

Application Note

Noise Histogram Analysis

by John Lis



INTRODUCTION

Many Analog-to-Digital Converters (ADC) are used to measure the level or magnitude of static signals. Applications include the measurement of weight, pressure, and temperature. These applications involve low-level signals which require high resolution and accuracy. An example is a weigh-scale that can handle up to a 5 kilogram load and yet resolve the measurement to 10 milligrams. The ratio of maximum load to lowest resolvable unit is five hundred thousand to one. This requires the weigh-scale's ADC to digitize a load cell's signal with a resolution of 500,000 counts.

When working with high resolution ADCs, an understanding of the error and noise associated with the conversion process is required. The goal for this application note is to show how histogram analysis is used to quantify static performance. Sample sets of data are collected and used to measure noise and offset. Statistical techniques are used to determine the "goodness" and confidence interval associated with the estimates. Averaging is addressed as a means of decreasing uncertainty and improving resolution.

In an ideal situation, the output of an ADC would be exact with no offset error, gain error, nor noise. However, the actual output from the ADC includes error and noise. Static testing methods can be used to evaluate the ADC's performance. A dc signal is applied to the ADC's input and the digital output words are collected. The signal's level is adjusted to measure offset and gain errors associated with deviations in the slope and intercept of the ADC's transfer function. Noise is measured as the variability of the output for a constant input.

Statistical techniques can be used to acquire performance measurements, assess the effects of noise, as well as compensate for the noise. An ADC's output varies for a constant input due to noise. The noise is defined by a Probability Den-

sity Function (PDF), which represents the probability of discrete events. Statistical parameters can be calculated from the PDF. The PDF's shape describes the certainty of the ADC's output and its noise characteristics.

Noise histogram analysis assumes that the noise is random with a Gaussian distribution. This means that the noise amplitude at a given instant is uncorrelated with the output amplitude at any other instant. A sample set of random noise produces a normal distribution which is used to estimate the PDF. If the noise is not random and does not have a normal distribution, the following histogram analysis would not apply. Examples of non-Gaussian noise include 1/f noise, clock coupling, switching power supply noise, and power line interference.

The ensuing sections discuss noise histogram analysis and the estimation of unknown parameters. The discussion addresses sampling requirements, statistics, and performance trade-offs. Statistical methods are used to determine "goodness" and confidence intervals of parameters estimated from sampled data. Goodness relates to how well the sample set parameters correlate to the actual system. Averaging is discussed as a method of reducing uncertainty and improving resolution. This paper will lead to an understanding of sampling issues and the trade-offs that can be made to improve performance and the consequences of the various choices among sample size, confidence level, and throughput.

NOISE HISTOGRAM DESCRIPTION

Usually the PDF describing the ADC noise is not specified. The PDF can be estimated by static testing. This estimated PDF is actually a histogram plot of the occurrences of a random variable versus the individual variations. For an ADC, the random variable is the resulting digital codes, so the frequency at which each code occurs is plotted against each discrete code.




Figure 1. Histogram for a Noiseless ADC




Figure 2. Histogram for an ADC with Noise




Figure 3. Histogram and Estimated PDF

In a noiseless ADC, the output code for a specific input voltage will always be the same value. The histogram plot for a noiseless converter is shown in Figure 1. If noise sources are present in the ADC, the histogram plot of the output codes contains more than one value. Figure 2 is a histogram plot of an ADC with internal noise. The histogram suggests that the output for a single input can be one of eleven possible codes.

Electrical noise resulting from random effects forms a Gaussian or normal distribution, which is a bell-shaped curve called a normal curve. The Gaussian PDF is continuous and completely determined with the specification of a mean (μ) and variance (σ^2). The Gaussian PDF is defined by the following equation:

$$p(x) = \frac{n}{\sqrt{2\pi} \sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

$$n = 1 \quad \text{for the actual PDF.}$$

Other values for n scale the PDF to fit a sample set. The data presented in Figure 2 can be used to estimate the PDF. Figure 3 plots the histogram of an ADC with noise along with its estimated PDF. The mean and variance were estimated from the sample set of data. From these PDF parameters, the ADC’s performance can be quantified. The mean is the expected or average value. It is used to measure offset errors. The variance describes the variability of the distribution about the mean. It is used as a measurement of uncertainty or noise. The square root of the variance is called the standard deviation (σ), and it is a measure of the effective or rms noise. The peak-to-peak noise can be determined from the rms noise value.

