www.czasopisma.pan.pl PAN www.journals.pan.pl

Metrol. Meas. Syst., Vol. XXII (2015), No. 1, pp. 89–100.

METROLOGY AND MEASUREMENT SYSTEMS

Index 330930, ISSN 0860-8229

www.metrology.pg.gda.pl

ON THE VALIDITY OF THE NOISE MODEL OF QUANTIZATION FOR THE FREQUENCY-DOMAIN AMPLITUDE ESTIMATION OF LOW-LEVEL SINE WAVES

Diego Bellan

Politecnico di Milano, Department of Electronics, Information and Bioengineering, piazza Leonardo da Vinci 32, 20133 Milan, Italy (✉diego.bellan@polimi.it, +39 02 2399 3708*)*

Abstract

This paper deals with the amplitude estimation in the frequency domain of low-level sine waves, i.e. sine waves spanning a small number of quantization steps of an analog-to-digital converter. This is a quite common condition for high-speed low-resolution converters. A digitized sine wave is transformed into the frequency domain through the discrete Fourier transform. The error in the amplitude estimate is treated as a random variable since the offset and the phase of the sine wave are usually unknown. Therefore, the estimate is characterized by its standard deviation. The proposed model evaluates properly such a standard deviation by treating the quantization with a Fourier series approach. On the other hand, it is shown that the conventional noise model of quantization would lead to a large underestimation of the error standard deviation. The effects of measurement parameters, such as the number of samples and a kind of the time window, are also investigated. Finally, a threshold for the additive noise is provided as the boundary for validity of the two quantization models.

Keywords: quantization, amplitude estimation, sine wave, discrete Fourier transform, additive noise.

1. Introduction

Low-level sine waves (i.e. sine waves with a low amplitude with respect to the measurement range of an analog-to-digital (A/D) converter) occur in many application fields ranging from aerospace to industrial processes (e.g. see [1–3]). The amplitude estimation of a sine wave can be performed either in the time domain or in the frequency domain. In many applications it is usually preferred the frequency-domain approach since a simple and widely investigated mathematical tool such as the discrete Fourier transform (DFT) is available, together with its fast implementation known as the fast Fourier transform (FFT). A measurement procedure, therefore, consists of an A/D conversion (i.e. sampling and quantization) of the waveform and a time-to-frequency transformation (i.e. DFT) of the digitized signal.

The amplitude estimation in the frequency domain is affected by many error sources which can be roughly classified in three groups. First, a non-ideal A/D conversion process including the quantization and other more specific issues, such as a jitter and non-linearity of the converter. Second, the choice of measurement parameters, such as the sampling frequency, the number of samples, and the kind of the time window against the spectral leakage. Third, the additive noise always present in practical applications. All the mentioned aspects have been deeply investigated in the related literature (e.g. see [4–6]). This paper, however, is specifically devoted to low-level sine waves, i.e. sine waves spanning only a small number of quantization steps in the A/D conversion process. Notice that this issue is quite common in high-speed A/D converters, usually characterized by a rough resolution.

Article history: received on May. 18, 2014; accepted on Sep. 27, 2014; available online on Mar. 15, 2015; DOI: 10.1515/mms-2015-0004.

Thus, the main aspect to be investigated is the quantization, under the assumption that a small number of quantization steps are spanned by the input waveform. This is not a minor point since it is well-known that the widely used noise model of quantization (also called the statistical model of quantization) [4] works properly only under the assumption that, roughly speaking, a large number of quantization steps is involved. In fact, only in this case, together with the assumption of a dynamic behavior of the waveform, the quantization can be treated as a uniform and white additive noise.

In the literature it has been shown that the noise model of quantization is not appropriate for low-level sine waves [4], [7–9]. Thus, the amplitude estimation of low-level sine waves in the frequency domain cannot be characterized in terms of the standard deviation (or uncertainty) through the simple noise model of quantization. In this paper, quantization effects are taken into account by following an analytical approach, i.e. by studying the harmonic distortion produced by the quantization on the sine wave. Indeed, the fundamental component in the Fourier series of the quantized waveform provides the estimation of the sine wave amplitude. Such an approach will be exploited to derive approximate expressions for the standard deviation of the error in the estimated amplitude by assuming a random offset and phase of the sine wave [10]. It is shown that the usage of the noise model of quantization would result in a large underestimation of such a standard deviation. Moreover, the impact of the additive noise is evaluated. In particular, it is shown that, while low noise levels do not impact on the derived analytical results, by increasing the level of the additive noise above a specific limit the standard deviation of the error estimate gradually approaches the prediction of the noise model of quantization. Finally, the effects of measurement parameters such as the selection of the time window and the number of samples, are investigated.

The paper is organized as follows. In Section 2 the frequency-domain amplitude distortion due to the sine wave quantization is analyzed. In Section 3 a statistical approach is developed by assuming the sine wave offset and phase as random variables. The behavior of the probability density function of the error in the amplitude estimate is shown, and approximate expressions of the standard deviation of the error are provided. In Section 4 a comparison between the proposed analytical model and the noise model of quantization is presented by means of a numerical simulation of the whole A/D conversion process. In particular, the underestimation provided by the noise model of quantization is shown. In Section 5 the impact of the additive Gaussian noise for increasing noise levels is investigated. Finally, concluding remarks are drawn in Section 6.

