The Sampling Theorem in Pi and Lambda Digital Frequency Dividers

Claudio E. Calosso
Division of Optics
INRIM
Torino, Italy
e-mail: c.calosso@inrim.it

Enrico Rubiola
Department of Time and Frequency
CNRS FEMTO-ST Institute
Besancon, France
e-mail: rubiola@femto-st.fr
web page http://rubiola.org

Abstract*— It is a common belief that a noise-free frequency divider by D scales down the input phase by a factor of $1/D$, thus the phase-noise power spectral density (PSD) by $1/D^2$. We prove that the behavior described *does not apply to digital dividers.

Instead, the digital divider scales the white phase-noise PSD down by $1/D$. Phase downsampling and aliasing, inherent in digital frequency division, is the reason. However the $1/D^2$ law holds asymptotically for flicker, where the aliases can be neglected.

We propose a new de-aliased divider, which scales the input phase-noise PSD by approximately $1/D^2$. The scheme is surprisingly simple and suitable to CPLD and FPGA implementation.

Keywords*— *aliasing; digital frequency divider; phase noise; white noise; flicker

I. Introduction

The noise-free frequency divider can be assimilated to an ideal gearbox that divides the input frequency by D, transferring the total jitter $x(t)$ – better referred to as `phase time' or `phase-time fluctuation' – from the input to the output. Accordingly, the output phase noise is $\varphi_{out}(t) = \varphi_{in}(t)/D$, and its power spectral density (PSD) $[S_\varphi(f)]_{out} = [S_\varphi(f)]_{in}/D^2$. So, a modulo-10 divider reduces the input phase noise by 20 dB.

Unfortunately, the output *white* noise of the classical digital divider (Π divider) is ruled by $[S_\varphi(f)]_{out} = [S_\varphi(f)]_{in}/D$. Since the denominator is D instead of D^2, a modulo-10 divider reduces the input noise by 10 dB instead of 20 dB. This is due to aliasing, thus it is inherent in the square-wave frequency division (Π divider), and independent of technology.

This article stands on [1], clarifies some issues, and introduces the Λ-type divider. This divider fixes the aliasing by adding shifted square waves with resistors, which results in a triangular-wave output (Fig. 2). Therefore, the noise of the Λ divider is ruled by $[S_\varphi(f)]_{out} = [S_\varphi(f)]_{in}/D^2$, which outperforms the phase noise of the Π-type divider by a factor of $1/D$.

The $1/D^2$ and $1/D$ laws result from simple mathematical properties of white noise. Of course thermal noise, shot noise, and flicker noise add up, and the noise due to the technology as well.

II. Sampling and Aliasing

It is well known that the spectrum around all multiples of the sampling frequency is down-converted to baseband, and overlaps to the main part of the spectrum. This is the essence of aliasing, fully explained by the sampling theorem. The same fact can be interpreted as a consequence of energy conservation, as the sampling process preserves the total noise power. The power is by virtue of the Parseval theorem equal to the noise PSD integrated over the full bandwidth. Therefore, band compression results into white-noise enhancement by the same factor.

Real-world white noise has constant PSD spanning from 0 to the bandwidth B of the system. In the presence of white noise, aliasing deteriorates the SNR by a factor of B/f_N, where f_N is the Nyquist frequency (half the sampling frequency).

By contrast, flicker results in little and generally negligible degradation of the SNR. This happens because the PSD is proportional to $1/f$, so the aliases are small as compared to the base band.

A. Input stage

In digital circuits, the input stage converts the sinusoidal clock input into a square wave. As a relevant consequence, phase noise shows up only at the rising and falling edges, while the signal is saturated in between. This means that phase noise is *sampled* at the frequency 2ν, hence the Nyquist frequency is $f_N = \nu$. Denoting the bandwidth of the input stage with B, aliasing increases $S_\varphi(f)$ by a factor of B/ν. For instance, a 100 MHz sinusoid processed by a 300 MHz logic circuit results into a degradation of the input noise by approximately 4 dB.

B. Π-type divider

The classical digital divider is called Π-type divider from the graphical similarity of the output pulses with the Greek letter Π.

In a modulo-D divider, one edge every D edges propagates from input to output (Fig. 1). This is a hard decimation process, with no low pass (anti-aliasing) filter. In this condition, aliasing increases the white noise by a factor of D. Since the pure gearbox scales $S_\varphi(f)$ down by $1/D^2$ and alias enhances it by a factor of D, overall the input white noise is governed by the $[S_\varphi(f)]_{out} = [S_\varphi(f)]_{in}/D$ law.

Oppositely, flicker ($1/f$) noise is well approximated by the $[S_\varphi(f)]_{out} = [S_\varphi(f)]_{in}/D^2$ law because the aliases are comparatively small.

