Phase Noise and Frequency Stability

Enrico Rubiola

About this text

Most of this material is published as the Chapter 2 of the book U. L. Rohde, E. Rubiola, J. C. Whitaker, *Microwave and Wireless Synthesizers* 2nd ed., Wiley 2021, ISBN 978-1-119-66600-4. Most figures in this text are the original artwork, which has modified and edited by Wiley.

Random phase is denoted with θ in the Wiley book, and with φ here, which is the mostly used symbol.

First and foremost

Before reading this text, it is strongly recommended to keep on hand the *Enrico's Chart of phase noise and two-sample*

variances (Rubiola E., The Enrico's chart of phase noise and two-sample variances, 2023), and its companion article (Rubiola & Vernotte, The Companion of Enrico's Chart for Phase Noise and Two-Sample Variances, 2023). They are available for free download from https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/records/7691686 and https://zenodo.org/document/10050257, respectively, under different flavors of the CC-BY 4.0 license. They can also be downloaded by flashing the QR codes below.





Download the <u>Companion</u> article.



Contents

1	Introduction to phase noise				
	1.1	The clock signal	6		
	1.2	The power spectral density (PSD)	9		
2	The origin of noise				
	2.1	Quantum noise	.15		
	2.2	Thermal (Johnson) noise	.15		
	2.3	Shot (Schottky) noise	.16		
	2.4	Flicker noise	.17		
	2.5	Spurs and other unwanted signals	.17		
	2.5.1	Noise factor, noise figure, and noise temperature	.19		
	2.5.2	The Measurement of the noise temperature	.25		
	2.6	Phase and frequency noise			
	2.6.1	The quantities $S\varphi(f)$, $Sx(f)$, and $S\alpha(f)$			
	2.6.2	The deprecated and misleading quantity $L(f)$			
	2.6.3	Heuristic derivation of Lf and $S\varphi f$ in the case of additive noise			
	2.6.4	Additive and parametric noise	32		
	2.7	The Time Channel (a new way to explain an old concept)			
	2.7.1	The polynomial law, or power law	34		
	2.7.2	Frequency Stability PSD	.39		
	2.7.3	The low-Fourier-frequency part of the phase-noise PSD	.40		
	2.7.4	The RF spectrum of the oscillator signal	.41		
3	The	Allan variance and other two-sample variances			
	3.1	Frequency counters			
	3.1.1	The Π frequency counter			
	3.1.2	The Λ frequency counter	.48		
	3.1.3	The Ω frequency counter.	.51		
	3.1.4	Comparison of the frequency counters	.52		
	3.2	The two-sample variances AVAR, MVAR and PVAR	.53		
	3.2.1	The Allan variance (AVAR)	.54		
	3 2 2	The modified Allan variance (MVAR)	56		
	323	The parabolic variance (PVAR)	.56		
	33	Conversion from spectra to two-sample variances	56		
	331	Comparison between AVAR MVAR and PVAR	60		
4	Phas	e noise in components	63		
•	4 1	Amplifiers	63		
	4.1.1	White and flicker phase noise	.63		
	412	How to choose a low PM noise amplifier	66		
	413	Isolation amplifiers	68		
	414	The Siccardi Isolation Amplifier	70		
	4 2	Frequency dividers	70		
	421	Digital frequency dividers	70		
	4 2 2	Phase noise scaling	71		
	4.2.2	Time-type and phase-type PM noise	74		
		The A divider	77		
	4.2.7 4.2.5	Analog frequency dividers	78		
	43	Frequency multipliers	81		
	44	Direct digital synthesizer (DDS)	80		
	т. т ДД1	Theory of operation	80		
	7.7.1	Theory of operation	.07		

	4.4.2	Signal to quantization ratio (SQR)	91
	4.4.3	Truncation spurs	92
	4.4.4	Phase noise	96
	4.4.5	Examples	98
	4.5	Phase detectors	105
	4.5.1	Noise in the phase-frequency detector	105
	4.6	Noise contribution from power supplies	110
5	Phase	noise in oscillators	115
	5.1 I	Modern view of the Leeson model	116
	5.1.1	The resonator and its impulse response	118
	5.1.2	The oscillator's phase-noise transfer function	122
	5.1.3	The phase noise of the complete oscillator	123
	5.1.4	Lessons from the examples	128
	5.2	Circumventing the resonator's thermal noise	130
	5.3	Oscillator hacking	132
	5.3.1	Moderate/low-Q (type A/C) oscillators	133
	5.3.2	High-Q (type B/D) oscillators	138
6	The n	neasurement of phase noise	143
	6.1 l	Double-balanced-mixer instruments	144
	6.1.1	The measurement of oscillators	147
	6.1.2	Background noise, spurs, and other experimental issues	150
	6.1.3	Asymmetric driving for low-power Signals	153
	6.1.4	Heterodyne measurement of oscillators	154
	6.1.5	The measurement of amplifiers and other two-port components	155
	6.1.6	The discriminator method	158
	6.2	The cross-spectrum method	162
	6.2.1	The rejection of background noise	164
	6.3 l	Digital instruments	168
	6.3.1	The Microchip family of phase noise and Allan deviation testers	172
	6.3.2	The Jackson Labs PhaseStation 53100A	175
	6.3.3	The NoiseXT DNA	180
	6.3.4	The Rohde & Schwarz FSWP family of phase noise analyzers	180
	6.4]	Pitfalls and limitations of the cross-spectrum measurements	185
	6.4.1	The effect of a disturbing signal	186
	6.4.2	Some concepts related to the measurement uncertainty	187
	6.4.3	Thermal energy in the input power divider	189
	6.4.4	The effect of AM noise	191
	6.5	The bridge (interferometric) method	193
	6.5.1	Phase-to-voltage gain and background noise	195
	6.5.2	Building your own system	196
	6.5.3	A practical example	197
_	6.6	Artifacts and oddities often found in real world	198
7	Refer	ences	203
8	Sugg	ested readings	213
	8.1 I	ower spectra and Fourier transform	213
	8.2 I	Electromagnetic Compatibility	213
	8.3 (Jeneral Aspects of Noise	213
	8.4 Phase Noise, Frequency Stability, and Measurements		
	8.5 Amplifiers		
	ð.0	requency Dividers	213

8.7	Frequency Multipliers	
8.8	DDS	
8.9	Phase-Frequency Detectors	
8.10	Oscillators	
8.11	Resonators	
8.12	Double Balanced Mixer	

1 Introduction to phase noise

The noise of an oscillator, of a synthesizer, or of other signal sources is an important parameter which requires a sophisticated framework to describe it appropriately. If the RF spectrum is measured at the output of an oscillator, a curve such as that of Figure 1 is observed. Rather than all the power being concentrated at the "exact" oscillator frequency, some is distributed in the sidebands. In principle, such sidebands are originated by phase or frequency fluctuations, and by amplitude fluctuations. A traditional spectrum analyzer cannot divide AM from PM or FM, and it cannot identify the noise correctly. Here, phase and frequency fluctuations are our main concern.

As noise is a form of stability, it is useful to characterize the frequency stability in the time domain in several areas. Such areas are referred to as short-term stability, medium-term stability, and long-term stability or aging. The interpretation of what is short, medium or long is a matter of context and jargon of specific communities. We can agree that short-term stability extends between a very small fraction of a second to 1 s, maybe under some considerations up to 1 minute, and the value for the stability between 1 s and 1 minute will be about the same. For longer time, we talk about longterm stability or aging. The aging is typically expressed in forms of how many parts in 10^{-10} or 10^{-11} per day the frequency changes. Lower values are found in atomic frequency standards, like the Hydrogen maser and the Cesium clock. This information is in the time domain. In the frequency domain, we find terms like *random walk*, *flicker*, and *white*, which describe the slope of spectral density on a log-log scale. The Fourier frequency, at times labeled f_m or just f when there is no ambiguity, is at times called sideband frequency, offset frequency, or modulation frequency. In this article, we will refer to it as Fourier frequency, describing the phase fluctuations of an oscillator at a certain frequency off the carrier frequency. The most common characterization of the phase noise of a source or of a component is the power spectral density L(f), or equivalently $S_{\omega}(f)$, which we will study extensively in the following pages.



Figure 1 – Example of the microwave spectrum of an oscillator. Courtesy of Yannick Gruson, FEMTO-ST Institute, France.

1.1 The clock signal

The clock signal is a highly pure sinusoidal signal which we can write as

$$v(t) = V_0[1 + \alpha(t)] \cos[2\pi f_0 t + \varphi(t)]$$
(1)

where V_0 is the peak amplitude, f_0 is the carrier frequency, $\alpha(t)$ is the random fractional amplitude, and $\varphi(t)$ is the random phase. The amplitude noise (AM noise) and phase noise (PM noise) features of such signal are sketched in Figure 2. In the presence of AM noise only, the peak amplitude changes at random by an amount equal to $V_0\alpha(t)$, while the phase is unperturbed, and consequently the zero crossings occur exactly when expected. Oppositely, a signal affected by PM noise only has peak amplitude exactly equal to V_0 , while the phase fluctuates at random, and consequently the zero crossings fluctuate. Both AM and PM are present in all real signals, albeit not in equal amount. The relevant notations and several basic concepts about phase noise and frequency stability presented in this Chapter are found in the milestone articles (Barnes & et al, 1971) (Rutman, 1978). A summary of our notation found on Table 1.



Figure 2 – Time-domain representation of AM and PM noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Quantity	Our notation	General time and frequency literature	Ch.2 cited (and PLL books)	Dimension
Carrier frequency	f_0	$ u_0 ext{ or } f_0$	f_0	Hz
Fourier frequency	f or f_m	f	f or f_m	Hz
Random phase	$\varphi(t)$	$\phi(t) \text{ or } \varphi(t) \qquad \theta(t)$		rad
	$S_{\varphi}(f)$	$S_{\phi}(f)$ or $S_{\varphi}(f)$	$S_{\theta}(f)$	rad²/Hz
PM noise PSD	L(f)	$\mathcal{L}(f)$	$\mathcal{L}(f)$	$10\log L(f)$
	$L(f) = \frac{1}{2}S_{\varphi}(f)$ by definition	$\mathcal{L}(f) = \frac{1}{2}S_{\varphi}(f)$ by definition	$\mathcal{L}(f) = \frac{1}{2}S_{\theta}(f)$ by definition	dBc/Hz
Random fractional amplitude	$\alpha(t)$	$lpha(t),\ \epsilon(t) ext{ or } \epsilon(t)$	$\alpha(t)$	dimensionless
AM noise PSD	$S_{\alpha}(f)$	$ \frac{S_{\alpha}(f),}{S_{\epsilon}(f) \text{ or } S_{\epsilon}(f)} $	$S_{\alpha}(f)$	Hz ⁻¹

Table 1 – Notation.

Another, minor, difference is that the normalized quantities x(t) and y(t), and some specific parameters, for example the coefficient b_n , k_n and h_n found in the polynomial law discussed later, are written in Sans Serif font instead of regular math font. In this way, we let the regular math-font alphabet free, available for general use.

A well-designed, high-quality oscillator exhibits high amplitude stability. For reference, we found that a short-term stability of 10^{-6} is rather common in high-end RF oscillators, with flicker power spectral density (Rubiola E., The measurement of AM noise of oscillators, 2005). The amplitude noise is generally considered a special topic, discarded in system analysis and design. For our purposes, we assume that V_0 is the best estimation of the amplitude, so that $\alpha(t)$ is very small and has a mean close to zero. In formulae, $|\alpha(t)| \ll 1$ and $\langle \alpha(t) \rangle \approx 0$.

By contrast, we cannot assess a boundary for $\varphi(t)$ because real oscillators are subject to ageing, drift and sensitivity to the environment. If we freeze the numerical value of f_0 , these phenomena go in $\varphi(t)$, which can accumulate a quite large number of cycles (2π rad). Thus, $\varphi(t)$ can be a divergent process.

The traditional analog phase detectors work in the range of $\pm \pi$, or $\pm \pi/2$ radians. When these detectors are used for phase-noise measurements, it is often necessary that $|\varphi(t)| \ll 1$ for the duration of the test. This can be achieved by slow phase locking, or with other techniques. Conversely, I-Q detection enables the phase measurement in unbound range, giving a valid result even if $\varphi(t)$ accumulates a large number of cycles. The unbound phase is often called *unwrapped phase*.

The quantity $\varphi(t)$ is not the one and only option to describe the oscillator fluctuations. Other physical quantities are often preferred, depending on purposes and applications, and sometimes on personal preferences. The clock signal carries a time reference. For this reason, it may be convenient to describe the phase noise as the *time fluctuation*, or *phase time* (fluctuation)

$$\mathbf{x}(t) = \frac{1}{2\pi f_0} \,\varphi(t) \tag{2}$$

which is $\varphi(t)$ converted into time and expressed in seconds. Of course, $\varphi(t)$ is allowed to exceed $\pm \pi$ radians, and the number of cycles accumulated is accounted for in x(*t*).

For the layman, x(t) is the time error of a clock driven by our oscillator through an appropriate gearbox. The quantity x(t) is independent of the carrier frequency, thus it is suitable to compare the timekeeping feature of watches in terms of daily or monthly error.

It is well known in radio engineering that the angular modulation can be formulated in two fully equivalent ways, as phase modulation (PM) or as frequency modulation (FM). In the same way, we can freeze the random phase and move the random fluctuation from $\varphi(t)$ to the carrier frequency f_0 by replacing

$$f_0 \to f_0 + (\Delta f_0)(t)$$

The quantity Δf_0 is the random frequency fluctuation or error, expressed in Hz. Accordingly, $\theta(t)$ is replaced with $2\pi \int (\Delta f_0)(t) dt$. The subscript "0" can be dropped if there is no risk of confusion. The notation $(\Delta f_0)(t)$ emphasizes the fact that Δf is a single variable, function of time. The clock signal becomes

$$v(t) = V_0 [1 + \alpha(t)] \cos[2\pi f_0 t + 2\pi \int (\Delta f)(t) dt]$$
(3)

The *fractional frequency fluctuation*, or for short *fractional frequency* y(t) is another quantity often used to describe the oscillator fluctuation, defined as

$$y(t) = \frac{1}{f_0} \, (\Delta f_0)(t) \tag{4}$$

Of course, y(t) is dimensionless. The main reason for y(t) is that it provides straightforward, fair comparison of oscillators at different frequencies, with no conversion factor. Very common specifications, like "0.1 ppm aging after one year" or "thermal drift of $10^{-9}/^{\circ}$ C," refer implicitly to the quantity y(t). It holds that

$$\mathbf{y}(t) = \frac{d}{dt} \mathbf{x}(t) \tag{5}$$

1.2 The power spectral density (PSD)

The PSD S(f) tells us how the power of a signal is distributed among frequencies, similarly to a prism which split the light in colors, or to a bank of filter which splits the input signal in bands (Figure 3). Taking the power in strict physical sense, the physical dimension of the PSD is W/Hz. The generalized power is often used, which is a squared quantity like a voltage, a current, or a phase. The physical dimension follows in a rather obvious way. For example, the PSD of a voltage v(t), denoted with $S_v(f)$, has the dimension of V²/Hz. Textbooks often say that there is a 1 Ω resistor implied in order to get a proper unit of power, but this is not necessary because the PSD is a mathematical tool. The PSD of the random phase $\varphi(t)$, denoted with $S_{\varphi}(f)$, has the dimension of rad²/Hz.



Figure 3 – Parallel spectrum analyzer. With the appropriate normalization, P_1/B_1 , $P_2/B_2...P_N/B_N$, the output is the PSD of the input signal. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Engineers, physicists and other experimentalists use the single-sided PSD, which is restricted to f > 0. There are two important, and straightforward properties associated with the PSD. The first is the *Parseval* theorem, which states that the power of a signal x(t) calculated in the time domain and in the frequency domain is the same. In formula

$$P = \lim_{T \to \infty} \frac{1}{T} \int_0^T |x(t)|^2 dt = \int_0^\infty S_x(f) df$$
 (6)

The second property states that the power of the signal x(t) after band pass filtering is calculated by integrating the PSD over the filter bandwidth. Denoting the band limits with a and b, the power of the filtered signal is

$$P = \int_{a}^{b} S_{x}(f) df$$

A more rigorous version of the above is formulated with the transfer function $|H(f)|^2$ of the filter as

$$P = \int_0^\infty S_x(f) \, |H(f)|^2 \, df$$
 (7)

General instruments calculate the PSD using the Fast Fourier Transform (FFT). This implies that the signal is sampled at a given frequency, digitized, and truncated in time. Of course, the input is a stream of real numbers, so the FFT has the usual symmetry properties. Using the uppercase X(f) for the Fourier transform, the subscript T for the truncation over the measurement time of duration T, and f_s for the sampling frequency, the PSD is

$$S_x(f) = \frac{2}{T} |X_T(f)|^2$$
 for $0 < f < f_s/2$ (8)

The factor of 2 is necessary for energy conservation. The energy of $X_T(f)$ is equally split between negative frequencies and positive frequencies, and the energy associated to the negative frequencies of $X_T(f)$ is folded to the positive frequencies in $S_x(f)$. The multiplication by 1/T comes from the mathematical development of $S_x(f)$, omitted here. However, the need for such factor is quite evident from physical dimensions. Letting x(t) be a voltage, a valid PSD must have the physical dimension of V²/Hz. Since the Fourier transform $X(f) = \int_{-\infty}^{\infty} x(t) e^{-j2\pi ft} dt$ has the dimension of Vs, or equivalently V/Hz, the quantity $|X_T(f)|^2$ has the dimension of V²/Hz². The multiplication by 1/T turns the unit into V²/Hz.

It is a common practice to improve the confidence of the measure, or the readability of the PSD plot, by averaging on a suitable number *m* of acquisitions. This is written as

$$\langle S_x(f) \rangle_m = 2 \frac{1}{T} \langle |X_T(f)|^2 \rangle_m \qquad \text{for } 0 < f < f_s/2 \tag{9}$$

The angle parentheses $\langle ... \rangle$ denote the average, and the subscript *m* the number of acquisitions.

The PSD has a resolution bandwidth (RBW) which results from the sampling frequency f_s , from the measurement time T, and from the window (tapering) function used in the FFT. The minimum RBW is 1/T, limited by the time-frequency indetermination theorem. The latter is usually written as $\Delta \omega \Delta t \ge 2\pi$. Smoother window functions result in smaller frequency-leakage and in better capability to identify correctly a dip between peaks, at the cost of broader RBW. The RBW may change with frequency because the full span is usually obtained by joining pieces sampled at different frequencies. The spectrum of a pure sinusoid of power P and frequency f_1 is a narrow line, ideally a Dirac delta function. However, a real instrument displays a line of bandwidth equal to the RBW, and PSD

$$S(f) = \frac{P}{\text{RBW}}$$
 at $f = f_1$, 0 elsewhere (10)

This can be misleading in the case of a smooth PSD affected by spurs. Oppositely, a regular RF and microwave spectrum analyzer displays the *Power Spectrum* (PS), which is denoted with G(f) and has the physical dimension of a power (W), or of a generalized power (V², A², etc.). In the presence of noise of PSD equal to N(f), such instrument displays

$$G(f) = N(f) \times \text{RBW}$$
(11)

Notice that the term *power spectrum* is sometimes used as an abridged form for *power spectrum density*, leading to a confusion between PSD and PS. Of course, we recommend great attention in all cases where the context may leave a doubt.

Power spectra and PSD are well established concepts. To the reader willing to know more, we recommend the books by Blackman and Tukey, Percival and Walden, and Brigham.

Example 2-1 – Power spectral density. Measuring a signal consisting of white noise plus spurs from the power grid, we get the PSD shown on Figure 4. We want to know the power and the RMS voltage of all the components of such signal. A straightforward calculation gives the results shown in Table 2. ■



Figure 4 – Example of PSD constituted of white noise and narrow spectral lines.

Signal component	Power (generalized)	RMS voltage
White noise	$P_n = 10^{-14} \frac{\mathrm{V}^2}{\mathrm{Hz}} \times 10^5 \mathrm{Hz} = 10^{-9} \mathrm{V}^2$	$e_n = 31.6 \ \mu V$
60 Hz	$P_1 = 10^{-8} \frac{V^2}{Hz} \times 0.125 \text{ Hz} = 1.25 \times 10^{-9} \text{ V}^2$	$e_1 = 35 \ \mu V$
120 Hz	$P_2 = 10^{-11.4} \frac{V^2}{Hz} \times 1.25 \text{ Hz} = 5 \times 10^{-12} \text{ V}^2$	$e_2 = 2.2 \ \mu V$
180 Hz	$P_3 = 10^{-10.2} \frac{V^2}{Hz} \times 1.25 \text{ Hz} = 7.9 \times 10^{-12} \text{ V}^2$	$e_3 = 8.9 \ \mu V$
Total	$P = P_n + P_1 + P_2 + P_3 = 2.26 \times 10^{-9} \mathrm{V}^2$	$e = 47.6 \mu\text{V}$

Table 2 – Interpretation of the spectrum of Figure 4.