2. The amplitude error of a quantized sine wave

Let us consider a sinusoidal waveform with an amplitude A and an offset B:

$$x(t) = A\sin(\omega t + \varphi) + B, \tag{1}$$

as the input of a continuous-time quantizer. In Fig. 1 a stretch of the input-output characteristic around the origin of a quantizer with a quantization step Δ is shown. A uniform quantizer is obtained when the transition levels $\{T_k\}$ are equally spaced by Δ. Notice that although the whole A/D conversion process foresees also the sampling of the quantized signal, the waveform distortion can be mainly ascribed to the quantization.

The quantized waveform can be expanded in a Fourier series in the explicit form and, in particular, it can be shown that the magnitude of the fundamental component is given by [8]:

$$A_q = \frac{2\Delta}{\pi}\sum_{k\in I}\sqrt{1-v_k^2}, \tag{2}$$

Metrol. Meas. Syst., Vol. XXII (2015), No. 1, pp. 89–100.

where

$$v_k = \frac{T_k - B}{A}, \tag{3}$$

and I is the set of the indexes of the transition levels crossed by the input sine wave.

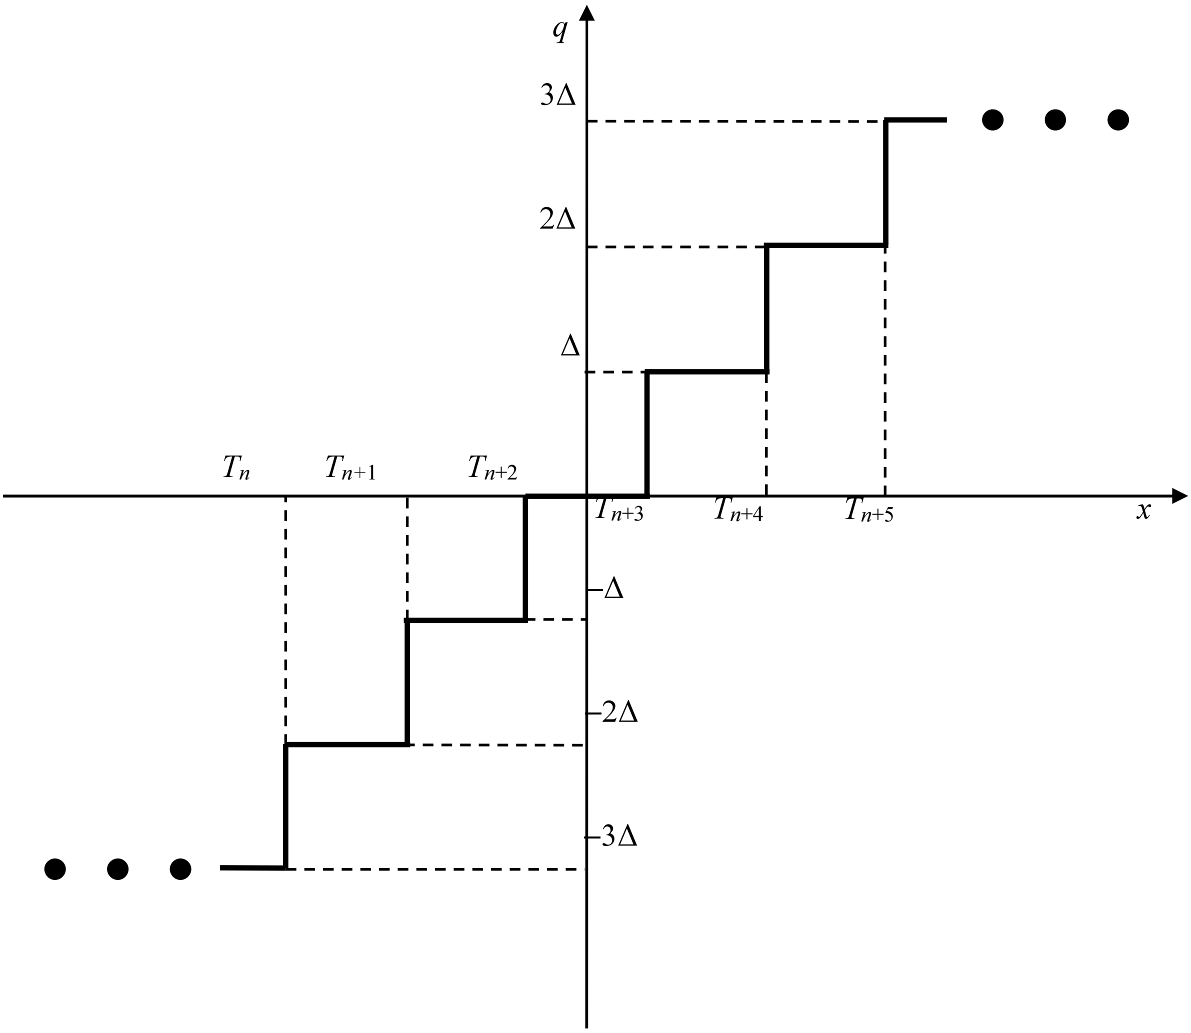

Fig. 1. The input-output characteristic around the origin of a quantizer.

