between the quantization-noise limited sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

None of the probability plots in Fig. 6 contain normally or Rayleigh-distributed data. Because the 20 kHz sinewave is sampled at 100 kHz, the upper left probability plot contains five samples per cycle of the distorted minus ideal sinewave data, $x[n] - v[n]$, which is referred to here as the difference data. The upper right probability plot contains five samples per cycle of the distorted minus fit sinewave data, $x[n] - x'[n]$, which is referred to here as the residual data. Because the sampling and sinewave frequencies are very well synchronized, the same points of the difference and residual sine waves are sampled each cycle, which explains why both probability plots contain many duplicates of only five different values of the sine wave. Because the fit can never perfectly reconstruct the ideal sinewave from the distorted sinewave, the values of the five different residuals in the upper right plot will in general never duplicate the five different values of the differences in the upper left plot. The fact that the range of the residuals is about a factor of two smaller than that of the differences is consistent with $\text{SYNAD}_x \approx 2\text{SINAD}$ in Fig.

4 B. Finally, because the phase of the sinewave relative to the sampling signal is 0.3 radians, for which π is not an integer multiple, none of the five points that are sampled on the sinewave have the same value. For special values of ϕ, some of the points being sampled will be duplicates, and the probability plot will contain fewer different values. For instance, only three values occur when $\phi = 0$.

The DFT noisefloor plotted in the lower left of Fig. 6 doesn't even look like a DFT noisefloor. Instead, it consists of an easily resolved signal at 40 kHz, which is the first harmonic of the 20 kHz sinewave. The true noisefloor is negligible at the scale of this figure. The fact that all of the energy in the noisefloor occurs at a harmonic of the sinewave frequency shows that quantization noise is caused by the non-linearity of the quantization process rather than a random-noise-like process.

A closer at look the $\sigma_N/A = 10^{-5}$ results: For this simulation, Fig. 7 presents the standard graphs.

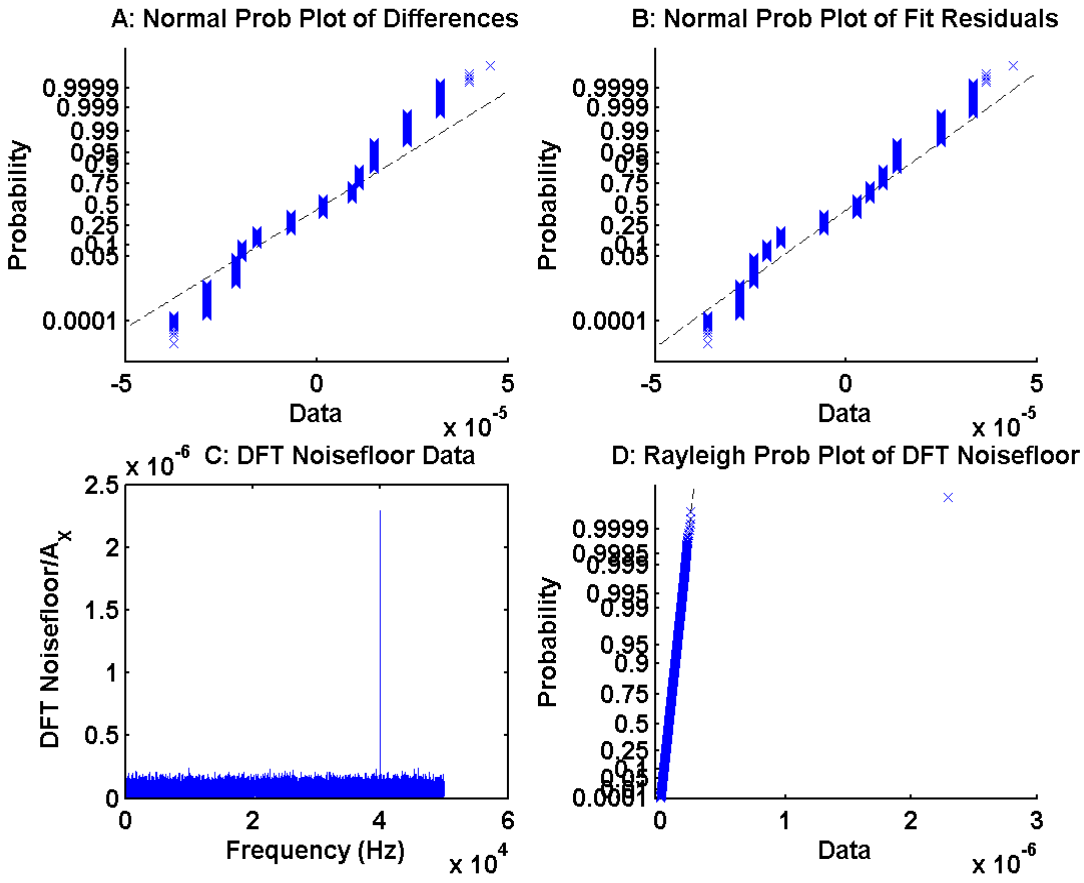

Figure 7: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

This simulation is on the quantization-noise side of the transition from quantization noise limited behavior to normally distributed, additive, random-noise-limited behavior. The probability plots in Fig. 7 still retain the qualitative characteristics of quantization noise, but the range has increased by about

a factor of 30, and 14 different values instead of just 5 are present in the differences and the residuals. The strong signal at the 40 kHz is still evident in the DFT noisefloor data, but unlike the case where $\sigma_N/A = 10^{-9}$, the true noisefloor, which is Rayleigh distributed, is resolved.

A closer at look the $\sigma_N/A = 10^{-4}$ results: For this simulation, Fig. 8 presents the standard graphs. This simulation is on the random-noise side of the transition from quantization noise limited behavior to normally distributed, additive, random-noise-limited behavior.

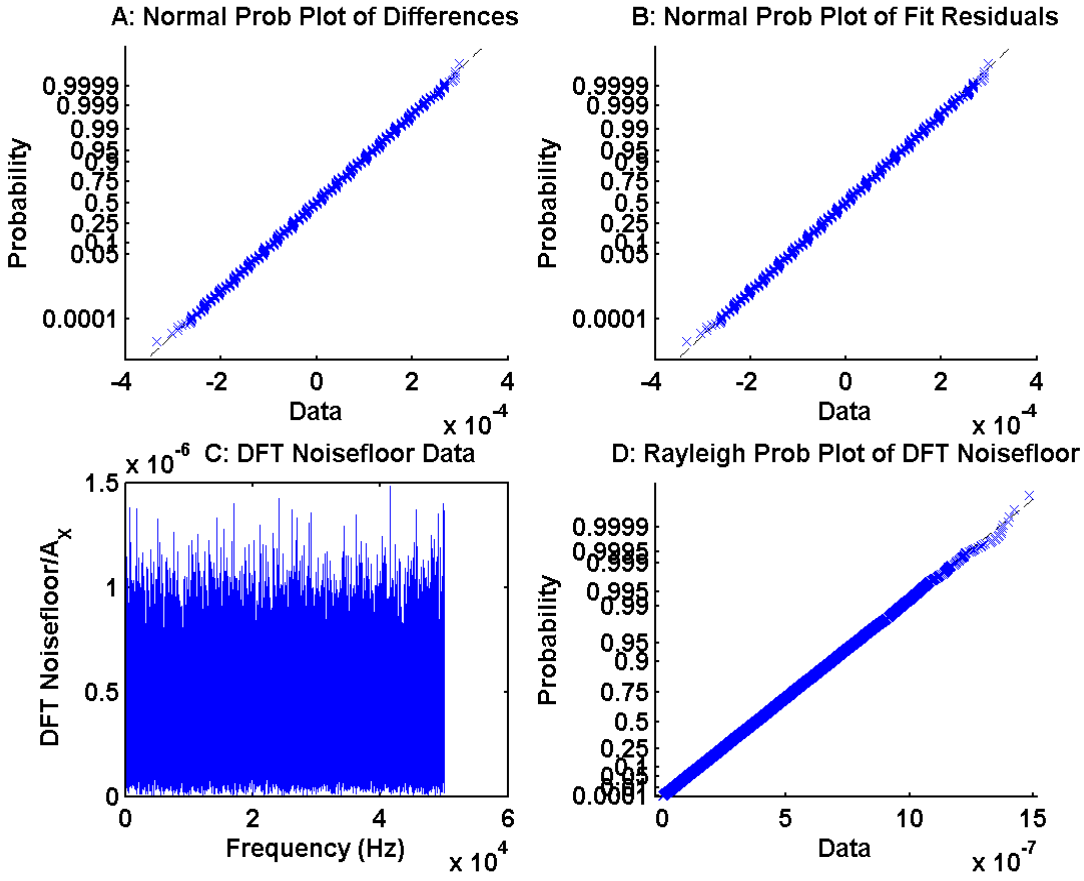

Figure 8: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

Increasing $\sigma_N/A = 10^{-5}$ by one order of magnitude to $\sigma_N/A = 10^{-4}$ completely changed the character of the probability plots as shown in Fig. 8. Not only do all four plots now bear a strong resemblance to the $\sigma_N/A = 1.0$ plots in Fig. 5, but the probability plots of the differences and of the residuals are identical for all practical purposes. Even the fine-grain structure of these two probability plots, which is the contribution of the quantization, looks very similar.

Before reporting the results of a similar simulation as a function of the jitter parameter σ_J/A with $\sigma_N/A = 0$, there is one more important point to consider. Fig. 4B and reference to Table 1 show that

$SINAD_x$ and $SINAD_X$ agree very well for $10^{-9} \leq \sigma_N/A \leq 1$ even though the shapes of the probability plots are very different because one is normally distributed and the other Rayleigh distributed. This is an approximate application of Parseval's energy theorem. It is not exact because different prefactors are used in the definitions of NAD_x and NAD_X. But the differences are quite small for large M or N. On the other hand, even though the probability plots of the differences between the distorted and ideal sinewaves and the residuals of the fit to the distorted sinewave look much more similar than do the probability plots of the DFT noisefloor and the residuals of the fit to the distorted sinewave, they agree very well only for additive noise-limited data, differing by almost a factor of two for quantization "noise"-limited data.

This space intentionally left blank .

SINAD versus jitter

Table 2 in Appendix B compares the $A_x/A, A_X/A, SINAD = A/\mathrm{NAD}$, $\mathrm{SINAD}_x = A_x/\mathrm{NAD}_x$, and $\mathrm{SINAD}_X = A_X/\mathrm{NAD}_X$ for times equal to 1 s, 10 s, and 100 s as a function of the relative jitter $\sigma_J/\Delta t$ from 10^{-9} to 0.1 with $b = 16$ and $\sigma_N/A = \epsilon_F = 0$. The excellent agreement between the corresponding quantities at the three different simulated times shows that NAD is independent of time for this type of

error, which is also true for relative additive noise and which justifies estimating the mean noise as the time average of NAD and the standard deviation uncertainty in the mean noise as NAD/M for jitter as well as random additive noise and also for their combination. However, no correction for clipping such as that given in Appendix A for random noise is required for jitter because jitter is not additive. Instead, it samples the signal at times that are randomly different from when it was supposed to be sampled, but doesn't change its value as additive noise does.

A closer look at the $\sigma_J/\Delta_t = 0.1$ results: For this simulation, Fig. 9 presents the standard graphs.