2 The origin of noise

2.1 Quantum noise

[I plan to add this section to a forthcoming version]

2.2 Thermal (Johnson) noise

This type of noise is a form of blackbody radiation confined in an electrical line, and originates from the thermal agitation of free charges in conductors. It has been observed experimentally by Johnson (Johnson, 1928) and explained theoretically by Nyquist (Nyquist, 1928) in 1928.

The Planck law for the blackbody radiation, in the case of a resistor *R* at the temperature *T* is usually written as

$$S(f) = \frac{hf}{e^{hf/kT} - 1} \qquad [W/Hz] \tag{12}$$

where $h = 6.626 \times 10^{-34}$ Js, or equivalently W/Hz, is the Planck constant, and $k = 1.381 \times 10^{-23}$ J/K is the Boltzmann constant. The quantity hf is the photon energy, and kT is the thermal energy of the free electrical charges in thermal equilibrium. Equation (12) can be interpreted as the available power spectral density, i.e., the power in 1 Hz bandwidth transferred from the resistor to the load of equal resistance R, ideally cold (T = 0 K). For the sake of rigorous treatise, one has to account the zero-point energy hf, which is a quantum-physics concept. The complete expression is

$$S(f) = hf + \frac{hf}{e^{hf/kT} - 1} \qquad [W/Hz]$$
(13)

The zero-point energy dominates over the blackbody radiation at low temperature and high frequencies. The cutoff frequency is given by

$$f_{\rm th} = \frac{kT}{h} \ln(2) \tag{14}$$

Beyond $f_{\rm th}$, the blackbody radiation rolls off because the average thermal energy kT is insufficient to produce a photon of energy hf. The exponential shape of the roll off function reflects the statistical nature of the thermal energy. For reference, $f_{\rm th}$ is of 4.3 THz at room temperature, and of 60.6 GHz at the liquid Helium temperature (4.2 K). So, in almost all practical applications it holds that $hf \ll kT \ln(2)$. Thus, $e^{hf/kT} \simeq 1 + hf/kT$, and (13) is approximated with white noise

$$N = kT \qquad [W/Hz] \tag{15}$$

where the S(f) is replaced with the symbol N to comply with usual notation.

The Planck law refers to the available power, in impedance matching conditions. If the conductor is left open, the electromotive force across its ends has PSD

$$S_{\nu}(f) = 4kTR \qquad [V^2/Hz] \tag{16}$$

The reference value

 $T_0 = 290 \text{ K}$ (17.2 °C or 66.3 °F)

is often used in electronics and radio engineering as a convenient approximation of the physical temperature found in most practical cases. The corresponding thermal energy is a round number

$$kT_0 = 4 \times 10^{-21} \text{ W/Hz.}$$
 i.e., -174 dBm/Hz (17)

At that temperature, the equivalent noise voltage kT_0/q is equal to 25 mV. In the jargon of semiconductors, this quantity is often denoted with V_T . The thermal emf in 1 Hz bandwidth is denoted with $e_n = \sqrt{4kT/R}$. Across a $R = 50 \Omega$ resistor we find

 $e_n = 0.9 \text{ nV}/\sqrt{\text{Hz}}$, or $-181 \text{ dB V}/\sqrt{\text{Hz}}$.

The thermal energy contributes to phase noise in two ways. The first and most important effect is the microwave noise added to the carrier. The second, and generally minor, is a phase modulation from the near-DC part of the thermal noise.

2.3 Shot (Schottky) noise

The shot noise in electrical circuits originates from the discrete nature of the electrical charge, which is an integer multiple of the electron charge $q = 1.602 \times 10^{-19}$ C. Its discovery is generally credited to Schottky (Schottky, 1918). Shot noise occurs in junctions, vacuum tubes, and other devices or physical experiments where electrons and holes appear as individual particles. The electrical current in regular conductors, like wires and resistors, is a field, which does not generate shot noise. The standard picture for the electrical current in vacuum is a stream of Φ electrons per second emitted at random time, with no memory and no space correlation. The average current is $\langle I \rangle = \Phi q$. Unlike thermal noise, shot noise is only present when electrical current flows, and it is independent of temperature and of the resistance of the electrical circuit. This type of noise is a Poisson process, which has uniform (white) PSD

$$S_I(f) = 2q\langle I \rangle \quad A^2/\text{Hz}$$
 (18)

The same formula holds in photodiodes, where $I = \eta P / h f_{opt}$ as the photocurrent, f_{opt} is the frequency of the light, and η is the quantum efficiency, i.e., the probability that a photon is captured and generates photoelectron or a photo-electron-hole pair.

The bandwidth cannot be infinite. The shot noise rolls off at the cutoff frequency

$$f_{\rm sh} = \frac{1}{2}\Phi = \frac{\langle I \rangle}{2q} \tag{19}$$

The reason is that the generalized power $S_I(f)f_{\rm sh}$ must be equal to the generalized DC power $\langle I^2 \rangle$. This is a consequence of the properties of the Poisson distribution. An alternate interpretation relies on the sampling theorem, which states that the maximum frequency is half the sampling frequency. The same holds for random sampling, just with a smooth cutoff. In this picture, the electrons are electrical pulses playing the role of the random samples.

The shot noise equals the thermal noise at the critical current

$$\langle I \rangle = \frac{kT}{2qR} \tag{20}$$

determined by $2q\langle I \rangle R = kT$. For reference, this critical current is of 250 µA at 290 K with $R = 50 \Omega$, and the associated power is of 3.1 µW (-25 dBm).

Like thermal noise, the shot noise contributes to phase noise as it adds to the carrier, and as a near-DC noise which modulates the carrier.

2.4 Flicker noise

Flicker noise is characterized by the PSD proportional to 1/f, or close to 1/f, in a wide range of frequency. The digression about whether flicker noise is fundamental or not, is more academic than pragmatic, and goes far beyond our scopes. The most interesting fact about flicker is its ubiquity (Milotti, 2002) (Levitin, Chordia, & Menon, 2012). After being discovered in carbon microphones (Christiansen & Pearson, 1937), it is found in geophysical phenomena, climatology, mechanics, optics, classical music, Internet traffic, and in a variety of other domains (Milotti, 2002), and of course electronics.

Flicker noise originates around DC. Flicker of phase and flicker of frequency are parametric noise types, generated by a near-DC process which modulates the phase or the frequency of a signal. Flicker is of paramount importance for us because it turns out to be a major limitation in the noise of synthesizers, and of oscillators as well.

2.5 Spurs and other unwanted signals

The generation of a clean microwave signal, free from spurs, interferences and other unwanted signals is a blend of engineering, experience and art. We all are used to the presence of unwanted signals at 60 Hz (50 Hz in Europe) and multiples, from the power grid. Such signals show up as spectral lines in phase noise, and as a hum sound in audio-frequency. They get in microwave circuits in several different ways, like the ripple of supply lines, unequal potential of different ground points, ground loops, and magnetic fields captured by loops and turned into emf. These signals can modulate the carrier in a various way, or are added at the VCO input of oscillators, and transposed to the carrier as parametric noise. Other interferences have similar behavior, like the ripple from switching power supplies, and the high-voltage raster signals from cathode ray tubes. Unshielded AC magnetic fields affect the magnetic permeability of ferrite cores, which modulate the phase of RF signals. Acoustic noise gets in microwave circuits via the sensitivity to acceleration. Most of such noise comes from fans, from the mechanical vibration of transformers, and again from unshielded magnetic fields via the AC attractive force on iron parts.

Disturbances from 50–60 Hz power grid usually extends up to approximately 1 kHz, becoming progressively smaller as the number of harmonic decreases. Odd-order harmonics are generally stronger than even-order harmonics. Acoustic noise is most present between 1 and 2 kHz, while switching power supplies and cathode-ray tubes are typically in the 10–100 kHz region.

Quartz oscillators and other electro-mechanical oscillators are highly sensitive to acoustic noise. Some are also sensitive to magnetic fields, mainly because of the presence of magnetic materials in packaging, and to springs. YIG materials are highly sensitive to magnetic fields, but packaged YIG oscillators are generally well shielded.

Digital circuits can be an annoying source of spurs and disturbances because of the variety of effects. Radiation occurs at the clock frequency, or at the bus frequency, which occurs from 100 MHz or less, up to 1 GHz. High peaks of current on supply lines or ground are driven by software in microprocessors and FPGAs, which sometimes cause a large number of transistors to switch with random, pseudo-random, or pseudo-periodical appearance. Spurs may be observed in a wide range of frequency from Hz to MHz and beyond. Impressively large spurs may be observed in digital phase-noise test equipment, if the user removes the post-processing filters that hide them. Digital circuits can also interfere with other parts of a system in another subtle way, via thermal fluctuations. The problem arises from modern VLSI integrated circuits, where high dissipated power per unit of Silicon surface is necessary to achieve the computing power. A dissipation of a few Watts is usual in FPGAs, DDSs, etc., in a small chip, and the surface is proportionally hot. If not appropriately shielded, temperature fluctuations show up generally below 10 Hz with a steep spectrum, of slope $1/f^5$ or higher. Every circuit is a special case, and the literature provides little or no help.

To complete the picture, the electromagnetic interference between different parts of a system is one of the earliest known forms of spurious signals, and however sometimes difficult to model and predict precisely. The electromagnetic interference impacts on systems as an additive disturbance, or through intermodulation in junctions.

Some classic reference books are available on this topic, by Goedbloed, Ott, Paul, and Perez.

2.5.1 Noise factor, noise figure, and noise temperature

The noise factor, or equivalently the noise figure, is likely the most used, if not overused, parameter to characterize the noise of two-port components or systems, like amplifiers, frequency converters, radio receivers etc. However often used interchangeably, one should prefer the term *noise factor* for the dimensionless quantity *F*, and *noise figure* (NF) for *F* expressed in dB

$$NF = 10 \log_{10} F [dB]$$
 (21)

A popular definition of *F*, given by Friis (Friis, 1944) in 1944, is

$$F = \frac{\text{SNR}_i}{\text{SNR}_o} \qquad (\text{Friis, obsolete definition}) \tag{22}$$

where SNR is the Signal-to-Noise Ratio, and the subscripts "*i*" and "*o*" stand for input and output. A substantially equivalent, and less known definition was given by North in 1942 (North, 1942), as the ratio of (1) the output noise power from the transducer to (2) the output noise power from an equivalent noise-free transducer.

The problem with (22), and also with the North's definition, is that the degradation to the SNR depends on the noise of the source that excites the device. This is seen by sending a signal of power P_i and thermal noise kT_i to a device of power gain A^2 . Denoting with N_{dev} the available noise contribution of the device in 1 Hz bandwidth, as observed at the output, Equation (22) gives

$$F = \frac{P_i}{kT_i} \frac{N_{dev} + A^2 kT_i}{A^2 P_i} = 1 + \frac{N_{dev}/A^2}{kT_i}$$

For the same device at the same temperature, N_{dev} is the same, but F is affected by the temperature T of the input termination.

The ambiguity of (22) is solved with the new definition proposed by the IRE (Haus & et al., 1960), and adopted by the NIST (at that time, NBS) (Arthur, 1974):

Definition (IRE): The **noise factor**, at a specified input frequency, is defined as the ratio of **(1)** the total noise power per unit bandwidth at a corresponding output frequency available at the output port when the noise temperature of the input termination is standard (290 K) to **(2)** that portion of (1) engendered at the input frequency by the input termination.

It is worth mentioning that the IRE definition differs from the Friis definition only in the use of the reference temperature $T_0 = 290$ K. For this reason, (22) is still found in the literature, with the reference temperature often implied. Both definitions are smart in that the input circuit is taken in its actual configuration, making no assumption on

impedance matching and on best noise impedance, if different from the characteristic impedance.

The *noise temperature*, denoted with T_e , is another widely used parameter to assess the noise, by analogy with the thermal noise.

The case of a source (Figure 5 A) is quite straightforward. We say that the source has equivalent noise temperature T_e at a specific frequency if the available noise power at its output in 1 Hz bandwidth is $N = kT_e$, the same of a resistor at temperature T_e .

The case of a two-port device is similar, and illustrated on Figure 5 B. The available output noise N_o is the same for the two cases, (1) when the real device is connected to a noise-free source, and (2) when a noise-free, and otherwise equal device is connected to a source at the temperature T_e .





Figure 5 – Noise temperature.

The noise temperature is related to the noise factor. Assuming impedance matching at the input, the device's noise contribution observed at the output is A^2kT_e . Accordingly, the noise factor (IRE definition) can be written as

$$F = \frac{A^{2}kT_{e} + A^{2}kT_{0}}{A^{2}kT_{e}}$$

$$F = 1 + \frac{T_{e}}{T_{0}}$$
(23)

thus

As we mentioned, the SNR degradation produced by a two-port device depends on the input noise, hence on the temperature of the source. Figure 6 shows an example.
All the plots match the Friis' definition
$$F = \text{SNR}_i/\text{SNR}_o$$
, but only the plot $T_i = 290$ K matches the IRE definition.

The noise factor of an attenuator is an interesting case because of the physical insight it provides. We first assume that the attenuator is impedance-matched at both ends, and that it receives the thermal noise kT_i from a resistor. The attenuator "amplifies" the input noise by a factor of $A^2 < 1$, as it does with any signal. Yet, the attenuator adds its own contribution, which we will calculate. For this purpose, we set the temperature of both input resistor and attenuator to T_0 . The output noise is the sum of (1) the input noise attenuated, that is, $N' = kT_0A^2$, and (2) the noise $N'' = A^2kT_e$ from the attenuator. Because the attenuator output is equal to $N_0 = kT_0$. Thus, $N'' = kT_0 - N'$, and consequently $N'' = kT_0(1 - A^2)$. Referring the attenuator noise to the input, we find $N_e = N''/A^2 = kT_0(1 - A^2)/A^2$. Finally, the equivalent temperature N''/k is $T_e = T_0(1 - A^2)/A^2$. Using $F = 1 + T_e/T_0$, we find $F = 1/A^2$. This proves the thumb rule that the noise factor of an attenuator is equal to the power attenuation.



Figure 6 – SNR degradation due to a two-port device of noise temperature T_e for various values of temperature T_i of the input source. From Fundamentals of RF and Microwave Noise Figure Measurements, September 2017, © 2010-2017 Keysight Technologies. (Keysight Technologies, September 2017). Used with permission, and adapted to the text.

The noise factor and the noise temperature are simplified representations of the reality. More accurate noise models, generally good for virtually all practical purposes, resort to the seminal paper by Rothe and Dahlke (Rothe & Dahlke, 1956) on the theory of linear four-poles. Following this approach, the noise is best described by adding a voltage generator e_n and a current generator i_n at the device input, as shown on Figure 7. Such generators are described in terms of their power spectral densities $S_e(f)$, $S_i(f)$ and $S_{ie}(f)$, or equivalently in terms of their variances $\langle |e_n^2| \rangle$, $\langle |i_n^2| \rangle$ and the covariance $\langle i_n e_n^2 \rangle$ in 1 Hz bandwidth. The most relevant fact is that e_n and i_n define an impedance which generally differs from the input impedance. Thus, the device has two "optimum" impedances, one for maximum power transfer, which refers to the usual conjugate matching, and one for lowest noise. Consequently, *the impedance of the generator impacts on the noise factor*, and low noise design requires a tradeoff between gain and noise.



Figure 7 – Generalized noisy device.

After doing the appropriate math, the noise factor is given by

$$F(\Gamma_g) = F_o + 4 \frac{R_n}{Z_0} \frac{\left|\Gamma_o - \Gamma_g\right|^2}{\left|1 + \Gamma_o\right|^2 \left(1 - \left|\Gamma_g\right|^2\right)}$$
(24)

where F_o is the minimum (optimum) noise factor, R_n is the noise resistance (the sensitivity of noise factor to source resistance changes), Z_0 is the nominal input, Γ_g is the reflection coefficient, and Γ_o is the value of Γ_g with which F_o is achieved. The parameters F_o , R_n , and Γ_o describe the noise of the device, and Γ_g is a free design choice. Represented on the Smith chart, the noise factor $F(\Gamma_g)$ looks like a set of equal-noise circles (see Figure 8 for an example).

Packaged amplifiers are often impedance-matched in a wide range of frequency, thus the noise factor is degraded because impedance matching is often privileged versus noise matching, and also because of the loss of the input circuit. Values of 1 dB to 4 dB are rather common. Conversely, the noise factor of a transistor can be quite low, yet at the cost of uncomfortable impedance matching (see the example on Table 3).



Figure 8 – Example of equal-noise circles of an amplifier, on the Smith chart. Edited, from Fundamentals of RF and Microwave Noise Figure Measurements, September 2017, © 2010-2017 Keysight Technologies (Keysight Technologies, September 2017). Used with permission and adapted to the text.

f_0 , GHz F_o , dB		$ \Gamma_o $	$\arg(\Gamma_o)$, deg	R_n, Ω	Gain, dB	
0.5 0.09		0.37	16.1	4.0	26.6	
0.7	0.12	0.37	28.5	3.5	24.6	
0.9	0.16	0.37	40.6	3.0	23.0	
1.0	0.18	0.37	46.6	3.0	22.3	
1.9	0.34	0.39	97.4	1.5	17.8	
2.0	0.35	0.39	102.7	1.5	17.4	
2.4	0.42	0.40	123.4	1.5	16.3	
3.0	0.53	0.41	152.5	1.5	14.9	
3.9	2.5	13.3				
5.0	0.89	0.45	-127.1	5.0	11.8	
5.8	1.03	0.46	-102.1	8.0	10.8	
6.0	1.06	0.47	-96.5	9.0	10.6	
Conditions: $V_{DS} = 4 \text{ V}$, $I_{DS} = 30 \text{ mA}$.						
Data are taken from the TAV-581+ data sheet (Mini Circuits, NY, USA)						

Table 3 - Typical noise parameters of the TAV-581+ transistor.

When several stages are cascaded, the noise factor of the chain is given by the Friis formula (Friis, 1944)

$$F = F_1 + \frac{F_2 - 1}{A_1^2} + \frac{F_3 - 1}{A_1^2 A_2^2} + \cdots$$
(25)

Accordingly, the first stage of a chain should have a low noise factor, while the noise factor requirement of subsequent stages is relaxed. The ideas underneath the Friis formula are simple. The overall noise factor F is referred to the input of the chain. The noise contribution of each amplifier is added as square voltage, power, or PSD, as appropriate, because the amplifier are separate devices, and their fluctuations are statistically independent. The contribution F_1 of the first stage is obvious. The second stage has no resistive load at the input, thus no kT_0 . There remains $(F_2 - 1)kT_0$, which becomes $(F_2 - 1)kT_0/A_1^2$ referred at the input of the chain, after dividing by the power gain A_1^2 of the first stage. Recursively, the noise $(F_3 - 1)kT_0$ of the third stage is divided by the power gain $A_1^2A_2^2$ of the two preceding stages, etc.

The Friis formula is an approximation based on impedance matching. A more accurate model should account for two main facts. First, impedance mismatch calls for a correction term which lowers the gains A_i^2 , based on the reflection coefficients between the *i*-th and the (i + 1)-th stage. Second, the noise factor F_1 should account for the source impedance, and likewise all the F_{i+1} should account for the output impedance of the preceding, *i*-th, stage.

2.5.2 The Measurement of the noise temperature

The equivalent temperature of a device is usually measured with the Y method shown on Figure 9. The method consists of switching two impedance-matched input sources, one at the temperature T_h (hot) and the other at the temperature T_l (cold). Asymptotically, if one can set $T_h \rightarrow \infty$ and $T_l \rightarrow 0$, the temperature T_h is the probe signal that enables the measurement of the power gain A^2 , and T_l gives the equivalent temperature T_e after taking away the gain A^2 . In actual experimental conditions, the output noise power is

$$P_h = A^2 k (T_h + T_e) B$$
$$P_l = A^2 k (T_c + T_e) B$$

where *B* is the bandwidth of a filter at the device output. The solution of the system is

$$T_e = \frac{T_h - YT_l}{Y - 1}$$
 with $Y = \frac{P_h}{P_l} = \frac{T_h + T_e}{T_l + T_e} > 1$ (26)

The main virtue of the *Y* method is that the factor A^2/B cancels in the evaluation of *Y*. This simplifies the calibration and results in improved accuracy because *A* generally

suffers from flatness defect, while *B* is the equivalent noise bandwidth, which results from integrating the transfer function.

It is worth mentioning that the equivalent noise temperature includes thermal noise in strict sense, the shot noise, and any other noise process. For this reason, people with a background in optics may find this concept particularly misleading. In fact, in optical systems there is no temperature and, in high SNR condition, the electrical noise at the detector output is chiefly shot noise.



Figure 9 – The Y method for the measurement of the equivalent noise temperature of a device.

2.6 Phase and frequency noise

2.6.1 The quantities $S_{\varphi}(f)$, $S_{x}(f)$, and $S_{\alpha}(f)$

The PSD of the random phase $\varphi(t)$, denoted with $S_{\varphi}(f)$ [rad²/Hz] is the obvious choice to characterize the phase noise in the frequency domain. Its use already appeared as $S_{\varphi}(f)$ in an article (Baghdady, Lincoln, & Nelin, 1965) presented at a NASA symposium (Chi, 1965) intended to clarify spectral purity and related problems.

