IMTC 2004 – Instrumentation and Measurement
Technology Conference
Como, Italy, 18-20 May 2004

Improved Determination of the Best Fitting Sine Wave in ADC Testing

István Kollár[1] and Jerome J. Blair[2]

[1]Dept. of Measurement and Information Systems
Budapest University of Technology and Economics
H-1521 Budapest, Magyar tudósok krt. 2., HUNGARY
Phone: +36 1 463-1774, fax: +36 1 463-4112
Email: kollar@mit.bme.hu

[2]Bechtel Nevada
P.O.Box 98521
Las Vegas, Nevada 89193, USA
Phone: +1 702 295-2647, fax: +1 702 295-3305
Email: blairjj@nv.doe.gov

Abstract *– The sine wave test of an ADC means to excite the ADC with a pure sine wave, look for the sine wave which best fits the output in least squares sense, and analyze the difference. This is described in the IEEE standards 1241-2000 and 1057-1994.*

Least squares is the 'best' fitting method most of us can imagine, and it yields very good results indeed. Its known properties are achieved when the error (the deviation of the samples from the true sine wave) is random, white (the error samples are all independent), with zero mean Gaussian distribution. Then the LS fit coincides with the maximum likelihood estimate of the parameters.

However, in sine wave testing of ADC's these assumptions are far from being true. The quantization error is partly deterministic, and the sample values are strongly interdependent. This makes the sine fit worse than expected, and since small changes in the sine wave affect the residuals significantly, especially close to the peaks, ADC error analysis may become misleading. Processing of the residuals (e.g. the calculation of the effective number of bits, ENOB) can exhibit serious errors.

This paper describes this phenomenon, analyses its consequences, and suggests modified processing of samples and residuals to reduce the errors to negligible level.

Keywords *– effective number of bits, ENOB, IEEE Standard 1241-2000, IEEE Standard 1057-1994, ADC test, analog-to-digital converter, sine wave fitting, least squares.*

I. INTRODUCTION

Sine wave testing of ADC's is perhaps the most popular method for the evaluation of ADC converters. The standard IEEE 1241-2000 strongly builds on this: the 3-parameter and 4-parameter methods perform sine wave fitting, and continue by the analysis of the residual errors.

Theoretically, when we apply a sine wave, we can know its parameters (amplitude, frequency, phase, and in addition, DC offset), therefore we can determine the properties of the ADC from these and the output of the device. However, it is impractical to make use of the parameters which are usually not known with sufficient accuracy. Therefore, we rather determine these from the ADC output data. Since these data are not only quantized by the ideal ADC characteristic, but also prone to (usually small) ADC errors, we *fit them* with a sine wave, by adjusting its parameters. The fit is performed by setting the parameters to minimize a global measure of the level of deviations of the output data from the corresponding sine wave samples. In the standard, the sum of the squares of the errors is minimized, that is, a *least squares* (LS) fit is performed.

$$\min_{A,B,C,\omega} \sum_{n=1}^{M} \left(x_n - A\cos(\omega t_n) - B\sin(\omega t_n) - C\right)^2 \qquad (1)$$

If the frequency is known, the model is linear in the parameters, so *linear least squares* is performed (three-parameter method), while if the frequency is unknown, *nonlinear least squares is executed* (four-parameter method).

Least squares is one of the most powerful methods since Gauss invented it. The fit has very nice properties, especially when the error sequence (the difference between the samples and the model) is random, zero-mean Gaussian and white. In this case the LS fit coincides with the *maximum likelihood estimate* of the parameters, unbiased and minimum variance for the linear LS case, and with asymptotic unbiasedness, and asymptotically minimum variance for the general case.

