Phase noise and amplitude noise in DDS

Claudio E. Calosso[‡], Yannick Gruson^{*}, and Enrico Rubiola^{*}

web page <http://rubiola.org>

[‡] INRIM, Torino, Italy

^{*} CNRS FEMTO-ST Institute, Besançon, France

Abstract—Their article reports on the measurement of phase noise and amplitude noise of direct digital synthesizers (DDS), ultimately intended for precision time and frequency applications. The DDS noise $S_\varphi(f)$ tends to scale down as $1/\nu_0^2$, until the noise hits the limit due to the output stage. The spurs, however disturbing in general, sink power from the white noise. Voltage noise can be more critical in the digital power supply than in the analog supply. Temperature fluctuations are an issue at $10^{-3} \dots 1$ Hz Fourier frequency. Passive stabilization (thermal mass) proves to be useful. Other paramours affect the phase noise, like the clock frequency and power. The amplitude $1/f$ noise is of the order of -110 dB(V²/V²)/Hz in some reference (typical) conditions.

Owing to the page and file size limitations, only a small part of the available data can be published here. An extended and free version of this article is available on <http://rubiola.org> and on <http://arxiv.org>.

I. INTRODUCTION

After the original article [1] published more than 40 years ago, the DDS is now a mature piece of technology. The development has been pushed by the semiconductor technology and by applications. The DDS owes its success to the frequency range (dc to GHz and beyond), to the high resolution (μ Hz), to the fast frequency switching (sub- μ s), to the small size and power, and to the suitability to almost-all-digital single-chip implementation. The newcomer can refer to a technical tutorial [2] and to two books [3], [4].

With the availability of single-chip DDSs in the 1980s, phase noise and spurs were on the stage (see for instance [5], [6], [7], [8]), and an exact and computationally efficient

solution for spurs was found [9]. However, a number of experimental issues of phase noise are not addressed properly in the literature, and amplitude noise is totally absent. We target filling this gap.

A. Notation

Following the general notation of frequency metrology (see for example [10], [11]), the clock signal is written as $v(t) = V_0 [1 + \alpha(t)] \cos [2\pi\nu_0 t + \varphi(t)]$, where $\alpha(t)$ is the random fractional amplitude, and $\varphi(t)$ is the random phase. As usual, we describe such noise in term of $S_\alpha(f)$ and $S_\varphi(f)$, i.e., the single-sided power spectral densities of $\alpha(t)$ and $\varphi(t)$. The spectra are approximated with the polynomial law, using $S_\alpha = \sum_i h_i f^i$ and $S_\varphi = \sum_i b_i f^i$.

II. DDS OPERATION

Referring to the block diagram of Fig. 1 and to the state representation of Fig. 2, the DDS is governed by

$$n_k = (n_{k-1} + \mathcal{N}) \bmod \mathcal{D} , \quad (1)$$

Figure 1. DDS scheme.

Figure 2. DDS state variable, and its relationships to the output phase.

where n_k is the state variable at the integer time, $\mathcal{D} = 2^m$ is the modulo, m is the number of bits of the phase accumulator (D register), $t = k/\nu_s$ is the time, and ν_s is the clock frequency. At each clock rising edge, the accumulator is incremented by the value $\mathcal{N}$ of the control word, and the carry is discarded when the accumulator overflows. This mechanism gives the output frequency

$$\nu_0 = \frac{\mathcal{N}}{\mathcal{D}} \nu_s. \quad (2)$$

The sampling theorem requires that $\nu_0 \leq \frac{1}{2}\nu_s$, or $\nu_0 \lesssim 0.4\nu_s$ accounting for the rolloff of a real lowpass. Higher frequency operation is allowed, selecting an alias with a bandpass.

The state variable n is converted into sinus and cosine by the look-up table, and further converted into the output analog signals by the two DACs. The look-up table is functionally a read-only memory, though the implementation is quite different. The design is generally determined by the high speed, made possible by the low normalized slew rate of the sinusoid (2π), which makes small the difference between contiguous data.

a) *Frequency resolution*: Since the control word $\mathcal{N}$ can be set with the resolution of one, the frequency resolution is

$$\nu_{\text{res}} = \frac{1}{\mathcal{D}} \nu_s. \quad (3)$$

For example, a 48 bit DDS has $\mathcal{D} = 2^m = 2.8 \times 10^{14}$, thus a resolution of 3.6 μHz at 1 GHz clock frequency. Such resolution is sufficient for virtually all practical applications, and also gives the opportunity to implement accurate phase or frequency modulation. Owing to the low cost and power of the digital circuits, there are little reasons to choose a significant smaller value of m . For instance, with $m = 32$ bits the resolution would be of 233 mHz at 1 GHz clock. So, most DDSs have $m = 48$ bits, while 32 bits is reserved to extremely-high frequency implementations, or then the power is a critical issue. Other values are seldom encountered.