Similarly, the phase time fluctuation can be characterized in the frequency domain in terms of $S_x(f)$, which is PSD of x(t). Because it holds that $x(t) = \varphi(t)/2\pi f_0$, the quantities $S_{\varphi}(f)$ and $S_x(f)$ are fully equivalent, and related by

$$S_{\rm x}(f) = \frac{1}{4\pi^2 f_0^2} S_{\varphi}(f) \tag{27}$$

Likewise, the quantity $S_{\alpha}(f)$ is the PSD of the random fractional amplitude $\alpha(t)$.

2.6.2 The deprecated and misleading quantity L(f)

The quantity L(f), sadly the most widely used one to measure phase noise, is deprecated because (i) it is incompatible with the International System of Units SI, and (ii) because it generates confusion. The reasons are detailed in Sections II-D and II-E of (Rubiola & Vernotte, The Companion of Enrico's Chart for Phase Noise and Two-Sample Variances, 2023), and summarized below. It is defined (Donley & et al., 2022) as

$$L(f) = \frac{1}{2}S_{\varphi}(f) \qquad \text{(definition, IEEE std 1139)}$$
(28)

is generally given in dBc/Hz using $10 \log_{10} L(f)$. Some authors include $10 \log_{10}$ in the definition of L(f). According to (28), L(f) and $S_{\varphi}(f)$ are fully equivalent, and differs only in the unit of angle. The unit is the radian for $S_{\varphi}(f)$, and $\sqrt{2}$ rad = $\sqrt{2} \times 180/\pi = 81.03^{\circ}$ for L(f).

If we started from the scratch now, we would use $S_{\theta}(f)$, and L(f) would not exist. The reason is that $S_{\varphi}(f)$ is a proper SI quantity, L(f) is not. The problem originates in the early attempts to measure the phase noise with a spectrum analyzer. At that time, the phase noise was measured as

Power in 1 Hz bandwidth at a

$$L(f) = \frac{\text{frequency } f \text{ off the carrier frequency } f_0}{\text{Carrier power}} \quad (\text{obsolete definition}) \quad (29)$$

The true measurement of phase noise became common in the 1970s (Walls, Stein, Gray, & Glaze, 1976), when the double balanced mixer was available as an off-the-shelf component, suitable to a wide range of carrier frequency.

Notice that the IEEE Standard 1139 replaces (29) with (28). This was done since the first edition published in 1988 (Hellwig & et al., 1988), and the second (Vig & et al., 1999) and third edition (Ferre-Pikal & al., 2009) of this Standard, published in 1999 and 2009 respectively, confirm this choice.

The obsolete definition (29) is *conceptually incorrect* and *experimentally incorrect*. Let us discuss why.

First and foremost, the sideband power originates any combination of amplitude noise and phase noise. The obvious consequence is that (29) is a conceptually wrong representation of phase noise. For example, there is a discrepancy of 3 dB in the actual phase fluctuation of two signals having the same spectrum, one affected by equal amount of AM and PM noise, and the other having negligible AM noise.

Second, phase noise is measured with a phase detector. Consequently, "the SSB power in 1 Hz bandwidth" does not match the operation of the instrument.

Third, phase noise is pure angular modulation, thus the total power is the same at any modulation index. The random nature of noise does not change this fundamental property. The definition (29) can be used only for small modulation, where most of the power is in the carrier, and the power associated to the sidebands is comparatively small. Slow phenomena, like frequency random walk and drift, yields to large phase swing, exceeding 1 rad² in 1 Hz bandwidth. In the presence of such phenomena, (29) gives nonsensical results. By contrast, the correct definition (28) is perfectly suitable to large phase swing.

Finally, the obsolete definition (29) suffers from several pathologies. What happens with an odd signal affected by strong AM noise, and small PM noise? What happens if a spur occurs only in the upper (or lower) sideband, in the PM noise measurement range? In both cases the sideband-to-carrier ratio gives a nonsensical picture of the phase fluctuations.

It is a common belief that L(f), or equivalently $S_{\varphi}(f)$, is a valid measure only for small angles. In reality, there is no reason for such limitation, and L(f) is valid even if $\varphi(t)$ accumulates a large number of cycles. In other words, there is no reason to restrict L(f) to values below 0 dBc/Hz. In optics, measuring lasers one may encounter values of +40 dBc/Hz or 60 dBc/Hz, which are theoretically and experimentally correct. Of course, the phase detector has to work correctly in this regime.



Figure 10 – Heuristic derivation of L(f) in the case of additive white noise. The resistance R is not shown. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

2.6.3 Heuristic derivation of L(f) and $S_{\varphi}(f)$ in the case of additive noise

It is instructive to derive the quantities L(f) and $S_{\varphi}(f)$ for the simple case of white noise having PSD equal to N [W/Hz] added to a sinusoidal signal of power P_0 [W].

Before proceeding, we have to make clear that the case described does not match the definition of L(f) and $S_{\varphi}(f)$, but approximates it for small N. The catch is that a true random PM keeps the total power constant, while the added noise does not. With this caveat, our heuristic derivation gives useful results.

Let us start with L(f), with the help of Figure 10. In the standard notation for microwave circuits, a sinusoidal signal $v(t) = V_0 \cos(\omega_0 t + \varphi)$ can be represented as the complex vector

$$V = V_0 e^{j\varphi} \tag{30}$$

The power of such signal is $P_0 = |V_0|^2/2R$, thus the vector length is $V_0 = \sqrt{2RP_0}$. Similarly, a narrow noise slot of bandwidth *B* centered at $f_0 + f_m$ can be represented as a vector

$$V = V_n e^{j2\pi f_m t} \tag{31}$$

of random amplitude $V_n(t)$ rotating at the frequency f_m and average absolute value $V_n = \sqrt{2RNB}$. Adding carrier and sideband under the approximation of small noise-to-signal ratio, we get

$$\varphi(t) = \frac{\sqrt{2RNB}}{\sqrt{2RP_0}} \sin(2\pi ft) = \sqrt{NB/P_0} \sin(2\pi ft)$$
(32)

The RMS value of $\varphi(t)$ for B = 1 Hz is $\sqrt{NB/2P_0}$. Accordingly, it holds that

$$L(f) = \frac{1}{2} \frac{N}{P_0}$$
(33)

The above formula can be rewritten in terms of equivalent noise temperature T_e or of the noise factor F as

$$L(f) = \frac{1}{2} \frac{k(T_e + T_0)}{P_0} \quad \text{or} \quad L(f) = \frac{1}{2} \frac{FkT}{P_0}$$
(34)



Figure 11 – Heuristic derivation of $S_{\varphi}(f)$ and of $S_{\alpha}(f)$ in the case of additive white noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

A similar procedure can be used to derive $S_{\varphi}(f)$, with the help of Figure 11. While the carrier is the same as above, now we have two symmetric narrow sideband slots of bandwidth *B* centered at $f_0 - f_m$ and at $f_0 + f_m$

$$V = -V_{LSB}e^{-j2\pi f_m t} + V_{USB}e^{+j2\pi f_m t}$$
(35)

These sidebands have random amplitude $V_{LSB}(t)$ and $V_{USB}(t)$. The power associated to each sideband is *NB*, equally split into AM and PM noise. Thus, the absolute value of the vectors that contribute to PM noise is $V_{USB} = V_{LSB} = \sqrt{RNB}$. Combining carrier and sidebands under the approximation of small noise-to-signal ratio, we get

$$\varphi(t) = \frac{2\sqrt{RNB}}{\sqrt{2RP_0}} \sin(2\pi f_m t) \tag{36}$$

and

$$\varphi(t) = \sqrt{\frac{2NB}{P_0}} \sin(2\pi f_m t) \tag{37}$$

The RMS value of $\theta(t)$ for B = 1 Hz is $\sqrt{NB/P_0}$. Thus

$$S_{\varphi}(f) = \frac{N}{P_0} \tag{38}$$

Using the equivalent noise temperature T_e or the noise factor F, e above formula becomes

$$S_{\varphi}(f) = \frac{kT_e}{P_0} \qquad \text{or} \qquad S_{\varphi}(f) = \frac{FkT}{P_0}$$
(39)

The same development can be used to derive the amplitude noise. In this case, the vector representing the LSB has opposite sign with respect to phase noise, hence the sum of the two sideband vectors is parallel to the carrier. The result is

$$S_{\alpha}(f) = \frac{N}{P_0} \tag{40}$$

Using Figure 10 instead, we notice that the old definition of L(f) based on the sideband-to-carrier ratio gives both amplitude fluctuations and phase fluctuations, of equal amount. After our digression on the reason why the definition of L(f) has be changed to $L(f) = \frac{1}{2}S_{\varphi}(f)$, this unpleasant fact does not come as a surprise. Should the reader have to face both PM and AM noise, we recommend the use of $S_{\varphi}(f)$ and $S_{\alpha}(f)$, or $L_{\varphi}(f)$ and $L_{\alpha}(f)$ with obvious meaning. The notation M(f) is sometimes encountered as the AM-noise counterpart of L(f).

The smallest amount of white noise for a source characterized by a resistance at a temperature *T* is

$$S_{\varphi}(f) = S_{\alpha}(f) = \frac{kT}{P_0}$$
(41)

Scaling the carrier frequency with an ideal noise-free synthesizer which delivers an output frequency $f_o = (n/d)f_i$, the quantities associated to phase noise scale according to the rules listed in Table 4. We want to draw the attention of the reader to the simple fact that the synthesizer scales up or down the input phase noise and the input frequency noise in the same way. By contrast, the amplitude at the output of the ideal synthesizer is determined by the output stage, rather being sensitive to the input amplitude.

Quantity	Time domain	Spectral domain		
Phase	$\varphi_{out}(t) = \frac{n}{d}\varphi_{in}(t)$	$S_{\varphi out}(f) = \left(\frac{n}{d}\right)^2 S_{\varphi in}(f)$		
Frequency	$(\Delta f_{out})(t) = \left(\frac{n}{d}\right) (\Delta f_{in})(t)$	$S_{\Delta f out}(f) = \left(\frac{n}{d}\right)^2 S_{\Delta f in}(f)$		
Phase time	$\mathbf{x}_{out}(t) = \mathbf{x}_{in}(t)$	$S_{x out}(f) = S_{x in}(f)$		
Fractional Frequency	$y_{out}(t) = y_{in}(t)$	$S_{yout}(f) = S_{yin}(f)$		

Table 4 – Scaling rules for a noise-free synthesizer delivering $f_{out} = (n/d)f_{in}$,

Example 2 – Noise factor. A system has a noise factor of 1.8 dB, and receives at the input a sinusoid of power $P_0 = 100 \ \mu\text{W}$ (-10 dBm). The phase noise L(f) is

$$L(f) = \frac{1}{2} \frac{FkT}{P_0} = \frac{1}{2} \frac{10^{1.8/10} \times 1.386 \times 10^{-23} \times 290}{10^{-4}}$$
$$= 3 \times 10^{-17} \quad (-165.2 \text{ dBc/Hz})$$

If the reader can think in dB, the above formula becomes

 $L_{dB} = -3 + 1.8 - 174 - (-10) = -165.2 \, \text{dBc/Hz}$

2.6.4 Additive and parametric noise

Experience shows that in all oscillators the sideband noise increases greatly as we observe very close to the carrier. Slowly tuning the oscillator to a different frequency, we are faced to the evidence that the noise sideband are attached to the carrier, and follow the carrier frequency. This behavior is incompatible with the noise model we have used in the last Section to derive $S_{\theta}(f)$. How could the additive noise have a sharp peak centered exactly at the carrier frequency, "know" when the oscillator is re-tuned or drifts, and track the carrier by shifting the peak to the new frequency? No way. Our derivation related to Figure 10 and Figure 11 is correct, but it does not explain this behavior. The answer is that there are two types of phase noise, and of amplitude noise as well, called *additive noise* and *parametric noise*. They already appeared in the seminal article (Baghdady, Lincoln, & Nelin, 1965), at that time called *additive* and *multiplicative* noise. The basic mechanisms are represented in Figure 12. Understanding the difference between these types of noise is of paramount importance to master phase noise. The *additive noise* is exactly what we have explained in the last Section when we derived $S_{\omega}(f)$ by adding carrier and sideband vectors. White, or nearly white noise is present in a wide radiofrequency and microwave spectrum, and it adds to the carrier. By contrast, the *parametric noise* originates from a near-DC noise which modulates the carrier phase, frequency or amplitude, or any combination of. The

noise spectrum of the near-DC process is transposed to the sidebands around the carrier with the appropriate rules and symmetry. In this way, it is perfectly sound that a noise pattern with spurs in the microwave spectrum is centered at the carrier, and it appears unchanged around the new frequency after tuning or drift.

Notice that the power conservation inherent in the angular modulation of any index or phase swing applies only to parametric phase or frequency noise. By contrast, *adding* a white noise process results in higher total power.

Two-port components show a similar behavior, with the difference that they have no frequency drift because the output frequency is rigidly determined by the input frequency.

The digression about additive and parametric noise and their difference deserves further attention, because the term *added noise* is sometimes encountered in the specs of components and of test equipment. This term denotes the phase noise, and sometimes also the amplitude noise, that the component "adds" to the incoming carrier. The choice of the term "added noise" is unfortunate because it is too easily mistaken for "additive noise," while it refers to both the additive noise and the parametric noise that the component "adds" to the incoming carrier.



Figure 12 – Additive and parametric noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

2.7 The Time Channel (a new way to explain an old concept)

[I plan to add this section to a forthcoming version]



2.7.1 The polynomial law, or power law

A model which is found useful to describe the phase noise of oscillators and components is the *polynomial law*, often also called *power law* (Barnes & et al, 1971), (Halford, General Mechanical Model for f^{α} Spectral Density Random Noise with Special Reference to Flicker Noise 1/f, 1968)

$$S_{\varphi}(f) = \sum_{n \le -4}^{0} \mathsf{b}_n f^n \tag{42}$$

where the coefficients b_n are the parameters which describe the corresponding noise process. Equivalently

$$L(f) = \frac{1}{2} \sum_{n \le -4}^{0} b_n f^n$$
(43)

The polynomial law for phase noise is shown in Figure 13 and Table 5. The latter is extended to the phase time and fractional frequency fluctuations, detailed later. For the reader found in mathematics, the polynomial law refers to a Laurent polynomial, which is the extension of the Taylor series to negative-exponent powers of the running variable. In the oscillator phase noise, we find f^0 (constant), 1/f, $1/f^2$ etc., and each of such processes takes a specific name. Some may be hidden underneath the neighboring terms.



Figure 13 – The polynomial law for phase noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Table 5 – Main noise processes of the polynomial law for PM and FM noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Noise type	Phase noise, $S_{\varphi}(f)$		Phase time, $S_{x}(f)$		Fractional f	requency, $S_{y}(f)$
	Law	unit	law	unit	Law	unit
Blue PM	$b_1 f$	rad ² /Hz ²	$k_0 f$	s²/Hz² (s ⁴)	$h_3 f^3$	Hz ²
White PM	b ₀	rad²/Hz	k _o	s²/Hz (s³)	$h_2 f^2$	Hz
Flicker PM	$\frac{b_{-1}}{f}$	rad ²	$\frac{k_{-1}}{f}$	s ²	$h_1 f$	dimensionless
White FM	$\frac{b_{-2}}{f^2}$	rad ² Hz	$\frac{k_{-2}}{f^2}$	s²Hz (s)	h _o	Hz ⁻¹
Flicker FM	$\frac{b_{-3}}{f^3}$	rad ² Hz ²	$\frac{k_{-3}}{f^3}$	s ² Hz ² (dimensionless)	$\frac{h_{-1}}{f}$	Hz ⁻²
Frequency RW	$\frac{b_{-4}}{f^4}$	rad ² Hz ³	$\frac{k_{-4}}{f^4}$	s ² Hz ³ (s ⁻¹)	$\frac{h_{-2}}{f^2}$	Hz⁻³
Integrated flicker FM	$\frac{b_{-5}}{f^5}$	rad ² Hz ⁴	$\frac{k_{-5}}{f^5}$	s ² Hz ⁴ (s ⁻²)	$\frac{h_{-3}}{f^3}$	Hz ⁻⁴
Integrated RW frequency	$\frac{b_{-6}}{f^6}$	rad ² Hz⁵	$\frac{k_{-6}}{f^6}$	s²Hz⁵ (s⁻³)	$\frac{h_{-4}}{f^4}$	Hz ⁻⁵

We have studied the additive noise extensively in Section 2.6.3. The additive noise is chiefly a **white PM** process, with at most some smooth irregularities. The reason is that we observe a narrow region $f_0 \pm f_m$ of the microwave spectrum, centered at the carrier frequency f_0 .

Surprisingly for some, the white noise is not necessarily all of additive origin. An amount of white noise in a modulation process can be present.

The *flicker PM* noise is a parametric noise process originated by near-DC flicker, whose PSD is proportional to 1/f, which modulates the phase of the microwave signal.

The **white frequency noise**, or white FM noise, is a parametric process originated by white noise which modulates the frequency of an oscillator. This can be due for example to the thermal fluctuations of a resonator, to white noise in the oscillator loop, or to the white noise of the signal at the VCO input of the oscillator. The phase noise PSD associated to a white FM noise process is proportional to $1/f^2$. The reason is the following. As the phase is the integral of a frequency, its fluctuation $\varphi(t)$ can be
expressed as the integral of the carrier fluctuation $(\Delta f_0)(t)$ because $\varphi(t) = 2\pi \int (\Delta f_0)(t) dt$. We exploit the property of the Fourier transform, that the time-domain integral maps into a multiplication by $1/j2\pi f$, i.e.,

$$\int x(t) dt \leftrightarrow \frac{1}{j2\pi f} X(f)$$
(44)

There follows that

$$S_{\varphi}(f) = \frac{1}{f^2} S_{\Delta f_0}(f) \text{ and equivalently } L(f) = \frac{1}{2f^2} S_{\Delta f_0}(f)$$
(45)

and equivalently

$$S_{\varphi}(f) = \frac{1}{f^2} S_{\Delta f_0}(f) \text{ and equivalently } L(f) = \frac{1}{2f^2} S_{\Delta f_0}(f)$$
(46)

Of course, in the presence of white frequency noise, $S_{\Delta f_0}(f)$ is a constant vs frequency, while $S_{\varphi}(f)$ and L(f) are proportional to $1/f^2$.

The *flicker frequency noise*, or *flicker FM noise*, is another type of noise very often found in oscillators and characterized by a phase noise PSD proportional to $1/f^3$. It originates from a flicker noise process which modulates the frequency of the oscillator. The same reasoning seen for the white PM noise yields the conclusion that $S_{\varphi}(f)$ and L(f) are proportional to $1/f^3$.

Textbooks of statistics teach us that random walk results from integrating a white noise process. But we have seen that the integral operator maps into the multiplication by a factor of $1/f^2$ in the PSD. As a consequence, in the presence of a **random walk of** *frequency*, or *frequency RW* for short, $S_{\varphi}(f)$ and L(f) are proportional to $1/f^4$. There are several physical reasons for the presence of frequency RW in oscillators, mainly related to the resonator's natural frequency changing with time or affected by environmental parameters.

There is no a-priori reason for the sum (42) to start from n = -4, and further negative terms can be added when needed. Oppositely, a 'true' b₁f term is not allowed because it results in too large power, if not infinite power, after integrating $S_{\varphi}(f)$. A '+1' slope is often found in optical fiber links and other applications involving a control loop. However, this behavior can only be local, that is, the left-hand side of a bump in the spectrum.

The *frequency drift* consists in a linear change of the oscillator frequency with time. In the presence of a drift D of the fractional frequency, we can replace the oscillator frequency f_0 as

$$f_0 \to f_{00}(1 + \mathrm{D}t)$$
 (47)

where f_{00} is the oscillator frequency at an appropriate origin of time t = 0. The reason for the frequency drift in oscillators are mainly related to the aging of the resonator.

The PSD is not a preferred tool to describe the drift because the curve is too steep for clear visual interpretation, and because the phase $\varphi(t)$ grows rapidly and gets too large for the dynamic range of most instruments. Time-domain techniques are more suitable. However, it is instructive to calculate $S_{\varphi}(f)$ in the presence of a frequency drift. We start from a frequency perturbation described by a Dirac $\delta(t)$ distribution. The Laplace transform of $\delta(t)$ is equal to 1. The integral of the $\delta(t)$ distribution is the Heaviside distribution u(t), and its Laplace transform is 1/s. Further integrating, we get a linear ramp tu(t) starting at t = 0, that is, a drift. Its Laplace transform is $1/s^2$. But the phase is the integral of the frequency, thus the Laplace transform is $1/s^3$. We can derive the PSD from the Laplace transform, first by converting the Laplace transform into the Fourier transform (replace $s \rightarrow j2\pi f$), and then by taking the absolute square value. In this way, we find $S_{\varphi} \propto 1/f^6$.