The LS fit exists and is reasonable even in cases when the above assumptions concerning the error sequence are not met, but the properties of the estimates of the parameters may not be as nice as above. In the testing of ADC's, this is definitely the case: the ideal quantization error and the ADC nonlinearity error often dominate over observation noise, therefore they will also dominate the properties of the results. LS fitting does still work, as it usually works, especially because the error values are more or less scattered, but since the error is deterministically related to the input and to the transfer characteristic of the ADC, the result is prone to possible strange errors. Maybe the most apparent one is that the calculated *effective number of bits* (ENOB) depends on the amplitude of the sine wave and the dc value, and the calculated *ENOB* can change by about 0.1 for smaller amplitudes [7]. This is not a very large error, and only appears when random noise is small compared to quantization errors. This is why it was not treated in the standard until now, but since we know about it, and it can be handled by proper means, it makes sense to deal with it, by slight modification of the algorithm.

The purpose of this paper is to analyze the causes of the error, explore the possibilities of correction, and furthermore suggest an improvement of the standard.

II. PRELIMINARIES

In the IEEE standards [1-2] the above problem is not mentioned. They simply assume that the LS fit works well. Practitioners did already observe the phenomenon of the uncertainty in the determination of the ENOB, or of the equivalent resolution, but up to now no systematic treatment was suggested.

A possibility would be to treat the errors random, and use an improved model of their distribution to develop a maximum likelihood estimate of the sine parameters and of the *ENOB*. Since numerical software is available for minimizing different cost functions [4], this seems to be attractive – however, since nonlinear errors and noise are unknown beforehand, this is not a viable approach.

An attempt to solve the problem was presented in [6], utilizing the special shape of the ideal quantization error of the converted sine wave, but it assumes ideal (error-free) ADC characteristics and no noise, so it cannot be applied in practical tests.

An improvement to the method in the standard was presented in [7], but the improved approach, which corrects for the noise-free case, proved inferior to the IEEE standard method in the presence of substantial noise, so it may not be automatically applied.

We need a robust algorithm which works also on true ADC data, both when the noise dominates, and when the quantization error dominates. We will develop such an algorithm below.

III. SOURCES OF THE ERRORS

As described above, the calculation of the *ENOB* consists of two steps:

- fitting of the sine to the samples,
- calculation of the RMS value of the residuals.

The error in the determination of the ENOB depends on both steps, therefore we will analyze both.

A. Imprecise sine fit

In order to understand what happens, let us plot the quantization error of a sine wave in an ideal quantizer. We use a simple, rough quantization case which however illustrates the problems.

We can observe that the quantization error is more or less sawtooth-like (that is, uniformly distributed) at most places, *except at the peaks*. There it is almost constant, as the sine wave is also almost constant, and the level of the almost-constant error depends on the relation of the sine amplitude to the closest quantization level.

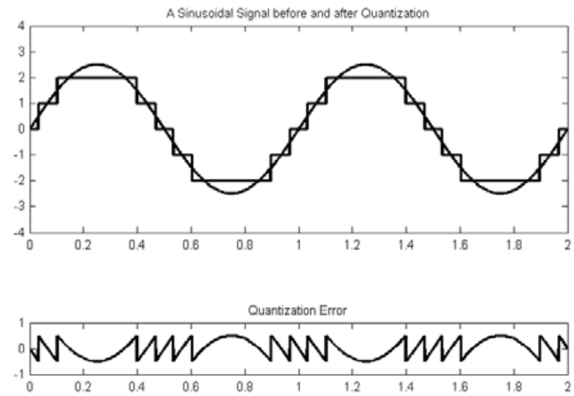

Fig. 1 Quantization of a sine wave (amplitude: 2.5Δ) in an ideal quantizer, with the dc offset is zero

According to the noise model of quantization, in the ideal case the PDF of the error should be uniform. However, as it can be seen from the probability density function, it is clearly not uniform: it has strong peaks, depending on the amplitude and the dc offset.