The Gaussian PDF can not be used to measure all types of noise. When using a normal or Gaussian distribution to estimate the PDF, the noise must be random. Figure 4 illustrates a histogram of non-random noise. Notice that the




Figure 4. Non-random Noise Histogram

histogram distribution does not possess the familiar bell-shape. This histogram is possibly the result of 60 Hz line interference or other type of sinusoidal noise. The PDF resembles that of a sine wave, which has a "cusp-shaped" distribution.

Figure 5 is another example of a PDF not having a Gaussian shape. Here, the reason is due to the ADC's poor differential nonlinearity (DNL). Poor DNL results in uneven code widths, which skew the distribution. Delta-Sigma and self-calibrating ADCs possess good DNL specifications. Good DNL is very important for applications which use averaging to improve resolution.

The histogram must possess a bell-shape distribution or the estimated Gaussian PDF will not correlate to the actual system. It is good practice to view the noise distribution and verify that random noise is being analyzed. If the noise is not random, the Gaussian PDF equation can not be used to model the histogram.

EXAMPLE

Figure 6 is a noise histogram of a 20-bit ADC with a grounded analog input, and ± 2.5 volt input range. For an ideal, noise free system, the expected output would always be zero. However, the experimental 1024-sample set ranges from

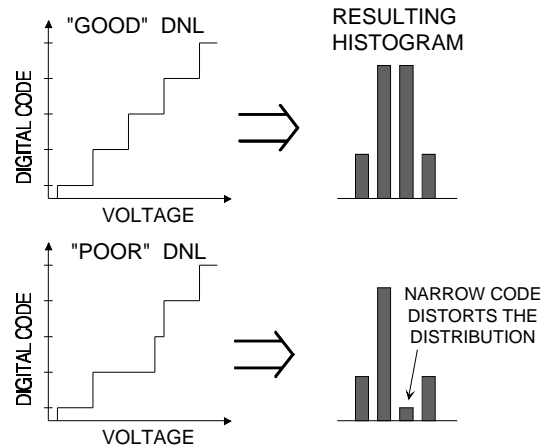


Figure 5. DNL Effects on a Noise Histogram

negative seven to positive five. Since the digital codes vary by more than one count, the system noise exceeds quantization error causing an uncertainty associated with the output. The range of code variations requires the histogram to contain at least thirteen discrete ranges or cells. Table 1 lists the number of occurrences of the fifteen cells used to create the histogram.

A PDF can be estimated by using statistical functions to characterize sampled data. Assuming the system noise is Gaussian, the noise can be measured by collecting a set of n samples from a normal population with mean (μ) and variance (σ^2) and calculating the sample mean (\bar{X}) and sample variance (S^2).

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$$

$$S^2 = \frac{\sum_{i=1}^n (X_i - \bar{X})^2}{n - 1}$$




Figure 6. Noise Histogram for a 20-bit ADC

CELL	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6
Occurrences	0	1	8	32	64	124	181	204	169	126	76	32	6	1	0

Table 1. Data for Histogram in Figure 6

\bar{X} and S^2 are estimators for the system μ and σ^2 . This information is used to create a mathematical model of the system's PDF. From the sample set of data used to create Figure 6, $\bar{X} = -0.98$ and $S^2 = 3.96$. The PDF for Figure 6 can be modeled by substituting n , \bar{X} , and S^2 into the scaled Gaussian PDF equation.

$$p(x) = \frac{n}{\sqrt{2\pi} \sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

Substituting \bar{X} for μ , S for σ , and S^2 for σ^2 .

$$p(x) = \frac{1024}{\sqrt{2\pi} \ 1.99} e^{-\frac{(x+0.98)^2}{(2 * 3.96)}}$$

Figure 7 overlays the estimated PDF over the histogram of 1024 samples shown in Figure 6. Notice that the PDF and histogram are highly correlated and the estimated PDF seems to be a good model of the actual system.