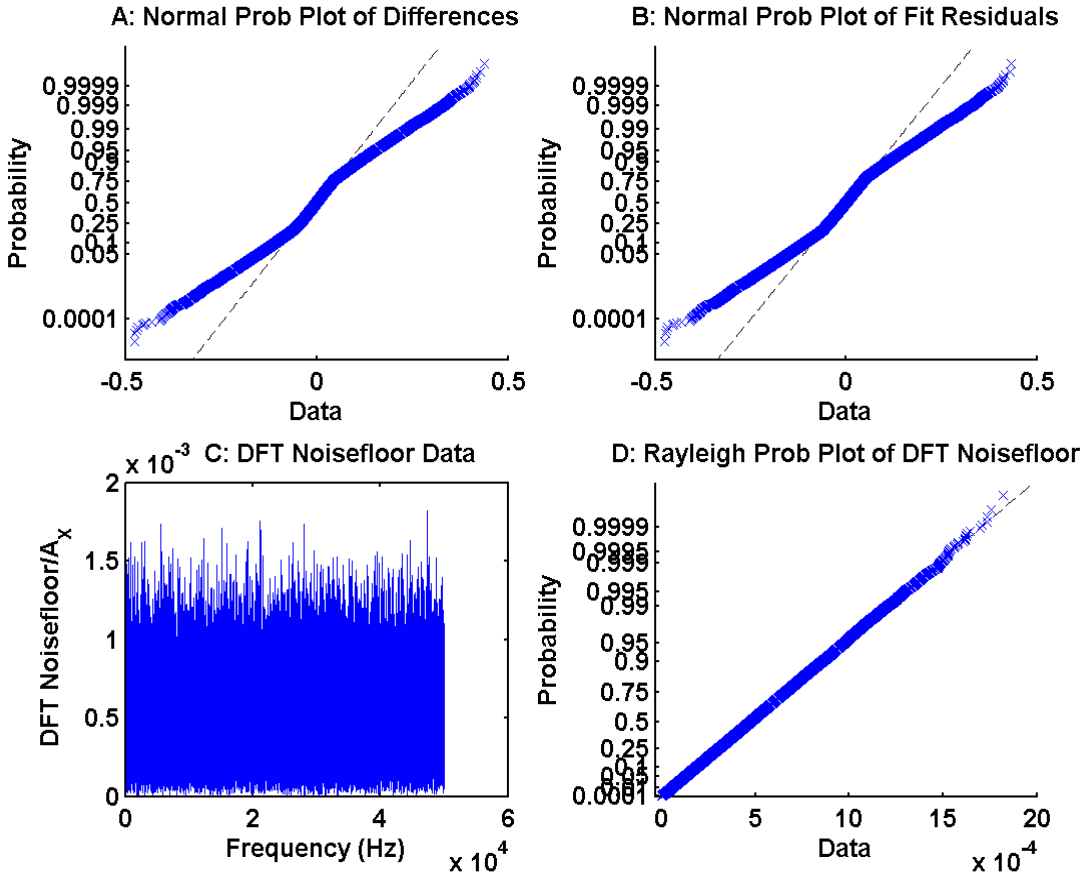

Figure 9: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

As was the case with the SINAD simulation at $\sigma_N/A = 1$, the probability plots of the set of differences $x[n] - v[n]$ and the residuals of the fit $x[n] - x'[n]$ for $\sigma_J/\Delta t = 0.1$ are very similar to each other. Furthermore, the probability plots of the noisefloors of the $\sigma_N/A = 1$ and $\sigma_J/\Delta t = 1$ DFTs are very similar. On the other hand, the probability plots of the SINAD(σ_J/Δ_t) and SINAD($\sigma_N/A = 1$) time-series data are quite different, the latter being normally distributed while the former has a small central region that is normally distributed between two nominally symmetric long-tailed distributions. This is a

result of the non-additive nature of jitter. Even though the sampling times t_n are normally distributed, the resulting errors in the signal amplitude are not normally distributed. Adding a random time to a sampling time produces a very different signal-amplitude error depending upon whether it occurs near a extremum or a zero crossing of the signal.

The other major difference between the simulation results for jitter and for additive noise is that above the transition region from quantization-dominated behavior, SINAD = $A/\text{NAD} \approx 0.8\Delta t/\sigma_J$ for the former and SINAD = $A/\text{NAD} \approx A/\sigma_N$ for the latter.

A closer look at the $\sigma_J/\Delta_t = 10^{-9}$: The results for $\sigma_N/\Delta_J = 10^{-9}$, cannot be distinguished from the results for $\sigma_N/A = 10^{-9}$ that are shown in Fig. 6 because the 16-bit quantization noise overwhelms the effects of $\sigma_N/\Delta_J = 10^{-9}$ in Fig. 4 (C) just as it overwhelms the effect of $\sigma_N/A = 10^{-9}$ in Fig. 4 (A).

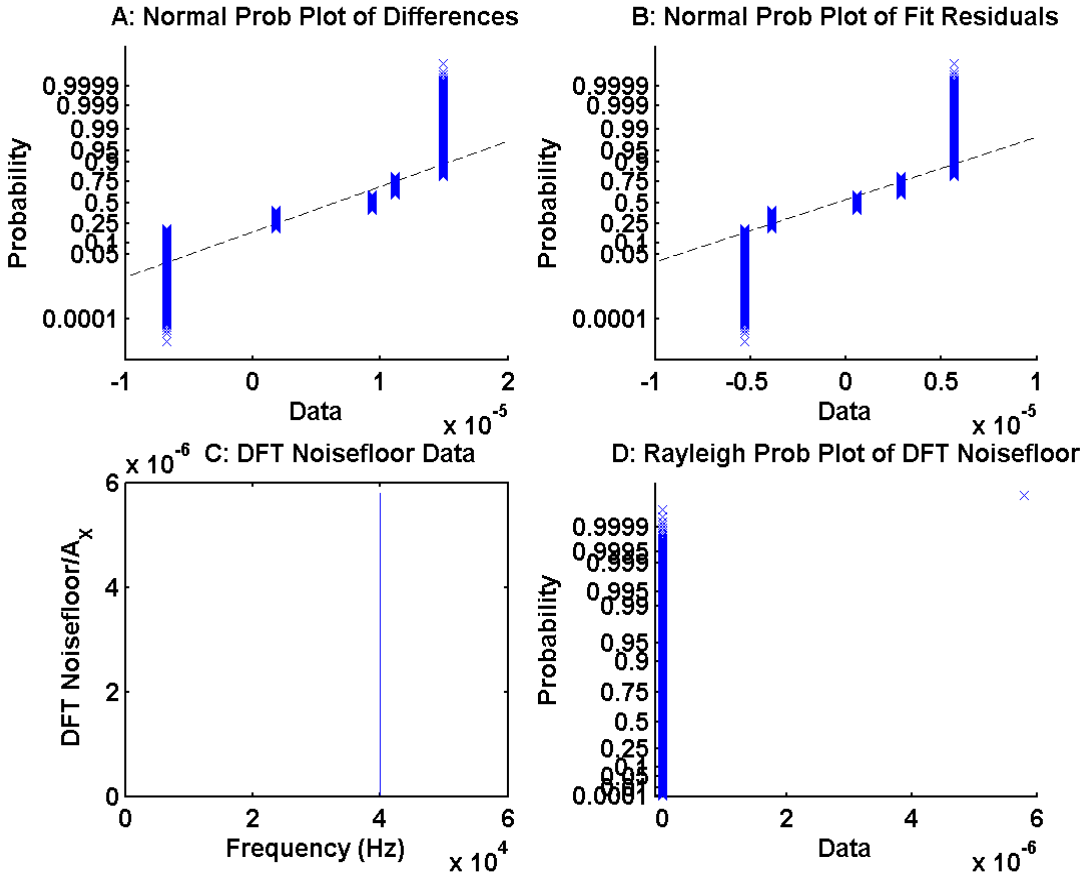

Figure 10: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

A closer look at the $\sigma_J/\Delta t = 10^{-5}$ results: For this simulation, Fig. 11 presents the standard graphs. This simulation is on the quantization-noise side of the transition from quantization noise limited behavior

to normally distributed, random, additive, noise-limited behavior.

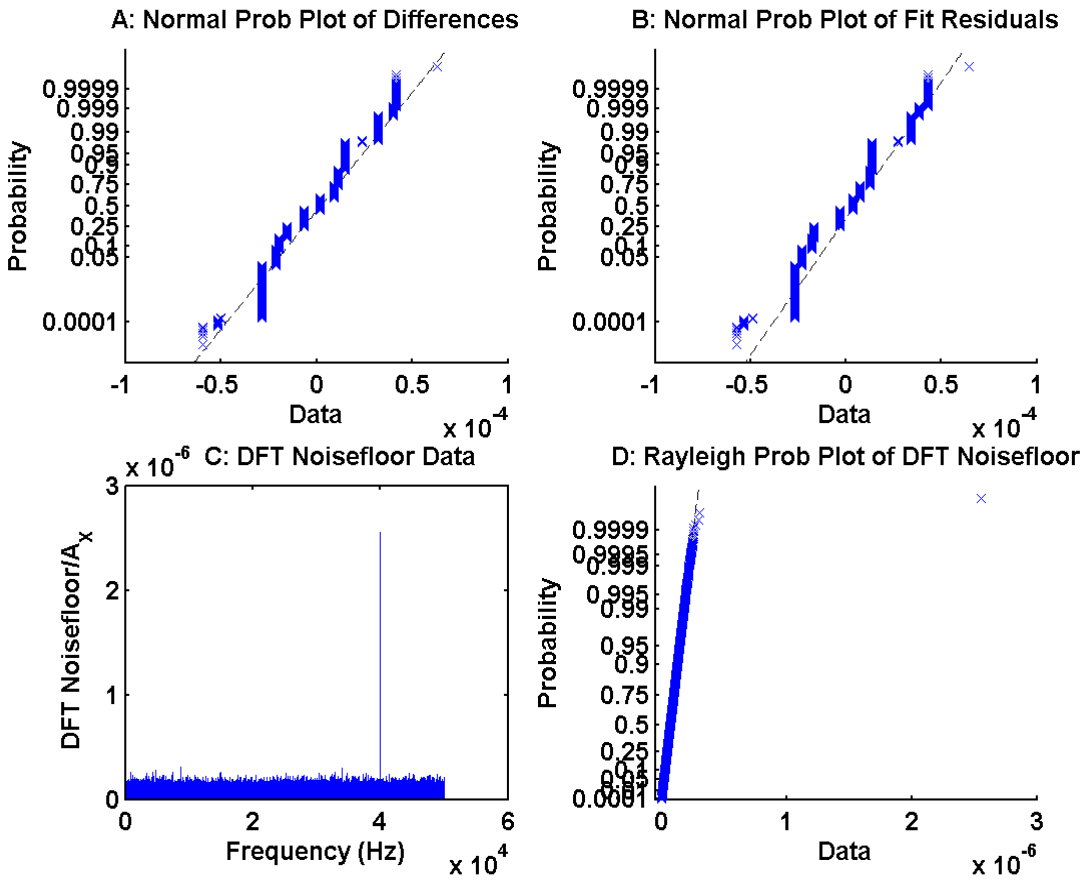

Figure 11: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

The probability plots in Fig. 11 have progressed a little further toward the random-noise limit than have those for σ_N/A in Fig. 7. This is also evident in Fig. 4 B versus Figs. 4 D. However, if these plots were shown without a description of how they were generated, one might erroneously conclude that Fig. 11 was created by additive noise and that Fig. 7 was created by jitter based on their similarities to Figs. 5 and 9, respectively.