The equation (2), therefore, provides the estimate A_q of the sine wave amplitude A after the quantization. Notice that such an estimate is also a function of the offset B, usually unknown in a measurement process. When the whole A/D conversion process is considered, the estimate (2) can be evaluated with the discrete Fourier transform (DFT) of the quantized samples.

The amplitude estimate (2) is affected by the error:

$$e = A_q - A, \tag{4}$$

which is clearly a rather complicated function of the sine wave parameters A and B. Notice that the phase φ of the input sine wave is not significant, since it does not affect the magnitude of the Fourier coefficient (2). In Fig. 2 the behavior of the error e is represented as a function of the sine wave amplitude A and offset B. The quantities are normalized with respect to the quantization step Δ, therefore they are given in units of the Least Significant Bit (LSB). In the figure, the range for the sine wave amplitude is bounded by 10 LSB since a distortion due to the quantization is more significant for a low-level signal. Indeed, the error e decreases as the amplitude A increases. Regarding the considered range for the offset B in Fig. 2, it is sufficient to consider a single quantization step since the behavior is periodic with the quantization step itself, provided that the quantizer is not overloaded.

The intricate behavior represented in Fig. 2 shows a need for a simpler mathematical representation of the quantization effects of an unknown amplitude and offset on the amplitude estimation of a low-level sine wave. For this aim, the sine wave amplitude A is treated as an independent variable, and for each specific value of A the worst case for the error e is evaluated with respect to the offset B taking all the values within one quantization step. The result is shown in Fig. 3 where both the maxima and minima of e are presented. In the figure, the envelopes of the two curves are also presented. It was found numerically that such envelopes can be approximated as [10]:

$$e_{\max} \cong \frac{0.17\Delta}{\sqrt{A/\Delta}}, \tag{5a}$$

$$e_{\min} \cong -\frac{0.37\Delta}{\sqrt{A/\Delta}}. \tag{5b}$$

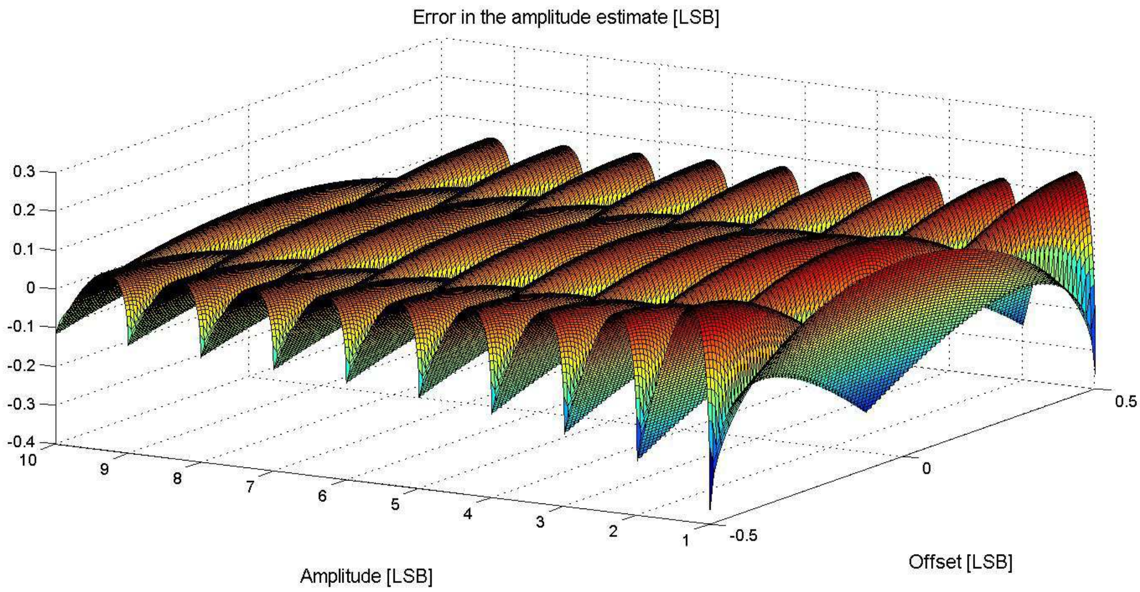

Fig. 2. The estimation error of the sine wave amplitude after the quantization as a function of the amplitude and the offset.

Metrol. Meas. Syst., Vol. XXII (2015), No. 1, pp. 89–100.