We recall that the phase time fluctuation is equivalent to the random phase after converting the unit from radians to seconds, i.e., $x(t) = \varphi(t)/2\pi f_0$. As an obvious consequence, the polynomial law applies to $S_x(f)$

$$S_{x}(f) = \sum_{n \le -4}^{0} k_{n} f^{n}$$
 with $k_{n} = \frac{1}{4\pi^{2} f_{0}^{2}} b_{n}$ (48)

The classification of noise from white PM to frequency RW applies to oscillators and to frequency sources in general. By contrast, only white and flicker phase noise are possible in two-port components, otherwise the input-output delay would diverge. Steeper terms of the polynomial law, for example $1/f^4$ or $1/f^5$, are often seen on phase noise plots. However, deeper analysis shows that this behavior is the right-hand side of a large bump in the spectrum due to environmental parameters or to other phenomena, and the full bump does not show up because its left-hand side occurs at too low frequencies, outside the measurement span.

The mean square phase and delay accumulated by a device, in the frequency span from f_1 to f_2 are given by

$$\langle \varphi^2 \rangle = \int_{f_1}^{f_2} S_{\varphi}(f) \, df \tag{49}$$

and

$$\langle \mathsf{x}^2 \rangle = \int_{f_1}^{f_2} S_\mathsf{x}(f) \, df \tag{50}$$

It is clear that, for $1/f^2$, $1/f^3$ and slower processes, the quantities $\varphi_{rms} = \sqrt{\langle \varphi^2 \rangle}$ and $x_{rms} = \sqrt{\langle x^2 \rangle}$, can get quite large as the lower boundary f_1 is small. Conversely, and surprisingly, flicker PM gives rise to small phase and delay, even if the PSD is integrated over a rather extreme frequency span. We will see this in the following example.

February 6, 2025

Example 3 – Flicker Noise. A two-port device used at the carrier frequency $f_0 = 100$ MHz, shows a flicker of -80 dBc/Hz extrapolated to 1 Hz (a rather poor value). Let us estimate the phase φ_{rms} and the delay x_{rms} .

For the purpose of convincing the reader that the flicker PM noise does not end in infinitely diverging phase and delay, we make an arbitrary and rather extreme choice of the frequency span. First, we agree that the lifetime of the device will not exceed 10^9 s (30 years), or we do not care about longer time. Accordingly, we take $f_1 = 10^{-9}$ Hz. Second, we agree that the bandwidth of the phase fluctuations is at most equal to the carrier frequency, thus we take $f_2 = 10^8$ Hz. From the statement of the problem, we find $b_{-1} = 2 \times 10^{-80/10}$, and $k_{-1} = \frac{b_{-1}}{4\pi^2 f_0^2} = 5 \times 10^{-26}$. Thus,

$$\langle \varphi^2 \rangle = \int_{f_1}^{f_2} \frac{\mathsf{b}_{-1}}{f} \, df = \mathsf{b}_{-1} \ln \frac{f_2}{f_1} = 2 \times 10^{-8} \, \ln \frac{10^8}{10^{-9}} = 7.8 \times 10^{-7} \, \mathrm{rad}^2$$

hence $\varphi_{rms} = 885 \,\mu rad$. Similarly

$$\langle x^2 \rangle = \int_{f_1}^{f_2} \frac{k_{-1}}{f} df = k_{-1} \ln \frac{f_2}{f_1} = 5 \times 10^{-26} \ln \frac{10^8}{10^{-9}} = 2 \times 10^{-24} s^2$$

thus $x_{rms} = 1.4$ ps.

The above results only mean that the internal delay of a device does not diverge in finite time. We encourage the reader to calculate φ_{rms} and x_{rms} in the most extreme conceivable case, where f_1 is the reciprocal of the age of the universe, and f_2 is the reciprocal of the Planck time. The result is surprisingly small. This does not change the fact, that flicker is a major concern for oscillators and synthesizers.

2.7.2 Frequency Stability PSD

It is well known that the angular modulation can be expressed as a phase modulation or as a frequency modulation, and that the two forms are equivalent. The same holds random phase and frequency modulation. The phase fluctuation associated to a frequency fluctuation is

$$(\Delta f_0)(t) = \frac{1}{2\pi} \frac{d\varphi(t)}{dt}$$
(51)

Using the property of the Fourier transform, that the time-domain derivative operator maps into a multiplication by $j2\pi f$,

$$dx(t)/dt \to j2\pi f X(f) \tag{52}$$

the PSD of $(\Delta f_0)(t)$ is given by

$$S_{\Delta f_0}(f) = f^2 S_{\varphi}(f) \tag{53}$$

Alternatively, we can use the fractional frequency fluctuation $y(t) = \frac{1}{f_0} (\Delta f_0)(t)$, and its PSD

$$S_{\gamma}(f) = \frac{f^2}{f_0^2} S_{\varphi}(f)$$
(54)

The polynomial laws, rewritten for $S_v(f)$, is (Figure 14)

$$S_{y}(f) = \sum_{n \le -2}^{2} h_{n} f^{n}$$
 with $h_{n} = \frac{1}{f_{0}^{2}} b_{n-2}$

The quantities $S_{\Delta f_0}(f)$ and $S_{\gamma}(f)$ are seldom used in in radio engineering. However, $S_{\gamma}(f)$ is an important step to assess the relationship between phase noise and the time-domain variances AVAR, MVAR, PVAR etc.



Figure 14 – Polynomial law for frequency noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

2.7.3 The low-Fourier-frequency part of the phase-noise PSD

We have seen that the polynomial law fits the phase noise PSD in a large number of cases, that the polynomial law extends to low frequencies, and that the higher negative-power terms, found on the left-hand side of the plot, reveal the slow frequency fluctuations.

A question arises, can $S_{\varphi}(f)$ be used to measure the long-term behavior of oscillators? Of course, $S_{\varphi}(f)$ is a mathematical tool based on the measurement of the physical quantity $\varphi(t)$, thus any valid mathematical manipulation yields correct results. However, in this case $S_{\varphi}(f)$ is not to a good tool, for the following reasons.

First, it is difficult to fit the experimental plot and to extract precisely the coefficients because the terms $b_i f^i$ associated to drift and other slow phenomena have steep negative slope. The reader can try with the frequency random walk b_{-4}/f^4 or with the frequency drift b_{-6}/f^6 .

Second, the measurement of a steep slope requires that the noise test set has a wide dynamic range. For example, the term b_{-4}/f^4 rises by 40 dB over a factor of 10 in frequency.

Third, the measurement time *T* needed to get $S_{\varphi}(f)$ with a resolution δf is governed by the time-frequency indetermination theorem, which states that $T \ \delta f \ge 1$, where T and δf are the RMS values. The equality holds for Gaussian distributions. The theory underneath is found in many textbooks on the Fourier integral, among which we prefer Papoulis, 1962. Actual instruments work with the acquisition time T_a , which is a finite and well identified quantity. The acquisition time is associated to the window (taper) function used in the FFT analysis, the most popular of which is the Hanning window. In practice, the resolution is governed by $T_a \ \delta f \ge C_1$, with C_1 is in most cases of at least 2–3. The lowest frequency of the FFT is $f_1 = 1/T$. However, the first points may not be plotted, or discarded, because of the poor resolution $\delta f/f$ and because of artifacts related to the window function. The consequence is that the minimum plotted frequency is ruled by $T_a f_{\min} \ge C$, with *C* of the order of 5-10. For example, the acquisition of a single spectrum down to $f_{min} = 10 \text{ mHz}$ gives $T_a = 600 \text{ s}$ (C = 6). If we decide to average on 12 spectra in order to get a comfortable confidence level, the measurement takes 2 hours. By contrast, the Allan variance and the other wavelet variances, described later, are way more efficient at estimating the long-term behavior of oscillators with a reasonably short data record. Interestingly, all these variances are easily calculated from a time series of phase data, the same used to calculate the PM noise PSD.

2.7.4 The RF spectrum of the oscillator signal

The oscillator signal, observed with a RF spectrum analyzer, looks like a rather narrow bell-shaped pattern, wobbling, wandering, and drifting. How does it relate to the PM noise PSD, which is seemingly unbounded at low Fourier frequencies?

However naïve the question may seem, the problem underneath is surprisingly complex. All difficulties start from the fact that the variance σ_{φ}^2 does not exist for flicker (1/f) and for steeper processes $(1/f^2, 1/f^3, \text{etc.})$, and from the fact that the variance of the truncated $S_{\varphi}(f)$, high-passed at a frequency f_H , diverges as $f_H \to 0$. Thus, the use of a high-pass filter does not help.

The problem of the line shape is addressed in (Godone, Micalizio, & Levi, 2008), (Brochard, Sudmeyer, & Schlit, 2017), and (Elliott, Roy, & Smith, 1982). These References rely on a rather difficult the mathematical framework, and provide a separate solution for each PM noise process. Reference (Elliott, Roy, & Smith, 1982) is particularly useful for optics, where the line broadening phenomena are amplified by beating two laser beams down to RF. However, the language used in these References may be unusual or difficult to microwave engineer. That said, we try to grab the main physical facts.

First of all, the angular modulation does not affect the total electrical power, and the same holds with phase noise. Consequently, the generalized PM power, given in rad², is allowed to be quite large, and even to diverge in the long run, without violating the energy conservation principle. The pattern seen on the display of the spectrum analyzer reveals the average electrical power. The averaging time is set by 1/RBW or by 1/VBW, which is longer (RBW is the resolution bandwidth, and VBW is the video filter bandwidth).

Second, the electrical power is spread in sidebands, governed by the infinite series of Bessel functions $J_n(\beta)$, where β is the modulation index. The sidebands are separate entities at small β , but they collapse in a single line at large β . Something happens in between.

Third, the spectrum analyzer has a limited RBW. In the traditional scanning spectrum analyzer, the transfer function associated to the RBW is a Lorentzian lineshape, determined by the IF filter. The displayed shape results from the convolution of the input spectrum and the IF frequency response. The behavior of modern spectrum analyzers, based on the FFT of a wide-band IF signal is rather similar, however more difficult to understand in rigorous mathematical terms. For the sake of clarity, let us say that the RBW is associated to a Lorentzian filter.

Now we apply the above concepts to the oscillator signal.

When the oscillator delivers a pure and stable signal, the RF spectrum is a clean and sharp line, narrower than the instrument RBW. The displayed spectrum is a Lorentzian pattern determined by IF filter. Such pattern wanders and drift slowly, following the oscillator frequency. The spectrum of white FM noise is a Lorentzian. Thus, if the dominant oscillator noise is white FM, wider than the RBW, the analyzer displays a Lorentzian determined by the oscillator FM noise. When flicker FM or FM random walk is dominant, the numerous random sidebands tend to cluster in Gaussian shape. However, if the width of the Gaussian is still comparable to the Lorentzian RBW, there results a Voigt distribution (Posener, 1959). Until now, we have implicitly assumed that the frequency reference inside the spectrum analyzer is stable and free from noise. Actually, the displayed pattern results from the PM noise of both the oscillator under test and the analyzer's internal oscillator, and the two contributions are indistinguishable.

3 The Allan variance and other two-sample variances

We study the frequency fluctuations of a real oscillator around the nominal value, as a function of the measurement time. Such fluctuations contain true random terms, plus other phenomena like drift and environmental effects. However the environmental effects are of systematic origin, they can only be described statistically for the part of the environment that escape from quantitative understanding. This goes under the term *influence quantities* in the formal language of metrology.

Before tackling the analysis of frequency fluctuations, we have to study the basic operation of the *frequency counters*, sometimes called *frequency-to-digital converters*. Other options are possible, chiefly time analyzers and phase meters. These instruments are highly specialized, and seldom found in general laboratories.

Modern frequency counters are rather complex, and their statistical properties may be difficult to understand. Reference (Kalisz, 2004) reviews the high resolution counting techniques, and References (Rubiola E., On the measurement of frequency and of its sample variance with high-resolution counters, 2005), (Rubiola E., The Ω Counter, a Frequency Counter Based on the Linear Regression, 2016) provide insight in the statistical processing techniques.

3.1 Frequency counters

3.1.1 The Π frequency counter

The II counter is an instrument that measures frequency or period by counting a number of events occurring during the gate time τ . The classical frequency counter, shown on Figure 15, is the simplest example. The instrument counts the integer number N_x of cycles of the input signal in the gate time τ generated by the reference clock. The classical frequency counter is seldom used because of the poor resolution at low input frequency. If $\tau = 1$ s, the number N_x is equal to the frequency expressed in Hz, thus the measurement of the 50 Hz frequency of the European power grid suffers from 2% quantization uncertainty. For this reason, the classical frequency counter is generally replaced with classical reciprocal counter. The role of the clock signal and of the input signal are interchanged, and the instrument measures the average period by counting the clock cycles in a suitable multiple of the input period that approximates τ . With a clock frequency of 10 MHz, the quantization s 10^{-7} in 1 s measurement time, regardless of the input frequency. More sophisticated instruments can measure a fraction of a clock cycle by interpolating between edges of the clock signal. Combining reciprocal counting and clock interpolation boosts the resolution up to 10 digits and

more, but the noise mechanism of the Π counter remains. In this type of instruments, the noise is determined by the time fluctuations found at the *start* and at the *stop* event that define the gate time τ , and the fluctuations occurring between start and stop do not contribute to the result. In practice, the time fluctuations originate from the quantization noise, from the interpolator, and from the noise of the input trigger. With sophisticated interpolators, the remaining noise from the input trigger is the dominant noise source. The contribution of the frequency reference is not accounted here, and it must be considered separately.



Figure 15 – Classical frequency counter. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Since the input frequency fluctuates, it is useful to replace f_0 with $f_0 + (\Delta f_0)(t)$. The average frequency is given by

$$\overline{f_0} = \frac{1}{\tau} \int_0^{\tau} f_0 + (\Delta f_0)(t) \, dt$$
(55)

Introducing the fractional frequency fluctuation $y(t) = (\Delta f_0)(t)/f_0$, we can replace the average (55) with the weighted average

$$\overline{\mathbf{y}} = \int_0^\infty \mathbf{y}(t) \, w_{\Pi}(t) \, dt \tag{56}$$

where

$$w_{\Pi} = \begin{cases} 1/\tau & 0 < t < \tau \\ 0 & \text{elsewhere} \end{cases}$$
(57)

The name Π counter was chosen because the Greek letter Π recalls shape of $w_{\Pi}(t)$, which is a rectangular pulse.

The value of our approach is that (56) also describes the other counters, just by replacing the weight function $w_{\Pi}(t)$ with a weight function of our choice. Of course, it is necessary that

$$\int_{-\infty}^{\infty} w(t)dt = 1$$
(58)

for w(t) to be a valid weight function.

Denoting with σ_x^2 the variance (time fluctuation) associated to each trigger event, the variance of the fractional frequency fluctuation is

$$\sigma_{\rm y}^2 = \frac{2\,\sigma_{\rm x}^2}{\tau^2} \tag{59}$$

The factor of 2 comes from the fact that σ_x^2 counts twice, start and stop being statistically independent events. It is to be made clear that (59) is the classical variance, as opposite to the two-sample variances which we will introduce later.



Figure 16 – Theory of operation of the Π (classical) frequency counter. Reprinted from E. Rubiola, "On the measurement of frequency and of its sample variance with high-resolution counters," *Rev. Sci. Instrum.*, vol. 76, no. 5, pp. 054703 1-6, May 2005 (Rubiola E. , On the measurement of frequency and of its sample variance with high-resolution counters, 2005), with the permission of AIP Publishing.

Example 4 – Frequency counter. Consider the measurement of a 10 MHz oscillator with a gate time of 100 ms. The instrument has a quantization $\delta t = 5$ ns, and the trigger noise is negligible in this case. The time deviation associated to the quantization noise is $\sigma_x = \delta t / \sqrt{12} = 1.44$ ns. The factor $1/\sqrt{12}$ accounts for the quantization uniformly distributed between $\pm \delta t / 2$. Using (59) with $\tau = 100$ ms, we find $\sigma_y = \sqrt{2} \times (2 \times 10^{-9} \text{ s})/(10^{-1} \text{ s}) = 2.5 \times 10^{-8}$. The same reasoning applies to the noise in the input trigger, just omitting the $1/\sqrt{12}$ factor, which is characteristic of the quantization.

3.1.2 The Λ frequency counter

Looking at Figure 16, there is a lot of information contained in the time fluctuations between the edges of the gate function, which can be exploited to reduce the background noise of the instrument. This requires a different type of hardware, capable of timing more transitions, ideally all the transitions between the two edges. With FPGA technology, such hardware is affordable even in moderate-cost instruments.

The first method to reduce the instrument noise is the Λ counter, whose principle is shown in Figure 17. The measurement process consists of averaging *n* highly overlapped measures of duration τ , spaced by $\tau_0 = \tau/n$. These *n* measures are the same as in the Π counter. The benefit of the Λ counter derives from the fact that most of the random fluctuations come from the input trigger, which is a wideband stage. Therefore, the white PM noise is dominant, and the samples of such noise, taken at t_0 , t_1 , t_2 etc., are statistically independent. Thus, the process of averaging on *n* independent values reduces $\sigma_{\rm Y}^2$ by a factor 1/n. This noise reduction comes at the cost of a longer measurement time, 2τ instead of τ . Ideally, *n* is maximized by taking a measure at each cycle of the input signal, that is, with $\tau_0 = 1/f_0$. However, τ_0 is limited by the data transfer inside the instrument and by the processing rate. In commercial equipment, we find values between a millisecond and a fraction of a microsecond.



Figure 17 – Theory of operation of the Λ frequency counter. Reprinted from E. Rubiola, "On the measurement of frequency and of its sample variance with high-resolution counters," *Rev. Sci. Instrum.*, vol. 76, no. 5, pp. 054703 1-6, May 2005, with the permission of AIP Publishing.

In the last example, the standard deviation is reduced from 2.5×10^{-8} to 2.5×10^{-9} if the instrument has a minimum τ_0 of 1 ms (n = 100), and to 2.5×10^{-10} if the minimum τ_0 is of 10 µs ($n = 10^4$).

The following points deserve attention before doing statistics with frequency counters.

- Some commercial instruments use the *Λ* averaging without saying. Of course, the statistical properties are totally different from the *Π* counter.
- The "gate time" shown on the front panel may be ambiguous and depend on the specific instrument. It can be identified either with *τ* or with the total measurement time, which is 2*τ* in our notation.
- The gate events can be contiguous, or exactly overlapped by one side of the triangle. The decimation rule is a good reason to choose the overlapped option. Suppose we have a stream of data f_1, f_2, f_3, \dots overlapped by one side. The rule

$$f_1' = \frac{1}{4}f_1 + \frac{1}{2}f_2 + \frac{1}{4}f_3$$
$$f_2' = \frac{1}{4}f_3 + \frac{1}{2}f_4 + \frac{1}{4}f_5$$
etc.

gives a new data stream with the same type, and overlap. The new n is twice the older, and τ is twice longer. By contrast, if the triangles are just contiguous, there is no decimation rule preserving the statistical properties associated to the measurement.

• Understanding the internal operation and the meaning of the gate time requires attention and some tests. It is useful to observe the counter readout on the presence of a phase-modulated RF signal. We can get information about the weight function by setting the modulation frequency for the readout to be constant, and for the readout to have the maximum fluctuation.

A more formal description of the Λ counter is the following. We start from a time series of *n* highly-overlapped measures f_i , of the same type of the Π counter. Each measure is shifted by $\tau_0 \ll \tau$ with respect to the previous one

$$f_i = \frac{1}{\tau} \int_{i\tau_0}^{i\tau_0 + \tau} f_0 + (\Delta f)(t) dt \qquad i = 0 \dots n - 1$$
(60)

Averaging on these n measures gives

$$\overline{f} = \frac{1}{n} \sum_{i=0}^{n-1} f_i \tag{61}$$

Replacing $f_0 + (\Delta f)(t)$ with the fractional frequency $\gamma(t)$, we write the weighted average

$$\overline{\mathbf{y}} = \int_0^\tau \mathbf{y}(t) \, w_A(t) \, dt \tag{62}$$

Taking the limit for large *n*, thus for $\tau_0/\tau \rightarrow 0$, the staircase function shown on Figure 17 becomes a triangular function

$$w_{\Lambda}(t) = \begin{cases} t/\tau^{2} & 0 < t < \tau \\ 2/\tau - t/\tau^{2} & \tau < t < 2\tau \\ 0 & \text{elsewhere} \end{cases}$$
(63)

It follows from our reasoning that the variance of the fractional frequency fluctuation is

$$\sigma_{\rm y}^2 = \frac{1}{n} \frac{2 \,\sigma_{\rm x}^2}{\tau^2} = \tau_0 \frac{2 \,\sigma_{\rm x}^2}{\tau^3} \tag{64}$$

This is the same as (59), but for a factor 1/n that accounts for the average on n measures.

As the reader may expect after our digression, the name Λ counter is due to the graphical analogy of the Greek letter Λ with the triangular step function shown on Figure 17.