A closer look at the $\sigma_J/\Delta t = 10^{-4}$ results: For this simulation, Fig. 12 presents the standard graphs. This simulation is on the random-noise side of the transition from quantization noise limited behavior to normally distributed, additive, random-noise-limited behavior. In contrast to the case when quantization "noise" is the dominant noise source, random additive noise can be distinguished from random jitter when one of the latter two is the only other source of noise.

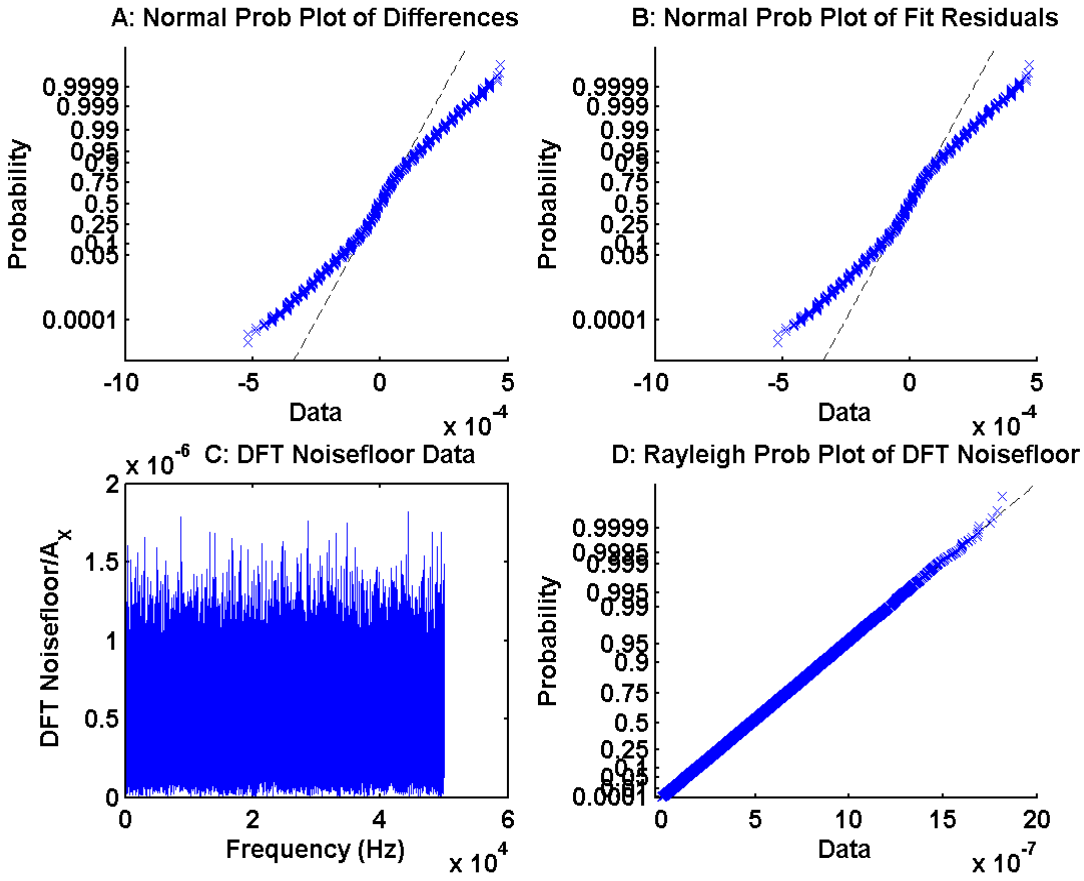

Figure 12: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

SINAD versus relative frequency offset

Table 3 in Appendix D compares the $A_x/A, A_X/A$, SINAD $= A/\text{NAD}$, $\text{SINAD}_x = A_x/\text{NAD}_x$, and $\text{SINAD}_X = A_X/\text{NAD}_X$ for times equal to 1 s, 10 s, and 100 s as a function of the relative frequency-offset error ϵ_F from 10^{-9} to 0.1 with $b = 16$ and $\sigma_N/A = \sigma_J/\Delta t = 0$. Examination of this table shows that $\text{SINAD}(\epsilon_F, t_1) \neq \text{SINAD}(\epsilon_F, t_2)$ when $t_2 \neq t_1$ for ϵ_F above the transition region, which means that errors of this type should not be thought of as noise and cannot be extrapolated over time by statistical means. However, no correction for clipping such as that given in Appendix A for random noise is required for frequency offset error because this source of distortion is not additive.

A closer look at the $\epsilon_F = 10^{-10}$ results: For this simulation, Fig. 13 presents the standard graphs.

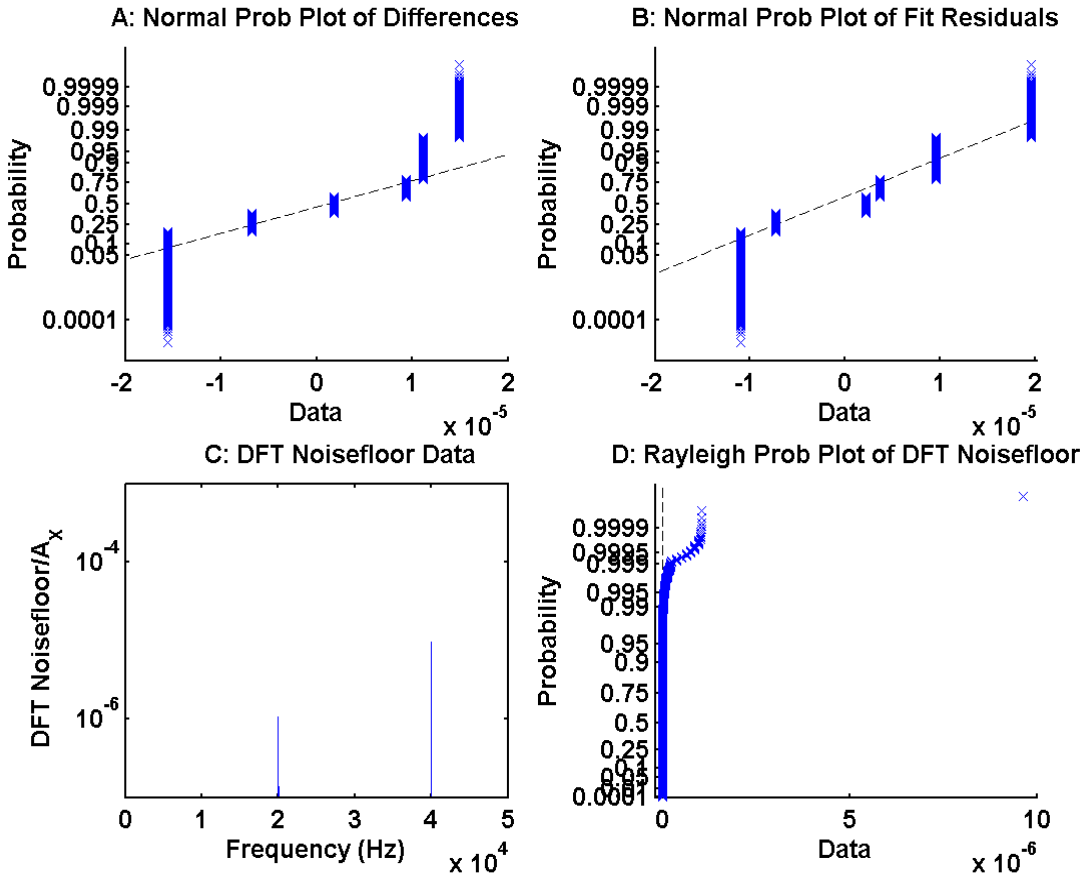

Figure 13: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

These results are clearly on the quantization noise side of the transition from quantization-noise-limited behavior to frequency-offset-limited noise behavior. These plots share some similarities with those for quantization-noise-limited behavior with $\sigma_N/A = 10^{-9}$ in Fig. 6, but the plots here also have

some noticeable differences. Both time-domain probability plots have only a few values that are repeated over multiple cycles, but unlike those in Fig. 6, which each have five values repeated in multiple cycles with a factor of two difference in range, these have six values that are repeated in multiple cycles with only a 30 % difference in range. Also, both noisefloor DFTs have large peaks at the first harmonic frequency, but the DFT noisefloor here also has double peaks (not resolvable at the resolution of the figure) at the fundamental frequency. These are caused by the frequency offset $\epsilon_F/\Delta t = 10^{-10}$, which is already large enough to put a non-negligible quantity of energy into the frequencies in the vicinity of the fundamental at k_A, as is shown more clearly in the next two figures and in Fig. 2. However, only $X[k_A]$ was set to zero when calculating NAD, which splits the fundamental into two peaks bracketing the missing fundamental in the plot of the DFT noisefloor.

A closer look at the $\epsilon_F = 10^{-8}$ results: For this simulation, Fig. 14 presents the standard graphs.

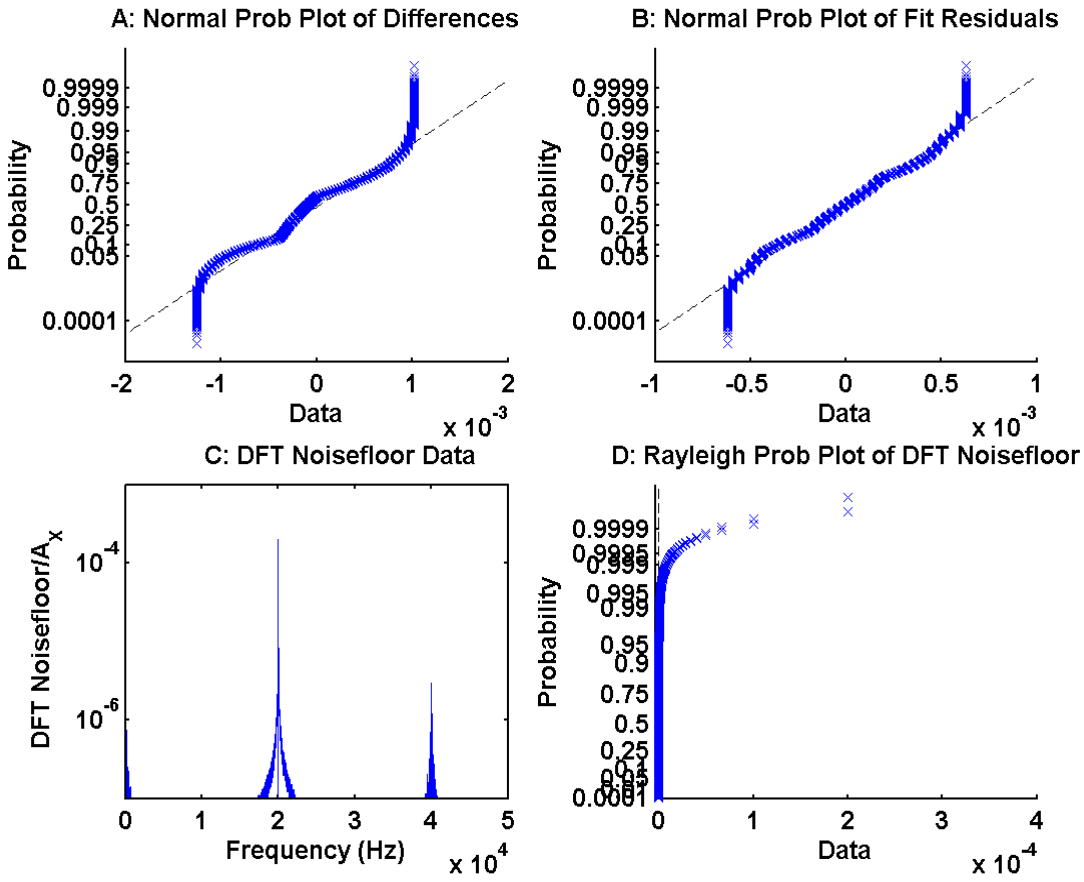

Figure 14: Top Left: Probability plot of the theoretical difference between the quantization-noise limited sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

These results are on the frequency-offset side of the transition from quantization-noise-limited behavior

to frequency-offset-limited noise behavior. These plots share some similarities with those for noise and jitter with $\sigma_N/A = \sigma_J/\delta t = 10^{-4}$ in Fig. 8. But there is a factor of two difference in the ratio of the range of the residuals to the range of the differences. Also the shape of the probability plots of the time domain data is very different from that for random noise limited data and that for random jitter limited data.