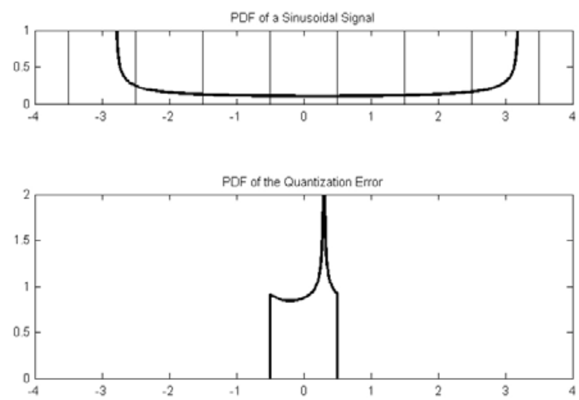

Fig. 2 Probability density function of a sinusoidal signal (amplitude: 3Δ), and that of the quantization error: n = Q(x) - x. The dc offset is dc = 0.2.

The LS fit tries to minimize the sum of the squares of the error samples. If there are dominant terms which can be reduced by modifying the DC level or the sine amplitude, the LS fit tends to decrease these terms, introducing a bias into the parameter estimates. The samples close to the peaks of the sine wave form such terms. For example, in Fig. 1 the corresponding quantized values are smaller than the input samples, that is, a slight decrease in the sine amplitude would bring the model closer to several quantized samples. Therefore, when this is the case, the estimated sine amplitude will be smaller than the true value. A similar, but amplitude-increasing case is when these samples are modified upwards by quantization. What is worse, it can well happen that the side bins are not totally filled with samples, only parts of them, as in Fig. 2, which makes their effect unpredictable.

It cannot be controlled how the sine peak is related to the quantization levels, thus these phenomena can happen at any time. Moreover, since this also depends on the dc level, it is even more uncontrollable.

In the following, until we specifically mention otherwise, all simulations are done with an ideal noiseless quantizer. This allows the investigation of the special case which, because of the deterministic pattern in the quantization errors, is most prone to the errors in question. When developing the modification, we will tacitly pay attention not to destroy the noisy case – afterwards, we will double-check that we caused no extra problem in the noisy/nonlinear cases.

Figure 3 illustrates the estimation of the amplitude as a function of the amplitude itself and of the dc level. We see that the error in the amplitude estimation can take any value between [-0.1,0.046].

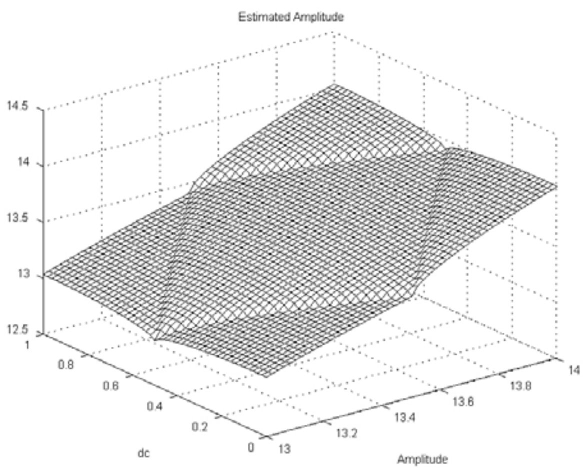

Fig. 3 Estimated value of the amplitude, as a function of the true value of the amplitude and of the dc. The deviation is between [-0.1,0.046].

Correction of this would be possible if we could make use of the pattern of the error: however, since ADC nonlinearities and noise destroy the nice pattern, this is impossible. It is, however, possible *to identify the samples which cause the problem* and eliminate them. The difficulty is that in practical cases the true input signal is not known, therefore we cannot rely on the input amplitude. However, as we are going to see, the output samples of the ADC provide enough information.

We want to eliminate the samples falling in the partly filled bins, but we have no direct information about the distribution in the bins, we only have the histogram values (the numbers of samples in each bin). Fortunately, we know that in the histogram of a sine wave the values increase towards each side. Therefore, the maximum values of the histogram at the two sides either represent the bins which contain the sine peaks and are only partly filled, or, if the quantization level is close to the sine peak, mark the bins *just below* the partly filled bins. Therefore, by looking for the maxima in the histogram (or, more directly, looking for the most often occurring ASDC output values close to the two extremes) we can determine an interval of the ADC output samples, which does not contain samples corresponding to the 'pathological' bins.