3.1.3 The Ω frequency counter

The total phase $\phi(t)$ of the clock signal the ever-growing ramp described as

$$\phi(t) = 2\pi \int f_0 + (\Delta f)(t) \, dt = 2\pi f_0 t + \varphi(t) \tag{65}$$

whose fluctuation is $\varphi(t) = 2\pi \int (\Delta f)(t) dt$ as the frequency fluctuates. Notice the difference between the total phase $\phi(t)$ and its fluctuation $\varphi(t)$. The instantaneous frequency is

$$f(t) = \frac{1}{2\pi} \frac{d\phi(t)}{dt}$$
(66)

The Ω counter measures the average frequency \overline{f} using the *linear regression* on the time series ϕ_0 , ϕ_1 , ϕ_2 ... sampled at regular intervals spaced by τ_0 , as shown in Figure 18. Denoting with t_i the time when the sample ϕ_i is taken, and with μ_{ϕ} and μ_t the averages of ϕ and t in the measurement time τ , we find

$$\overline{f} = \frac{1}{2\pi} \frac{\sum_{i=0}^{n-1} (\phi_i - \mu_\phi) (t_i - \mu_t)}{\sum_{i=0}^{n-1} (t_i - \mu_t)^2}$$
(67)

Replacing $f_0 + (\Delta f)(t)$ with the fractional frequency y(t), we write the weighted average

$$\overline{\mathbf{y}} = \int_0^\tau \mathbf{y}(t) \, w_\Omega(t) \, dt \tag{68}$$

The mathematics underneath is a little more complex than in the case of the Π and Λ counters, and the details are found in Reference [23]. The weight function has the "cap" parabolic shape

$$w_{\Omega}(t) = \begin{cases} \frac{6t}{\tau^2} \left[1 - \frac{t}{\tau} \right], & 0 < t < \tau \\ 0, & \text{elsewhere} \end{cases}$$
(69)

The name Ω has been chosen by the graphical analogy of $w_{\Omega}(t)$ with the Greek letter Ω , in the continuity of the Π and Λ counters. The theory indicates that the Ω counter exhibits the highest rejection of white PM noise, which is exactly what we want to reduce the effect of the quantization and of the trigger noise.

The variance of the fractional frequency is

$$\sigma_{\rm y}^2 = \frac{1}{n} \frac{12 \, \sigma_{\rm x}^2}{\tau^2} = \tau_0 \frac{12 \, \sigma_{\rm x}^2}{\tau^3} \tag{70}$$



Figure 18 – Operation of the Ω frequency counter, which measures the average frequency using the linear regression on phase data. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

3.1.4 Comparison of the frequency counters

The performance of the frequency counters is described by the variance σ_y^2 , that is (59), (64), or (70). These equations give the background noise of the instrument in the presence of white PM noise (input trigger and quantization). The frequency reference is still not included.

We will learn in the next Section that the measurement time τ is a more complex issue, and why the two-sample variances are necessary. That said, the classical variance is perfectly suitable to describe the white phase noise. Thus, we can draw some conclusions about the time required to test an oscillator, with obvious implications on industrial production. It goes without saying that the instrument background noise must be smaller than the noise of the oscillator under test. This sets a minimum measurement time needed to average out the instrument background.

First, the instrument noise is averaged out proportionally to $1/\tau^2$ in the Π counter, and proportionally to $1/\tau^3$ in the Λ counter and in the Ω counter. Thus, the Π counter is clearly slower in all cases. The two other counters are similar in speed, being governed by the $1/\tau^3$ law.

Second, a comparison between (64) and (70) gives the false impression that the Ω counter suffers from more noise than the Λ counter, by a factor of six. The catch is that

the Λ counter takes a time 2τ to deliver a value, while the Ω counter takes τ . For fair comparison, we have to replace $\tau \rightarrow \tau/2$ in (64), so that the measurement time is the same. This gives $\sigma_y^2 = 16 \tau_0 \sigma_x^2/\tau^3$ for the Λ counter. The conclusion is that the Ω counter is superior to the Λ counter, to the extent that it exhibits lower noise by a factor of $\frac{3}{4}$ (1.25 dB) for the same measurement time. However modest this result may seem, it can be a great choice for large volume tests.

Most old frequency counters work in \varPi mode.

Major manufacturers sell Λ counters, often without saying. Eventually, Reference (Rubiola E., On the measurement of frequency and of its sample variance with highresolution counters, 2005) originated from the attempt to sort out inconsistent results when we measured some oscillators using an Agilent (now Keysight) counter. We could interpret correctly these results only after setting the formal framework that describes the Π and Λ counters, and after reverse engineering the counter we had. Lange Electronics manufactures a counter specifically intended for the measurement of the two-sample variances, which can be programmed to operate in Π and Λ mode. Pendulum (later acquired by Philips, to the best of our knowledge) was arguably the first Company to produce a counter using the linear regression (Johansson, 2005). The statistical properties of the Ω counters were tackled independently in (Rubiola E. , The Ω Counter, a Frequency Counter Based on the Linear Regression, 2016), (Vernotte, Lenczner, Bourgeois, & Rubiola, 2016), and (Benkler, Lisdat, & Sterr, 2015). Recently, Carmel Instruments (formerly Brilliant Instruments) manufactures some time analyzers implemented as PXI Express modules, advertised as capable to operate in Π , Λ and Ω modes.

3.2 The two-sample variances AVAR, MVAR and PVAR

Until now, we have implicitly used the experimental variance found in many textbooks, defined as

$$\sigma_{\rm Y}^2 = \frac{1}{N-1} \sum_{k=1}^{N} ({\rm y}_k - \mu)^2 \tag{71}$$

where *N* is the number of samples f_k , and $\mu = \frac{1}{N} \sum_{k=1}^{N} \gamma_i$ is the average. The problem is that (71) is useful only for certain types of fluctuations, and of course it works perfectly with the white PM noise we were really concerned about. In such cases (71) converges to the "ideal" value called *mathematical expectation*, and the confidence improves progressively as *N* increases.

The problem is that real oscillators suffer from frequency flicker, random walk, and drift, and other processes. These processes make (71) depend on both N and τ , and the mathematical expectation does not exist. That something goes wrong with N, is clearly

seen with linear drift (constant aging). Try yourself to calculate the deviation σ_y in the case of a 100 MHz oscillator drifting by +1 mHz per second, measured with s gate time of 1 s. After taking $N = 10, 10^2, 10^3$, etc., we agree that the experimental variance is not a good choice.

The problem with τ is a little bit subtler. Let us take the distance from two points on a sheet of sandpaper as a simple example. A human-sized ruler may indicate a distance quite close to 6 inches, which we perceive as "correct" at our scale, and that's it. By contrast, a tiny insect experiences a much longer path of consisting of high obstacles which it can surmount only thanks to its incredible agility. What happens? The answer is that the distance depends on the size of the ruler used to measure it. Not surprisingly, something similar occurs with the oscillator fluctuations y(t) and the measurement time τ . Still on the example of the 100 MHz oscillator drifting at +1 mHz/s, try yourself to calculate σ_y^2 using (71), after switching the gate time to 0.1 s, 1 s, and 10 s. The difference between results should convince the reader that (71) is to be replaced with a more appropriate tool. Here the family of two-sample (Allan and Allan-like) variances gets on the stage. The key points are

- The measurement time τ is made explicit
- The variance is calculated using a simple and perfectly defined pattern.

3.2.1 The Allan variance (AVAR)

Let us first introduce the average fractional frequency $\overline{\mathbf{y}}_k(\tau)$, defined as

$$\overline{\mathbf{y}}_{k}(\tau) = \int_{(k-1)\tau}^{k\tau} \mathbf{y}(t) \, dt \tag{72}$$

This is the fractional frequency y we are familiar with, measured with a Π counter where the gate time has duration τ starting at $t = (k - 1)\tau$ and ending at $t = k\tau$. So, $\overline{y}_1(\tau)$ is averaged between 0 and τ , $\overline{y}_2(\tau)$ is averaged between τ and 2τ , etc. It is important that the measurements are exactly contiguous, with no dead time in between, otherwise the final result will differ from the Allan variance. This requirement is not a real limitation because the speed of nowadays digital hardware is generally sufficient to avoid such dead time. Understanding what happens in the presence of a dead time is a special topic for experts, out of our scope..

Having defined τ , we solve the dependence on *N* by setting N = 2 in (71). Notice that N = 2 is the smallest value that gives a valid variance. Rewriting (71) for N = 2

$$\sigma_{y}^{2}(\tau) = \left[\bar{y}_{2}(\tau) - \mu\right]^{2} + \left[\bar{y}_{1}(\tau) - \mu\right]^{2}$$
(73)

and expanding using $\mu = (\bar{y}_2 + \bar{y}_1)/2$, we find the *two-sample variance*

$$\sigma_{\gamma}^{2}(\tau) = \frac{1}{2} \left[\bar{\gamma}_{2}(\tau) - \bar{\gamma}_{1}(\tau) \right]^{2}$$
(74)

This variance is also called *Allan variance* (Barnes & et al, 1971) or AVAR. The obvious notation ADEV is often used for $\sigma_{\rm v}(\tau)$. The above (74) should be intended in the sense of the mathematical expectation, that is, the ideal average on infinite values of $\sigma_{\rm v}^2(\tau)$. In practice, we use a finite time series of *M* contiguous values of $\overline{y}_k(\tau)$, $k = 1 \dots M$. Consequently, the average on M - 1 values of $[\overline{y}_{k+1}(\tau) - \overline{y}_k(\tau)]^2$ gives

$$\sigma_{\mathbf{y}}^{2}(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M} \left[\bar{\mathbf{y}}_{k+1}(\tau) - \bar{\mathbf{y}}_{k}(\tau) \right]^{2}$$
(75)

Equation (75) is the usual formula for the Allan variance. Figure 19 shows the most common noise types, as they appear on the AVAR plot.



Figure 19 – Allan variance, plotted for the most common noise terms of the polynomial law. The constant *c* results from a cutoff frequency f_H , necessary to limit the bandwidth of white phase noise. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

It is a common practice to use a counter with a gate time τ_0 , and to combine 2^m contiguous measurements to get $\tau = 2^m \tau_0$. In this case, the terms $[\bar{y}_{k+1}(\tau) - \bar{y}_k(\tau)]$ used in (75) can be highly overlapped, and spaced by τ_0 . Strictly speaking, this is the *overlapped Allan variance*. For the same noise process, overlapped AVAR and non-overlapped AVAR converge to the same value. The overlapped AVAR is generally used because it exhibits superior confidence level for finite *M*.

Some confusion may arise with the notation $\sigma_y^2(\tau)$ because the number of samples is implied. A more rigorous notation would be $\sigma_y^2(N,\tau)$ for the *N*-sample variance, and consequently $\sigma_y^2(2,\tau)$ for AVAR, but this notation is seldom used. Notice that the Allan variance is always represented as a function of τ , while it is rather uncommon to find τ in the classical variance (our Section 3.1 is a necessary exception). Another confusion is around the corner, between AVAR and other similar variances, some of which are described in the next Section.

3.2.2 The modified Allan variance (MVAR)

The Modified Allan Variance (Snyder, 1981)- (Allan & Barnes, 1981), denoted with mod $\sigma_y^2(\tau)$ or MVAR, is similar to the Allan variance and uses the same formula (75), but the weight function used to calculate the average \overline{y} is replaced with $w_{\Lambda}(t)$. In practice, we can use a Λ counter instead of the Π counter, or implement the Λ average in some other ways. The latter option is let to the experts. Using the Λ counter, we should make sure that averaging is done correctly:

- The time τ is interpreted as in Figure 17. Accordingly, the measurement of a single $\overline{y}_k(\tau)$ takes a time 2τ
- Two "contiguous" measures overlap by τ . Thus, the (k + 1)-th measure starts exactly in the middle of the k-th measure. In other words, the falling side of the k-th measure overlaps exactly to the rising side of the (k + 1)-th measure. Consequently, the measurement of $\bar{y}_{k+1}(\tau) \bar{y}_k(\tau)$ takes 3τ .

In the domain of oscillators, frequency synthesis and telecom, MVAR should be preferred to AVAR because of its superior capability to divide the fast noise processes, namely white PM and flicker PM. More details are given in the review article (Bregni, 2016).

3.2.3 The parabolic variance (PVAR)

The Parabolic Variance, denoted with "par $\sigma_y^2(\tau)$ " or PVAR, relies the same formula (75) already used for AVAR and PVAR, but the stream of values $\overline{y}_k(\tau)$ results from Ω averaging (68). Of course, the measurements are exactly contiguous. Consequently, replacing the Π counter with an Ω counter, we obtain PVAR.

3.3 Conversion from spectra to two-sample variances

We have seen that the PSD of frequency noise and the two-sample variances are both valid measures of the frequency stability. The obvious question arises, about how to convert data between these two worlds. The question is best answered after rewriting the general formula (74) for the two-sample variance as

$$\sigma_{\mathbf{y}}^{2}(\tau) = \left[\int_{0}^{\infty} \mathbf{y}(t) w(t) dt\right]^{2}$$
(76)

where the weight function w(t) is either $w_A(t)$, $w_M(t)$ or $w_P(t)$, plotted in Figure 20. As one may expect, the subscripts A, M and P stand for Allan, Modified Allan, and Parabolic. Hence, the two-sample variance can be seen as the filter described by

$$\sigma_{\gamma}^{2}(\tau) = \int_{0}^{\infty} S_{\gamma}(f) |H(f)|^{2} df$$
(77)

where the transfer function H(f) is the Fourier transform of the impulse response h(t), which is equal to w(-t), and the trivial subscripts A, M and P apply. After some boring math, it can be proved that

$$|H_A(f)|^2 = 2 \frac{\sin^4(\pi \tau f)}{(\pi \tau f)^2}$$
(78)

$$|H_M(f)|^2 = 2 \frac{\sin^6(\pi \tau f)}{(\pi \tau f)^4} \quad \text{for } \tau \gg \tau_0$$
 (79)

$$|H_P(f)|^2 = 9 \frac{[2\sin^2(\pi\tau f) - \pi\tau f \sin(2\pi\tau f)]^2}{2(\pi\tau f)^6} \quad \text{for } \tau \gg \tau_0$$
(80)

Should the reader wish to calculate $|H(f)|^2$ by hand, the exercise is simpler if $w_A(t)$, $w_M(t)$ and $w_P(t)$ are shifted and centered at t = 0. Of course, the phase introduced by this time shift has no effect on $|H(f)|^2$.



Figure 20 – Weight function used in AVAR, MVAR and PVAR. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

We can see on Figure 21 that $|H_A(f)|^2$, $|H_M(f)|^2$ and $|H_P(f)|^2$ are band pass filters with a bandwidth of approximately one octave centered at $f \approx 0.4/\tau \dots 0.45/\tau$. Center frequency and bandwidth are not the same for the three functions. For high frequency, $|H_A(f)|^2$ rolls off as $1/f^2$. For this reason, an additional lowpass filter at a suitable cutoff frequency f_H is necessary if the fluctuations contain white and flicker phase noise. An interesting treatise is given in (Calosso, Clivati, & Micalizio, Avoiding Aliasing in Allan Variance: An Application to Fiber Link Analysis, 2016). However, the language may be more difficult because (Calosso, Clivati, & Micalizio, Avoiding Aliasing in Allan Variance: An Application to Fiber Link Analysis, 2016) is about an application in optical fibers. At first sight, the low-pass filter does not seem necessary for MVAR and PVAR because $|H_M(f)|^2$ and $|H_P(f)|^2$ roll off as $1/f^4$. The catch is that MVAR and PVAR result from sampling y(t) at the rate $1/\tau_0$, thus a lowpass filter at the cutoff frequency $1/(2\tau_0)$ is needed.



Figure 21 – Frequency response associated to AVAR, MVAR and PVAR. Reprinted, with permission, from F. Vernotte, M. Lenczner, P-Y Bourgeois, E. Rubiola, "The Parabolic Variance (PVAR), a Wavelet Variance Based on the Least-Square Fit," *IEEE Transact Ultrason. Ferroelec. Frequency Control* vol. 63 no. pp. 611-623, April 2016 (Vernotte, Lenczner, Bourgeois, & Rubiola, 2016).

Equations (78), (79) and (80) enable the conversion from $S_{\gamma}(f)$ to $\sigma_{\gamma}^2(\tau)$. The reverse conversion, from $\sigma_{\gamma}^2(\tau)$ to $S_{\gamma}(f)$, is not possible in the general case. The reason is that the frequency responses (Figure 21) have side lobes at high frequency. Hence, a point on the $\sigma_{\gamma}(\tau)$ plot represents the FM noise in the main lobe of the bandpass, plus the FM noise captured by the side lobes. The side lobes are smaller for MVAR and PVAR.

Table 6 provides all the formulae to convert from PM or FM noise PSD to the three variances described.

Table 6 – Noise response of the two-sample variances. The formulae for MVAR and PVAR hold for $\tau \gg \tau_0$, being τ_0 he sampling interval. Reproduced from the 2019

noise type	$S_{arphi}(f)$	$S_{y}(f)$	$S_\varphi\leftrightarrow S_y$	AVAR $\sigma_y^2(\tau)$	MVAR mod $\sigma_y^2(\tau)$	$\mathrm{PVAR}~\sigma_{\mathrm{y}}^{2}(\tau)$
white PM	þ	$h_2 f^2$	$h_2 = \frac{b_0}{\nu_0^2}$	$rac{3f_H}{4\pi^2}rac{{\sf h}_2}{ au^2}$ 0.0760 $f_H{\sf h}_2/ au^2$	$\frac{3}{8\pi^2} \frac{h_2}{\tau^3} \\ 0.0380 h_2/\tau^3$	$\frac{3}{2\pi^2} \frac{h_2}{\tau^3} \\ 0.1520 h_2/\tau^3$
flicker PM	$b_{-1}f^{-1}$	h_1f	$h_1 = \frac{b_{-1}}{\nu_0^2}$	$\frac{3\gamma - \ln 2 + 3\ln(2\pi f_H \tau)}{4\pi^2} \frac{h_1}{\tau^2}$ $3\gamma - \ln 2 + 3\ln 2\pi]/4\pi^2 = 0.166$	$rac{(24\ln 2 - 9\ln 3)}{8\pi^2} rac{ extsf{h_1}}{ au^2} rac{ extsf{2}}{ au^2}$	$rac{3\left[\ln(16)-1 ight] }{2\pi^{2}}rac{h_{1}}{ au^{2}} rac{1}{ au^{2}} rac{1}{ au^{2}}$
white FM	$b_{-2}f^{-2}$	h ₀	$h_0 = \frac{b_{-2}}{\nu_0^2}$	$rac{1}{2}rac{h_0}{ au}$	$\frac{1}{4} \frac{h_0}{\tau}$	$\frac{3}{5} \frac{h_0}{\tau}$
flicker FM	$b_{-3}f^{-3}$	$h_{-1}f^{-1}$	$h_{-1} = \frac{b_{-3}}{\nu_0^2}$	$2 \ln(2) h_{-1}$ 1.3863 h_1	$\frac{27\ln 3 - 32\ln 2}{8} h_{-1}$ 0.9352 h_1	$\frac{2 \left[7 - \ln(16)\right]}{5} h_{-1}$ 1.691 h_1
random walk FM	$b_{-4}f^{-4}$	$h_{-2}f^{-2}$	$h_{-2} = \frac{b_{-4}}{\nu_0^2}$	$\frac{2\pi^2}{3}h_{-2} au$ 6.5797 h_2 au	$\frac{11\pi^2}{20} h_{-2} \tau$ 5.4283 h_{-2} \tau	$\frac{26\pi^2}{35}h_{-2}\tau$ 7.3317 h_2 $ au$
linear freq	uency drif	t Dy		$rac{1}{2}{ m D}_y^2 au^2$	$rac{1}{2}D_{y}^{2} au^{2}$	$rac{1}{2}D_{y}^{2} au^{2}$
MVAR an	d PVAR f	ormular ne	sed $\tau > 1/f_H$, v	where $f_H < 1/2 au_0$ is the cutoff f	requency, and τ_0 is the	e sampling interval.

Enrico's Chart of Phase Noise and Two-Sample Variances (older version), http://rubiola.org.

3.3.1 Comparison between AVAR, MVAR, and PVAR

A few more questions arise. First, why MVAR and PVAR are used, after the original Allan variance? Second, if these alternate variances are more modern and efficient, why

February 6, 2025

they do not have replaced the Allan variance? And third, why all these variances give different values (Table 6) for the same noise process?

The original Allan variance was invented at NBS (the ancient name of NIST) to describe the stability of the atomic clocks used in the time scale. Monitoring the long-term performance is of paramount importance in this application. When a long data record is available (say, 6 years) we wish to plot $\sigma_{y}(\tau)$ for τ as large as possible, as close as possible to half of the data record (3 years). AVAR is still the best we have for this purpose, and still the standard tool used by National Laboratories to monitor their primary clocks. Additionally, AVAR has been studied for longer time, and many technical publications are available. It is therefore not a surprise that the engineers and practitioners tend to believe that the Allan deviation is what they should use.