A closer look at the $\epsilon_F = 10^{-6}$ results: For this simulation, Fig. 15 presents the standard graphs.

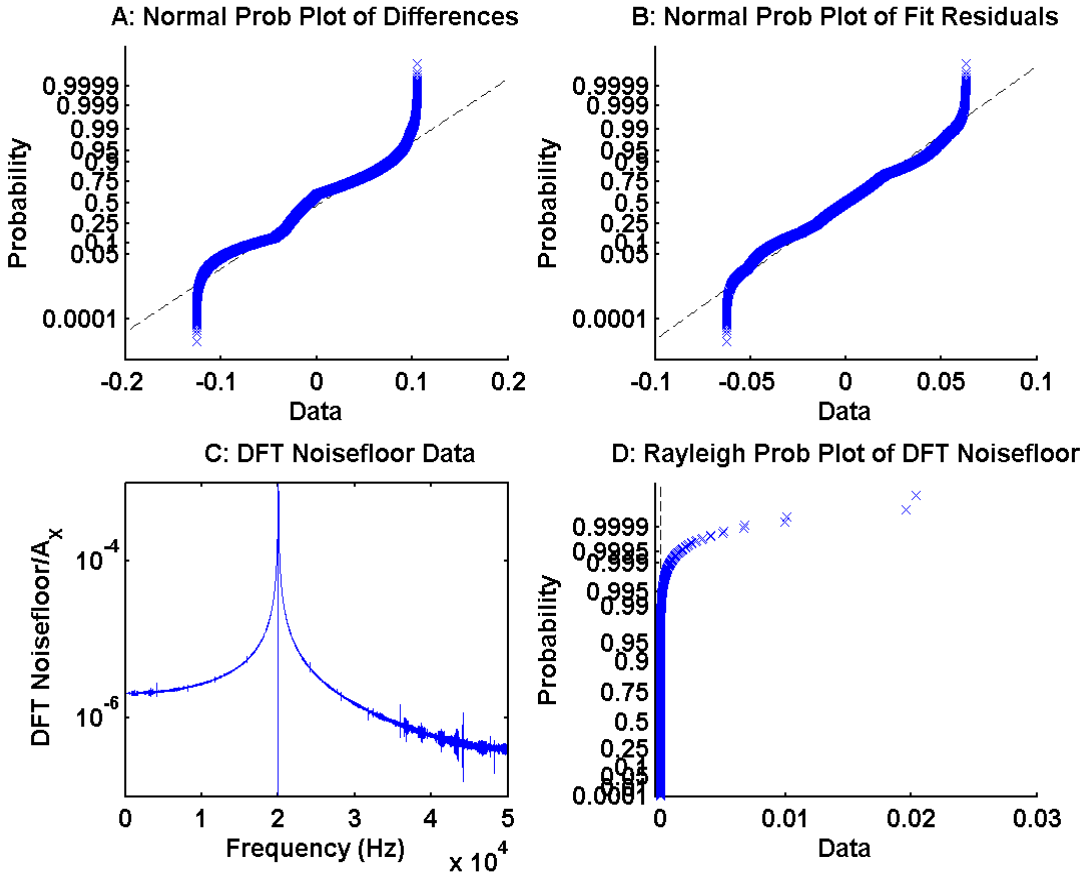

Figure 15: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

These results, which are for the same parameters used in Fig. 1, and which are strongly on the frequency-offset side of the transition from quantization-noise-limited behavior to frequency-offset-limited noise behavior, continue the trends apparent in Fig. 14. The contribution to the first harmonic from frequency offset is now larger than that from the quantization noise, so the latter cannot be detected in the noisefloor DFT.

The probability plot of the residuals of the fit of a sinewave to the distorted sinewave data is nominally identical to the probability plot of the differences between the distorted and undistorted sinewave.

Furthermore the data in these probability plots appears consistent with random distributions that are approximately normal near their means but that become extremely short-tailed at their extremes. However, this does not mean that these results can be extrapolated to different size samples collected over different times as described above with reference to the simulation results in Table 3.

No closer look at the $\epsilon_F = 10^{-16}, 10^{-14}$, and 10^{-12} results are shown because they are similar to the $\sigma_N/A = \sigma_J/\Delta = 10^{-9}$ quantization-noise-limited results in Fig. 13.

This space intentionally left blank .

SINAD versus Available Bits

Table 4 in Appendix E compares the $A_x/A, A_X/A$, SINAD $= A/\text{NAD}$, $\text{SINAD}_x = A_x/\text{NAD}_x$, and $\text{SINAD}_X = A_X/\text{NAD}_X$ for times equal to 1 s, 10 s, and 100 s as a function of the available number of bits from $b = 8$ to $b = 64$ and $\sigma_N/A = \sigma_J/\Delta t = \epsilon_F = 0$. Examination of this table shows that $\text{SINAD}_x(b, t_1) \neq \text{SINAD}(b, t_2)$ $\text{SINAD}_x(b, t_1) \neq \text{SINAD}(b, t_2)$ when $\neq t_1 \neq t_2$, which means that errors of this type should not be thought of as noise and cannot be extrapolated over time by statistical means. However, no correction for clipping such as that given in Appendix A for random noise is required for frequency offset error because this source of distortion is not additive.

A closer look at the $b = 8$ results: For this simulation, Fig. 16 presents the standard graphs.

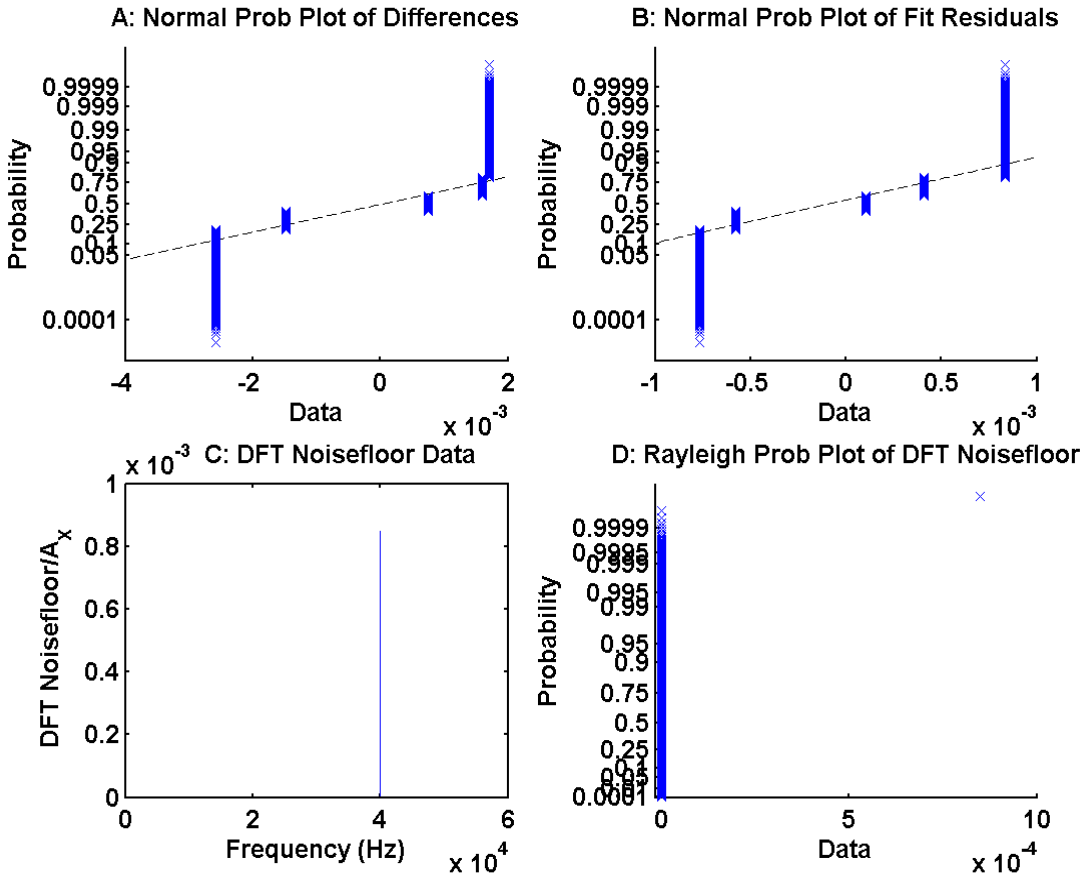

Figure 16: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

These results are qualitatively similar to those in Fig. 6 for $\sigma_N/A = 10^{-9}$. Specifically, the probability plot of the difference between the ideal $v[n]$ and the simulated distorted data $x[n]$ has only five values that are repeated synchronously over the data set, and the probability plot of the residuals of the fit of

the sinewave $x'[n]$ to the $x[n]$ data is similar to the probability plot of $x[n] - v[n]$, also having five values repeated over many cycles, but with a reduced range. Finally, the noise floor DFT has only one spike, which shows that the effect of quantization is synchronous with the ideal sinewave $v[n]$.

A closer look at the $b = 16$ results: For this simulation, Fig. 17 presents the standard graphs.