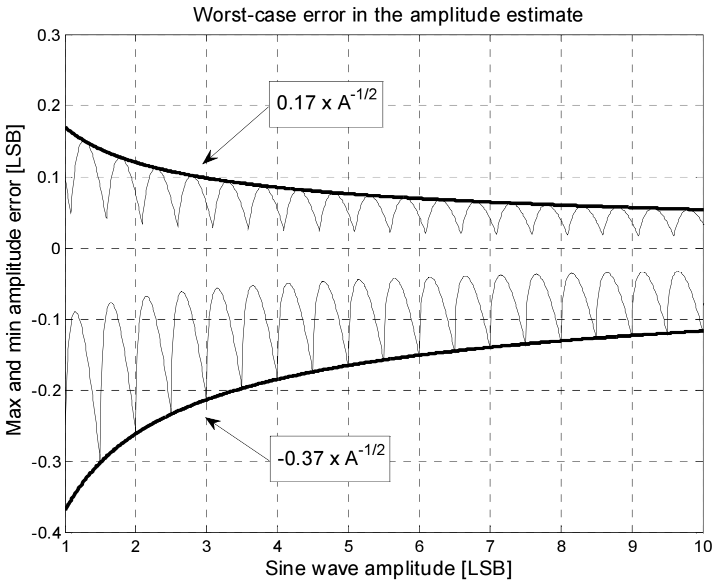

Fig. 3. The maximum and minimum error *e* in the frequency-domain estimation of the sine wave amplitude after the quantization. Thick lines represent the error envelopes.

3. The statistical characterization

A more complete characterization of the accuracy of the sine-wave amplitude estimation can be attained by using a statistical approach. By treating the offset *B* as a random variable (RV) with a uniform distribution within the interval (–Δ/2, Δ/2), and the amplitude *A* as a parameter, the estimation error *e* defined in (4) becomes an RV whose probability density function (PDF) can be obtained numerically by means of repeated-run simulations. As an example, Fig. 4 shows the behavior of the numerical PDF of the estimation error *e* assuming the parameter A/Δ=5.20. The range of *e* (approximately between –0.05 and 0.065 LSB) is coherent with the bounds given by the oscillating curves in Fig. 3. The presence of two peaks is a common characteristic for the PDFs of *e* corresponding to most of the values of the parameter *A*.

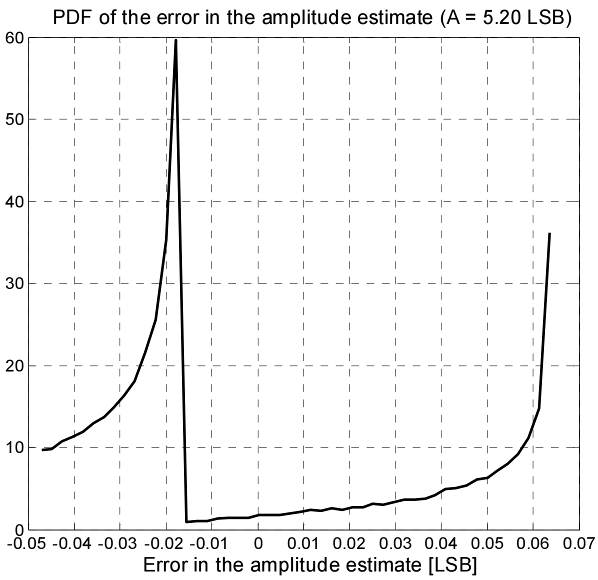

Fig. 4. The PDF of the error in the amplitude estimation for a specific sine wave amplitude, i.e. *A*=5.20LSB.

Fig. 5 shows several PDFs of e for different values of the parameter A within the range between 5 and 5.5 LSB (the behavior is almost periodic with respect to a half quantization step). The location of the peaks changes with A. The relative location of such peaks is essential in determining the numerical value of the variance of the RV e. Indeed, it is expected that the minimum variance (and therefore its square root, i.e. the standard deviation) is around A=5.10 LSB where the two peaks are very close. On the contrary, it is expected that the maximum variance is reached around A=5.40 LSB where the strong peaks are located at the two edges of the range of e.

Such remarks are confirmed by Fig. 6, showing the behavior of the standard deviation of e for A ranging in the interval (5, 5.5). In the same figure, the numerical standard deviation evaluated in the whole A/D conversion process is shown for validation purposes. In this case, the number of samples N_S=1024 was taken, with the sampling frequency f_S=1GHz. Coherent sampling was implemented by acquiring the integer number of sine wave periods N_p=101 and assuming the sine wave frequency f_0=98.63MHz. Each sample was then quantized by rounding its value, and the DFT was calculated. The magnitude of the spectral line corresponding to the sine wave was evaluated and compared with the actual sine wave amplitude to obtain the estimation error. The simulation was repeated 10^4 times for each sine wave amplitude from 5 to 5.5 by selecting, at each simulation run, random values for the offset and the phase of the sine wave, such that the standard deviation was calculated for each sine wave amplitude.