The modified Allan variance came in the 1980s from the field of optics as a better alternative to AVAR for the measurement of fast phenomena. The main features of MVAR are the superior rejection of the instrument noise, related to the Λ counter, and the capability to resolve between flicker PM and white PM, which impossible with AVAR. For this reason, MVAR is highly recommended for all applications we are concerned with, like telecom, oscillators and frequency synthesis. By contrast, MVAR is not a good choice to measure the long-term instability of atomic clocks. The reason is that w_M spans over 3τ , instead of 2τ for w_A . For example, with a data record of 6 years, the maximum τ is close to 3 years for AVAR, and close to 2 years for MVAR.

The parabolic variance, par $\sigma_y^2(\tau)$ or PVAR, is much less known because it has been proposed only in 2015. Compared to MVAR, PVAR is superior in all cases and for all types of noise, to the extent that it enables the measurement of noise processes with a shorter data record. Compared to AVAR, we notice that PVAR is a better choice in most cases, but it is slightly less suitable to the measurement of the slowest phenomena, like frequency random walk and drift of atomic clocks, where one wants the result with the shortest data record. PVAR is superior to MVAR in all cases for the detection of random processes (Vernotte, Lenczner, Bourgeois, & Rubiola, 2016), thus we can expect that it will replace MVAR. However, at the time of writing, PVAR is still not present in standards, in recommendations, and in all the software packages.

The three tools start from the same formula $\sigma_y^2(\tau) = [\bar{y}_2(\tau) - \bar{y}_1(\tau)]^2/2$, but the averages $\bar{y}(\tau)$ are different. Thus, it is not a surprise that we end up with different statistical properties. We see on Table 6 that the response to the linear drift is the same. This results from a choice of normalization (scale factor), which has historical roots. By contrast, an engineer with a background in signal processing would have probably scaled MVAR. The different coefficients found in Table 6 do not mean that one variance

is better than another, no more than inches are a better or a worse unit than centimeters.

Example 5 – Two-sample variances. The phase noise spectrum of a 100 MHz oscillator has the following noise terms, extrapolated to 1 Hz. Calculate the two-sample deviations ADEV, MDEV, and PDEV. For ADEV, use $f_H = 500$ Hz.

Noise type	RW FM	Flicker FM	White FM	Flicker PM	White PM
Slope	$1/f^4$	$1/f^{3}$	$1/f^{2}$	1/f	constant
I(f)	-99	-134	$(\mathbf{N}\mathbf{A})$	-164	-180
	dBc/Hz	dBc/Hz	(NA)	dBc/Hz	dBc/Hz
at <i>f</i> =	10 Hz	100 Hz	(NA)	1 kHz	10 kHz

Let us start with the FM random walk, which rolls off as $1/f^4$, i.e., -40 dB/decade. The value of -99 dBc/Hz at 10 Hz is equivalent to -59 dBc/Hz extrapolated at 1 Hz. The corresponding term of $S_{\varphi}(f)$ is b_{-4}/f^4 , with $b_{-4} = 10^{(59+3)/10} = 2.5 \times 10^{-6} \text{ rad}^2/\text{Hz}^3$. We convert $S_{\varphi}(f)$ into $S_{\gamma}(f)$ using $S_{\gamma}(f) = (f/f_0)^2 S_{\varphi}(f)$. The FM random walk of $S_{\gamma}(f)$ is h_{-2}/f^2 , with $h_{-2} = b_{-4}/f_0^2 = 2.5 \times 10^{-6}/(10^8)^2 = 2.5 \times 10^{-22} \text{ Hz}^{-3}$.

Then, we follow the same steps with the flicker FM (-30 dB/decade). We calculate the coefficient b_{-3} of $S_{\varphi}(f)$, and in turn the term h_{-1} of $S_{\gamma}(f)$. And we repeat for all the other noise terms.

Noise type	ADEV	MDEV	PDEV
White PM	$152 \times 10^{-15} / \tau$	$2.76 \times 10^{-18} / \tau \sqrt{\tau}$	$5.51 \times 10^{-18} / \tau \sqrt{\tau}$
Flicker PM	1.52 × 10 / 1	$8.25 \times 10^{-16} / \tau$	$1.46 \times 10^{-15} / \tau$
White FM	(NA)	(NA)	(NA)
Flicker FM	3.32×10^{-12}	2.73×10^{-12}	3.70×10^{-12}
RW FM	$4.07 \times 10^{-11} \sqrt{\tau}$	$3.70 \times 10^{-11} \sqrt{\tau}$	$4.30 \times 10^{-11} \sqrt{\tau}$

4 Phase noise in components

4.1 Amplifiers

The amplifiers are one of the most ubiquitous building-block, serving to different purposes. Amplification is necessary for example to raise the power after the loss associated to frequency conversion, or after splitting a signal into several branches. Second, it is an absolute necessity to provide isolation between a VCO and the following circuits (e.g., a divider chain or others) to minimize any feedback or noise contribution because of periodic loading. Similarly, high isolation amplifier is necessary at the output of a reference oscillator. Then, power amplification must be provided at the output of practical instruments. Different design rules follow, depending on whether the engineer designs with transistors, or at system level by assembling blocks. The latter approach is privileged here. A wider treatise of phase noise in amplifiers is given in Reference (Boudot & Rubiola, Phase noise in RF and microwave amplifiers, 2012). We restrict our attention to amplifiers at room temperature, within a reasonable range. Hightemperature and cryogenic electronics are highly specialized topics, beyond our scope.

4.1.1 White and flicker phase noise

If we assume that we generate our signal in a noise-free environment, or at least start off with the theoretical minimum of – 174 dBm/Hz set by the thermal noise generated by a resistor at room temperature, the signal will be degraded by the amplifier that follows. In virtually all cases, the amplifier can be described in terms of the noise factor *F*. For reference, the noise factor of most microwave amplifiers is of 1–3 dB. A value of 0.4–0.5 dB is found in some special narrowband amplifiers, while power amplifiers often have a noise factor of 6 dB and more.

The white phase noise of a signal of power P_0 is

$$S_{\varphi}(f) = \frac{FkT_0}{P_0} \tag{81}$$

This is the term b_0 of the polynomial law we have seen in the previous Sections.

The white noise floor can be observed only beyond a cutoff offset frequency. Various noise sources must be taken into consideration, affecting low offset frequencies. As mentioned previously, flicker noise is the major reason for the noise degradation, and its contribution is device-dependent and can range from a few hundred hertz to 1 MHz. It is caused by internal near-DC noise modulating the phase of the passing signal, and the input and output impedances of the amplifier. References (Walls, Ferre-Pikal, & Jefferts, Origin of 1/f PM and AM Noise in Bipolar Transistor Amplifiers, 1997) (FerrePikal, Walls, & Nelson, 1997) (Howe & Hati, 2005) provide a more detailed analysis and guidelines for bipolar transistors.

It has already been observed long time ago (Halford, Wainwraight, & Barnes, Flicker Noise of Phase in RF Amplifiers and Frequency Multipliers, Characterization, Cause, and Cure, 1968) that the amplifier's flicker of phase, that is, the term b_{-1} of the polynomial law, is rather independent of the carrier power P_0 in a relatively wide power range from very small signals to moderate clipping, and it is also rather independent of carrier frequency in the amplifier bandwidth. Moreover, the spread of values for amplifiers of similar technology and different manufacturers is surprisingly small. Figure 22 shows an example of amplifier PM noise, measured at different levels of input power.



Figure 22 – Example of amplifier phase noise. Courtesy of Vincent Giordano, FEMTO-ST Institute.

The coefficient b_{-1} can be considered a characteristic parameter of the amplifier. We can take as a reference the following values: – 100 to – 110 dB rad²/Hz for microwave amplifiers, up to 20 GHz; – 120 dB rad²/Hz for microwave amplifiers in SiGe technology (which unfortunately have a high noise factor, typically 5 to 8 dB); and – 130 to – 140 dB rad²/Hz for bipolar RF amplifiers, up to 1 GHz.

There is very little that one can do to reduce flicker PM noise. Some negative feedback at low frequency proved to be useful, such as an emitter resistor with no bypass capacitor. Useful tips are negative feedback at RF to stabilize the transconductance, and a careful design of the amplifier for a low noise factor at low frequency. Kuleshow and Boldyreva (Kuleshov & Boldyreva, 1997) suggest that there is a magic bias which reduces the 1/f noise in common-emitter amplifiers.

Figure 22 also shows what happens combining white and flicker PM noise. An interesting fact is that is that the cutoff frequency f_c , where white PM noise equals flicker PM noise, is proportional to P_0 . This happens because flicker PM noise is constant vs power, while white PM noise gets lower at higher power. Computer simulation programs like SPICE may use f_c as the noise parameter. This choice is unfortunate and misleading because the user unaware of the mechanism may not understand the need to update f_c after changing any parameter affecting P_0 .

Cascading amplifiers yields a surprising result. The general rule, that the noise of the chain is essentially determined by the noise of the first stage, does not apply to flicker. Since the coefficient b_{-1} is in a first approximation independent of the carrier power, the flicker PM noise is

$$S_{\varphi}(f) = \frac{b_{-1}}{f}$$
 with $b_{-1} = [b_{-1}]_1 + [b_{-1}]_2 + \cdots$ (82)

Most people are surprised to learn that the result is independent of the order of the amplifiers in the chain.

The flicker of two or more amplifier connected in parallel is another amazing consequence of the property that b_{-1} is independent of the carrier power. Splitting the input into two amplifiers gives the same amount of PM noise in each branch, despite of the smaller power. Combining the two outputs, the PM noise is half of the noise of one amplifier because the carrier adds up coherently, and the PM noise adds up statistically. To this extent, push-pull stages, distributed amplifiers, and other configuration where the output power results from adding the contribution of multiple transistors are equivalent to parallel amplifiers. The idea of achieving low flicker by paralleling multiple amplifiers is exploited in some commercial special-purpose amplifiers (Boudot & Rubiola, Phase noise in RF and microwave amplifiers, 2012). By contrast, the white PM noise cannot be improved in this way because the PM noise of each branch scales up as the power is reduced. Actually, the loss associated to power splitting degrades the white noise.

Example 6 – Amplifier noise. A clean noise-free 100 μ W signal is amplified by two cascaded amplifiers detailed in the Table below. We calculate the phase noise PSD at the output, considering the A-B order (lower noise amplifier first), and the B-A order.

Amplifier	Ga	Gain <i>A</i> Noise Factor <i>F</i>		Flicker coeffic	ient b_{-1}	
А	3.98	12 dB	1.41	1.5 dB	10^{-13} rad ² /Hz	-130 dB rad ² /Hz
В	3.16	10 dB	2.51	4 dB	$1.58 \times 10^{-13} \text{ rad}^2/\text{Hz}$	–128 dB rad ² /Hz

Considering the A-B order first, the noise factor is

$$F_{AB} = F_A + \frac{F_B - 1}{A_A^2} = 1.41 + \frac{2.55 - 1}{3.98^2} = 1.51$$
 (1.8 dB)

Thus, the white PM noise is

$$S_{\theta}(f) = \frac{F_{AB}kT_0}{P_0} = \frac{1.51 \times 1.38 \times 10^{-23} \times 290}{10^{-4}} = 6.06 \times 10^{-17}$$

that is, $-162.1\ dBrad^2/Hz$, or $-165.1\ dBc/Hz$

When the amplifiers are interchanged, we get

$$F_{BA} = F_B + \frac{F_A - 1}{A_B^2} = 2.51 + \frac{1.41 - 1}{3.16^2} = 2.55$$
 (4.1 dB)

and

$$S_{\theta}(f) = \frac{F_{BA}kT_0}{P_0} = \frac{2.55 \times 1.38 \times 10^{-23} \times 290}{10^{-4}} = 1.03 \times 10^{-16}$$

that is, $-159.9 \text{ dBrad}^2/\text{Hz}$, or -162.9 dBc/Hz. As expected, the noise of the two configurations is different, with $\approx 2 \text{ dB}$ higher PM noise when the noisier amplifier comes first.

The flicker noise is the same for the two configurations (remember that b_{-1} has the physical dimension of rad² because $S_{\theta} = b_{-1}/f$ must have the dimension of rad²/Hz)

 $b_{-1} = b_{-1}|_A + b_{-1}|_B = 10^{-13} + 1.58 \times 10^{-13} = 2.58 \times 10^{-13}$ rad² that is, -125.9 dBrad, or -128.9 dBc/Hz extrapolated to 1 Hz. ■

4.1.2 How to choose a low PM noise amplifier

The noise factor *F* is generally well documented, and describes the white noise rather precisely. Thus, the amplifier with the lowest *F* also exhibits the lowest white PM noise according to $S_{\varphi} = FkT/P_0$. By contrast, flicker PM noise is seldom mentioned in the specs, and even when it is present, the documentation is generally poor. Generally, we find a plot, or at most a small number of plots referring to practical cases measured one sample at the Company site. No typical/max values are given, and no rule for extrapolation to different conditions. Thus, a discrepancy between actual phase noise and the datasheet is not a surprise.

Figure 23 shows the flicker of the Analog Devices HMC-072 amplifier taken from the datasheet. This amplifier was chosen because the phase noise is better documented than that we find in similar devices. Designing a complex oscillator intended to provide ultimate low phase noise at that frequency, we were looking for the lowest flicker. We will see later in this Chapter that the flicker of the sustaining amplifier impacts on the $1/f^3$ noise (flicker of frequency) of the oscillator. The 1/f noise found in the datasheet (Figure 23) is equivalent to -132 to -135 dBc/Hz extrapolated to 1 Hz. The flicker we measured is in fair agreement, being -131 dBc/Hz at 1 Hz (Figure 23). The

measurement was done with a double balanced mixer saturated at both inputs, following the classical scheme of Figure 72. The mixer is from Marki, and the line stretcher is from Arra.



Figure 23 – Phase noise of the Analog Devices HMC-C072 amplifier at $f_0 = 10$ GHz. The spectra are from the HMC-072 data sheet, © Analog Devices, reproduced with permission. Graphical adaptation is ours.



Figure 24 – Phase noise of two Analog Devices HMC-072 amplifiers at $f_0 = 10$ GHz, measured with a saturated double-balanced mixer. Each amplifier receives 7 V bias and 0 dBm input power. Subtract 3 dB for the noise of one amplifier, and add 3 dB to convert dBrad²/Hz to dBc/Hz. Courtesy of Yannick Gruson, FEMTO-ST Institute.

Having said all that, we can get useful tips by combining physical insight with some other parameters apparently not related to flicker PM noise. This is a good way to avoid numerous blind tests, but measurements are always necessary when flicker PM noise is a critical parameter.

The most important fact about flicker PM noise is that it comes from up conversion of near-DC flicker, that is, the 1/f fluctuation of the bias current, internal capacitances, or other parameters. Because up conversion is a nonlinear phenomenon, high linearity is a desired feature. High intercept power IP2 and IP3 and low harmonic distortion are good signs for a potentially good low-flicker amplifier. These parameters must be interpreted together with the amplifier topology. For example, a push-pull configuration cancels the second harmonics thanks to symmetry, but symmetry does not cancel the PM flicker because the near-DC processes inside the two transistors are independent.

Internal flickering in components is a microscopic phenomenon. The evidence of this is that the probability density distribution of the internal voltage or current, and therefore of random phase, is Gaussian. This distribution can only be originated by large number of small random fluctuation via the central limit theorem. Consequently, we expect lower flicker from devices having larger volume of the gain region and not-too-sharp junctions. If we design at transistor level, these considerations point to transistors with not-too-high gain β (or h_{FE}) and not-too-high transition frequency f_T . Higher V_{CB} max reveals that the collector-base junction is not too thin, thus the C_{BE} capacitance is more stable versus the fluctuations of the bias current. Germanium transistors may exhibit low 1/f noise, paradoxically for the same reason why are not as good as Silicon transistors in more general terms: the base region is thicker, the current gain is lower, etc. With modern technology, low 1/f noise is found in SiGe transistors (Ashburn, 2003). Unfortunately, often this comes at the cost of higher noise factor.

4.1.3 Isolation amplifiers

Isolation amplifiers are special devices optimized for high forward-to-reverse gain ratio, and intended to keep a precision oscillator or a frequency standard free from interferences brought inside from its output. The typical problem is that a resonator accumulates the incoming energy over a time equal to the relaxation time $Q/\pi f_0$, and it is not rare than even a power of picowatts or even femtowatts reaching the core of a precision oscillator impacts on its stability and noise in unpredictable way, or phase lock it. Isolation amplifiers are also needed in frequency distribution systems, where one input signal has to be delivered to multiple destinations. High output-to-input and output-to-output isolation is the main feature, and also isolation from load change.

At microwave frequencies, the isolation of regular microwave amplifiers can be increased by using ferrite isolators. One of these devices provides 20 dB typical isolation, and a few can be cascaded. In the radiofrequency region, below some GHz, the isolators are not an option, being the ferrite too large and impractical. At these frequencies, the grounded-base (or common-base) bipolar transistor is the most common choice. The idea is that the base region is a shield, which breaks the capacitive coupling between emitter (input) and collector (output). The grounded-base stage often goes with a common emitter stage to form a cascode amplifier. Increased isolation can be achieved with two or more grounded-base amplifiers stacked on the collector of the first common-emitter stage (Nelson, Walls, Siccardi, & De Marchi, 1994), as shown on Figure 25. A problem of this configuration is the higher supply voltage (20–30 V), uncomfortable for regular electronic design. A differential pair is an alternate option, with the input on the base of the first transistor, and the base of the second transistor connected to ground as the shield.



Figure 25 – Isolation amplifier. The use of two or more common-base transistors results in enhanced output-to-input isolation.

No matter whether the electrical circuit exploits circulators or grounded-base transistors, high isolation, say 100 dB or more, requires that enclosure and shielding are accounted for. Preventing radiation and blocking the signal propagating on the outer surface of shielded cables and enclosures is of paramount importance. In the HF-VHF

region, a common-mode ferrite filter on the input and output coaxial cables is often necessary.

4.1.4 The Siccardi Isolation Amplifier

[I plan to add this section to a forthcoming version]

4.2 Frequency dividers

4.2.1 Digital frequency dividers

Frequency division is one of the early applications of digital circuits. The reader can refer to the Egan article (Egan, 1990) for a review based on available data about the classic TTL and ECL families. There are two basic types, asynchronous dividers and synchronous dividers (Figure 26). The main difference is that in the asynchronous divider the jitter of each flip-flop adds up as the signal propagates along the chain, while in the synchronous dividers all the flip-flops switch simultaneously on the edges of the input clock. The output of an asynchronous divider can be synchronized to the input clock with a D-type flip-flop.



Figure 26 – Asynchronous and synchronous frequency dividers.

In accordance with Figure 27, the practical noise limit is of the order of -170 dBc/Hz for the old TTL dividers and -155 dBc/Hz for the ECL dividers. Yet, these values are not easy to achieve in practice. Small-scale integration CMOS dividers, up to an input frequency of 10 MHz, are similar in phase noise to the TTL devices. However, the close-in noise or noise between 1 and 10 Hz off the carrier is higher than that of TTL devices. TTL devices require higher shielding and better power supply decoupling to prevent external crosstalk between the various stages, which otherwise results in unwanted spurious and sidebands. Programmable logic devices offer new design perspectives, but they are only usable up to 1 GHz. Commercial dividers up to 12–26 GHz input frequency are available from Analog Devices, Hittite (now with Analog Devices), Keysight Technologies, Microsemi, ON Semiconductor, Pasternack, and other manufacturers.

Most of these dividers have another unpleasant effect, in the form of internal crosstalk. Crosstalk is defined as the amount of input signal appearing at the output. In high-performance synthesizers, it is necessary to use a low-pass filter after the reference or the programmable divider and a pulse shaper to translate the resulting sine wave back into a square wave for appropriate suppression of the crosstalk.



Figure 27 – Phase noise of some dividers as. All data refer to the output frequency in the left-hand plot, and are scaled up to 10 GHz in the right-hand plot. Reprinted from U. L. Rohde, Microwave and Wireless Synthesizers, 1st ed, Wiley 1997.

4.2.2 Phase noise scaling

In a noise-free divider, the device divides the input phase by *N* for the same reason it divides by *N* the input frequency

$$f_o = \frac{f_i}{N}$$
 and $\varphi_o(t) = \frac{\varphi_i(t)}{N}$ (83)

Naively, we may be inclined to extend this result to $S_{\varphi o}(f) = S_{\varphi i}(f)/N^2$. Reality is more complex because *aliasing* strikes on some types of signals, but not on others. Some relevant cases will be analyzed below.

Let us start from white noise. In digital circuits, phase noise is sampled at the frequency $2f_0$ for the simple reason that the phase fluctuations exist only during the rising and falling edges of the signal, while the pulse level has no effect on the output. Thus, the *bandwidth* of phase noise is equal to f_0 , which is half the sampling frequency. Some circuits switch only on one active edge, either rising of falling. In such case the sampling process takes place at f_0 , and the noise bandwidth is $f_0/2$. We take f_0 as the bandwidth, leaving to the reader the extension to $f_0/2$.