Let us thus preprocess the ADC output values themselves before any kind of fit, by eliminating the most often occurring positive value and the values above it, and, similarly, by eliminating the most often occurring negative value and the values below it. By this, we will certainly eliminate the pathological part of the sine wave (maybe, in addition, one more bin, but this will not significantly deteriorate the estimated parameters of the sine). Since Eq. (1) just fits the available samples and does not utilize the continuous nature of the record, we can simply drop the not wanted samples and the corresponding t_n values from the

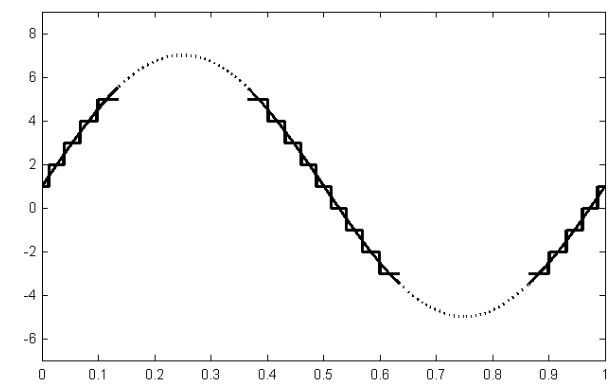

Fig. 4 Elimination of the samples corresponding to 'pathological' bins

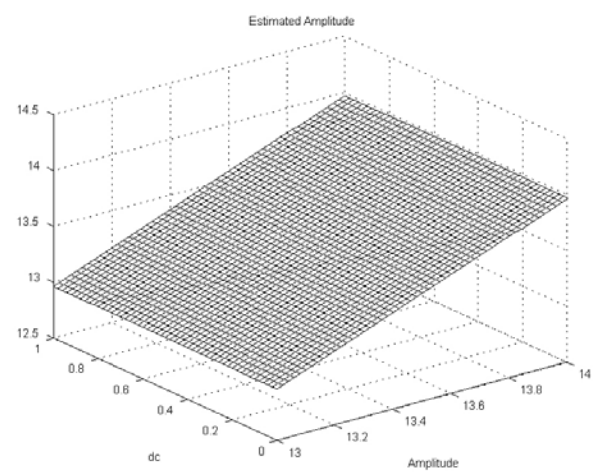

Fig. 5 Estimated value of the amplitude after elimination of the 'pathological' bins, as a function of the true values of the sine amplitude and of the dc. The deviation is between [-0.045,-0.033], significantly decreased by the modification

formula. This principle can be also explained as follows: the samples between the 'pathological' bins contain enough information concerning the sine peak, so elimination of the 'pathological' samples makes no harm.

By this manipulation, the amplitude error significantly decreases.It is interesting to observe that by elimination of the 'pathological' bins, the fit yields amplitudes which are much more accurate than without this manipulation, but are systematically underestimated. The reason is simple: since the sine wave is concave at the positive wave half, and convex at the negative wave half, we have more values rounded towards zero than ones rounded in the 'outside' direction. The cost function therefore overemphasizes the former ones, and minimization decreases them a little by setting the amplitude somewhat smaller than the ideal value. This can also be compensated for, but in the case of the *ENOB* we will see that we need not further improve the amplitude estimate.

B. Imprecise ENOB calculation

The definition of the effective number of bits can be given in several forms. Maybe the most useful is

$$ENOB = B - \log_2\left(\frac{RMS}{\Delta/\sqrt{12}}\right) \quad (2)$$

where B is the nominal number of bits, RMS is the measured rms value of the residuals, and Δ is the nominal quantum size (LSB size) of the ADC. For an ideal noiseless quantizer, $ENOB = B$.