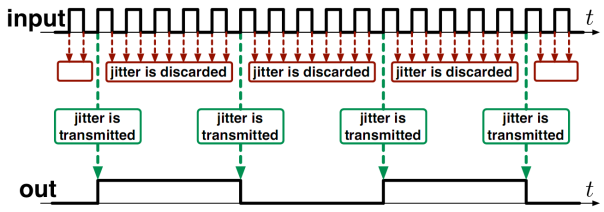

Fig. 1. Output sampling mechanism in Π divider.

C. Λ-type divider

The Λ divider generates a triangular wave signal by combining D square-wave signals shifted by half cycle of the input clock. Figure 2 shows an example of implementation for $D = 10$. Alternate topologies are also possible.

In Fig. 2, two ÷10 Π dividers are used, one driven by the rising edge of the input clock, and one driven by the falling edge. The outputs are delayed by 5-bit shift register, so that 10 phases of the intermediate square wave at $\nu/10$ are available, spanning over 180°. The sum of these intermediate signals is a triangular-wave like staircase (Fig. 3). Of course, the Greek letter Λ in the name derives from the graphical analogy with the triangular waveform.

The relevant feature of the Λ divider is that the output is sampled at the frequency 2ν, i.e., twice the input frequency, instead of $2\nu/D$, which fixes the alias of the Π divider. An alternate interpretation is that the white random signals at the intermediate outputs are statistically independent. Adding D of them results in facror-of-1/D lower phase noise. As a consequence, the Λ divider is governed by the $1/D^2$ law, which is the theoretical limit for the noise-free divider.

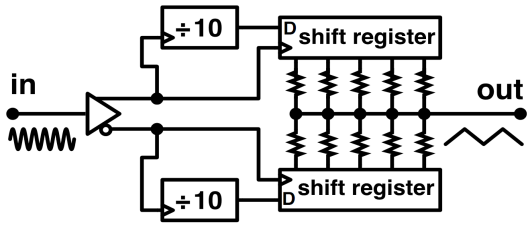

Fig. 2. Λ divider block diagram.

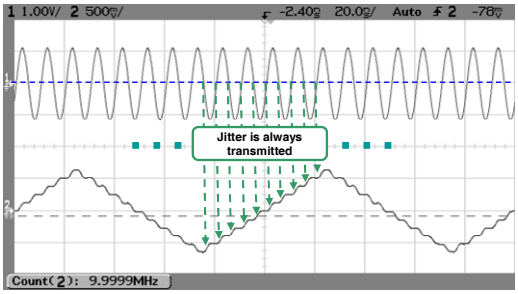

Fig. 3. Waveform at the input and at the output of a Λ divider.

III. Experiments and Results

The Divider Under Test (DUT) is implemented in a CPLD Altera EPM3064A [2], a simple 64-cell device with 10 ns speed grade, kind of the `ancestor' of the FPGA, which for our purposes can be programmed in three ways:

- The simple Π-type divider,
- A multi-buffer Π-type, which is the scheme of Fig. 2 with the `shift registers' re-configured as buffers, i.e., with no shift at all,
- The Λ divider, as it is in Fig. 2.

The resistor network is a SMD array mounted close to the CPLD.

The phase-noise test instrument [3] implements the cross-spectrum technique described in [4].

A. Aliasing

The first experiment is the demonstration of the aliasing phenomenon in the Π and Λ dividers, and of the benefit of the Λ topology. In order to isolate the aliasing effect from other phenomena in the divider, we work with an artificially – and unrealistically – high level of white noise. This is done by attenuating and re-amplifying the clock signal (Fig. 4).

Flicker noise cannot be obtained in the same way. So, we use the phase-modulation feature of a commercial synthesizer driven by an audio-frequency flicker-noise generator.

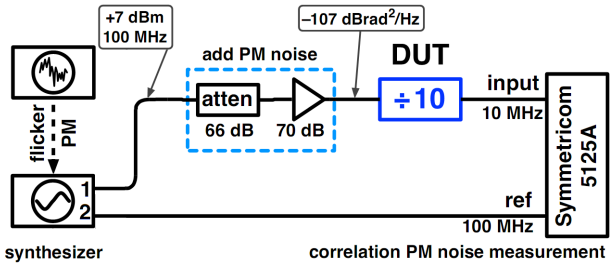

Fig. 4. Experimental setup.

The results of this first experiment are shown in Fig. 5. After waveform shaping, the phase noise is 4 dB higher. This is exactly what expected from the bandwidth of the front end.