III. NOISE AND SPURS

A. The Egan model for phase noise

Figure 3 shows a model for phase noise in frequency synthesis, inspired to the Egan's article on digital dividers [12]. Assuming that the synthesizer core is noise free, the input phase-time $\times$ — i.e., the time jitter — is transferred from the input to the output as it is. Since the unit of angle scales up as $\mathcal{D}/\mathcal{N}$, $S_\varphi(f)$ scales down as $(\mathcal{N}/\mathcal{D})^2$. Of course, this also applies to the noise of the input stage. Below a given ν_0 , the input noise scaled down hits the phase noise of the output stage, set by the SNR (white) and by the up-conversion of near-dc noise (flicker and temperature fluctuations).

However, there is a significant difference between the simple digital divider and the DDS. The divider samples the phase noise at the rising edges of the output, at the rate $\nu_0 = \nu_s/\mathcal{D}$ (the divider has $\mathcal{N} = 1$). In the cases of white noise this increases $S_\varphi(f)$ by a factor of $\mathcal{D}$, hence overall $S_\varphi(f)$ scales down as $1/\mathcal{D}$ instead of $1/\mathcal{D}^2$. Conversely, the DDS samples

Figure 3. Phase noise model for a frequency synthesizer.

the output at the full clock rate ν_0 . Therefore $S_\varphi(f)$ scales down as $1/\mathcal{D}^2$.

B. Quantization noise

Assessing the resolution, we notice that the number q of DAC bits is the most severe technical limitation, related to the sampling frequency. For reference, at 1 GHz maximum clock we find 14 bit converters (Analog device AD9912).

The voltage associated to the least significant bit is $V_{\text{LSB}} = V_{\text{FSR}}/2^q$, where FSR is the (peak-to-peak) dynamic range. The symbols V_{LSB} and V_{FSR} are consistent with the technical literature. If q is large enough (at least 8 bits) and the signal uses the dynamic range efficiently, the quantization noise has rectangular distribution in $\pm \frac{1}{2}V_{\text{LSB}}$. This is granted by the Wiener-Khinchine theorem for stationary ergodic systems, which states that the statistical properties can be calculated interchangeably in the ensemble or in the time domain. The variance is $\sigma^2 = \frac{1}{12}V_{\text{LSB}}^2 = V_{\text{FSR}}^2/(12 \times 2^{2q})$. Assuming that the quantization noise is a true random process with no structure, by virtue of the sampling theorem the noise spectrum is uniformly distributed from 0 to the Nyquist frequency $B = \frac{1}{2}\nu_s$. The Parseval theorem states that the power calculated in the time domain and in the spectrum is the same. Denoting the white noise level with N , the Parseval theorem gives

$$N = \frac{V_{\text{FSR}}^2}{6 \times 2^{2q} \nu_s}. \quad (4)$$

The signal power is $\frac{1}{8}V_{\text{FSR}}^2$. Thus, the white phase noise $b_0 = N/P_0$ turns into

$$b_0 = \frac{4}{3} \frac{1}{2^{2q} \nu_s} \text{ rad}^2/\text{Hz}. \quad (5)$$

The results discussed in this paragraph derive from a seminal article by Bennet [13], adapted to phase noise. For example, a 14 bit coverer at the sampling frequency $\nu_s = 400 \text{ MHz}$ yields a white noise of $-169 \text{ dBrad}^2/\text{Hz}$.

C. DAC resolution

Given the 32–48 bit resolution of the phase accumulator, it is obvious that the DAC resolution cannot be that high. The DAC has a number q of bits that is determined by the frequency and by the available technology. Accordingly, the

Figure 4. A simulation shows a reduction in the noise floor related to the presence of some spurs.

state variable must be truncated by discarding the appropriate number $m - p$ of bits. The minimum number p of bits at the input of the look-up table is determined by the need of preserving the DAC resolution, and estimated using the fact that the slew rate of the sinusoid, normalized to ν_s and to the full-scale range V_{FSR} is equal to π . Hence, the digital resolution at the look-up table input must be $p_{\min} = q + \log_2(\pi)$, rounded to $p_{\min} = q + \lceil \log_2(\pi) \rceil = q + 2$. The value $p = q + 5$ is often encountered.

b) *Truncation and spurs*: Deriving Eq. (4), we assumed that the quantization noise is white, as it happens in most data acquisition systems. The DDS is more complex because the truncation of the state variable ($p < m$) introduces a deterministic error governed by the arithmetic of reminders in Eq. (1).