RMS is nothing else than the minimized cost function in Eq. (1), divided by the number of samples M. Usually $ENOB<B$. However, when the measured RMS is smaller than $\Delta/\sqrt{12}$, the measured $ENOB$ is larger. This is the case for an ideal, noiseless quantizer when the deterministic pattern of

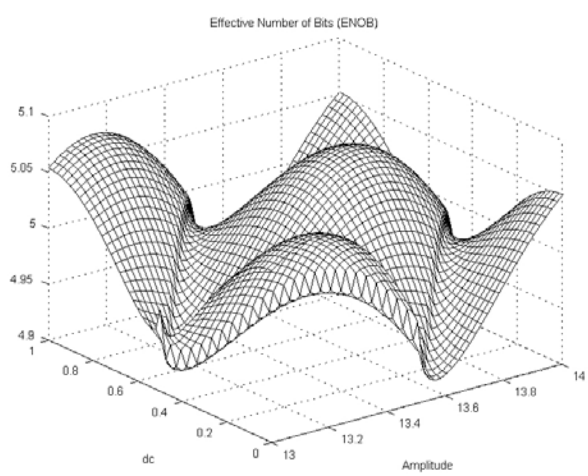

Fig. 6 *ENOB* value calculated according to the standards as a function of the true values of the sine amplitude and of the dc. Limits of the ENOB values: [4.92,5. 07]

the error samples allows to make a 'too good' fit by selecting a smaller amplitude.

We know from earlier results [7] that, especially for small amplitudes, the calculated *ENOB* can be erroneous. E.g. for a 5-bit quantizer this value seriously depends on the amplitude and on the dc value.

The cause is again the effect of the 'pathological' bins. The asymmetrically filled-in bin(s) may cause serious deviations. The remedy is the same as above: eliminate the undesired bins.

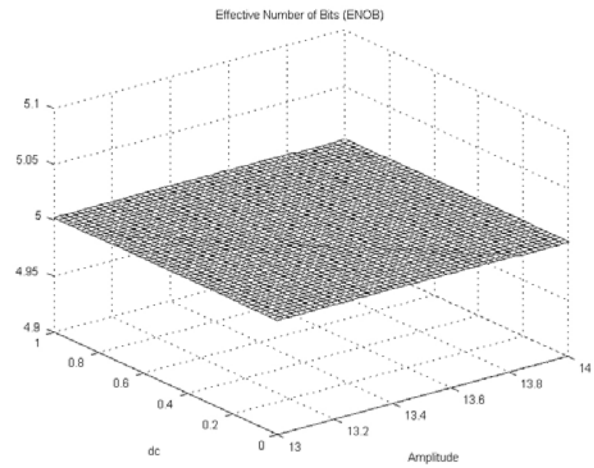

Fig. 7 *ENOB* value as a function of the true value of the amplitude and of the dc, after elimination of the 'pathological' bins. Limits: [5.001,5.003]

As we observe in Fig. 7, the *ENOB* is precisely measured. With no quantization error or noise, it is somewhat larger than the theoretical value – this is because of the somewhat small estimated sine amplitude. However, the observable deviation is already negligible.

Knowing the cause of the problem, and having the remedy, we can pose the question: Which mechanism is the *ENOB* most influenced by: the improper setting of the amplitude, or the improper estimation of the rms. Therefore, a last calculation has been executed: a (maybe distorted) sine amplitude estimation was made on all samples, and the 'pathological' bins were eliminated for the *ENOB* calculation only.

We can state that the main reason of the error is the effect through *ENOB* calculation, but the effect of the imprecise amplitude calculation is also not negligible. To have a real good *ENOB* value even for small numbers of bits, we can suggest to modify the algorithm by dropping the samples corresponding to the two maximum height bins and those which are outside these in amplitude.