Continuing with the assumption that the data in Figure 7 has a normal distribution, the performance of the ADC can be quantified using the measurements based on estimates for the mean and standard deviation.

The sample mean in Figure 7 is -0.98 counts. A perfect ADC, operating in the same mode with its input grounded, has a mean of zero. Such a mean deviation from ideal is called an offset. In terms of voltage, the ADC's Zero Offset in Figure 7 is -0.98 counts or -4.67 μV

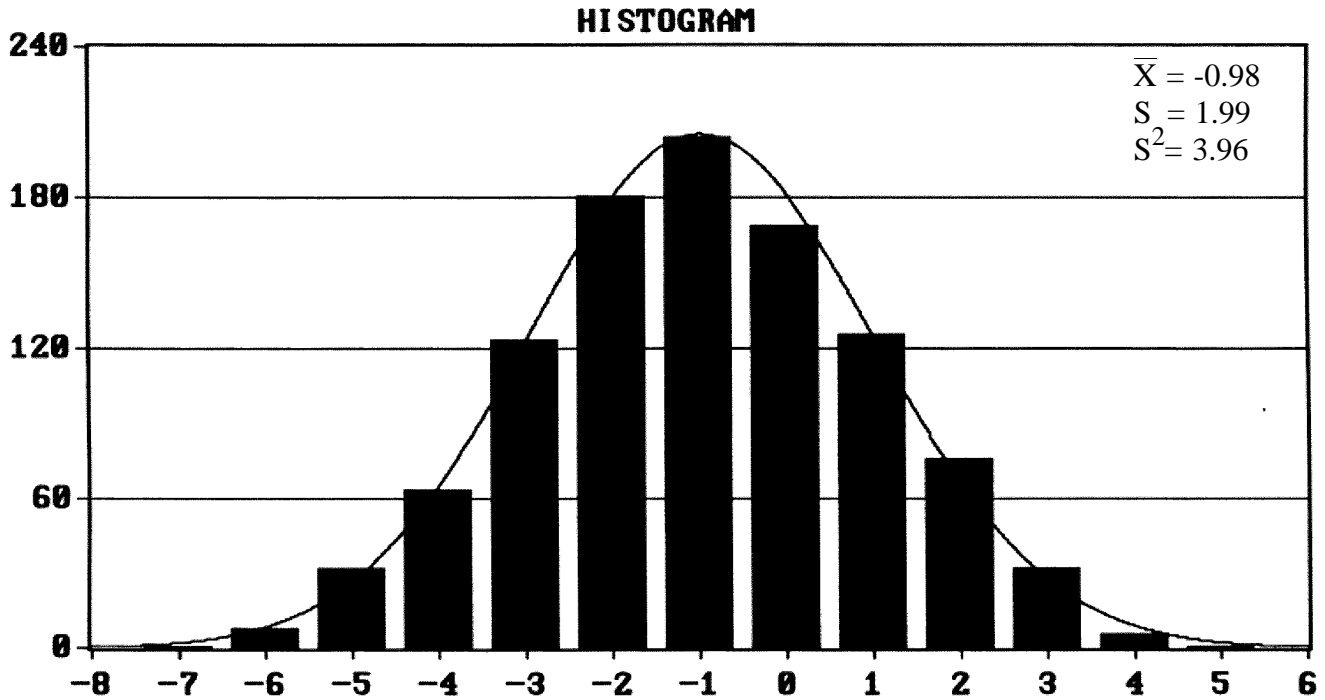


Figure 7. Histogram and Estimated PDF

(1 count = 5 volts / 2²⁰). The Full Scale Error is measured in a similar manner, by calculating a mean for a full-scale input signal.