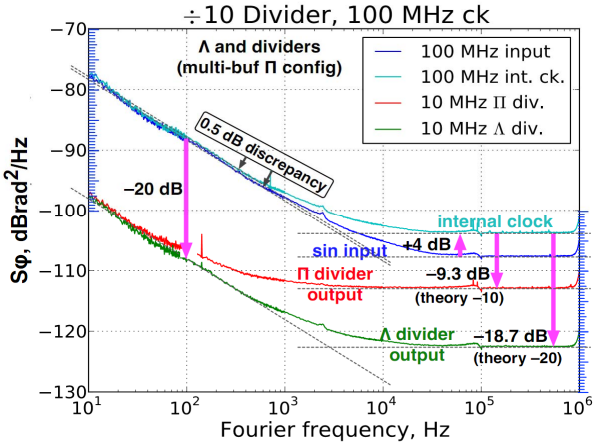

Fig. 5. Result on aliasing in case of flicker and white phase noise.

The flicker noise is scaled down by 20 dB in all cases. This confirms the theoretical prediction that the aliases of $1/f$ noise are too low to be visible on the output spectrum. At 100-1000 Hz, the spectrum differs by 0.5 dB from the straight-line approximation. This is still not explained, and left for future investigation.

The Π divider reduces the phase noise by 9.3 dB, which is close to the expected value of 10 dB. Similarly, the Λ divider reduces the input phase noise by 18.7 dB, close to the value of 20 dB predicted by the theory.

B. Residual Phase Noise of the Dividers

The second experiment is the measurement of the divider phase noise in real-life conditions, in order to demonstrate the practical usefulness of our solutions. The synthesizer is now replaced with a low-noise OCXO, manufactured by Wenzel Inc. [5]. The phase noise spectrum is shown in Fig. 6.

The 1-Hz phase-noise coefficient is $b_{-1}=10^{-12}$ rad^2/Hz (–120 dB rad^2/Hz) for the Π divider, and $b_{-1}=9\times10^{-14}$ rad^2/Hz (–130.5 dB rad^2/Hz) for the Λ divider. Beyond 10 Hz, the flicker noise is 2–3 dB higher. This discrepancy could be an artifact, however the phenomenon is still unexplained.

The white noise is $b_0=2.5\times10^{-16}$ rad^2/Hz (–156 dB rad^2/Hz) for the Π divider, and $b_0=3.2\times10^{-17}$ rad^2/Hz (–165 dB rad^2/Hz) for the Λ divider.

The Allan deviation of the Λ divider is of 2×10^{-14} at 1s. This is equivalent to that of a DDS [6], yet at a way lower cost, complexity, size, and electrical power.

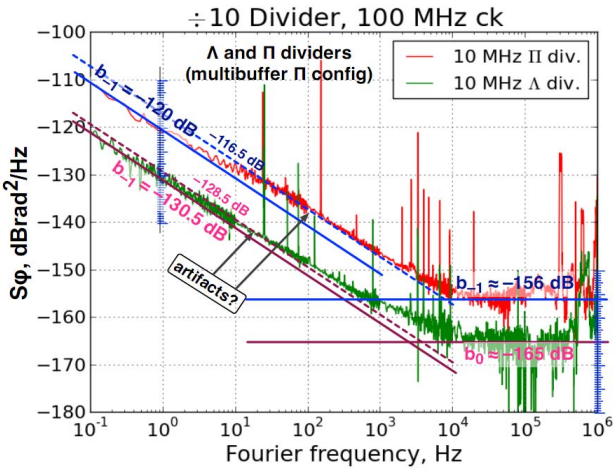

Fig. 6. Residual phase noise of the Π and Λ dividers we tested.

IV. Conclusions

We demonstrated theoretically and experimentally that the white noise of traditional digital frequency dividers (Π topology) is affected by aliasing, and that the new Λ topology fixes the problem.

The phase noise of the Λ divider sets the state of the art for digital dividers. Only the regenerative divider [7] performs lower noise. However, the Λ divider is a fully digital circuit, therefore simple, reproducible, and requires no tuning or adjustment. Thus it is suitable to straightforward integration in larger blocks.

Acknowledgments

This work is partially funded by the ANR (Oscillator IMP project and First-TF network) and by the Région Franche Comté.

We thank Vincent Giordano, Marco Siccardi, Jean Michel Friedt, Marcello Chiaberge for fruitful discussions.

References

[1] W. F. Egan, "Modeling phase noise in frequency dividers", IEEE T UFFC 37(4), July 1990.

[2] http://www.altera.com

[3] http://www.symmetricom.com

[4] E. Rubiola, F. Vernotte, "The cross-spectrum experimental method", arXiv:1003.0113

[5] http://www.wenzel.com

[6] C. E. Calosso, Y. Gruson, E. Rubiola, "Phase noise in DDS", Proc. 2012 IFCS p. 777-782

[7] A. Hati, C. W. Nelson, C. Barnes, D. Lirette, J. DeSalvo, D. A. Howe, "Ultra-low-noise regenerative frequency divider", IEEE T UFFC 59(11), November 2012