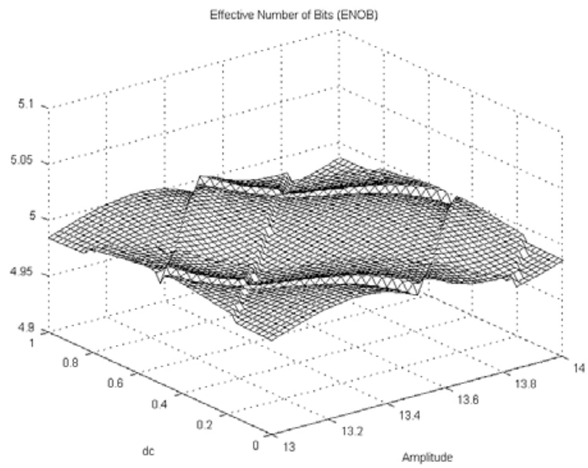

Fig. 8 *ENOB* value as a function of the true value of the amplitude and of the dc, after elimination of the 'pathological' bins for the *ENOB* calculation only. Limits: [4.97,5.002].

IV. THE EFFECT OF NOISE

Until now, we have investigated the noiseless ideal quantizer, although we did not make use of the ideal nature of it. But, to be sure that this method can be equally well used for true ADC's, we need to try what happens when noisy ADC's are used.

First of all, let us mention that additive noise at the input acts like dither. This means that moments of the quantized samples (the ADC output samples) are less biased than without: we expect that even for non-manipulated processing the calculated amplitude and *ENOB* values are closer to the true values. Therefore, manipulation may not be necessary, at least when the noise is large enough. However, if we can show that manipulation makes no harm, we may suggest the modification for all cases.

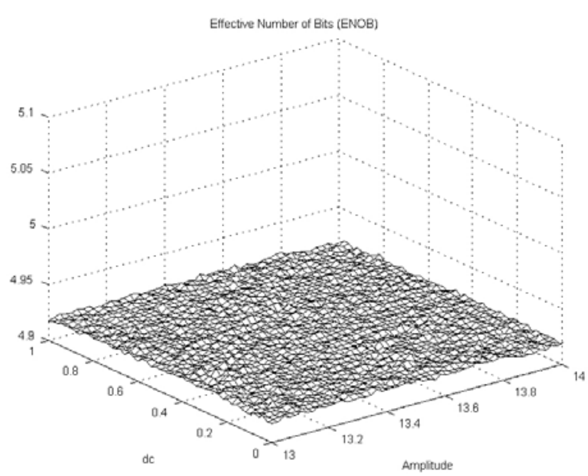

Fig. 9 Modified algorithm with additive Gaussian noise, $\sigma = 0.1\Delta$, $M = 10^5$. The theoretical *ENOB* is 4.92, the limits are [4.914,4.924].

We can observe in Fig. 9 that no additional error occurs by applying the algorithm to noisy ADC data.

Finally, let us explore what happens in the case of larger noise.

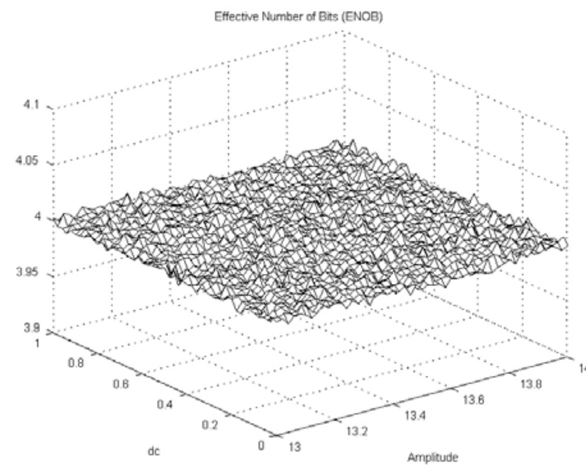

Fig. 10 Original algorithm, sine with additive Gaussian noise, $\sigma_n = 0.5\Delta$ ($\sigma_n = 1.7\sigma_q$), $M = 10^5$. The theoretical *ENOB* is 4.0, the limits are [3.99,4.011].

In Fig. 10 the error caused by the noise dominates. The modification of the algorithm only slightly changes much the calculated *ENOB*. In both cases the noise dominates the error.