The PDF's shape, which is used in noise calculations, is defined by the sample variance or its positive square root called the sample standard deviation (S). The standard deviation of a noise distribution is a measure of the rms or effective noise. In Figure 7, the rms noise is 1.99 counts or 9.49 μV. Additionally, the peak-to-peak noise can be calculated by using the standard deviation. Peak-to-peak noise is defined as the interval which contains six standard deviations. For a normal distribution, this interval represents 99.6% of the occurrences. In Figure 7, the peak-to-peak noise is 11.94 counts or 56.93 μV.

CONFIDENCE INTERVAL ESTIMATE

In the preceding section, the performance of an ADC was quantified from a sample set of data. For the data in Figure 7, the offset is calculated at -4.67 μV, and the noise at 9.49 μV rms and

56.93 μV peak-to-peak. Since these values are calculated from sampled data, a degree of uncertainty is associated with the estimated performance.

Using statistics, confidence intervals can be calculated for the estimated performance values. First a confidence interval for σ² is described. The accuracy of the model derived from sampled data depends upon how well the sample set resembles the actual system. The sample distribution alone is not useful in determining how well the sample variance correlates to the actual variance. However, a confidence interval can be obtained from the sample data. If the degrees of freedom (ν) and the actual system variance (σ²) are included with S², then the random variable [(n-1)S²/σ²] has a chi-squared distribution with ν degrees of freedom, where ν = n-1 for n number of samples.

$$\chi^2 = (n-1) \frac{S^2}{\sigma^2}$$

By treating the variance as a chi-squared variable, the variance PDF becomes a function of the degrees of freedom. Now the chi-squared distribution can be used to determine the accuracy of S^2 and state σ^2 as a confidence interval.

Chi-squared percentiles can be obtained from statistics tables. Statistic tables may not be available for large sample sizes. Fortunately, as the number of samples or degrees of freedom increases, a chi-squared distribution approaches a normal distribution. A good approximation of a chi-squared percentile, $\chi^2(\alpha;v)$, in terms of a standard normal percentile $[z(\alpha)]$ is given by:

$$\chi^2(\alpha;v) \approx 0.5(z(\alpha) + \sqrt{2v-1})^2, \quad v > 100$$

In the above formula, χ^2 is a function of α = area under the left tail of the standard normal curve and v = degrees of freedom.

For example, a 99% confidence interval is determined for χ^2 when $n=1024$:

Covering 99% of the total area leaves 1% uncovered, which is divided equally between the left and right tails. Thus the desired percentiles are $\alpha = 0.005$ and $\alpha = 0.995$. From a table for a Normal Distribution curve, the corresponding percentile points $z(\alpha)$ are -2.575 and 2.575 respectively. Thus 99% of the area under the standard Normal Distribution lies between -2.575 and 2.575.

Substituting values for $z(\alpha)$ and v in the $\chi^2(\alpha;v)$ formula, a range of χ^2 values is calculated.

$$\chi^2(0.005; 1023) = 909.37$$

$$\chi^2(0.995; 1023) = 1142.26$$

From these numbers, it can be stated that χ^2 has a 0.5% possibility of being less than 909.37 and 99.5% possibility of being less than 1142.26. Combining these two conditions, χ^2 has a 99% possibility of being greater than 909.37 and less than 1142.26.

Finally, a 99% confidence interval for σ^2 is obtained by substituting $(n-1)s^2/\sigma^2$ for χ^2 . Here the values $n = 1024$ and $S^2 = 3.96$ are used to further illustrate Figure 7.

$$909.37 < \chi^2 < 1142.26$$

$$909.37 < (n-1)s^2/\sigma^2 \text{ and } (n-1)s^2/\sigma^2 < 1142.26$$

$$\sigma^2 < (n-1) \frac{S^2}{909.37} \text{ and } (n-1) \frac{S^2}{1142.26} < \sigma^2$$

$$3.55 < \sigma^2 < 4.45$$

The calculated sample variance (S^2) is 3.96 LSB^2 s. There is a 99% confidence that the actual system variance (σ^2) is between 3.55 and 4.45 LSB^2 s. The uncertainty associated with 3.96 is less than a half LSB^2 ($4.45 - 3.96 = 0.49$). That is to say the maximum error is 0.49, with 99% confidence.

