Improved Determination of the Best Fitting Sine Wave in ADC Testing

István Kollár, *Fellow, IEEE,* and Jerome J. Blair, *Fellow, IEEE*

***Abstract*—The sine wave test of an analog-to-digital converter (ADC) means to excite the ADC with a pure sine wave, look for the sine wave which best fits the output in least squares (LS) 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 ADCs, these assumptions are far from being true. The quantization error is partly deterministic, and the sample values are strongly interdependent. For sine waves covering less than, say, 20 quantum levels, this makes the sine wave 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, analyzes its consequences, and suggests modified processing of samples and residuals to reduce the errors to negligible level.**

***Index Terms*—Analog-to-digital converter (ADC) test, analog-to-digital converter, effective number of bits (ENOB), four-parameter method, IEEE Standard 1057–1994, IEEE Standard 1241–2000, least squares (LS), sine wave fitting, three-parameter method.**

I. Introduction

SINE wave testing of analog-to-digital converters (ADCs) is perhaps the most popular method for the evaluation of ADC converters. IEEE Standard 1241–2000 strongly builds on this: the three-parameter and four-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 would 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

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

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 well behaved as above. In the testing of ADCs, 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 small amplitudes (below 15–20 quantum levels) [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. Background

In the IEEE standards [1], [2], the above problem is not mentioned. These simply assume that the LS fit works well. Practitioners have already observed the phenomenon of the

Manuscript received June 15, 2004; revised June 30, 2005.

I. Kollár is with the Department of Measurement and Information Systems, Budapest University of Technology and Economics, H-1521 Budapest, Hungary (e-mail: kollar@mit.bme.hu).

J. J. Blair is with Bechtel Nevada, Las Vegas, NV 89193 USA (e-mail: blairjj@nv.doe.gov).

Digital Object Identifier 10.1109/TIM.2005.855082

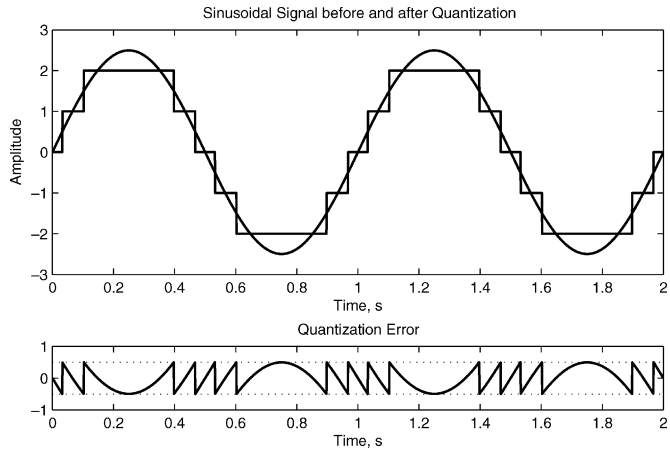

Fig. 1. Quantization of a sine wave (amplitude: 2.5Δ) in an ideal quantizer, with the dc offset is zero, and $\Delta = 1$.

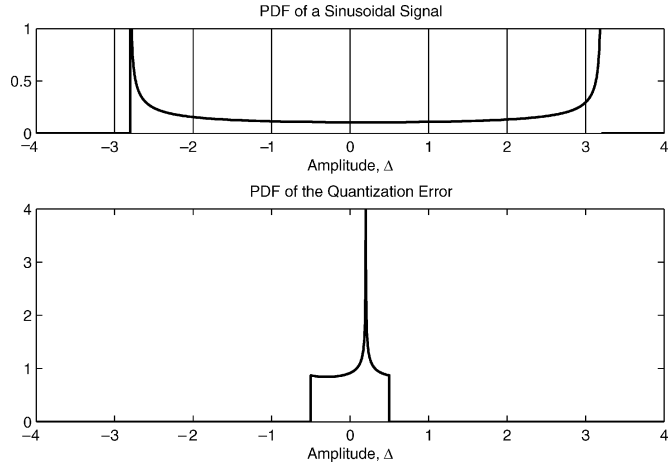

Fig. 2. Probability density function of a sinusoidal signal (amplitude: 3Δ), and that of the quantization error: $n = Q(x) - x$. The dc offset is dc $= 0.2$.

uncertainty in the determination of the ENOB, or of the equivalent resolution [7], but up to now, no systematic treatment was suggested.