For the sake of heuristic reasoning, we assume that the phase error is determined only by p , letting aside the relationship between p and q . For example, a 48 bit DDS ($m = 48$) may have $p = 16$. Accordingly, two contiguous values of the analog output are separated by $2^{m-p} - 1 = 2^{48-16} - 1 \simeq 4.3 \times 10^9$ invisible values of n .

The rotating vector (angle θ_k of Fig. 2) gives a sawtooth-like phase error distributed in $\pm \frac{\pi}{2^{m-p}}$, and sampled at the frequency ν_s . Such sawtooth contains harmonics beyond the Nyquist frequency, which are aliased down and brought into the output bandwidth. This phenomenon gives rise to a bunch of spurs which change suddenly as N is changed. The structure introduced by the spurs breaks the hypothesis that the quantization noise is white. The most amazing consequence is that the power of such spurs comes at expenses of the floor, and therefore the floor is *lower* than the value given by Eq. (5). This concept is confirmed by the simulation shown in Fig. 4, based on the parameters of the AD9854 in the same configuration of some experiments.

Interestingly, the spurs are generally not seen at the low resolution of a regular spectrum analyzer, but they show up when observed at the high frequency-resolution usually achieved in phase noise measurements. The sawtooth-like

Figure 5. AD9854 phase noise measured at different output frequencies.

Figure 6. AD9912 phase noise measured at different output frequencies.

phase error has a long period referred to as the grand repetition rate $GRR = \frac{D}{\gcd(N, D)} \frac{1}{\nu_c}$.

D. Nonlinearity

The DAC nonlinearity generates harmonics, integer multiples of ν_0 . The harmonics that exceed the Nyquist frequency $\frac{1}{2}\nu_s$ are brought to the output bandwidth by aliasing.

IV. PHASE NOISE

A. Experimental method

We experimented on three samples, an AD9854 demo board, an AD9912 demo board, and a home made AD9854 board. Our board gives flexibility in the configuration of the output stage. We used a traditional saturated mixer, and also the Symmetricom TSC5120A and TSC5125A test sets. The latters can compare the phase of two signals at different frequencies, and rejects the background noise by correlation and averaging on multiple spectra.

B. Phase noise

Figures 5 and 6 show the phase noise of the two DDSs, measured at different output frequencies. We see that the

Figure 7. Phase noise of the output stage, measured by comparing the phase difference between the I output and the Q output.

Figure 8. AD9854 phase noise measured at different output amplitudes

flicker noise scales down as the output frequency (-6 dB per factor-of-two), until it hits the noise of the output stage. Unfortunately, both figures suffer from noise in the supply line, which was fixed later.

The AD9854 has two outputs in quadrature, which we exploit to measure the phase noise of the output stage. In fact, the $I - Q$ phase noise contains only the noise of the two DACs, including their own clock circuits, and the output buffer. Looking at Fig. 7, the flicker is generally 10 dB lower than that of the whole DDS, but it hits the same value at low output frequency. This confirms our conclusion that the lowest flicker in Figures 5 and 6 is that of the output stage.

Unfortunately, this experiment does not reveal the DDS white noise. This happens because the measurement was done with an AD8002 output buffer instead of a RF amplifier. This choice was made for different purposes, because the AD8002 provides kHz output, at the cost of large white noise.

Figure 8 shows the AD9854 phase noise measured at different output amplitude levels. While the flicker is constant,

Figure 9. Phase noise spectra with a small change in the control word.

Figure 10. Power-supply rejection of the AD9912 1.8 V line, measured by switching the supply voltage by $\pm 5\%$ around the nominal value.

the white noise scales up by 6 dB per factor-of-two as the amplitude decreases. This is a consequence of the AD9854 specific architecture, which uses digital multiplication to set the output level. So, a lower level is obtained by sending smaller integers to the DAC, which is equivalent to using less bits.

Most of the spectra we measured show steep slope at Fourier frequencies below 1 Hz. This is due to thermal fluctuations, discussed further in this article.