If the confidence interval or uncertainty is unacceptable, adjustments can be made. Notice that the interval width is effected by two variables $z(\alpha)$ and n . The interval width can be reduced by either increasing the sample set size or tolerating more uncertainty.

The following calculations using peak-to-peak noise show how $z(\alpha)$ and n affect the confidence interval. Peak-to-peak noise is the uncertainty associated with a single sample. Above, the 99% confidence interval was calculated for the variance. In Figure 7, the estimated variance is between 3.55 and 4.45 with 99% confidence.

Confidence	Number of Samples	α	$z(\alpha)$	χ^2	σ^2 Range	Peak-to-Peak Noise Range
99%	1024	0.005 0.995	$z(0.005) = -2.575$ $z(0.995) = 2.575$	909.37 1142.26	$3.55 < \sigma^2 < 4.45$	11.30 to 12.66
95%	1024	0.025 0.975	$z(0.025) = -1.96$ $z(0.975) = 1.96$	935.79 1113.06	$3.64 < \sigma^2 < 4.33$	11.45 to 12.49
99%	2048	0.005 0.995	$z(0.005) = -2.575$ $z(0.995) = 2.575$	1885.08 2214.55	$3.66 < \sigma^2 < 4.30$	11.48 to 12.44
95%	2048	0.025 0.975	$z(0.025) = -1.96$ $z(0.975) = 1.96$	1923.03 2173.81	$3.73 < \sigma^2 < 4.22$	11.59 to 12.33

Table 2. Degree of Confidence and Sample Size Effects on the Confidence Interval

This interval can be translated for peak-to-peak noise:

$$\begin{aligned} \text{peak-to-peak noise} &= 3\sigma + 3\sigma = 6\sigma \\ 6\sqrt{3.55} &< \text{peak-to-peak noise} < 6\sqrt{4.45} \\ 11.30 &< \text{peak-to-peak noise} < 12.66 \\ & \text{(99\% Confidence Interval)} \end{aligned}$$

The width of this confidence interval is $12.66 - 11.30 = 1.36$ counts. If the confidence is relaxed to 95%, $\chi^2(\alpha; \nu)$ is recalculated using $\alpha = 0.025$ and $\alpha = 0.975$. This results in the confidence interval width for peak-to-peak noise being reduced by 0.32 counts:

$$\begin{aligned} 11.45 &< \text{peak-to-peak noise} < 12.49 \\ & \text{(95\% Confidence Interval)} \end{aligned}$$

If the set is increased to 2048 samples and the sample variance remains at 3.96, $\chi^2(\alpha; \nu)$ is recalculated using $\nu = 2047$ and the 99% confidence interval becomes:

$$\begin{aligned} 11.48 &< \text{peak-to-peak noise} < 12.44 \\ & \text{(99\% Confidence Interval)} \end{aligned}$$

The original 99% confidence interval is reduced from 1.36 to 0.96 counts by increasing the sample set. As seen, thousands of samples may be required to reduce the peak-to-peak noise confi-

dence interval to an acceptable width. The size of the sample set depends upon system capabilities and performance requirements. Large data sets are affected by memory size and processing capabilities of the data collection equipment, and ADC throughput. Table 2 indicates how the degree of confidence and number of samples affect the confidence interval.

AVERAGING

Once the ADC noise has been characterized, the effect of averaging can be analyzed. When samples are collected, the average or sample mean (\bar{X}) is an estimator of the population mean (μ). \bar{X} is likely to estimate μ very closely when the sample size is large.