The behavior shown in Fig. 6 is characteristic for each half quantization step. Of course, by increasing A it is expected that the standard deviation decreases, according to the error bounds presented in Fig. 2. Fig. 7 (the solid lines) shows the behavior of the maximum and the minimum standard deviations of the RV e as functions of the sine wave amplitude. Such evaluations have been performed for sine wave amplitudes corresponding to the maximum and the minimum in each half quantization step (e.g. 5.1 and 5.4 in Fig. 6). Approximate expressions can be obtained by means of a numerical analysis, leading to

$$\sigma_{\max} \cong \frac{0.15\Delta}{\sqrt{A/\Delta}}, \tag{6a}$$

$$\sigma_{\min} \cong \frac{4.5\times10^{-2}\Delta}{\sqrt{A/\Delta}}. \tag{6b}$$

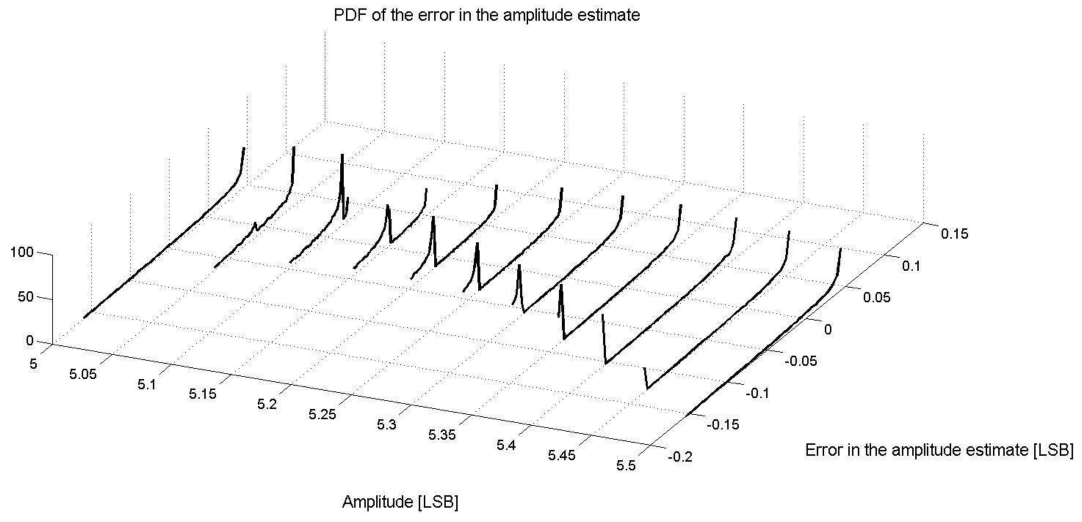

Fig. 5. The behavior of the PDF of the error in the amplitude estimate for different sine wave amplitude values.

Metrol. Meas. Syst., Vol. XXII (2015), No. 1, pp. 89–100.

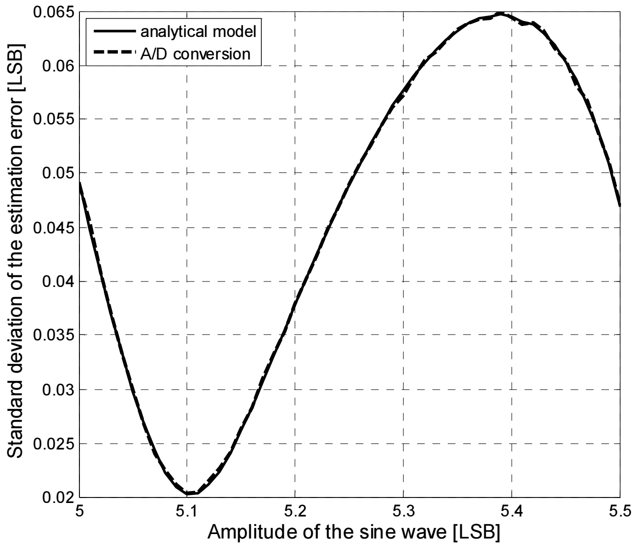

Fig. 6. The behavior of the standard deviation of the error in the amplitude estimate as a function of the sine wave amplitude ranging within a half quantization step. The solid line is derived by a repeated-run analysis of (4), while the dashed line is obtained by a repeated-run evaluation of the DFT of quantized samples.

4. A comparison with the noise model of quantization

According to the well-known noise model, the quantization is commonly modeled as a zero-mean white additive noise, uniformly distributed within a quantization step, with the variance:

$$\sigma_q^2 = \frac{\Delta^2}{12}. \tag{7}$$

Sampling is performed by acquiring a given number N_S of samples, such that the noise power in each DFT frequency bin of a one-sided spectrum is [5]

$$\sigma_{DFT}^2 = \frac{\Delta^2}{12 \cdot N_S/2} = \frac{\Delta^2}{6N_S}. \tag{8}$$