Figure 9 shows two phase noise spectra observed with a small change in the control word. As expected, the flicker is the same and the spurs change dramatically. The floor, lower than the prediction of Eq. (5) relates to the fact that the noise is not white, and the spurs introduce a structure.

C. Power supply rejection

The AD9912 has two supply voltages, 1.8 V for digital circuits and 3.3 V for analog circuits. We measured a coefficient of 15 ps/% for the 1.8 V fluctuations (Fig. 10). The constant time shift, independent of the output frequency, is the signature of a threshold issue. Converted in radians, this effect becomes dramatic at high output frequencies. For instance, 15 ps/% is equivalent to approximately 0.1 rad/V for the 1.8 supply at 25 MHz output. The effect of the fluctuations of the 3.3 V supply is a factor-of-100 lower, thus negligible.

Figure 11. The supply-voltage noise impacts on the DDS phase noise.

Experience suggests that voltage regulators commonly used for digital circuits are rather noisy. In some experiments we used a LM1117MP-1.8 1.8 V regulator with a noise of 250 nV/ $\sqrt{\text{Hz}}$, a bump at 30 kHz, and rolling off beyond. The effect, shown in Fig. 11, shows up at about 1 kHz, and becomes dominant at 10 kHz and above. We also come in trouble with Ta capacitors. The AD9854 has similar poor immunity to supply-voltage noise. Optimal filtering of the supply lines is of paramount importance for low phase noise applications, and the digital circuits are surprisingly more critical than the analog circuits.

D. Thermal effects on the output phase

We measured the temperature sensitivity of the AD9912 in two ways, by heating the card in a temperature controlled oven, and through the self-heating induced by a square wave added to the supply voltage. The phase step is due to the voltage fluctuation already discussed. Two thermal constants are clearly identified in the transient, 10 s and 1 m, likely due to the DDS chip and to the heat sink. After the transient, the phase shift is due to the induced temperature change. The reason is that the oven heats the whole board and cables, while the self-heating acts only on the DDS chip. The results are in a fair agreement for the two methods. Therefore, the additional sensitivity due to the card and cables is smaller than that of the DDS chip.

The AD9912 temperature coefficient (Fig. 12) is constant vs. frequency (-2 ps/K) above 20 MHz, and shows a $1/\nu_0^3$ slope below. This is the signature of the fluctuation of a digital threshold at high frequency, and likely of some kind of drift in the analog electronics at low frequency.

The thermal resistance and capacitance of the heat sink impact on the temperature fluctuations below 1 Hz, and in turn on phase noise. Figure 13 shows the phase noise of a AD9854 with two different heat sinks, a 0.64 cm³ Al cube, and 150 cm³ fin heat sink originally intended for Peltier cells. The combined effect of the 160 g mass and the low gradient due to the conductance stabilizes the junction temperature, and reduces the phase fluctuations. Below 2 Hz, and until approximately 1 mHz, the large heat sink reduces the phase

Figure 12. Thermal coefficient of the AD9912 vs. the output frequency.

Figure 13. Effect of the heat sink on the phase noise spectrum.

noise by some 10 dB/dec or more. Below 1 mHz, both heat sinks are driven by the room temperature and the stabilization effect vanishes.

E. Clock frequency and power.

The sampling theorem suggests that the white phase noise is proportional to the reciprocal of the clock frequency, as stated by Eq. (4) and (5). Besides, the integrated-circuit technology sets a bound to the minimum and maximum clock frequency. For example, the AD9912 must operated in the 200–1000 MHz range. A significant degradation in the flicker noise shows up at 50 MHz and 100 MHz clock. Conversely, the AD9912 seems to be tolerant to the clock amplitude in a range of at least 15 dB, with no degradation of phase noise (Fig. 14). This exceeds most practical needs.

V. AMPLITUDE NOISE

We measured the AM noise with the scheme of Fig. 15, giving special attention to flicker. The method resorts to two previous articles [14], [15]. The former describes the problems specific to the measurement of AM noise, and the latter reveals secrets and subtleties of the cross-spectrum method. In short, since for small fluctuations the fractional amplitude noise is equal to half the power fluctuation, the Schottky-diode power detector proves to be a suitable transducer. The cross-spectrum method is necessary because, unlike for PM noise, we cannot

Figure 14. AD9912 phase noise spectrum measured at 400 MHz clock for a few values of the clock amplitude. The amplitude is given in dBm on a 50 Ω resistance, 0 dBm is 224 mV.