Again a confidence interval describes how closely \bar{X} estimates μ , and the sample size governs the width of the interval. The distribution of \bar{X} is Gaussian with a mean μ and variance $\frac{\sigma^2}{n}$. In the case where the ADC output is not Gaussian, the distribution of \bar{X} will approach the above gaussian distribution as n gets large, by the Central Limit Theorem. The confidence interval for μ is:

$$\bar{X} - z(\alpha) \frac{\sigma}{\sqrt{n}} < \mu < \bar{X} + z(\alpha) \frac{\sigma}{\sqrt{n}}$$

where α is set by the confidence interval.

Note that $z(\alpha) \sigma$ is the peak noise and $2 z(\alpha) \sigma$ is the peak-to-peak noise value. Restated, the actual mean differs from the sample mean by a range of the peak-to-peak noise divided by the square root of the number of samples. Thus averaging multiple samples reduces the error by $1/\sqrt{n}$.

$$\mu = \bar{X} \pm \frac{\text{peak-noise}}{\sqrt{n}}$$

The peak-to-peak noise of the sample set for Figure 7 is 11.94 counts. If one sample is taken, the 99.6% confidence interval is 11.94 counts or ± 5.97 counts. If all 1024 samples are averaged, the actual population mean is between -1.17 and -0.79 with a 99.6% confidence. The uncertainty is reduced to ± 0.19 counts. Note that the quantization error for an ideal ADC produces an error of ± 0.5 counts. Averaging 1024 samples improves this noisy 20-bit ADC's accuracy to better than 21 bits!

As shown above, averaging can reduce the effects of Gaussian distributed noise as well as quantization error. However, the tradeoff is in reduced throughput. To get the confidence interval to less than one count, \sqrt{n} has to be greater than the peak-to-peak noise. For the sample set of data plotted in Figure 7, 143 samples (11.94^2) need to be acquired and averaged. Over 36,496 samples are required to create a 24-bit ADC with less than one count of peak-to-peak noise (reduce the uncertainty of a 20-bit converter to

$\pm 1/32$ counts). This would reduce a 100kHz, ADC to an effective throughput of 2.74 Hz. Averaging sacrifices throughput for improved resolution and reduced uncertainty.

CONCLUSION

Statistical methods are available to measure the performance of an ADC. The testing involves inputting a noise free, accurate DC signal to the ADC and collecting a sample set of data points. The sample set is used to calculate estimators for the mean and standard deviation. More statistics are used to decide the "goodness" and confidence level associated with the estimates. Averaging was introduced for reducing uncertainty and improving resolution. However, averaging reduces the ADC's effective throughput. Figure 8 illustrates the tradeoff between reducing uncertainty and lowering the effective throughput.


The same methods used to measure an ADC's performance can be used to measure the performance of an entire system which includes additional components containing multiple noise sources and offsets. During system integration or production test, tests can be performed as subsystems are added. This can be used to measure the performance of individual subsystems or isolate problems to a subsystem or component. The results can then be used with compensation techniques to improve system performance or to determine corrective actions.

TABLE OF VARIABLES


μ mean
 σ standard deviation
 σ^2 variance
 $p(x)$ probability density function

\bar{X} sample mean
 S sample standard deviation
 S^2 sample variance

χ^2 chi-squared variable
 n number of samples
 ν degrees of freedom.
 α area under the normalized curve
 z standard normal distribution



PDF for a Gaussian random variable. The area $1 - \alpha$ is the confidence interval



PDF for a Chi-Square variable. The area $1 - \alpha$ is the confidence interval




Figure 8. Tradeoff between Uncertainty and throughput

REFERENCES

- [1] John Neter, William Wasserman, & G.A. Whitmore, "Applied Statistics" Allyn and Bacon, Inc.

- [2] Ronald E. Walpole & Raymond H. Myers, "Probability and Statistics for Engineers and Scientists", Macmillan Publishing Co., Inc, New York, 1978.

- [3] Richard H. Williams, "Electrical Engineering Probability", West Publishing Company.

- [4] Ferrel G. Stremmer, "Introduction to Communications Systems", Addison-Wesley Publishing Company, Inc. 1982.

• **Notes** •