If a time window is used against the spectral leakage [11], the noise power (8) must be corrected with the so-called Equivalent Noise Bandwidth (ENBW) of the selected window, such that the resulting standard deviation of the DFT estimate of the sine wave amplitude is given by [5]

$$\sigma_{DFT} = \Delta\sqrt{\frac{ENBW}{6N_S}}. \tag{9}$$

The standard deviation (9), derived from the noise model of quantization, should be compared with (6a) (i.e. the worst case derived from the analytical approach used in Sections 2 and 3). It can be readily observed that for low-level sine waves (e.g. $A/\Delta<10$), for common values of the number of samples (i.e. N_S of the order of 10^3), and for the time windows usually employed in most applications (i.e. ENBW between 1 and 2), the standard deviation (6a) is much larger than (9). It means that, in such conditions, the noise model of quantization leads to an underestimation of the standard deviation of the measured sine wave amplitude. This is shown in Fig. 7 where the standard deviation (9) is also presented (the dashed line) for $N_S=2^{10}=1024$ and for the Hanning window (ENBW=1.5) which is very commonly used in many applications (also called the Hann window). The corresponding value, $\sigma_{DFT}=0.0156$, is much lower than the maximum values (6a) predicted by the analytical approach. On the contrary, as it was expected, the noise model of quantization improves its validity as the sine wave amplitude increases. This phenomenon has been already observed in practical experiments [12] and correctly ascribed to deterministic effects of quantization, but a detailed theoretical explanation has not been provided.

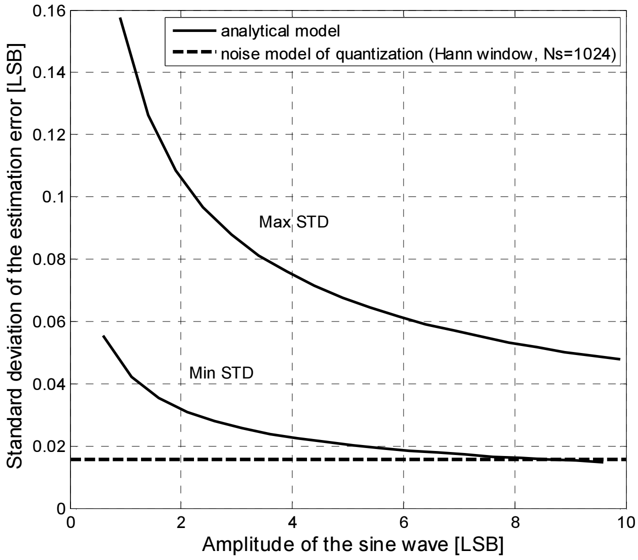

Fig. 7. The maximum and minimum standard deviation of the error in the amplitude estimation, given by (6a)-(6b) (solid lines). The dashed line represents the standard deviation provided by the noise model of quantization (9).

5. The impact of the additive noise

The additive noise is always present in measurements, therefore it is important to test the above analytical results against different noise levels added to the pure sine wave (1). The results presented in this Section have been all obtained by a numerical simulation of the whole A/D conversion process of noisy sine waves. Sampling conditions are the same as reported in Section 3, i.e. coherent sampling with N_S=1024, f_S=1GHz, N_p=101, and f_0=98.63MHz. The numerical results corresponding to the special case of zero noise must be compared with the analytical results obtained in the previous Sections for the validation.

When an independent white Gaussian noise with the standard deviation σ_n is added to the input sine wave, the effects on the standard deviation of the measured sine-wave amplitude are represented in Fig. 8 where only the maximum standard deviation is considered. The solid black curve corresponds to the noiseless case, and it is in a very good agreement with the maximum standard deviation presented in Fig. 7. The other curves correspond to increasing noise levels in LSB units. A clear behaviour is shown in the figure, i.e. the maximum standard deviation decreases as the noise level increases. This phenomenon can be explained with a smoothing action of the additive noise which mitigates the sharp effects of the deterministic noiseless quantization. A well known technique, called dithering, has been extensively

Metrol. Meas. Syst., Vol. XXII (2015), No. 1, pp. 89–100.

investigated in the literature to show an averaging action of the additive noise on some general properties of digitized waveforms [13]. The results presented in this paper, therefore, can be read as a further confirmation of this theory, with a new emphasis on the specific problem of the measurement of low-level sine waves. Moreover, it is interesting to observe from Fig. 8 that, as the noise level increases, the behaviour of the standard deviation of the amplitude error reaches a sort of an inversion point, where it stops to decrease and starts to increase according to the noise model of quantization. The inversion point can be located around σ_n=0.4 LSB. At this noise level, the behaviour of the total standard deviation is almost flat (the black dashed line) and its value (i.e. 0.027 LSB) is in a good agreement with the total standard deviation that can be obtained by combining the two noise sources as an independent additive noise [5]:

$$\sigma_{tot} = \sqrt{\left(\sigma_q^2 + \sigma_n^2\right)\frac{2ENBW}{N_S}}\,. \tag{10}$$

At σ_n=0.5 LSB (the brown dotted line) the standard deviation is flat and larger than the previous level. Its numerical value (i.e. 0.031 LSB) is in a good agreement with (10).