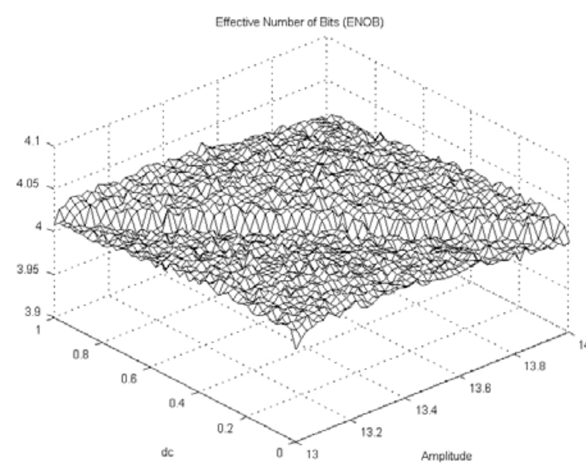

Fig. 11 Modified algorithm, sine with additive Gaussian noise, $\sigma_n = 0.5\Delta$ ($\sigma_n = 1.7\sigma_q$), $M = 10^5$. The theoretical *ENOB* is 4.0, the limits are [3.999,4.04].

V. GLOBAL DESCRIPTION OF AN ADC

We have seen that both for noiseless and for noisy ADC's the modified algorithm works well. A true ADC usually differs from these primarily in one important aspect: nonlinearity effects. Is this a disturbing factor?

Considering the above thoughts we can speculate that elimination of the ‘pathological’ bins causes no harm for real ADC’s, on the contrary, it eliminates the arbitrary weight of the side bins. However, understanding this, we can observe another problem. For uniformly sampled sinusoidal data, the larger amplitudes are overemphasized by using here far more samples. Therefore, the calculated *ENOB* refers primarily to the errors and noise *at both sides of the* characteristic, rather than giving an overall picture. Consequently, it makes sense for *ENOB* calculation, after elimination of the ‘pathological’ bins, to weight the residuals during rms calculation by the reciprocal of the PDF of the fitted sine. By this, a better overall characterization of the ADC can be obtained, instead of one referring primarily to the larger amplitudes. The details of this weighting will be discussed in another paper.

VI. SUMMARY

We have identified two error sources in *ENOB* calculation, based on a sine fit:

a) bias in the determined sine amplitude, caused by ‘pathological’ bins,
b) effect of the errors miscalculated around the sine amplitudes.

We have suggested a method which deals with these problems, and, consequently, suggested a modification of the algorithms in the standards:

- eliminate the contents of the maximum-value histogram bins at the edges, along with all the values surpassing these,
- make a sine fit on the pre-treated data,
- use the error samples in the pre-treated data for the calculation of the *ENOB*,
- for an improved overall description of an ADC, calculate the *ENOB* using the inverse sinusoidal weight.

REFERENCES

[1] IEEE Standard 1241-2000, “IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters.”

[2] IEEE Standard 1057-1994, “IEEE Standard Digitizing Waveform Recorders.”

[3] Blair, J., “Sine-fitting Software for IEEE Standards 1057 and 1241,” Proc. IEEE Instrumentation and Measurement Technology Conference, IMTC/1999, Venice, May 24-26, 1999. pp. 1504-6.

[4] van den Bos, A., “Estimation of Fourier Coefficients,” IEEE Trans. on Instrumentation and Measurement, Vol. 38, No. 5, Oct. 1989, pp. pp. 1005-1007.

[5] Kollár, I., “Evaluation of Sine Wave Tests of ADCs from Windowed Data.” Computer Standards and Interfaces, Vol. 22, pp. 261-68, 2000.

[6] Hejn, K., and A. Pacut, “Effective Resolution of Analog to Digital Converters – Evolution of Accuracy,” IEEE Instrumentation and Measurement Magazine, Vol. 6, No. 3, pp. 48-55, Sep. 2003.

[7] Blair, J., and T. Linnenbrink, “Corrected RMS error and effective number of bits for sinewave ADC tests,” Computer Standards and Interfaces, Vol. 26, No. 1, pp. 43-49. Jan. 2004.