The easiest way to understand aliasing is to derive the spectrum from the time fluctuation x(t) and its mean square value $\langle x^2 \rangle$. In principle, a noise-free divider transfers the time fluctuation from the input to the output. By virtue of the Parseval identity, it holds that

$$\langle \mathbf{x}^2 \rangle = S_{\mathbf{x}}(f)B \tag{84}$$

where the bandwidth *B* is equal to f_i at the input, and to $f_o = f_i/N$ at the output. The consequence is that

$$S_{xo}(f) = NS_{xi}(f)$$
 (with aliasing) (85)

Thus, the bandwidth reduction by 1/N is compensated by increasing the noise PSD by the same factor *N*. The time fluctuation can be converted into phase fluctuation using $\varphi(t) = 2\pi f_0 \mathbf{x}(t)$. This applies to the input frequency f_i and to the output frequency f_o . There results

$$S_{\varphi o}(f) = \frac{1}{N} S_{\varphi i}(f) \quad \text{(with aliasing)} \tag{86}$$

By contrast, aliasing produces no detectable effect on flicker noise because the aliases are attenuated by a factor $1/2f_i$, $1/3f_i$, etc. Thus, the output phase noise is described by

$$S_{xo}(f) = S_{xi}(f)$$
 (no aliasing) (87)

and

$$S_{\varphi o}(f) = \frac{1}{N^2} S_{\varphi i}(f) \qquad \text{(no aliasing)} \tag{88}$$

A further phenomenon occurs when an analog (sinusoidal) is sent to the input, instead of a clean digital signal from the same logic family. The phase noise PSD in the input stage increases because of non-linearity and other phenomena.

Figure 28 shows the interplay of noise levels in a digital divider. The input PM noise increases by a few dB in the first stage, and then it is scaled down by a factor of 1/N or $1/N^2$, depending on the noise type. As an unpleasant consequence, a ÷10 divider may
scale down the input white PM noise by a mere 5–6 dB instead of the 20 dB of the $1/N^2$ rule.



Figure 28 – Phase noise scaling in a digital frequency divider.

The case of a signal with a pedestal (Figure 29) is so common in PLLs and frequency synthesizers, that it deserves special attention. For this purpose, we can divide the input phase fluctuation $\varphi(t)$ into pedestal and floor, denoted with the subscripts p and f. Additionally, pedestal and floor have different origins, thus they are statistically independent. The pedestal is rather narrow as compared to the carrier, but it contains most of the mean square fluctuation $\langle \varphi^2 \rangle$. In most practical cases, N is not large enough to shrink the bandwidth of the pedestal ($N < f_i/f_p$), thus there is no aliasing. This is why $S_{\varphi}(f)$ follows the $1/N^2$ rule. Further dividing, at some point the pedestal shrinks to f_i , aliasing takes place, and the scaling rule changes from 1/N to $1/N^2$. By contrast, aliasing strikes on the floor, which is wide and uniform, and $S_{\varphi}(f)$ scales as 1/N. The latter result is of paramount importance in PLLs.





4.2.3 Time-type and phase-type PM noise

For the purpose of frequency synthesis, we identify two basic types of phase noise, called *phase-type noise*, or φ -type noise for short, and *time-type noise*, or x-type noise. These types of noise, first introduced in (Calosso & Rubiola, Phase Noise and Jitter in Digital Electronics, 2016.), differ in how the noise is scaled with the carrier frequency f_0 .

The *phase-type noise* is characterized by the statistics of φ independent of f_0 . Then, x follows the rule

$$\mathbf{x}(t) = \frac{\varphi(t)}{2\pi f_0}$$
 and $S_{\mathbf{x}}(f) = \frac{1}{4\pi^2 f_0^2} S_{\varphi}(f)$ (89)

This is the case of an input signal v(t) crossing a fluctuating threshold n(t). The time fluctuation is x(t) = n(t)/SR, where SR is the *slew rate*. In the case of a sinusoidal signal of peak amplitude V_0 , the slew rate is $SR = 2\pi V_0 f_0$. Thus $\varphi(t) = n(t)/V_0$, and consequently $\langle \varphi^2 \rangle = \langle n^2 \rangle / V_0^2$. The PSD $S_n(f)$ of the threshold fluctuation has a bandwidth *B* equal to the full analog bandwidth at the device input, while $S_{\varphi}(f)$ has a bandwidth equal to the input frequency f_0 . Thus,

$$S_{\varphi}(f) = \frac{\langle n^2 \rangle}{f_0 V_0^2} = \frac{S_n(f)B}{f_0 V_0^2}$$
(90)

Having to guess about a digital circuit with no specific information on hand, one can start from the following values: $\sqrt{S_n(f)} = 10 \text{ nV}/\sqrt{\text{Hz}}$ for the analog white noise at the device input, and $B = 4f_{max}$, four times the maximum switching frequency.

Oppositely, *the time-type noise* is characterized by the statistics of time fluctuation x independent of f_0 , and φ given by

$$\varphi(t) = 2\pi f_0 \mathbf{x}(t)$$
 and $S_{\varphi}(f) = 4\pi^2 f_0^2 S_{\mathbf{x}}(f)$ (91)

This is the case of the signals propagating inside an integrated circuit. Such signals have full amplitude and slew rate, thus the delay and its fluctuations result from the sum of the contribution of all the individual gates or cells along the path. Crosstalk is not considered in this simplified description, but it can be a major nuisance in practical cases.

The noise-scaling rule, either 1/N or $1/N^2$, is not sufficient to describe the noise in a divider. The phase noise cannot be reduced arbitrarily by dividing a high input frequency by a large N because at some point the scaled-down phase noise hits the phase noise of the output stage. When this happens, the phase noise cannot be reduced by further frequency division.

Figure 30 shows an example of a digital circuit when input frequency changes in a wide range. This example is the clock line in a FPGA, including the input an output circuits. Since $f_o = f_i$, there is no 1/N or $1/N^2$ scaling. On the right-hand side of the plot, where the phase noise is of the white type, we identify clearly the phase-type with aliasing, ruled by $S_{\varphi}(f) = S_n(f)B/f_0V_0^2$. Changing f_0 in powers of 2, the change in the white PM noise is close to the 3 dB expected. With $V_0 = 1$ V and B = 2.5 GHz (four times the maximum input frequency), we infer that the threshold fluctuation is of $11 \text{ nV}/\sqrt{\text{Hz}}$, rather independent of f_0 . Focusing on flicker noise, we see that at high f_0 , 100 MHz and beyond, the noise is clearly of the time type. Converting $S_{\varphi}(f)$ into $S_x(f)$, we find that the 1/f coefficient k_{-1} of the polynomial law is rather constant vs f_0 . The associated rms time fluctuation is 21 fs. Finally, thermal effects appear at f < 10 mHz (not visible on Figure 30).



Figure 30 – Example of noise in a digital circuit (Cyclone III FPGA). For the sake of proving concepts, the FPGA is programmed to replicate the input, so that we observe the input and output stages and the internal clock distribution. Reprinted from C. E. Calosso and E. Rubiola, arXiv:1701.00094 [physics.ins-det], 2016 (Calosso & Rubiola, Phase Noise and Jitter in Digital Electronics, 2016.).

Unfortunately, the data sheets seldom provide more than one example of phase noise, thus it is difficult to extract the noise parameters, and measurements are always necessary for critical applications. Figure 31 shows an example of phase noise PSD measured on a microwave divider. A flicker of -110 dBc/Hz extrapolated to 1 Hz, and a floor of -150 dBc/Hz can be expected from similar dividers.



Figure 31 – Example of phase noise in a microwave divider. Courtesy of V. Giordano, CNRS FEMTO-ST Institute, Besancon, France.

4.2.4 The Λ divider

The Λ divider (Calosso & Rubiola, The Sampling Theorem in Pi and Lambda Digital Frequency Dividers, 2013), shown in Figure 32, is a simple trick to circumvent the aliasing phenomenon inherent in digital dividers. The output signal is a triangular wave obtained by adding N square waves shifted by an integer number of clock half-periods. In the triangular wave, the power associated n-th harmonics is proportional to $1/n^4$, instead of $1/n^2$ for the square wave. No aliasing takes place because the output signal is sampled at $2f_i$, which is the same sampling frequency of the input. Thus, the phase noise follows the rule $S_{\varphi o}(f) = S_{\varphi i}(f)/N^2$. The triangular wave can be easily cleaned by filtering out the harmonics. The name " Λ divider" derives from the similarity of the Greek letter Λ with a triangle, and for the same reason the regular digital divider is called Π divider for the similarity of the Greek letter Π to the pulse of a square wave. Reference (Calosso & Rubiola, The Sampling Theorem in Pi and Lambda Digital Frequency Dividers, 2013) shows an example of \div 10 divider exhibiting a phase noise of -165 dBrad²/Hz (white) and -130 dBrad² (flicker at 1 Hz), measured at the 10 MHz output frequency.



Figure 32 – The Lambda divider. Reprinted with permission from C. E. Calosso and E. Rubiola, "The Sampling Theorem in Pi and Lambda Digital Frequency Dividers," *Proc Internat Frequency Control Symp*, © 2013 IEEE (Calosso & Rubiola, The Sampling Theorem in Pi and Lambda Digital Frequency Dividers, 2013).

4.2.5 Analog frequency dividers

Analog frequency dividers may be a good choice in some special cases, for example for extremely low noise applications or when the input frequency is too high and out of reach for digital technology. Unfortunately, these dividers are complex to design, work in a narrow range of input frequency, and a quantitative understanding of phase noise is difficult.

The first of such dividers is the *injection locked oscillator*, or ILO (Adler, 1946), (Kurokawa, 1973), (Rategh & Lee, Superharmonic Injection-Locked Frequency Dividers, 1999). Figure 33 shows an example of integrated implementation (Rategh, Samavati, & Lee, A CMOS Frequency Synthesizer with an Injection-Locked Frequency Divider for a 5-GHz Wireless LAN Receiver, 2000). The idea is quite old, and exploits the fact that the nonlinearity in the oscillator loop, naturally, generates harmonics and beats notes. So, a signal injected in the oscillator, having frequency f_i which is sufficiently close to an integer multiple N of the oscillation frequency range, which depends on the amplitude. The injection locked oscillator is broadly equivalent to a PLL without the integration effect inherent in the VCO.



Figure 33 – Injection locked frequency divider.

The second option is the *regenerative frequency divider*, or *Miller divider* (Miller, 1939), (Rubiola, Olivier, & Groslambert, 1992), shown in Figure 34. To start, let us assume that a signal at the frequency f_i/N is present at the mixer output. Such signal crosses the bandpass filter tuned at f_i/N , is amplified, frequency-multiplied by N - 1, and sent back to the mixer. The mixer beats the input f_i with $(N - 1)f_o$, and regenerates f_i/N at its output. The feedback equation, as seen at the mixer output, is

$$f_i - (N - 1)f_o = f_o (92)$$

This gives $f_o = f_i/N$, as expected. Self-starting operation requires that regeneration starts from noise or from the power-up transient. The challenge is to design a frequency multiplier with no threshold, still effective at very low input power. With N = 2, self starting is rather easy to achieve because there is no multiplier. The mixer works well at low power at one input (the feedback signal) if the other input (the divider input) is saturated. The phase noise is limited by the amplifier, in the ideal case multiplied by a factor of $1/N^2$ inherent in the phase feedback operation. The full noise theory and some practical examples are found in reference (Driscoll, 1990). This reference also explains an optimization method which exploits the third harmonic generated by saturation in the mixer, to reduce the noise of a divider by two. The regenerative divider is broadly similar to the injection locked oscillator. The relevant difference appears in the absence of the input signal. The divider delivers no output, while the injection locked oscillator oscillates at its own free-running frequency.



Figure 34 – Regenerative frequency divider.

The third option is the *parametric frequency divider* (Penfield & Rafuse, 1962) (Driscoll, 1990). The idea is that the current flowing in the varactor contains two frequencies, f_i and $f_o = f_i/2$. The varactor is a nonlinear capacitance, thus it beat signals with high power efficiency because the capacitance has small energy loss in the charge-discharge cycle. Consider the divider-by-two operation shown in Figure 35. The input circuit is a *LC* which transfers the power to the varactor at the frequency f_i , and has high attenuation at f_o . Similarly, the output circuit extracts the power at the frequency $f_o = f_i/2$, and at f_i it isolates the varactor from the output. Dividers by N > 2 make use of a series of idlers (resonators) tuned at the sub-harmonics of the input frequency f_i/N . The parametric divider has low phase noise because it can manage relatively high signal level (20 dBm and more) and has low dissipation.



Figure 35 – Example of parametric frequency divider ÷2.

Example 7 – PLL synthesizer. The frequency synthesizer shown in Figure 36 is used to generate a 5-GHz signal with 1 kHz resolution. A 1-kHz reference signal is obtained from a 5-MHz reference oscillator (M = 5000), which is specified to have a single-sideband noise power of –140 dBc/Hz at f = 500 Hz off the carrier. If the loop bandwidth is assumed to be approximately 1 kHz, the noise from the reference oscillator will not be reduced by the low-pass filtering of the PLL. Although the divider by N will reduce the noise power by the factor $1/N^2$, the approximate loop transfer function is

$$\varphi_o(s) = \frac{\varphi_r[K_v F(s)/s]}{1 + [K_v F(s)/sN]} = N\varphi_i(s)$$

for reference frequencies below the loop bandwidth of 1 kHz. The net effect is that the output phase noise is the reference oscillator PM noise multiplied by N^2/M^2 . Of course, we need $N = 5 \times 10^6$ to scale the output 5 GHz frequency down to the 1 kHz frequency for the phase detector. The output PM noise due to the reference oscillator, at 500 Hz offset, is

$$L(f) = -140 \text{ dBc/Hz} + 10 \log_{10} \left(\frac{5 \times 10^6}{5 \times 10^3}\right)^2 = -80 \text{ dBc/Hz}$$

This example illustrates a problem inherent in PLL frequency synthesizers used to generate an output frequency much higher than the reference oscillator frequency, with high resolution. Although the reference oscillator noise power may be small, the same noise power appears on the output signal amplified by a large factor. Notice that the phase detector (Section 2.4.5), not accounted here, may introduce even higher the phase noise. ■



Figure 36 – A 5-GHz YIG oscillator harmonic stabilized to a 5-MHz reference.

4.3 Frequency multipliers

Generally, the long-term frequency stability available in an instrument comes from the internal 10 MHz quartz oscillator, or from an external reference. The external

frequency of 100 MHz is sometimes preferred in high-end instruments, where the lowest phase noise is of paramount importance. In synthesizers there is a frequent requirement for internal signals at higher frequencies, having low phase noise and the stability of the main reference. We have different options for frequency multiplication, which fall in two classes:

- A locked oscillator:
 - $\circ~$ A PLL with a divider in the feedback, or with harmonic sampling.
 - An injection-locked oscillator which exploits its internal nonlinearity to generate harmonics of the input frequency, and locks to one of it.
- A dedicated frequency multiplier.
 - o Transistor, optionally in push-push or push-pull configuration
 - Step-recovery diodes or snap-off diodes.
 - Varactors
 - Nonlinear delay lines (sometimes referred to as nonlinear transmission lines)
 - Diode (rectifier) networks

The choice is a matter of frequency, target PM noise, electrical power constraints, and complexity. High order frequency multipliers are very difficult to build. It is often better to use a cascade of multipliers of lower order than a single high order multiplier. Applications requiring fixed or variable frequency yield different solutions.

The noise-free wide-band frequency multiplier $\times N$ transfers the input time fluctuation to the output, $x_o(t) = x_i(t)$, thus it scales up the input PM noise according to

$$\varphi_o(t) = N\varphi_i(t)$$
 and $S_{\varphi o}(f) = N^2 S_{\varphi i}(f)$ (93)

This is +20 dB per factor-of-ten. A reduction in the phase noise bandwidth applies in some cases, typically the PLL multiplier and the injection-locked oscillator, due to the narrower bandwidth of the loop.

The conversion efficiency is defined as the output-to-input power ratio

$$\eta = \frac{P_o}{P_i} \tag{94}$$

This definition is most relevant for passive multipliers, where $\eta < 1$, because the output power may limit the phase noise. By contrast, the conversion efficiency makes little sense in active multipliers, where the power is provided externally, and $\eta > 1$ is allowed. Of course, the conversion efficiency is not a useful parameter to describe PLLs and injection-locked oscillators.

Designing a PLL multiplier, harmonic sampling is generally used for upper microwave frequencies, where there are no convenient dividers. However, nowadays commercial microwave prescalers work up to at least 26 GHz. Attempts to use tunnel diodes for this purpose, or parametric effects in tunnel diodes, show up in the literature from time to time but they have failed to show reliable performance, due to complexity and component tolerances in temperature.

A harmonic sampler is typically a balanced modulator that uses hot carrier diodes, driven from a pulse or comb generator with extremely high harmonic contents. A typical application for such a circuit is in spectrum analyzers, where the input frequency and the YIG oscillator can be locked together. A similar application is where a harmonic comb is generated from a 1-MHz reference, and locking can occur every 1 MHz up to several GHz. These circuits require a pre-tuning mechanism to make sure that the desired harmonic is selected, preventing false locking. This type of multiplication is used in systems where the frequency of the VCO is changed frequently and low spurious contents and high signal-to-noise ratio are required.

For higher frequency ranges, IMPATT diodes or other exotic devices can generate the necessary frequencies, and some of these multipliers are also built as injection-lock oscillators. An injection-lock oscillator can be considered as a frequency multiplier within a certain pulling range, where the oscillator locks up with the reference frequency. These are highly nonlinear phenomena, described in the literature from time to time. Explanations and mathematical models are built primarily around experimental data, not always reliable for extrapolation or new design. Low-frequency injection locking is a very convenient way of combining extremely high stability in certain types of crystal oscillators, which are being used as a reference for extremely low noise crystal oscillators operating at the same frequency.

For single-frequency applications, we find synthesizer loops using high-frequency crystal oscillators at discrete frequencies between 70 and 150 MHz, locked to a frequency standard with a narrow-band loop, so that the output PM noise depends only on the VHF oscillator, rather than on the input frequency. These loops have bandwidth of a few Hz or less, and therefore compensate only for temperature effects and aging.

The step recovery diode (SRD) generates a comb of sharp pulses at the transition between the on and the off region, and in turn a comb spectrum extending to high harmonics. Such diodes may be useful to generate frequencies up to tens of GHz, and generally require a power of +20 dBm or higher. A problem with the step recovery diodes is the relatively high phase noise, inherent in abrupt change of capacitance used to generate the sharp pulse. The typical design of a high-order multiplier may include a low-noise pre-multiplier, followed by a step recovery diode. E. Rubiola

February 6, 2025

Example 8 – Frequency multiplication. The frequency multiplier from a stable 5 MHz oscillator to 9.18 GHz for a Cs frequency standard¹ can be implemented with a \times 18 low-noise multiplier, followed by a \times 102 SRD multiplier. If the floor of the quartz oscillator is -152 dBc/Hz, a noise-free multiplier would rise such floor to -127 dBc/Hz after multiplication \times 18 (+25 dB), and to -87 dBc/Hz after further multiplication by 102 (+40 dB). The -127 dBc/Hz at output of the \times 18 multiplier sets relaxed specs for the PM noise of the \times 102 SRD multiplier at approximately -130 dBc/Hz (3 dB less than the PM noise at its input). Introducing the contribution of a low-noise first stage, the result may differ by $+1 \dots + 2$ dB. This reinforces the conclusion that a low-noise premultiplier is useful to relax the noise specs of a high-order second stage.

For fixed-frequency applications, tuned frequency multipliers with transistors work well up to the order of a GHz. Saturation should be avoided because the charge stored in the forward-biased CB junction slows down the transistor and introduces phase noise. Class-C push-push configurations can be used to improve the efficiency by reinforcing the even harmonics, and canceling the odd harmonics. Likewise, the pushpull configuration reinforces the odd harmonics, and cancels the even harmonics. The differential pair is an appealing option fo a switch because it generates a square wave with very sharp edges, making efficient use of the transistors bandwidth and keeping the transistors in the linear region. Baugh (Baugh, 1972) suggested the use of an inductor to turn the edges into sharp pulses for efficient low-noise multiplication, but this idea is seldom seen in practical applications.

Varactors (Penfield & Rafuse, 1962) are a good choice for high power efficiency and low noise because they exploit smooth nonlinear capacitance. The efficiency of an ideal reactive nonlinear reactance, i.e. loss-free varactor, is $\eta = 1$, regardless of the order of multiplication. In practice, η is limited by the small series resistance of varactors. A problem with the varactor multiplier is that it works at fixed frequency or in a rather narrow bandwidth. The design can be rather complex because it requires resonant circuits at the input and at the output, and idlers at all intermediate harmonics if N > 2. Complexity may limit N to a convenient maximum of 3–5. However, efficiency can be so high (up to $-2 \dots -1$ dB) that the varactor multiplier is an option for the output stage in small radio transmitters, instead of an active multiplier.