A possibility would be to treat the errors as random errors, and use an improved model of their distribution to develop a maximum likelihood estimate of the sine wave parameters and of the ENOB. Since numerical software is available for minimizing different cost functions [4], this seems to be attractive. However, since the exact properties of the nonlinear errors and the noise levels 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 in this paper.

III. Sources of the Errors

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

- fitting of the sine wave 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 Wave Fit

In order to understand what happens, let us plot the quantization error of a sine wave in an ideal quantizer as shown in Fig. 1. We use a simple, coarse quantization case which 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 wave amplitude to the closest quantization level.

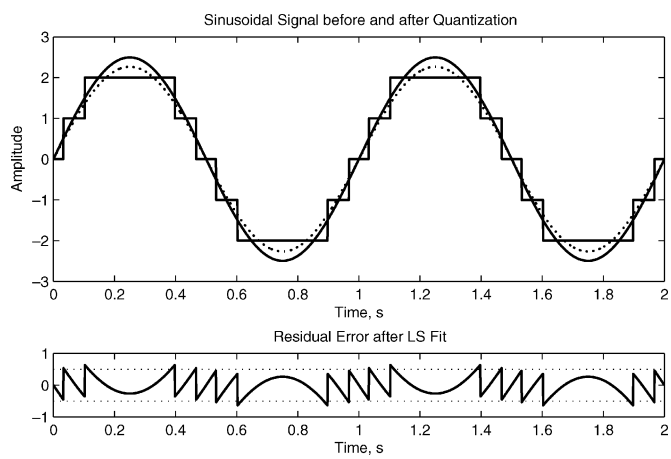

Fig. 3. Fitted sine wave (dotted line) and residuals in an ideal quantizer. $\Delta = 1$.

According to the noise model of quantization, in the ideal case, the probability density function (pdf) of the error should be uniform. However, as it can be seen from the pdf shown in Fig. 2, it is clearly not uniform: it has strong peaks, depending on the amplitude and the dc offset.

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 wave 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 wave amplitude would bring the model closer to several quantized samples. Therefore, when this is the case, the estimated sine wave amplitude will be smaller than the true value (Fig. 3).

A similar, but amplitude-increasing case is when these samples are modified upwards by quantization. What is worse, the side bins maybe only partially filled with samples, as in Fig. 2