Figure 15. Scheme for the measurement of AM noise.

assess the background noise by removing the DUT because relevant noise phenomena would vanish when the input signal was set to zero. The cross spectrum provides accurate results with easy implementation and calibration because under normal conditions the measured AM noise is not influenced by the PM noise.

Figure 16 shows the AM noise measured on a sample of AD9854. This Figure shows a smooth flicker process spanning on 2–3 decades, with a coefficient h_{-1} between -105 and -115 dB rad^2/Hz . A minimum is present around 20 MHz. We believe that this odd behavior is an outcome of the converter technology, possibly related to some change in the operation mode from low frequencies to high frequencies.

The core of a Spectra Dynamics LNFS 100 synthesizer is an AD9852. The AM noise of this synthesizer was measured at European Gravitational Observatory. The AM noise is of the order of -110 dB rad^2/Hz . Yet, we could not identify a frequency range where the flicker noise is lower, as it happens with the AD9854. However, general experience suggests that the AM noise of amplifiers is significantly lower. Accordingly, the observed AM flicker is ascribed to the DDS.

Acknowledgments

We thank Jean-Michel Friedt (Senséor, France) and Marco Siccaldi (SKK, Italy) for fruitful discussions on digital systems; Federico Paoletti (INFN, Italy) shared his results on AM noise. Symmetricom Inc. borrowed us a TSC5125A for a long time. E.R. is indebted to Giovanni Costanzo and Filippo Levi for inviting at the INRIM in 2011; and to Vincent Giordano for discussions and support over 15 years.

Figure 16. AM noise of the AD9854.

This work is partially funded and as a part of the Uliss program (<http://uliss-st.com>), and under the Contrat-programme LNE/DRST 11 7 002.

REFERENCES

- [1] J. Tierney, C. M. Rader, and B. Gold, "A digital frequency synthesizer," *IEEE Trans. Audio Electroacoust.*, vol. 19, pp. 48–56, Mar. 1971.
- [2] Analog Devices, Inc., *A Technical Tutorial on Digital Signal Synthesis*, 1999. Free online from <http://analog.com>.
- [3] B.-G. Goldberg, *Digital Frequency Synthesis Demystified*. Eagle Rock, VA, USA: LLH Technology Publishing, 1999.
- [4] V. F. Kroupa, ed., *Direct Digital Frequency Synthesizers*. New York: IEEE Press, 1999.
- [5] H. T. Nicholas, III and H. Samuelli, "An analysis of the output spectrum of direct digital frequency synthesizers in the presence of phase-accumulator truncation," in *Proc. Freq. Control Symp.*, (Philadelphia, PA, USA), pp. 495–502, May 27–29, 1987.
- [6] E. M. Mattison and L. M. Coyle, "Phase noise in direct digital synthesis," in *Proc. Freq. Control Symp.*, pp. 352–356, June 1–3, 1988.
- [7] V. F. Kroupa, "Phase and amplitude disturbances in direct digital frequency synthesizers," *IEEE Trans. Ultras. Ferroelec. and Freq. Contr.*, vol. 46, pp. 481–486, May 1999.
- [8] V. F. Kroupa, "Spurious signals in direct digital frequency synthesis due to phase truncation," *IEEE Trans. Ultras. Ferroelec. and Freq. Contr.*, vol. 47, pp. 1166–1172, Sept. 2000.
- [9] A. Torosyan and A. N. Willson, Jr., "Exact analysis of DDS spurs and SNR due to phase truncation and arbitrary phase-to-amplitude errors," in *Proc. Int. Freq. Control Symp.*, (Vancouver, BC, Canada), pp. 50–58, Aug. 29–31, 2005.
- [10] H. G. Kimball, ed., *Handbook of selection and use of precise frequency and time systems*. ITU, 1997.
- [11] J. R. Vig (chair.), *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities (IEEE Standard 1139-1999)*. IEEE, New York, 1999.
- [12] W. F. Egan, "Modeling phase noise in frequency dividers," *IEEE Trans. Ultras. Ferroelec. and Freq. Contr.*, vol. 37, pp. 307–315, July 1990.
- [13] W. R. Bennett, "Spectra of quantized signals," *Bell Syst. Techn. J.*, vol. 27, July 1948.
- [14] E. Rubiola, "The measurement of AM noise of oscillators," <http://arxiv.org>, document arXiv:physics/0512082, Dec. 2005.
- [15] E. Rubiola and F. Vernotte, "The cross-spectrum experimental method," arXiv:1004.5539 [physics.ins-det], Apr. 2010.