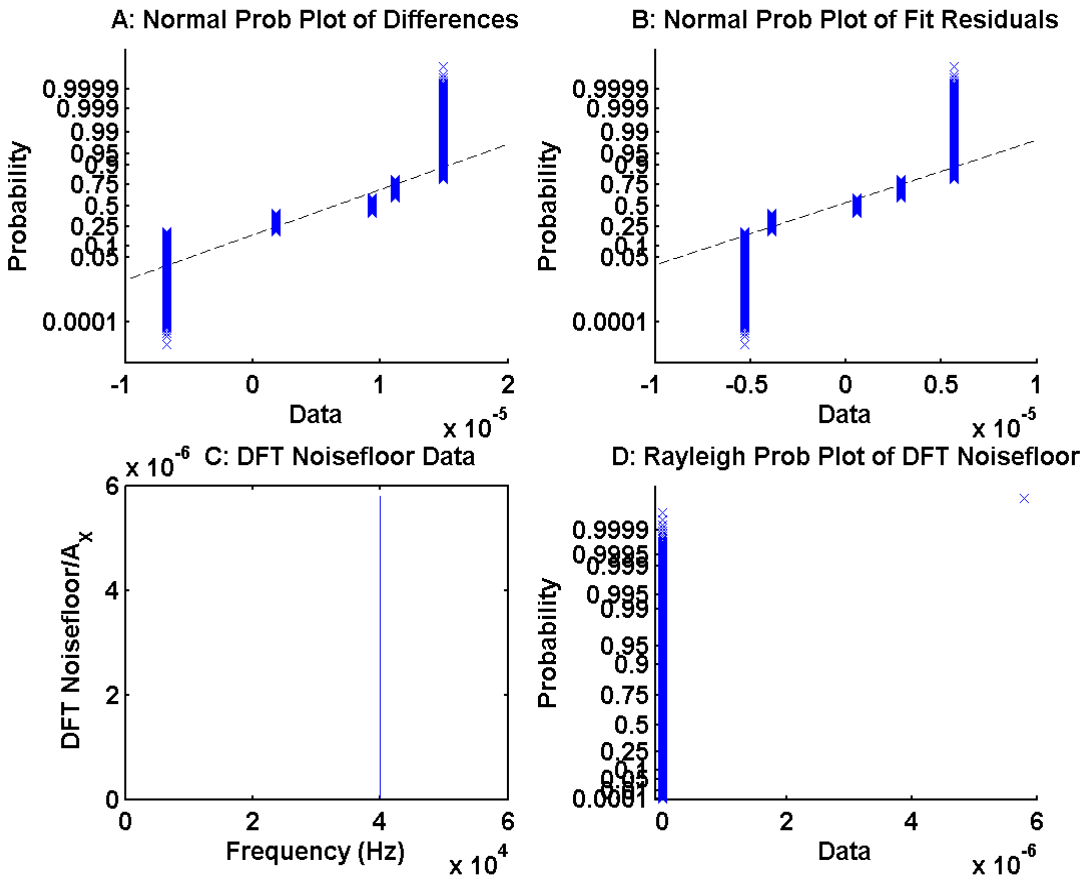

Figure 17: Top Left: Probability plot of the theoretical difference between the quantization-noise limited sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor obtained from $X[k]$ following removal of the fundamental and the DC component. Bottom Right: Probability plot of the DFT noise floor.

These results are qualitatively similar to those for 8 bits, but the ratio of the range of the residuals to that of the differences is reduced.

A closer look at the $b = 24$ results: For this simulation, Fig. 18 presents the standard graphs.

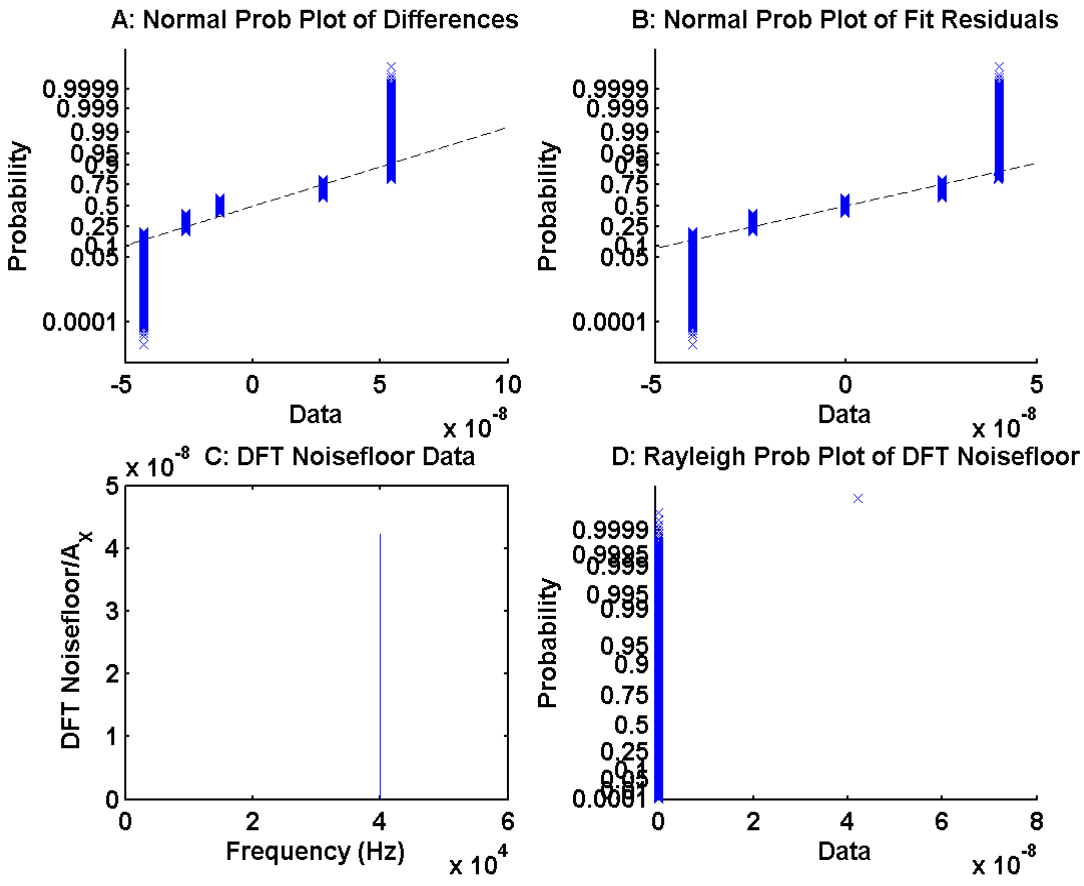

Figure 18: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

These results are qualitatively similar to those for 8 and 16 bits, but the ratio of the range of the residuals to that of the differences is further reduced.

A closer look at the $b = 32$ results: For this simulation, Fig. 19 presents the standard graphs.

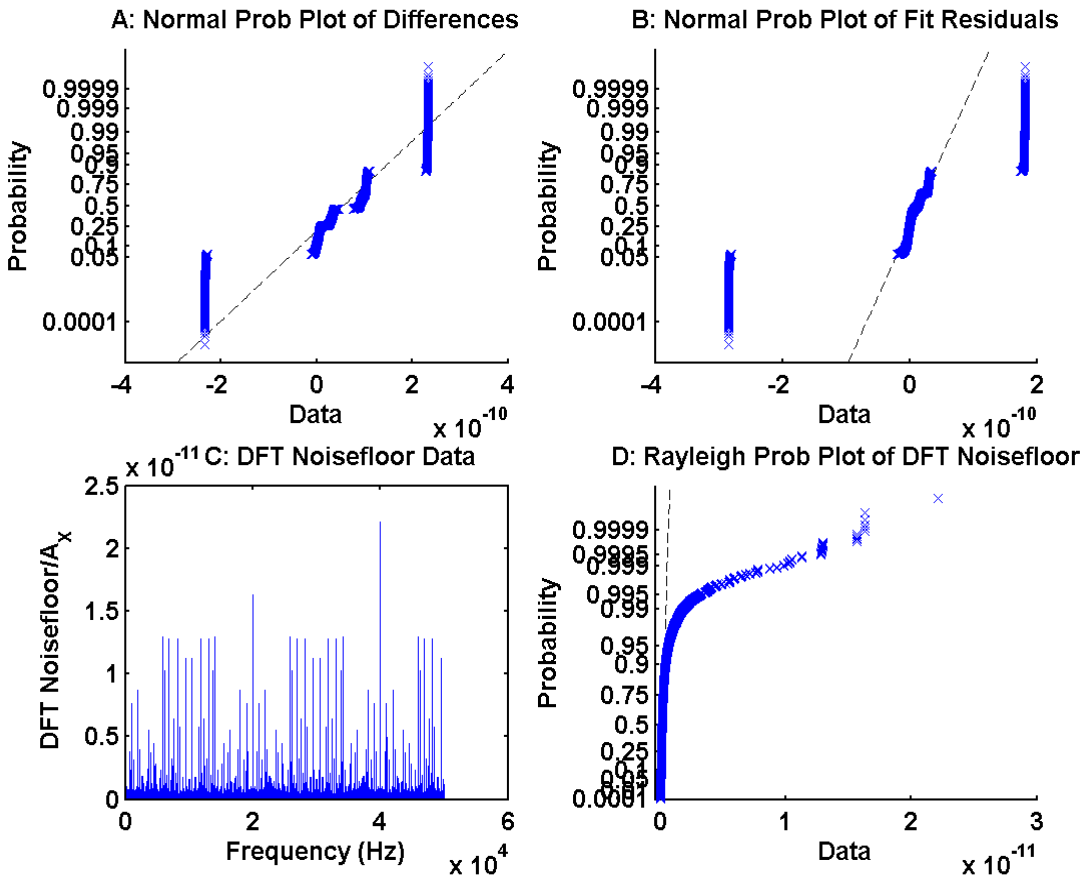

Figure 19: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

These results are qualitatively different from those for 8, 16, and 24 bits. First, the ratio of the range of the residuals to that of the differences is nominally 1. More significantly, the effect of the quantization error is no longer synchronous with the sinewave as is evident both in the probability plots in the top left and top right and in the noise floor DFT in in the lower left. This is surprising. Presumably, it is caused by the interaction of the rounding of $x[n]$ at 32 bits which was programmed into the simulation with the rounding at 51 bits intrinsic to the MATLAB double (precision) arithmetic implemented according to the IEEE 754 standard for double precision. But there is no obvious reason why the synchronization that produced five discrete values for $b = 8, 16$, and 24 bits should be broken. Therefore, these and the following results must be viewed with suspicion.

A closer look at the $b = 40$ results: For this simulation, Fig. 20 presents the standard graphs.

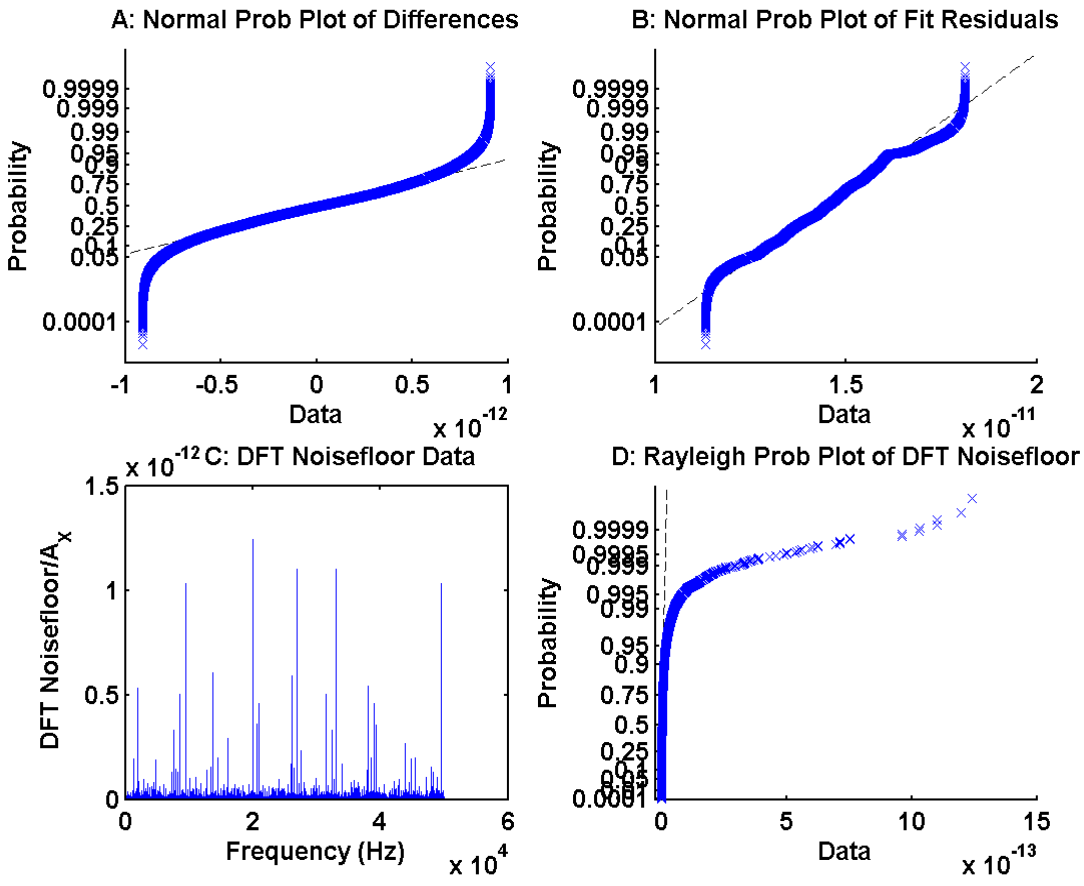

Figure 20: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

These results are qualitatively different from those for 32 bits as well as those for 8, 16, and 24 bits. The shape of the probability plot of the residuals is similar to that of the differences, but the former is much more asymmetric with respect to zero than any of the residual probability plots for 8 to 32 bits. The shape of all of the probability plots are remarkably similar to those for $\epsilon_F = 10^{-8}$ and $\epsilon_F = 10^{-6}$, for no obvious reason.