Figs. 9 and 10 show the effects of a different time window. In fact, while in Fig. 8 the Hanning window (ENBW=1.5) was used, in Figs. 9 and 10 the rectangular (ENBW=1) and the minimum 4-term Blackman-Harris window (ENBW=2) were used, respectively. It is apparent from the figures that the impact is negligible for low noise levels and low sine-wave amplitudes, for which the curves corresponding to $\sigma_n \le 0.3$ LSB are very close to each other. For larger noise levels the noise model of quantization comes into play and the total standard deviation level changes according to (10) with respect to the window parameter ENBW. Therefore, the brown dotted line in Fig. 9 (the rectangular window) has a lower level than the corresponding line in Fig. 10 (the Blackman-Harris window).

Finally, Fig. 11 shows the effect of the number of samples. The Hanning window was used as in Fig. 8, but the number of samples was 4096 instead of 1024. Also in this case, the effect is negligible on the curves related to a low noise level. On the contrary, for noise levels for which the noise model of quantization comes into play, the effect of the number of samples agrees with (10). As an example, the brown dotted line in Fig. 11 takes on a half of the corresponding curve value in Fig. 8.

6. Conclusion

The amplitude estimation of low-level sine waves in the frequency domain has been characterized in terms of the standard deviation of the estimated error. It was shown that the well-known noise model of quantization provides a large underestimation of such a standard deviation in the case of a low-level additive noise. This could be the case of a high-speed low-resolution A/D converter. In this case it was shown that a kind of the time window and the number of samples have a negligible impact. By increasing a noise level, noise model of quantization is gradually approached, and in this case the effects of the time window and the number of samples are readily explained by treating the quantization as the conventional additive noise.

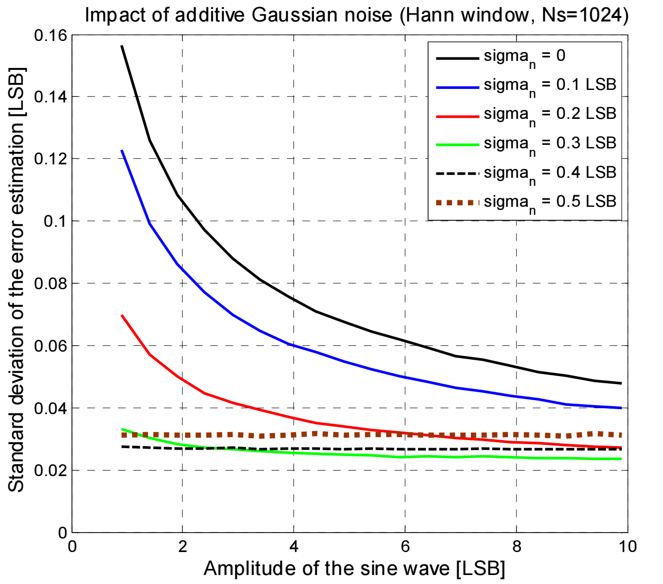

Fig. 8. The maximum standard deviation of the error in the amplitude estimation for different levels of the Gaussian additive noise after DFT. The noiseless case (the black solid line) must correspond to the max STD line in Fig. 7.

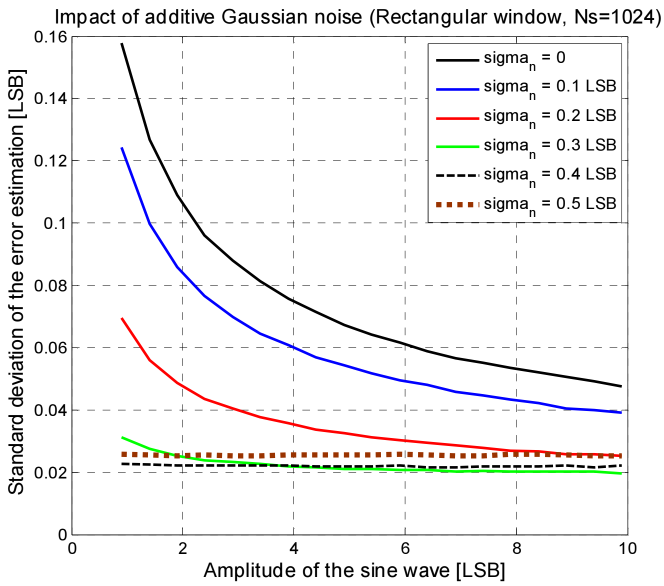

Fig. 9. The same as Fig. 8, but with the rectangular instead of the Hanning window.

Metrol. Meas. Syst., Vol. XXII (2015), No. 1, pp. 89–100.