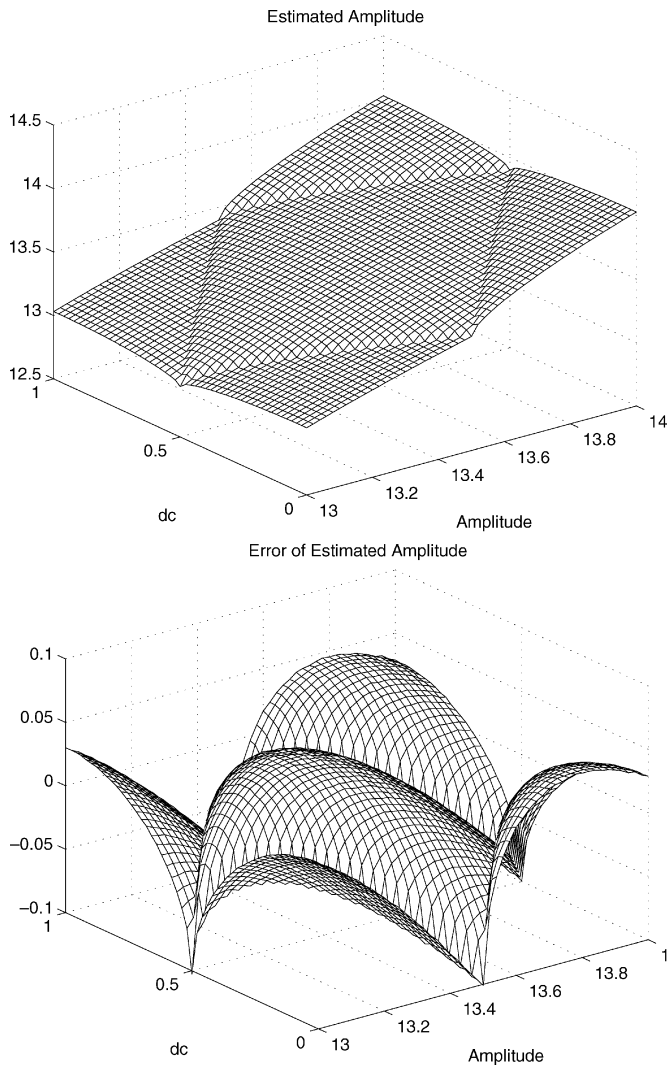

Fig. 4. Estimated value of the amplitude, as a function of the true value of the amplitude and of the dc. The deviation is between [−0.1, 0.046].

(the amplitudes of the sine wave do not reach the sides of the largest and smallest bins, respectively), which makes their effect unpredictable.

How the sine wave peak relates to the quantization levels cannot be controlled. Thus, these phenomena can occur at any time. Moreover, since these phenomena also depend on the dc level, their influence is even more uncontrollable.

In the following discussion, until we specifically mention otherwise, all simulations are done with an ideal, noiseless quantizer. This approach 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. The amplitude of the sine wave is changed in [13,14], the dc offset in [0,1].

When developing the modification, we will tacitly pay attention not to destroy the noisy case. Later, we will double-check that we have caused no extra problem in the noisy/nonlinear cases.

Fig. 4 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].

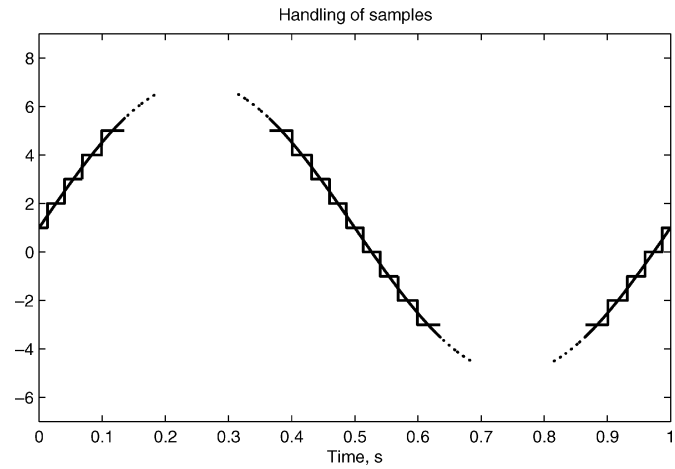

Fig. 5. Elimination of the samples corresponding to "pathological" bins.

Correction of this error 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 its input amplitude. However, as we are going to see, the output samples of the ADC provide enough information concerning this.

B. Suggested Algorithm

We want to eliminate the samples falling in the partly filled bins (see Fig. 2), 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 histogram values increase toward each side. Therefore, the maximum values of the histogram at the two sides:

- either represent the bins which contain the sine wave peaks and are only partly filled;
- or, if the quantization level is close to the sine wave peak, nearer to the zero than to the 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 ADC output values close to the two extremes) we can determine the intervals of the ADC output samples, which may contain samples corresponding to the "pathological" bins, and exclude them from further calculations.

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 as shown in Fig. 5. By this method, we will certainly eliminate the pathological part of the sine wave plus, perhaps one more bin on each side. This procedure will not significantly deteriorate the estimated parameters of the sine wave as shown in Fig. 6. Since (1) just fits the available samples and does not utilize the continuous nature of the record, we can simply drop the undesirable samples and the corresponding t_n values from the formula. This principle can be also explained as follows: the samples, we