A closer look at the $b = 48$ results: For this simulation, Fig. 21 presents the standard graphs.

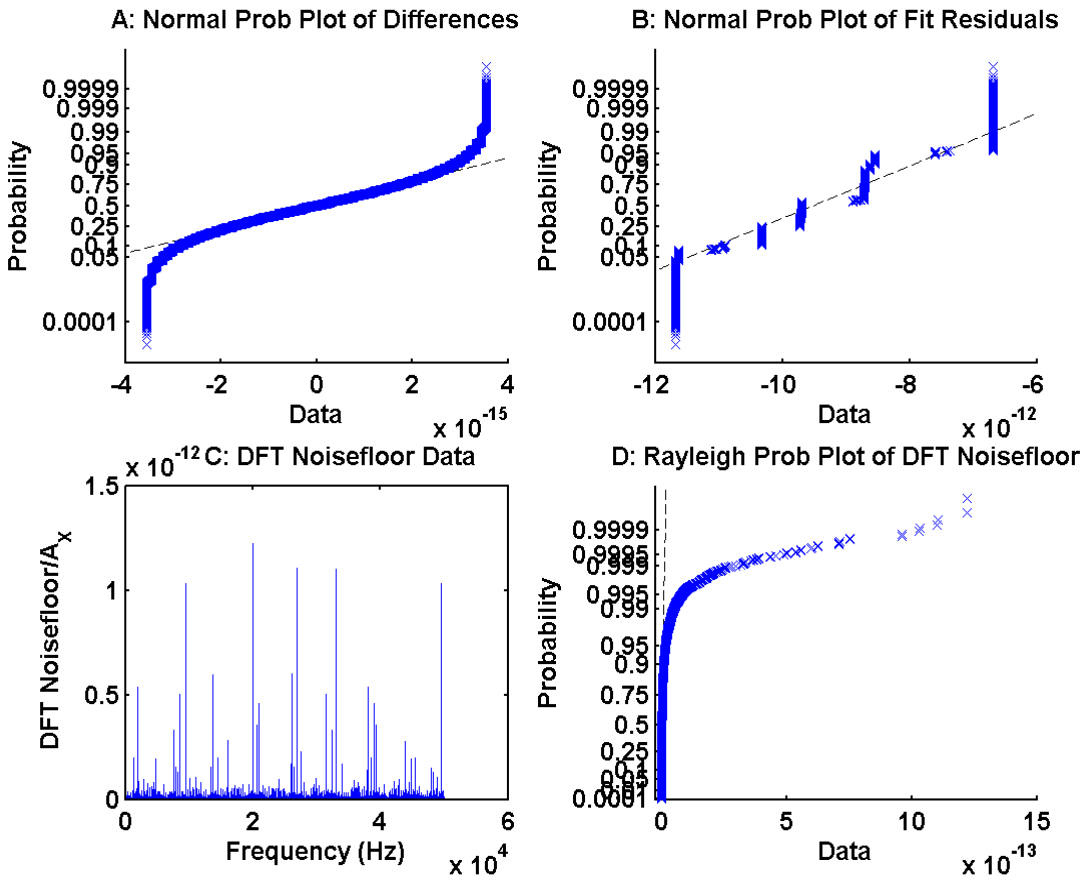

Figure 21: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

Three of the plots are qualitatively similar to the corresponding three for 40 bits, but the probability plot of the residuals has changed dramatically as discussed in a little more detail below.

A closer look at the $b = 56$ results: For this simulation, Fig. 22 presents the standard graphs.

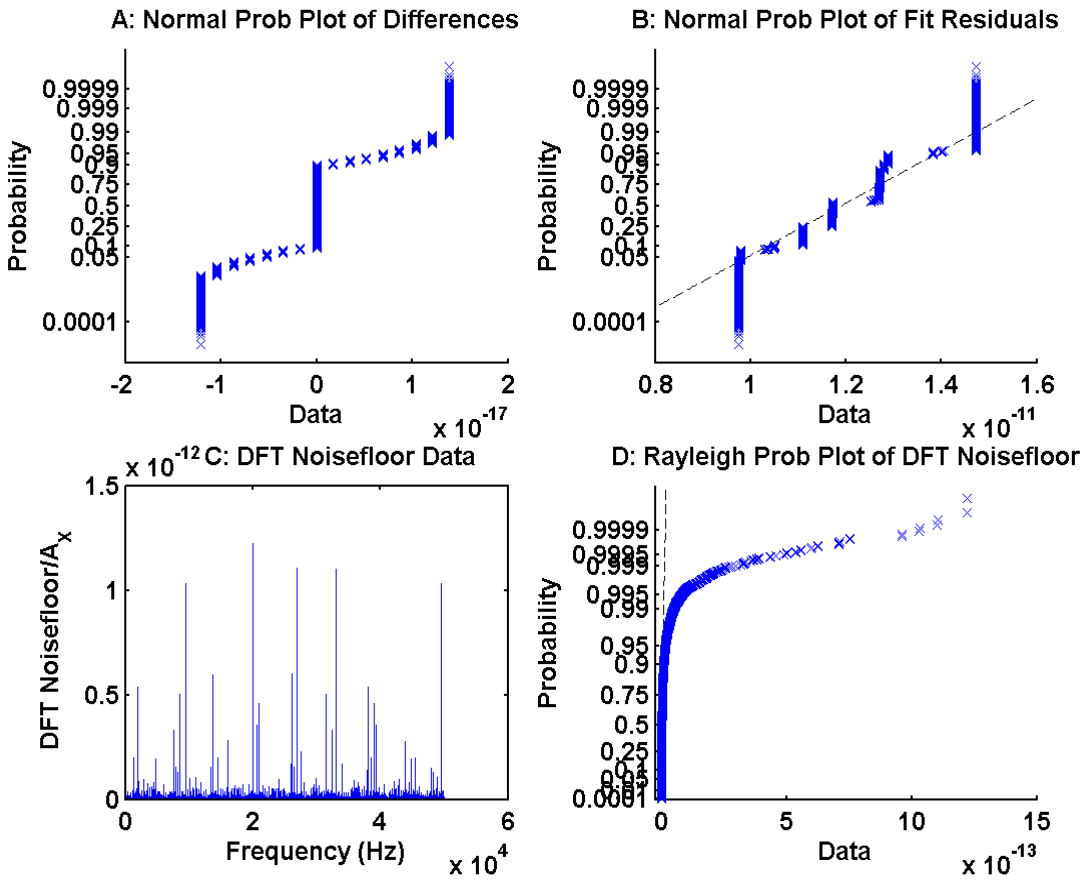

Figure 22: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

The shape of the probability plot of the residuals appears almost identical to that for 48 bits, which suggests that the residuals of the sinewave fit are smaller than a threshold in the program that carries out the fit. The shape of the probability plot of the differences is entirely new.

A closer look at the $b = 64$ results: For this simulation, Fig. 23 presents the standard graphs.

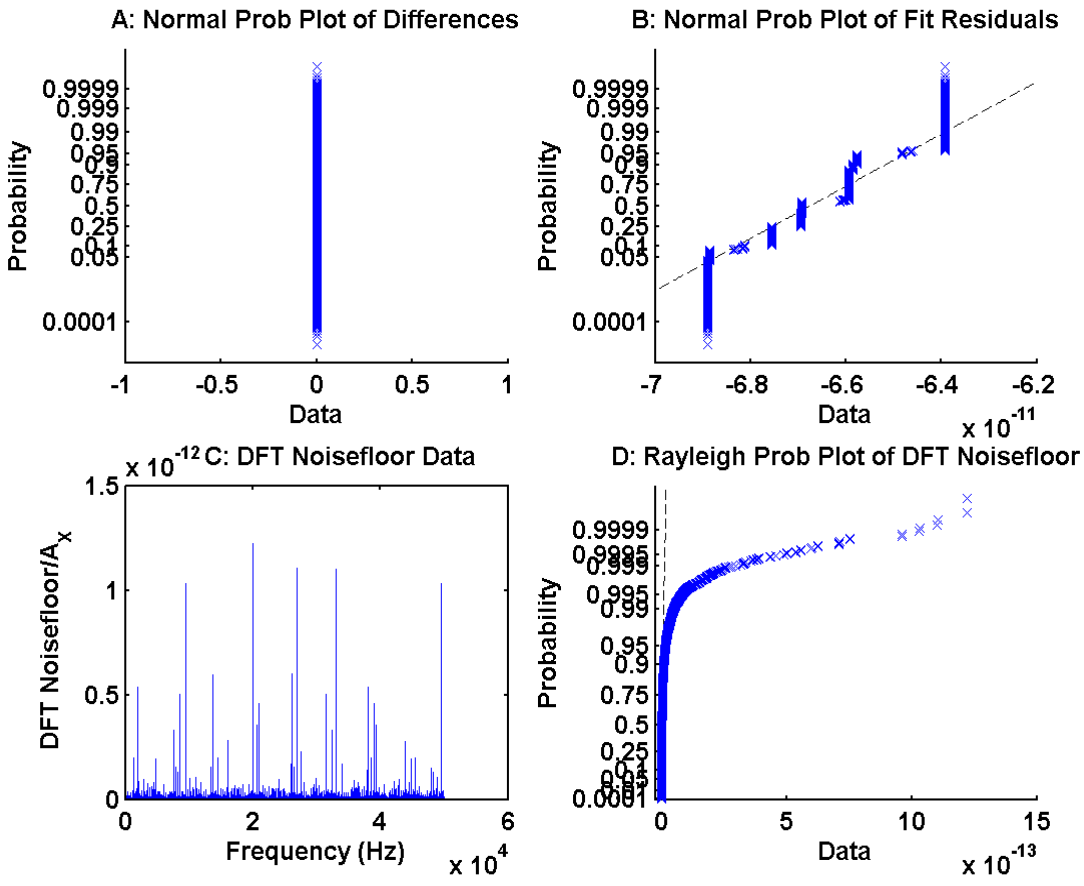

Figure 23: Top Left: Probability plot of the theoretical difference between the noisy sinewave $x[n]$ and the near perfect sinewave $v[n]$. Top Right: Probability plot of the residual after a least squares fit of $x'[n]$ to $x[n]$. Bottom Left: Standard plot of the DFT noisefloor. Bottom Right: Probability plot of the DFT noise floor obtained from $X[k]$ following removal of the fundamental and the DC component.