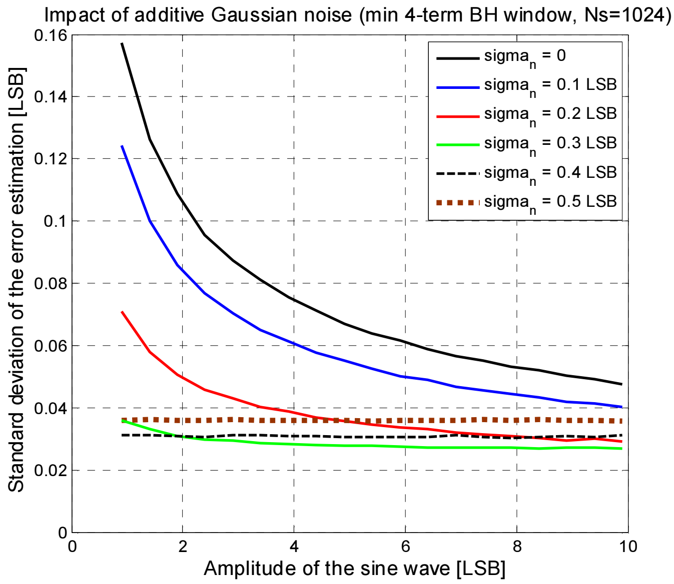

Fig. 10. The same as Figs. 8 and 9, but with the minimum 4-term Blackman-Harris window.

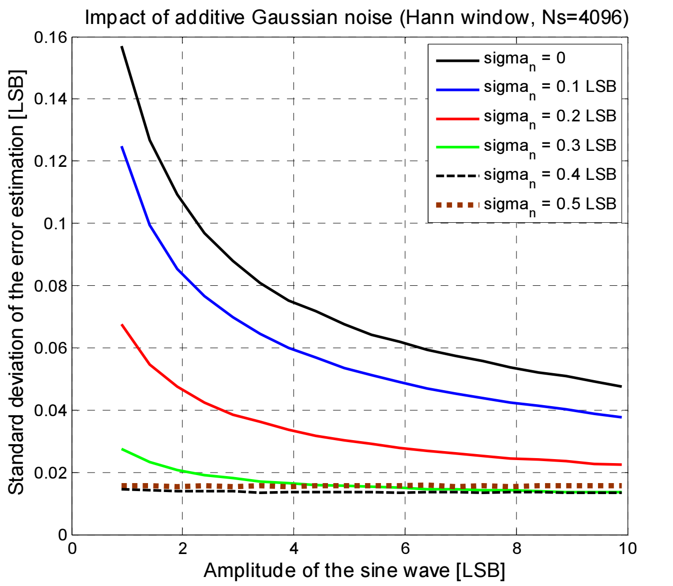

Fig. 11. The same as Fig. 8, but with 4096 samples instead of 1024.

The results presented in the paper are useful to assign a proper uncertainty value to low-level sine waves estimated in the frequency domain. Future work will be devoted to extending the approach to low-level non-sinusoidal waveforms.

References

[1] So, H. C., Chan, Y. T., Ma, Q., Ching, P. C. (1999). Comparison of various periodograms for sinusoid detection and frequency estimation. *IEEE Trans. Aerospace and Electronic Systems*, 35(3), 945–952.

[2] Moschitta, A., Carbone, P. (2007). Cramer-Rao Lower Bound for Parametric Estimation of Quantized Sinewaves. *IEEE Trans. on Instrum. Meas.*, 56(3), 975–982.

[3] Belega, D., Dallet, D., Slepicka, D. (2010). Accurate Amplitude Estimation of Harmonic Components of Incoherently Sampled Signals in the Frequency Domain. *IEEE Trans. Instrum. Meas.*, 59(5), 1158–1166.

[4] Gray, R.M. (1990). Quantization noise spectra. *IEEE Trans. Information Theory*, 36(6), 1220–1244.

[5] Petri, D. (2002). Frequency-domain testing of waveform digitizers. *IEEE Trans. Instrum. Meas.*, 51(3), 445–453.

[6] Solomon, O. M. (1992). The effects of windowing and quantization error on the amplitude of frequency-domain functions. *IEEE Trans. Instrum. Meas*., 41(6), 932–937.

[7] Bellan, D., Brandolini, A., Gandelli, A. (1999). Quantization Theory - a Deterministic Approach. *IEEE Trans. Instrum. and Measurement*, 48(1), 18–25.

[8] Bellan, D. (2000). Model for the Spectral Effects of ADC Nonlinearity. *Measurement*, 28(2), 65–76.

[9] Hejn, K., Pacut, A. (2003). Effective Resolution of Analog to Digital Converters. *IEEE Instrum. & Meas. Magazine*, 9, 48–55.

[10] Bellan, D. (2013). Detection and Estimation of Weak Sine Waves with Random Offset and Additive Noise. In *Proc. of IEEE International Conference on Instrumentation, Communication, Information Technology and Biomedical Engineering (ICICI-BME 2013)*. Bandung, Indonesia, 1–6.

[11] Harris, F. J. (1978). On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1), 51–83.

[12] Gupta, S., Paliakara, V. (October 2011). *Understanding Low-Amplitude Behavior of 11-bit ADCs*, Texas Instruments, Application Report SBOA133.

[13] Carbone, P., Petri, D. (1994). Effect of additive dither on the resolution of ideal quantizers. *IEEE Trans. Instrum. Meas.*, 43(3), 389–396.