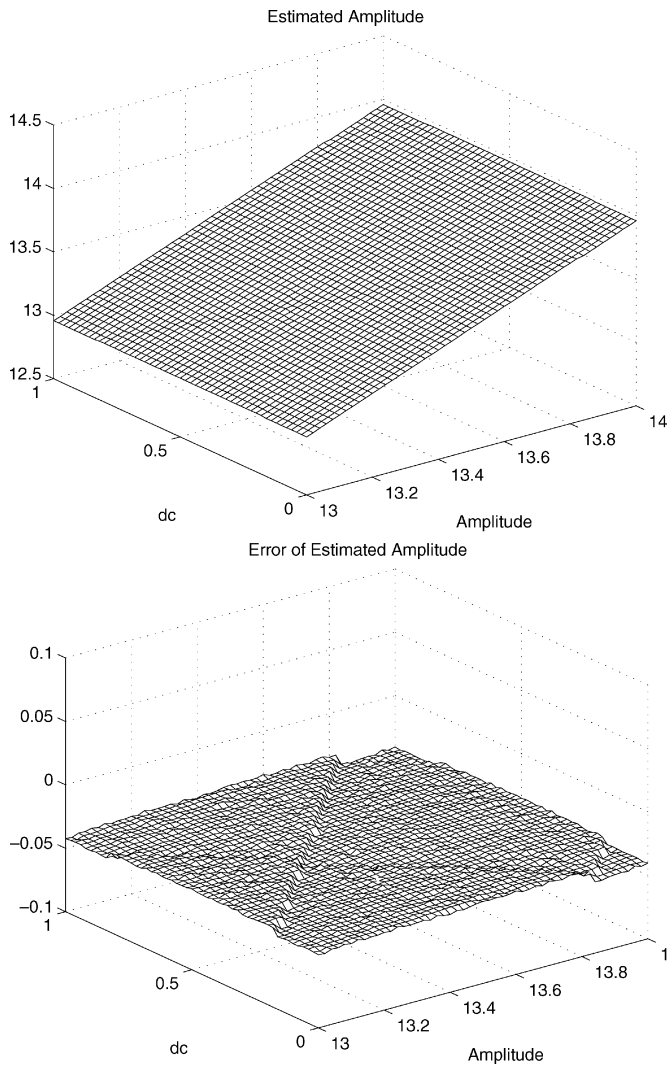

Fig. 6. Estimated value of the amplitude after elimination of the "pathological" bins, as a function of the true values of the sine wave amplitude and of the dc. The deviation is between $[-0.045, -0.033]$, significantly decreased by the modification.

rely upon (see Fig. 5), contain sufficient information concerning the peak of the sine wave, so that eliminating the "pathological" samples does not significantly degrade the fit.

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 toward 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.

We should realize, however, that the above described manipulation is only necessary when the sine wave covers a few quantization intervals only. When the sine wave is larger, the distorting effect becomes negligible.

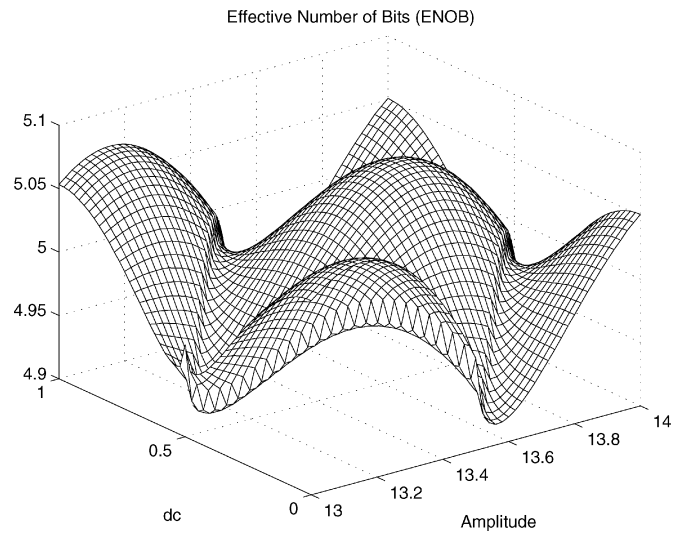

Fig. 7. ENOB value calculated according to the standards as a function of the true values of the sine wave amplitude and of the dc. Limits of the ENOB values: [4.92, 5. 07].