The shape of the probability plot of the residuals appears almost identical to that for 48 and 56 bits, which lends support to the idea that the shape for 48, 56, and 64 bits is determined by a threshold in the sinewave fitting program. All of the differences are now zero because the simulation program is rounding the simulated distorted data to a number of bits that is greater than the 51 bits in the mantissa of a MATLAB double (precision) number. In retrospect, it is clear that the differences $x[n] - v[n]$ and the residuals $x[n] - x'[n]$ were approaching this situation as the number of bits was increased from 8 to 64 in the simulation, but a computation-terminating threshold in the fitting function prevented the residuals from reaching zero. However, this explanation is at best incomplete because it does not explain why the differences were not all zero for the 56 bit simulation, which also exceeds the 51 bit length of the mantissa in the MATLAB double (precision) numbers, nor does it explain the detailed changes in shape observed when adding 8 bits to the simulated number of available bits to represent the distorted sinewave data.

Conclusion

This report described results of simulating the effect of random noise σ_N, random jitter σ_J, the number of bits b, and frequency-offset error ϵ_F between the sampling frequency and the sinewave frequency on the time and frequency domain estimates of the amplitude A and the signal-to-noise-and-distortion metric SINAD of a sinewave sampled by a 16-bit DAS. For these simulations, the theoretical value of SINAD was defined as the standard deviation of the difference $x[n] - v[n]$ between an ideal sinewave $v[n]$ and a distorted version $x[n]$ of the ideal sinewave.

A suitably normalized discrete Fourier transform (DFT) was used to determine the frequency domain estimates A_X and SINAD_X of A and SINAD, respectively, where SINAD_X was defined as the ratio of A_X to the standard deviation of the noisefloor of the DFT of $x[n]$, taking proper account of the loss of degrees of freedom when calculating the noisefloor from the DFT.

Fitting a three parameter sinewave $x'[n]$ to $x[n]$ was used to determine the time domain estimates A_x and SINAD_x of A and SINAD, respectively, where SINAD_x was defined as the standard deviation of the residuals of the fit, taking proper account of the loss of degrees of freedom when fitting the sinewave $x'[n]$ to the distorted sinewave $x[n]$.

It was also shown that these tools work best when the sinewave and the sampling frequency are integers, and that SINAD_x and SINAD_X are extremely sensitive to violations of this requirement, which may be much more stringent than what is required in many applications. In this case, a four parameter non-linear fit (which may be implemented iteratively with three-parameter linear fits) is required in the time domain and suitably windowed DFTs are required in the frequency domain. A future entry in this series will describe the relations among SINAD, SINAD_x, and SINAD_X when analyzed with these more sophisticated tools.

It was shown that SINAD_x and SINAD_X, as well as A_x and A_X, are identical for all practical purposes. It was also shown that

- A_x and A_X are excellent approximations to A for all cases studied,
- $\text{SINAD}_X \approx \text{SINAD}_x \approx 1.9$ SINAD when quantization is the dominant source of error,
- Quantization error and frequency-offset error cannot be treated as random errors by standard statistical tools,
- Random additive noise and random jitter can be treated as random errors by standard statistical tools, and for these sources of error:
  - SINAD_x and SINAD_X are excellent estimates of SINAD for $\sigma_N/A = \sigma_J/\Delta t \geq 10^{-4}$,
  - $\text{SINAD}_X(\sigma_N/A) \approx \text{SINAD}_x(\sigma_N/A)$ and $\text{SINAD}_X(\sigma_J/\Delta t) \approx \text{SINAD}_x(\sigma_J/\Delta t)$.

It was further shown that probability plots of the residuals of the fits of three parameter sinewaves to distorted sinewaves have distinctively different shapes for normally distributed random additive noise, normally distributed random jitter noise, and frequency-offset error if the effect of quantization-induced error is negligible compared to the effect of these types of distortion. This suggests that it may be possible to quantify the separate contribution from each of these types of distortion in the outputs of data acquisition systems acquiring sinewave data.

Finally, it was shown that this type of analysis can also be applied to 24-bit, and possibly 32-bit DAS, but that some anomalous behavior starting at 32 bits has the potential to produce errors of unknown magnitude in the estimates of SINAD.

References

[1] IEEE Std 1057^{tm}-2007 Standard for Digitizing Waveform Recorders, The Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, New York, NY 10016-5997, USA (2008).

[2] Certain commercial equipment, instruments, or materials are identified in this report in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

[3] Reference [1], Annex A, page 112.

[4] Reference [1], eq. 66 on page 77.

[5] Reference [1], page 21.

[6] Reference [1], page 25.

[7] Reference [1], eq. 68 on page 78.

[8] Nien Fan Zhang, Calculation of the uncertainty of the mean of autocorrelated measurements, Metrologia 43 (2006) S276S281, doi.org/10.1088/0026-1394/43/4/S15

[9] Nien Fan Zhang,The batched moving averages of measurement data and their applications in data treatment, Measurement 39 (2006) 864875, www.elsevier.com/locate/measurement

This appendix describes how to account for signal clipping in the simulation and in practice. The goal is to derive a conservative effective peak amplitude $A_E < A2^{b-1}$ to be used to assure that the probability that no $x[n]$ in eq. 4 for $n = 1, \ldots, M$ will exceed A is less than some user defined threshold Q. Assume that $A = 2^{b-1}$ is the peak amplitude of the maximum noise-free sinewave that can be measured without clipping. Further assume that the dominant source of noise is normally distributed, random, additive noise characterized by σ_N. To start, approximate Q by

$$\begin{aligned}
\mathcal{Q} &= \sum_{n=1}^{M} \mathcal{Q}_n \quad A_E v[n] + r_N[n]\,\sigma_N \; - A\big) \\
&= \sum_{n=1}^{M} \mathcal{Q}_n\left(\; r_N[n] \; < \frac{A - \; A_E v[n]}{\sigma_N} \right) \\
&\quad \sum_{n=1}^{M} \frac{1}{\sqrt{2\pi}} \int_{\frac{A - \; A_E v[n]}{\sigma_N}}^{\infty} \exp(-t^2/2) dt \\
&= \sum_{n=1}^{M} \frac{1}{2} \mathrm{erfc}\left(\frac{A - \; A_E v[n]}{\sqrt{2}\sigma_N} \right) \ll 1,
\end{aligned} \tag{21}$$

where $v[n]$, which is given in eq. 7, is the expectation value of $x[n]$, $\mathcal{Q}_n$ is the probability that its argument is negative, $|\ldots|$ denotes the absolute value function, and the $\ll$ symbol emphasizes that this approximation is accurate only when $\mathcal{Q} \ll 1$. Finally, let $\hat{A}_E$ be the maximum value of A_E for which eq. 21 is true.

Unfortunately, eq. 21 is very difficult to solve exactly for $\hat{A}_E$, but a very conservative estimate is given by

$$\mathcal{Q} = \sum_{n=1}^{M} \frac{1}{2} \mathrm{erfc}\left(\frac{A - A_E}{\sigma_N} \right) = \frac{M}{2} \mathrm{erfc}\left(\frac{A - A_E}{\sqrt{2}\sigma_N} \right), \tag{22}$$

which is readily solved as

$$\frac{A - A_E}{\sqrt{2}\sigma_N} = \mathrm{inverfc}\left(\frac{2\mathcal{Q}}{M} \right), \tag{23}$$

which in turn gives

$$A_E = A - \sqrt{2}\sigma_N \mathrm{inverfc}\left(\frac{2\mathcal{Q}}{M} \right) \leq 6\sqrt{2}\sigma_N \mathrm{inverfc}\left(\frac{2\mathcal{Q}}{M} \right). \tag{24}$$

As an example, let $M = 10^5$, $\sigma_N/A = 10^{-3}$, and $\mathcal{Q} = 0.1$, which means that in a set of 100 measurements of a sinewave, each of which consists of M samples of the sinewave, 10 measurements containing one or more values of $x[n] > A$ will be considered acceptable. In this case $A_E/A = 0.9952$. If Q is reduced to 0.001, A_E/A decreases only slightly to 0.9943. Note that these very favorable results are crucially dependent upon the assumption of normally distributed noise. Deviations for normality will change the results drastically. On the other hand, eq. 24 is very conservative. The net result is that for real measurements $A_E/A = 1 - 6\sigma_N$ for $\sigma_N/A \leq 10^{-2}$ is probably safe from noise-derived clipping, and less limiting than other restrictions of A_E for other sources of error such as gain non-linearity well below the clipping threshold.

In the absence of clipping $\mathrm{SINAD}_x(A_E) < \mathrm{SINAD}_x(A)$ if $A_E < A$. However, this relation may not be true in the presence of noise induced clipping. In this case an iteration based on eq. 24 will quickly converge to the value of A_E that maximizes $\mathrm{SINAD}_x(A_E)$.

Table 1: Amplitudes (A, A_x, and A_X), and SINADs (A/NAD, A_x/NAD$_x$,and A_X/NAD$_X$) as a function of σ_N and time.

```
SINAD_vs_Time_2015_07_10 run at 2015-4-22T22:41
Target sampling frequency:         fs = 100000.000000 Hz
Target sinewave frequency:         f = 20000.000000 Hz
Target sinewave phase:             phi = 0.300000 rad
Sampling time:                     time = 1.000000 s
Relative Std. Dev. of jitter:      sigJ = 0.000000e+00
Number of bits in ADC convertor:   bits = 16 bits
Peak full-scale signal amplitude:  A = 65536.000000 counts
Relative frequency offset:         epsF = 0.000000e+00
Relative additive noise:           sigN = 0.000000e+00
        SigN         Ax/A         AX/A        A/NAD      Ax/NADx      AX/NADX
Sampling time:                     time = 1.000000 s
1.000000e-09 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-08 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-07 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-06 1.000004e+00 1.000004e+00 7.858388e+04 9.634392e+04 9.634289e+04
1.000000e-05 1.000001e+00 1.000001e+00 5.865573e+04 5.920136e+04 5.920065e+04
1.000000e-04 1.000000e+00 1.000000e+00 9.956720e+03 9.956699e+03 9.956558e+03
1.000000e-03 1.000004e+00 1.000004e+00 9.993095e+02 9.993117e+02 9.992981e+02
1.000000e-02 9.999738e-01 9.999738e-01 1.000393e+02 1.000364e+02 1.000350e+02
1.000000e-01 1.000136e+00 1.000136e+00 1.000255e+01 1.000391e+01 1.000385e+01
1.000000e+00 1.002425e+00 1.002425e+00 1.002177e+00 1.004601e+00 1.004592e+00
Sampling time:                     time = 10.000000 s
1.000000e-09 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-08 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-07 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-06 1.000004e+00 1.000004e+00 7.861966e+04 9.643945e+04 9.628322e+04
1.000000e-05 1.000001e+00 1.000001e+00 5.861457e+04 5.915715e+04 5.900965e+04
1.000000e-04 1.000000e+00 1.000001e+00 9.911960e+03 9.911952e+03 9.912022e+03
1.000000e-03 1.000000e+00 1.000002e+00 1.000280e+03 1.000280e+03 1.000298e+03
1.000000e-02 1.000001e+00 1.000019e+00 9.996103e+01 9.996105e+01 1.001045e+02
1.000000e-01 9.999952e-01 9.996796e-01 9.989817e+00 9.989760e+00 9.951990e+00
1.000000e+00 1.000582e+00 1.002793e+00 9.991335e-01 9.997149e-01 1.000885e+00
Sampling time:                     time = 100.000000 s
1.000000e-09 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
1.000000e-08 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
1.000000e-07 1.000008e+00 1.000008e+00 9.168098e+04 1.726392e+05 1.726348e+05
1.000000e-06 1.000004e+00 1.000004e+00 7.865024e+04 9.652275e+04 9.607787e+04
1.000000e-05 1.000001e+00 1.000002e+00 5.866650e+04 5.921463e+04 5.930748e+04
1.000000e-04 1.000000e+00 1.000000e+00 9.926256e+03 9.926255e+03 9.933983e+03
1.000000e-03 9.999997e-01 9.999979e-01 9.999839e+02 9.999835e+02 9.994964e+02
1.000000e-02 9.999988e-01 9.999535e-01 9.997583e+01 9.997570e+01 9.980950e+01
1.000000e-01 9.999636e-01 1.000171e+00 1.000233e+01 1.000197e+01 9.996901e+00
1.000000e+00 1.000203e+00 1.004682e+00 9.998174e-01 1.000020e+00 1.004140e+00
```