C. Imprecise ENOB Calculation

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

$$\mathrm{ENOB} = B - \log_2\left(\frac{\mathrm{rms}}{\frac{\Delta}{\sqrt{12}}}\right) \tag{2}$$

where B is the nominal number of bits, rms is the measured root-mean-square value of the residuals, and Δ is the nominal quantum size (LSB size) of the ADC. For an ideal noiseless quantizer, ENOB should be equal to B.

The rms is nothing else than the minimized cost function in (1), divided by the number of samples M. Usually ENOB $< B$, since the noise increases the rms, decreasing ENOB by this. 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 the error samples enable a "too good" fit by selecting smaller amplitude (see Fig. 7).

We know from earlier results [7] that, especially for small amplitudes, the calculated ENOB can be erroneous. For example, for a 5-bit quantizer, this value significantly 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.

As we observe in Fig. 8, when eliminating the "pathological" bins, 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 wave amplitude. However, the observable deviation is already negligible.

Knowing the cause of the problem, and having the remedy, we can pose the question: By which mechanism is the calculated value of the ENOB most influenced: by the improper setting of the amplitude, or by the improper estimation of the rms. Therefore, a last calculation has been executed: *a (perhaps distorted) sine wave-amplitude estimation was made on all samples, and the "pathological" bins were eliminated for the ENOB calculation only (see Fig. 9).*

We can state that the dominant cause of the error is inclusion of the "pathological" bins into the ENOB calculation, but their

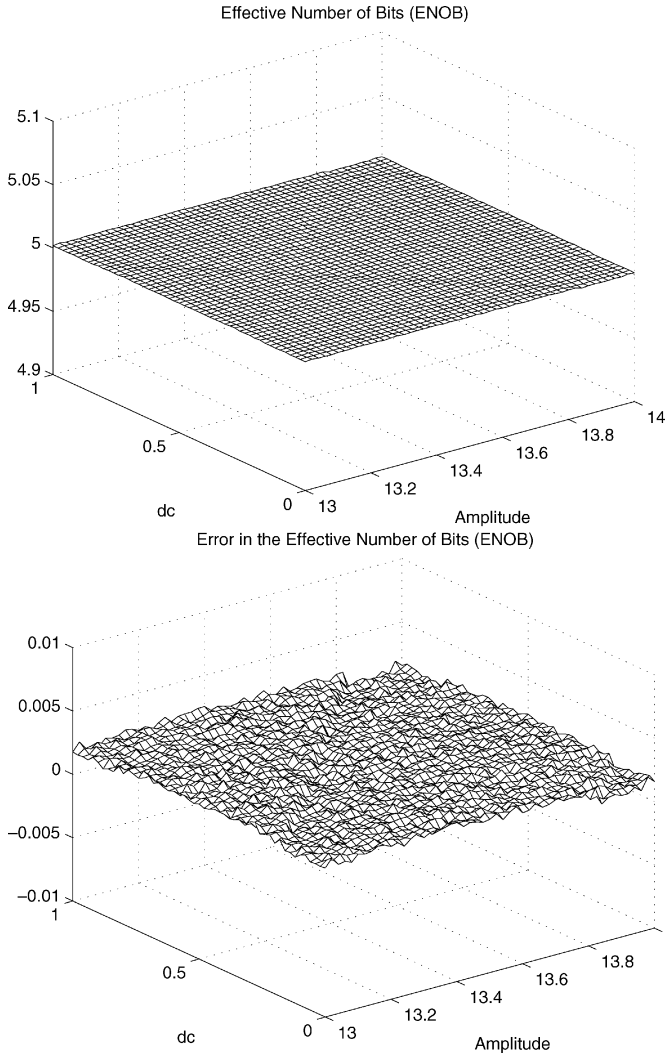

Fig. 8. ENOB value as a function of the true value of the amplitude and of the dc, after elimination of the "pathological" bins. Limits: [5.001, 5.003].

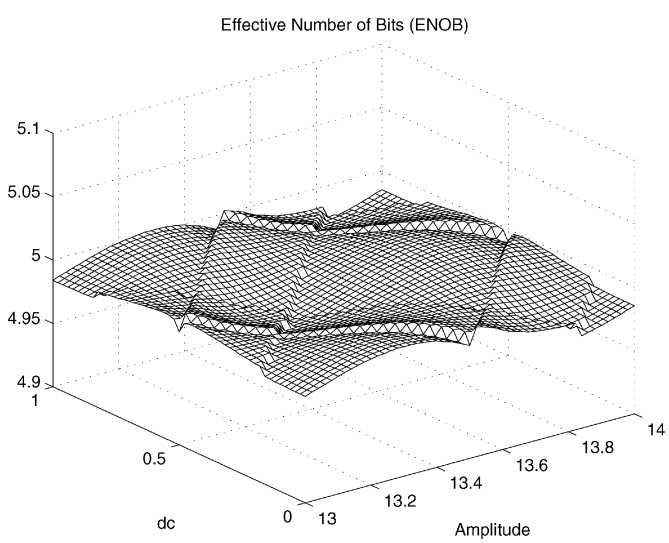

Fig. 9. ENOB value as a function of the true value of the amplitude and of the dc, after elimination of the "pathological" bins for the ENOB calculation only. Limits: [4.97, 5.002].

effect on amplitude calculation is also noticeable. To have a really good ENOB value even for small numbers of bits, *we suggest modifying the algorithm by disregarding the samples which fall in the two maximum height bins and those which are larger in absolute value than these, as it is suggested in Section III-B.*

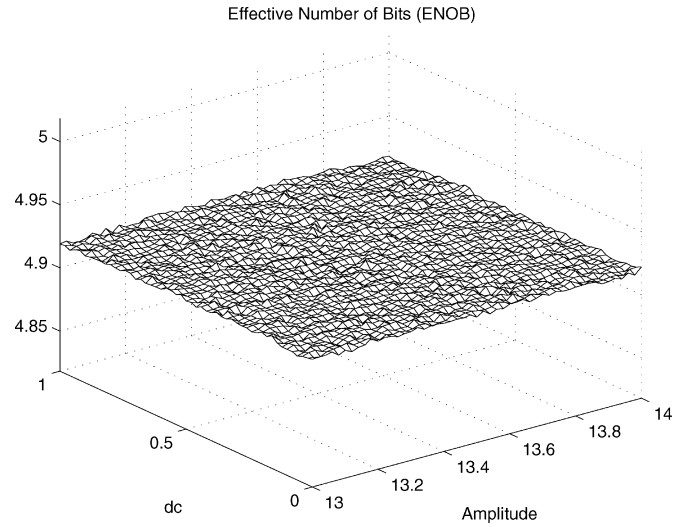

Fig. 10. Modified algorithm with additive Gaussian noise, $\sigma = 0.1\Delta$, $M = 10^5$. The theoretical ENOB is 4.92, the limits are [4.914, 4.924].

IV. Effect of Noise

Until now, we have investigated the noiseless ideal quantizer, although we did not make use of the ideal nature of it. However, to be sure that this method can be equally well used for real ADCs, we need to see what happens when noisy ADCs are used.

First of all, let us mention that additive noise at the input acts like dither [9]. This means that moments of the quantized samples (the ADC output samples) are less biased with noise than without it: we expect that even for nonmanipulated 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 this manipulation does not deteriorate the result, we may suggest the modification for all cases.

The argument for this is that dither in general decreases bias, thus additive noise works in the right direction, while disregarding a small fraction of the samples only slightly increases the variance. Indeed, as it is illustrated by Fig. 10, no additional error occurs by applying the algorithm to noisy ADC data.

In Fig. 11, the noise is larger than in Fig. 10; therefore, the error caused by the noise dominates. The modification of the algorithm only slightly changes the calculated ENOB (Fig. 12). In both cases, the noise dominates the error.

V. Global Description of a True ADC

We have seen that both for noiseless and for noisy ADCs, 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 ADCs. On the contrary, it eliminates the arbitrary weight of the side bins, and gives a good average measure of the ADC errors by the ENOB.