Table 2: Amplitudes (A, A_x, and A_X), and SINADs (A/NAD, A_x/NAD$_x$,and A_X/NAD$_X$) as a function of σ_J and time.

```
SINAD_vs_Time_2015_07_14_sigJ run at 2015-7-14T10:47
Target sampling frequency:          fs = 100000.000000 Hz
Target sinewave frequency:          f = 20000.000000 Hz
Target sinewave phase:              phi = 0.300000 rad
Sampling time:                      time = 1.000000 s
Relative Std. Dev. of jitter:       sigJ = 0.000000e+00
Number of bits in ADC convertor:    bits = 16 bits
Peak full-scale signal amplitude:   A = 65536.000000 counts
Relative frequency offset:          epsF = 0.000000e+00
Relative additive noise:            sigN = 1.000000e-20
        SigJ         Ax/A         AX/A        A/NAD      Ax/NADx      AX/NADX
Sampling time:                      time = 1.000000 s
1.000000e-09 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-08 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-07 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-06 1.000006e+00 1.000006e+00 8.284128e+04 1.103313e+05 1.103297e+05
1.000000e-05 1.000003e+00 1.000003e+00 5.206340e+04 5.269339e+04 5.269293e+04
1.000000e-04 9.999997e-01 9.999997e-01 7.915009e+03 7.915255e+03 7.915188e+03
1.000000e-03 9.999923e-01 9.999923e-01 7.925129e+02 7.925277e+02 7.925259e+02
1.000000e-02 9.999237e-01 9.999237e-01 7.960788e+01 7.960329e+01 7.960213e+01
1.000000e-01 9.919425e-01 9.919425e-01 7.974859e+00 7.926913e+00 7.926838e+00
Sampling time:                      time = 10.000000 s
1.000000e-09 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-08 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-07 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-06 1.000006e+00 1.000006e+00 8.277603e+04 1.100729e+05 1.104472e+05
1.000000e-05 1.000003e+00 1.000003e+00 5.220942e+04 5.283870e+04 5.287260e+04
1.000000e-04 1.000000e+00 1.000000e+00 7.922270e+03 7.922266e+03 7.919426e+03
1.000000e-03 1.000000e+00 1.000002e+00 7.962786e+02 7.962778e+02 7.956851e+02
1.000000e-02 9.999256e-01 9.999166e-01 7.964489e+01 7.964037e+01 7.976714e+01
1.000000e-01 9.920834e-01 9.921038e-01 7.971927e+00 7.924611e+00 7.944234e+00
Sampling time:                      time = 100.000000 s
1.000000e-09 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
1.000000e-08 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
1.000000e-07 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
1.000000e-06 1.000006e+00 1.000006e+00 8.278869e+04 1.101255e+05 1.101523e+05
1.000000e-05 1.000003e+00 1.000003e+00 5.225219e+04 5.287804e+04 5.271312e+04
1.000000e-04 1.000000e+00 1.000000e+00 7.922765e+03 7.922766e+03 7.925654e+03
1.000000e-03 9.999993e-01 1.000000e+00 7.959639e+02 7.959634e+02 7.962958e+02
1.000000e-02 9.999196e-01 9.999474e-01 7.958379e+01 7.957901e+01 7.964762e+01
1.000000e-01 9.921077e-01 9.925145e-01 7.970749e+00 7.923536e+00 7.861342e+00
```

Table 3: Amplitudes (A, A_x, and A_X), and SINADs (A/NAD, A_x/NAD$_x$,and A_X/NAD$_X$) as a function of ϵ_F and time.

```
SINAD_vs_Time_2015_07_14_epsF run at 2015-7-14T13:1
Target sampling frequency:         fs = 100000.000000 Hz
Target sinewave frequency:         f = 20000.000000 Hz
Target sinewave phase:             phi = 0.300000 rad
Sampling time:                     time = 1.000000 s
Relative Std. Dev. of jitter:      sigJ = 0.000000e+00
Number of bits in ADC convertor:   bits = 16 bits
Peak full-scale signal amplitude:  A = 65536.000000 counts
Relative frequency offset:         epsF = 0.000000e+00
Relative additive noise:           sigN = 1.000000e-20
        epsF         Ax/A         AX/A        A/NAD      Ax/NADx      AX/NADX
Sampling time:                     time = 1.000000 s
1.000000e-16 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-14 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
1.000000e-12 1.000008e+00 1.000008e+00 9.168113e+04 1.726340e+05 1.726348e+05
1.000000e-10 9.999971e-01 9.999971e-01 7.067603e+04 8.690616e+04 9.020434e+04
1.000000e-08 1.000000e+00 9.999999e-01 1.378113e+03 8.036879e+04 2.755136e+03
1.000000e-06 1.000000e+00 9.993427e-01 1.378864e+01 8.026442e+04 2.755476e+01
Sampling time:                     time = 10.000000 s
1.000000e-16 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
1.000000e-14 1.000008e+00 1.000008e+00 9.168168e+04 1.726397e+05 1.726348e+05
1.000000e-12 1.000006e+00 1.000008e+00 8.562573e+04 1.234753e+05 1.726348e+05
1.000000e-10 9.999994e-01 9.999971e-01 1.346802e+04 8.183719e+04 9.020434e+04
1.000000e-08 1.000000e+00 9.999999e-01 1.378325e+02 8.026292e+04 2.755136e+03
1.000000e-06 1.000000e+00 9.993427e-01 1.433928e+00 8.031384e+04 2.755476e+01
Sampling time:                     time = 100.000000 s
1.000000e-16 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
1.000000e-14 1.000008e+00 1.000008e+00 9.168104e+04 1.726367e+05 1.726348e+05
1.000000e-12 9.999971e-01 1.000008e+00 7.067575e+04 8.690867e+04 1.726348e+05
1.000000e-10 1.000000e+00 9.999971e-01 1.378116e+03 8.037023e+04 9.020434e+04
1.000000e-08 1.000000e+00 9.999999e-01 1.378866e+01 8.026561e+04 2.755136e+03
1.000000e-06 1.000000e+00 9.993427e-01 7.071067e-01 8.031513e+04 2.755476e+01
```

Table 4: Amplitudes (A, A_x, and A_X), and SINADs (A/NAD, A_x/NAD$_x$,and A_X/NAD$_X$) as a function of Number of Bits and time.

```
SINAD_vs_Time_2015_07_14_sigJ run at 2015-7-14T18:17
Target sampling frequency:         fs = 100000.000000 Hz
Target sinewave frequency:         f = 20000.000000 Hz
Target sinewave phase:             phi = 0.300000 rad
Sampling time:                     time = 1.000000 s
Relative Std. Dev. of jitter:      sigJ = 0.000000e+00
Number of bits in ADC convertor:   bits = 16 bits
Peak full-scale signal amplitude:  A = 65536.000000 counts
Relative frequency offset:         epsF = 0.000000e+00
Relative additive noise:           sigN = 1.000000e-20
        Bits         Ax/A         AX/A        A/NAD      Ax/NADx      AX/NADX
Sampling time:                     time = 1.000000 s
8.000000e+00 9.978583e-01 9.978583e-01 4.107682e+02 1.178420e+03 1.178402e+03
1.600000e+01 1.000008e+00 1.000008e+00 9.168113e+04 1.726374e+05 1.726348e+05
2.400000e+01 1.000000e+00 1.000000e+00 1.974972e+07 2.369203e+07 2.369167e+07
3.200000e+01 1.000000e+00 1.000000e+00 6.851407e+09 7.318231e+09 7.286117e+09
4.000000e+01 1.000000e+00 1.000000e+00 1.358524e+12 5.974162e+11 1.886386e+11
4.800000e+01 1.000000e+00 1.000000e+00 3.447107e+14 6.599051e+11 1.902051e+11
5.600000e+01 1.000000e+00 1.000000e+00 1.967085e+17 6.599082e+11 1.902053e+11
6.400000e+01 1.000000e+00 1.000000e+00      Inf     6.599082e+11 1.902053e+11
Sampling time:                     time = 10.000000 s
8.000000e+00 9.978583e-01 9.978583e-01 4.107700e+02 1.178436e+03 1.178402e+03
1.600000e+01 1.000008e+00 1.000008e+00 9.168168e+04 1.726398e+05 1.726348e+05
2.400000e+01 1.000000e+00 1.000000e+00 1.974595e+07 2.369233e+07 2.369167e+07
3.200000e+01 1.000000e+00 1.000000e+00 5.567443e+09 6.404409e+09 6.832991e+09
4.000000e+01 1.000000e+00 1.000000e+00 1.347659e+12 3.663300e+10 1.741727e+11
4.800000e+01 1.000000e+00 1.000000e+00 3.447085e+14 3.664569e+10 1.752944e+11
5.600000e+01 1.000000e+00 1.000000e+00 1.967561e+17 3.664569e+10 1.752945e+11
6.400000e+01 1.000000e+00 1.000000e+00      Inf     3.664569e+10 1.752945e+11
Sampling time:                     time = 100.000000 s
8.000000e+00 9.978583e-01 9.978583e-01 4.107702e+02 1.178437e+03 1.178402e+03
1.600000e+01 1.000008e+00 1.000008e+00 9.168104e+04 1.726400e+05 1.726348e+05
2.400000e+01 1.000000e+00 1.000000e+00 1.976388e+07 2.368886e+07 2.369167e+07
3.200000e+01 1.000000e+00 1.000000e+00 5.100547e+09 2.144688e+09 6.832991e+09
4.000000e+01 1.000000e+00 1.000000e+00 1.346720e+12 2.379477e+09 1.741727e+11
4.800000e+01 1.000000e+00 1.000000e+00 3.446705e+14 2.379480e+09 1.752944e+11
5.600000e+01 1.000000e+00 1.000000e+00 1.968036e+17 2.379480e+09 1.752945e+11
6.400000e+01 1.000000e+00 1.000000e+00      Inf     2.379480e+09 1.752945e+11
```