VI. Further Error Sources

Having done the above modification, we can observe a secondary error source. For uniformly sampled sinusoidal data,

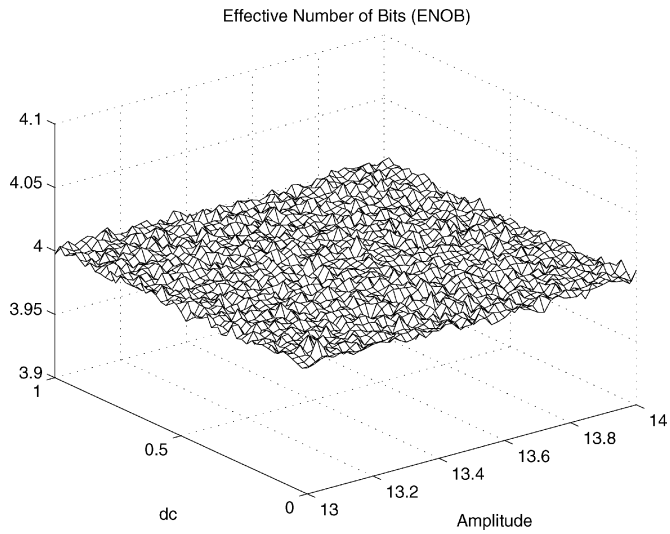

Fig. 11. Original algorithm, sine wave with additive Gaussian noise, $\sigma_n = 0.5\Delta$ ($\sigma_n = 1.7\sigma_q$), $M = 10^5$. The theoretical ENOB is 4.0, the limits are [3.99, 4.011].

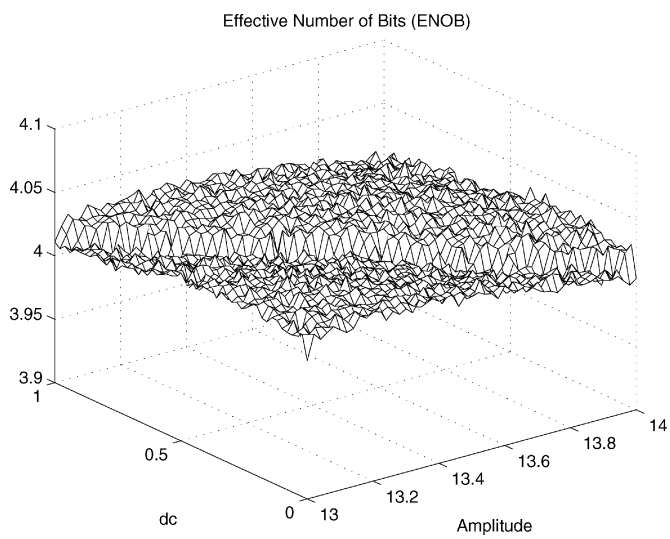

Fig. 12. Modified algorithm, sine wave with additive Gaussian noise, $\sigma_n = 0.5\Delta$ ($\sigma_n = 1.7\sigma_q$), $M = 10^5$. The theoretical ENOB is 4.0, the limits are [3.999, 4.04].

larger samples are overemphasized in each bin by using more samples around them than elsewhere. Therefore, the calculated ENOB is still somewhat distorted. Consequently, after elimination of the "pathological" bins, we can weight the residuals during rms calculation in each bin by the reciprocal of the pdf in the bin. However, this modification brings little additional benefit [10], so here we do not go into any more details.

VII. Summary

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

1) bias in the determined sine wave amplitude, caused by "pathological" bins;
2) effect of the errors overrepresented around the sine wave 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 wave fit on the pretreated data,
- use the error samples in the pretreated data for the calculation of the ENOB.

References

[1] *IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters*, IEEE Standard 1241–2000, 2000.
[2] *IEEE Standard Digitizing Waveform Recorders*, IEEE Standard 1057–1994, 1994.
[3] J. Blair, "Sine-fitting software for IEEE standards 1057 and 1241," in *Proc. IEEE Instrumentation Measurement Technology Conf.*, Venice, Italy, May 24–26, 1999, pp. 1504–6.
[4] A. van den Bos, "Estimation of Fourier coefficients," *IEEE Trans. Instrum. Meas.*, vol. 38, no. 5, pp. 1005–1007, Oct. 1989.
[5] I. Kollár, "Evaluation of sine wave tests of ADC's from windowed data," *Comput. Stand. Interfaces*, vol. 22, pp. 261–268, 2000.
[6] K. Hejn and A. Pacut, "Effective resolution of analog to digital converters—Evolution of accuracy," *IEEE Instrum. Meas. Mag.*, vol. 6, no. 3, pp. 48–55, Sep. 2003.
[7] J. Blair and T. Linnenbrink, "Corrected rms error and effective number of bits for sinewave ADC tests," *Comput. Stand. Interfaces*, vol. 26, no. 1, pp. 43–49, Jan. 2004.
[8] T. Z. Bilau, T. Megyeri, A. Sárhegyi, J. Márkus, and I. Kollár, "Four-parameter fitting of sine wave testing result: Iteration and convergence," *Comput. Stand. Interfaces*, vol. 26, no. 1, pp. 51–56, Jan. 2004.
[9] R. A. Wannamaker, S. P. Lipshitz, J. Vanderkooy, and J. N. Wright, "A theory of nonsubtractive dither," *IEEE Trans. Signal Process.*, vol. 48, no. 2, pp. 499–516, Feb. 2000.
[10] I. Kollár, "Improved residual analysis in ADC testing," in *Proc. 13th Int. Symp. Measurements Research Industry Applications, 9th Workshop ADC Modeling Testing*, Athens, Greece, Sep. 29–Oct. 1, 2004, pp. 869–874.

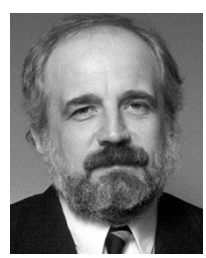

István Kollár (M'87–SM'93–F'97) was born in Budapest, Hungary, in 1954. He received the M.S. degree in electrical engineering from the Technical University of Budapest, Budapest, in 1977, and the "Candidate of Sciences" and "Doctor of the Academy" degrees from the Hungarian Academy of Sciences, Budapest, in 1985 and 1998, respectively.

From 1989 to 1990, and during spring of 2005, he was a Visiting Scientist at the Vrije Universiteit Brussel, Brussels, Belgium. From 1993 to 1995, he was a Fulbright Scholar and Visiting Associate Professor in the Department of Electrical Engineering, Stanford University, Stanford, CA. Currently, he is a Professor of electrical engineering at the Budapest University of Technology and Economics, Budapest, Hungary. His research interests are in the areas of digital and analog signal processing, measurement theory, and system identification. He has published about 110 scientific papers, and is the coauthor of the book *Technology of Electrical Measurements* (New York: Wiley, 1993). He authored the Frequency Domain System Identification Toolbox for Matlab. He was Editor of *Periodica Polytechnica Ser. Electrical Engineering* from 1988 to 1991, then Chief Editor of *Periodica Polytechnica* from 1991 to 1997. He was Technical Co-Chairman of IMTC/2001.

Prof. Kollár is a Member of the Administrative Committee of the Instrumentation and Measurement Society. He is also an active member of EUPAS in IMEKO.

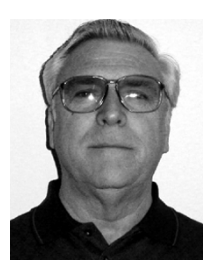

Jerome J. Blair (M'89–SM'94–F'04) received the B.S. degree in engineering physics in 1966 and the Ph.D. degree in applied mathematics in 1970, both from the University of California, Berkeley.

Since 1970, he has worked for EG&G Energy Measurements and Bechtel Nevada, Las Vegas, NV, designing, characterizing, and evaluating complex measurement systems and their software.

Dr. Blair has been an active Member of the IEEE Instrumentation and Measurement Society's TC-10, the Waveform Generation, Measurement, and Analysis Committee since 1989. He is an Associate Editor for the IEEE Transactions on Instrumentation and Measurement.