¹ The clock atomic transition is, by definition, 9.192631770 GHz for the unperturbed ¹³³Cs atom. An offset of the order of 1–2 Hz applies, due to the magnetic C-field needed to align the magnetic momentum of all atoms. Therefore, a high-resolution frequency of \approx 12.6 MHz must be added to the 9.18 GHz carrier. More modern schemes start from 9.2 GHz, and subtract a high-resolution 7.4 MHz frequency.

The nonlinear transmission line (NLTL) is another application of varactors suitable to high-order frequency multiplication (Ikezi, Woitowicz, Waltz, & Baker, 1988). It exhibits low phase noise, and relatively wide input bandwidth (one octave). Such device is a ladder LC network with series inductors, and the parallel capacitors replaced with varactors. The mechanism derives from studies on soliton waves. The input signal is turned into a pulse comb, and progressively shrunk and sharpened as the signal propagates through the line. Commercial devices are available, delivering multiplication up to \times 100 at a maximum output frequency up to tens of GHz. The typical input power is of +20 dBm, and in some cases up to +27 dBm. The loss is of 10-40 dB depending on the order of multiplication, on power and on other parameters. These components are often difficult to use, and the phase noise is highly dependent on the input power (Francois, Calosso, Danet, & Boudot, 2014). Figure 37 shows the phase noise of a pair of multiplication chains from 100 MHz to 9.4 GHz implemented with nonlinear delay lines. This chain is a part of a miniature atomic clock. The multiplication \times 94, inherently, increases the phase noise by 39.5 dB (20 log₁₀(94) = 39.5). So, a floor of approximately –120 dBc/Hz at 9.4 GHz output is equivalent to –160 dBc/Hz at the 100 MHz input. For frequency synthesis, the critical region is beyond about 100 Hz. A flicker PM of -80 dBc/Hz at 1 Hz is extremely low for a 10 GHz signal, and exceeds practical needs. For comparison, the PM noise of the best 10 GHz cryogenic oscillators is of the order of -90 dBc/Hz, which gives a frequency stability of 10^{-15} (ADEV at 1 s).



Figure 37 – Phase noise of a pair of multiplication chains from 100 MHz to 9.4 GHz (Boudot, Guerandel, & de Clercq, Simple-Design Low-Noise NLTL-Based Frequency Synthesizers for a CPT Cs Clock, 2009). The noise is shown for two different values of the input power between +16 dBm and +17 dBm. The phase noise spectrum is © 2009 IEEE, reprinted with permission from R. Boudot et al, "Simple-Design Low-Noise NLTL-Based Frequency Synthesizers for a CPT Cs Clock," *IEEE Transact. Instrum. Meas.*, vol. 58, no. 10, pp. 3659-3665, October 2009 (Boudot, Guerandel, & de Clercq, Simple-Design Low-Noise NLTL-Based Frequency Synthesizers for a CPT Cs Clock, 2009).

Schottky rectifiers are an appealing option for low-order multiplication (usually 2–3, but up to 7) because low-noise packaged components are available, requiring only an external filter. Figure 38 shows some examples of such multiplier. Theoretical efficiency is limited to

$$\eta \propto 1/N^2 \tag{95}$$

but practical efficiency is lower. However, the low efficiency is partially compensated by the $\times N^2$ increase in PM noise, inherent in the frequency multiplication. In the end, these multipliers may be a good option as the first stages of high-order multiplication. An example will be provided.



Figure 38 – Two examples of Schottky diode multiplier. With the lower scheme, frequency doubling (N = 2) is preferred.

The Schottky diode multiplier deserves attention for special applications, where the thermal and long-term stability of the phase is of paramount importance (timekeeping for space application and for radio navigation). The multiplier can be implemented using only \times 2 stages, frequency converters, low-pass filters and notch filters. The point is that bandpass filters are to be avoided because they suffer from a phase drift determined by the drift of the internal components multiplied by the quality factor *Q*. Oppositely, low-pass filters and notch filters are immune from this phenomenon if the main signal falls in a region where the frequency response is flat. Accordingly, a multiplier \times 4 from 100 MHz can be implemented with two \times 2 stages followed by a low-pass at 250 MHz and 500 MHz. These filters are stable because the cutoff frequency of each is far from the carrier frequency. A multiplier \times 5 starts from the same \times 4, followed by a double balanced mixer which adds 100 MHz. Since the mixer delivers 400 ± 100 MHz as the main products, a notch removes the unwanted 300 MHz, and a 550 MHz low-pass cleans the output. The notch filter may suffer from thermal effects at 300 MHz, yielding minor changes in the spur rejection, but the phase of the 500 MHz signal is stable.

Example 9 – Frequency Multiplication. We analyze the frequency doubling of a lownoise 100 MHz OCXO using a Mini Circuits LK3000+ frequency doubler. The OCXO has a PM noise floor of -175 dBc/Hz and output power P = +14 dBm. The doubler has a loss of 10.5 dB. It is followed by a 250 MHz low-pass filter (0.5 dB loss at 200 MHz) and an amplifier (noise factor F = 2 dB). Let's evaluate the phase noise.

First, the thermal noise at the oscillator output is kT/P = -174 - 14 = -188dBrad²/Hz, thus -191 dBc/Hz. This is 16 dB lower than the oscillator floor, thus guessing 2 dB noise factor for the multiplier will not change the result. *Second*, the power at the amplifier input is +3 dBm (+14 - 10.5 - 0.5 = 3). Hence, the PM noise of the amplifier is FkT/P = +2 - 174 - 3 = -176 dBrad²/Hz, thus -179 dBc/Hz. *Third*, the oscillator noise scaled up to 200 MHz is -169 dBc/Hz (-175 + 6 = -169), which is 10 dB higher than the amplifier noise. The *conclusion* is that the noise is set by the $S_{\varphi o}(f) = N^2 S_{\varphi i}(f)$ rule, and that the overall noise of the multiplier is negligible. Should we want *further multiplication*, we start from -169 dBc/Hz instead of -175 dBc/Hz. Therefore, the second stage will be more tolerant to the noise of the components, and it will be easier to keep with the × N^2 law.

For extremely high frequencies, hundreds of GHz or some THz, it is necessary to start from a clean microwave oscillator, optionally locked to a stable HF/VHF reference. Conversely, the direct multiplication of a HF oscillator will probably fail because of an intrinsic property of frequency multiplication. Phase modulation is ruled by the Bessel $J(\beta)$ functions, where β is the modulation index. Thus, J_0 for the carrier, J_1 for the firstorder sidebands, J_2 for the second-order sidebands, etc. The $\times N^2$ scaling rule is an approximation which holds for small angle modulation, where only carrier and firstorder sidebands are considered. Since the total RF power is independent of the modulation index, energy conservation requires that the sideband power comes at expenses of the carrier. When the modulation index approaches 2.4, $J_0(\beta)$ nulls, and the carrier sinks abruptly in the noise pedestal. Unlike the simple case of sinusoidal modulation, where the carrier re-appears at higher modulation index following the oscillating behavior of $J_0(\beta)$, the carrier is lost because of the statistical nature of the many spectral component which constitute the random phase modulation. This phenomenon, called *carrier* collapse, challenged the early attempts to design high frequency synthesizers starting from the stable 5 or 10 MHz OCXOs (Stewart, 1954) (Walls & De Marchi, RF Spectrum of a Signal after Frequency Multiplication, Measurement and Comparison with a Simple Calculation, 1975).

Two technologies are in competition for THz frequency multipliers, the Heterostructure Barier Varactor, and the Schottky diode (Maestrini, et al., 2010). The former is suited for the generation of odd harmonics due to internal symmetry. The latter is the simplest devices, and indeed probably the best for high efficiency. For these reasons it will probably be the preferred option for future applications. Anyway, the THz region is beyond our scope, and the reader should refer to the literature.

4.4 Direct digital synthesizer (DDS)

The DDS is such an important block in modern frequency synthesis that deserves special attention. This Section describes the general principles, and the phase noise of commercial components. Most of the material is based on our earlier article (Calosso, Gruson, & Rubiola, Phase Noise in DDS, 2012). The reader may learn more about the DDS from several references listed in the Suggested Readings at the end of this Chapter.

4.4.1 Theory of operation

The principle of operation follows immediately from the block diagram shown on Figure 39. The register is a *m*-bit D-type flip-flop called *phase accumulator*. The accumulator content at the discrete time *k* is the integer number x_k , which takes a value from 0 to $2^m - 1$. At each clock cycle, *x* is incremented by *N* modulo 2^m . This means that, when *x* reaches or exceeds 2^m , the overflow is ignored and counting goes on from the reminder. In formula

$$x_{k+1} = (x_k + N) \mod 2^m$$
(96)

The modulo- 2^m register is the hardware implementation of a finite field, which we find most convenient to represent as a set of 2^m points equally spaced on the circle $z = \exp(j2\pi x/2^m)$ in the complex plane (Figure 40). At each point x_k is associated the complex number z_k that has absolute value equal to one and phase $\vartheta_k = \arg(z_k) = 2\pi x_k/2^m$. The accumulator content x describes a discrete sawtooth waveform, which is converted into a sinusoidal signal by the Look-Up Table (LUT) and the digital-to-analog converter that follows.

For the layman, the complex-plane representation of the phase accumulator is similar to a watch dial, with the trivial difference that our "dial" has 2^m points instead of 60 seconds, and the hand jumps forward (counterclockwise) by *N* points at each tick of the clock signal, starting from the origin at "3 o'clock." Far beyond our concerns, the circular representation of a finite field is a serious branch of number theory, which has roots in the ancient Greek problem of the cyclotomy, that is, dividing the circle into a given number of equal angles and constructing regular polygons.

The output frequency is

$$f_o = \frac{N}{D} f_{ck}, \quad D = 2^m \tag{97}$$

By changing $N \rightarrow N + 1$, we find that the frequency resolution is given by

$$\Delta f_o = \frac{1}{D} f_{ck} \tag{98}$$



Figure 39 – Principle of operation of the Direct Digital Synthesizer.



Figure 40 – Complex-plane representation of the phase accumulator operation.

The term Numerically Controlled Oscillator (NCO) was originally used instead of DDS. In current jargon, the term NCO refers to the scheme of Figure 39, without the output DAC. The NCO is found as a useful building block in FPGAs, and in advanced

digital components. Oddly, some recent DACs have a built-in NCO. This is the case of the AD9144 (quad DAC, 2.8 GS/s, 16 bit, with internal 48-bit NCO) and of the AD9161/9162 DAC (6 GHz RF synthesis, 16 bit, 11 bit ENOB, internal 48-bit NCO). Modern DDS chips integrate complex functions like modulation, sweep, amplitude and phase control etc., which may not be implemented in the DACs with embedded NCO.

Our graphical and mathematical introduction describes a DDS free from noise and spurs. Practical DDSs have a phase accumulator with m = 24 to 48 bits, an output DAC with 10 to 16 bits, and operate at clock frequency up to a few GHz. Noise and spurs are a particularly complex topic. A simplified digression is given below.

4.4.2 Signal to quantization ratio (SQR)

The quantization noise is easy to derive using the methods given in the seminal article (Bennett, 27). A sinusoidal current swinging over the full-scale range I_{FSR} flowing through a resistor R results in a power

$$P = \frac{1}{8}RI_{FSR}^2 \tag{99}$$

The quantization step of a *n*-bit converter is $I_q = I_{FSR}/2^n$. Assuming that the quantization results in a random error with rectangular probability function uniformly distributed between $\pm I_q/2$, the associated noise power is

$$\sigma_q^2 = \frac{1}{12} R I_q^2 = \frac{1}{12} \frac{R I_{FSR}^2}{2^{2n}}$$
(100)

The Signal to Quantization Ratio SQR = P/σ_q^2 is given by

$$SQR = \frac{3}{2} 2^{2n}$$
(101)

or

$$SQR = 1.76 + 6.02 n dBc$$
 (102)

Using a fraction a < 1 of the full-scale range, the SQR decreases by a factor a^2 . The hypothesis that the quantization is random in amplitude fits well the observation. However, the SQR relates only to the total power of the quantization error, not to the spectrum. Because the quantization applied to deterministic signals is not random, a fraction of the quantization noise may be organized in harmonic distortion and spurs.

Analog components inside the DAC and at the output contribute a term σ_a^2 . Thus, the total noise is $\sigma^2 = \sigma_q^2 + \sigma_a^2$. Interestingly, the noise σ_a^2 of actual components is close to the limits of the technology, and also close to fundamental limits. By contrast, the quantization noise σ_q^2 can be reduced by increasing the number of bits, at least within certain limits. Increasing the number of bits has moderate impact on complexity and cost. The critical number n_c of bits, where $\sigma_q^2 = \sigma_a^2$, is of paramount importance in the

design of converters, and also in designing with converters. It is wise to have $n > n_c$, so that the total noise is chiefly limited by σ_a^2 . With 2-4 bits in excess, the quantization noise is 12-24 dB smaller than the analog noise. This choice also results in significantly reduced distortion and spurs because these artifacts originate from the non-random nature of the quantization.

The ENOB (Equivalent Number Of Bits) results from an attempt of simplification, describing total noise $\sigma^2 = \sigma_q^2 + \sigma_a^2$ with a single parameter. So, the engineer uses the formula

$$\sigma^2 = \frac{1}{12} \frac{R I_{FSR}^2}{2^2 \text{ ENOB}}$$
(103)

and that's it. The catch is that the ENOB hides the difference between analog noise and quantization noise, and their statistical proprieties. The ENOB is suitable to signals with sufficiently good random characteristics, like audio communication. Conversely, synthesizers deliver highly coherent signals. In this case, the ENOB fails to describe the quality of the output because the quantization noise yields artifacts, harmonics and spurs, while the analog noise does not.

4.4.3 Truncation spurs

The output DAC has a finite number *n* of bits, which in turn determines the number *p* of address bits that gives distinct values at the LUT output. Higher number of bits results in duplicated codes at the LUT output, and in no improvement. The value p = n + 2 or p = n + 3 is often found. The full *m*-bit word of the accumulator represents the exact instantaneous phase, given by $x_{k+1} = x_k + N \mod 2^m$. However, the voltage delivered to the output is determined only by the higher *p* bits. In other words, the accumulator defines 2^m possible phases, or states (all the dots on the circle of Figure 40), but only 2^p distinct phases (the thick dots on the circle) can be delivered to the output. The resolution $\Delta x = 2^{m-p}$ results in a round off phase error distributed from 0 to $\Delta \varphi$, where $\Delta \varphi = 2\pi/2^{m-p}$ rad. Such error is of pseudo-random nature because it results from the fully deterministic operation of the accumulator (the DAC analog noise is not counted here). Pseudo-randomness, as opposite as true randomness, originates spurs. This can be understood by analyzing carefully the Accumulator and the Frequency Control Word (Figure 41).

m	output phase	Accumulator	phase error	
MSE	AH		AL	LSB
	× _Н		×L	
m–1	I	m—p m—p—1		0
	p bits		b = m–p bits	
		m bits		



Figure 41 – Accumulator and Frequency Word.

For easier interpretation, it is useful to divide the accumulator into two sections, AH and AL (Accumulator High and Low), thus $x = x_H + x_L$. AH has the size of p bits, and generates the output phase. AL has the size of m - p bits, and generates the truncation spurs. Similarly, we divide the Frequency Word (FW) in two parts, FWH (the higher p bits) and FWL (the lower m - p bits. Notice that the MSB of FW must be zero, otherwise f_o exceeds $f_{ck}/2$ and the output frequency results from aliasing.

An important case is FWH containing an odd number, and FWL containing all zeros

FWH	FWL			
0xxxxx1	000000000000000000000000000000000000000			

Accordingly, the value x_k is incremented in steps equal to Δx starting from zero. Thus, at each clock cycle x_k jumps to the next thick dots on the circle, and AL contains all zeros. There is no phase truncation, and no truncation spurs. The case of FWH containing at least one "1," and FWL containing all zeros is an obvious extension. Now x_k walks through the thick dots in steps multiple of Δx , the step being determined by the position of the rightmost "1," and there is no truncation.

Another important case is FWL containing all zeros, but the MSB is equal to one

FWH FWL

0xxxxxx 100000000000000

The accumulator content x increments in steps exactly equal to $\Delta x/2$, hence the truncation error is a square wave of peak-to-peak amplitude $\Delta x/2$. This is the condition of maximum spurs. The maximum spur-to-carrier ratio is 2^{-2p} , that is, -6.02 p dBc/Hz.

The accumulator content x is periodic. It starts from zero and it first returns to zero after a number of clock cycles called Grand Repetition Period² (GRP). The maximum GRP is equal to 2^m clock periods, thus $2^m/f_{ck}$ seconds, and occurs when N is an odd number (the FW has the LSB equal one). The GRP can be rather long. For example, with $f_{ck} = 1$ GHz and m = 48 bits, the GRP is $2^{48} = 2.81 \times 10^{14}$ clock periods, thus 2.81×10^5 s (3.25 days). If the FW contains r trailing zeros

the GRP is given by

$$GRP = 2^{m-r} \tag{104}$$

This is rather obvious because x is incremented in steps odd multiples of 2^r , thus the lowest r bits of the accumulator will always be zero. We can see this as a smaller DDS, where the accumulator has m - r bits.

The simplest way to understand the truncation spurs is to interpret AL as an accumulator, which behaves in the same way as the full accumulator (Figure 41). In fact, the content x_L increments in steps and overflows, exactly as x does. The trivial difference, that the overflow of AL goes to AH while the overflow of AH is discarded, is not relevant here.

Since AL has a number b = m - p of bits, its operation is described by the equation $x_{L,k+1} = x_{L,k} + N_E \mod 2^b$ (105)

where the increment

$$N_E = \begin{cases} N_L & \text{for } N_L < 2^{b-1} \\ 2^b - N_L & \text{for } N_L \ge 2^{b-1} \end{cases}$$
(106)

is the content of the Equivalent Frequency Word (EFW). N_E results from the following reasoning. If the MSB of the FWL is zero ($N_L < 2^{b-1}$), N_L is a valid frequency word, and $N_E = N_L$. By contrast, if the MSB is equal to one ($N_L \ge 2^{b-1}$), the frequency exceeds half the clock clock frequency. The frequency observed is in the first Nyquist zone, determined by $N_E = 2^b - N_L$.

² The GRP is a period in a strict sense. However the term GRR (Grand Repetition Rate) is sometimes used, which is misleading because the word "rate" refers to a frequency.

The accumulator content x_L is periodic. It starts from zero and it first returns to zero when x falls on one of the thick dots in the circle. This occurs after a number of clock cycles called Truncation Grand Period

$$TGP = 2^{b-r} \tag{107}$$

The value x_L describes a sawtooth (Figure 42) whose period is the Sawtooth Repetition Period

$$SRP = \frac{2^b}{N_E}$$
(108)

The spectrum of the sawtooth waveform contains all the harmonics multiple of 1/SRP, with amplitude proportional to the reciprocal of the order. The sawtooth is sampled, so all the harmonics exceeding $f_s/2$ are remapped to the first Nyquist zone. The number of samples of x_L in the grand period is equal to TGP. Hence, the discrete transform has $TGP/2 = 2^{b-r-1}$ frequencies. This completely describes the spectrum of the truncation spurs. Unfortunately, the theory is of limited usefulness because the spurs spectrum depends on N_L , and changes abruptly changing the frequency word.



Figure 42 – Phase error due to the truncation of the accumulator content.

Modified and updated, from Chapter 2 of U.L. Rohde, E. Rubiola, J.C. Whitaker, *Microwave and Wireless Synthesizers* 2nd Ed., Wiley 2021. **Draft version, limited circulation for review only.** Page 95/219 Our approach to the truncation spurs derives from Reference (Gentile, 1999). Among the suggested readings, we recommend a technical tutorial published by Analog devices, an article by Torosyan, the two seminal articles by Nicholas and Samueli, and a book by Widrow and Kollar. The latter only to the reader willing to tackle an extensive mathematical treatise on the quantization noise.

Other types of spurs arise from harmonic distortion in the output DAC and in the analog electronics which follows. Of course, aliasing is always present, so all the harmonics exceeding $f_s/2$ are remapped to the first Nyquist zone.



Figure 43 – Simplified noise model of a synthesizer. Reproduced with permission from C. E. Calosso, Y. Gruson, and E. Rubiola, Phase noise and amplitude noise in DDS, Proc. 2012 International Frequency Control Symposium pp. 777-782 (Calosso, Gruson, & Rubiola, Phase Noise in DDS, 2012), © 2012 IEEE.

4.4.4 Phase noise

Figure 43 shows a simplified noise model of a DDS, and highlighs the two basic mechanisms. The noise of the input clock is scaled down according to

$$S_{\varphi o}(f) = \left(\frac{N}{D}\right)^2 S_{\varphi ck}(f) \tag{109}$$

because the noise-free synthesizer transfer the time fluctuation x(t) from the input to the output. Unlike in frequency dividers, aliasing has no practical effect on noise scaling because the output DAC samples always at the clock frequency, regardless of the output frequency. The same rule applies to the noise of the input stage, and to the noise of the clock distribution as well, both of the time type. This is raher obvious because the input frequency has a fixed value, thus the time fluctuation at the input and all long the clock distribution does not depend on the output frequency. The output stage adds its own noise. The latter is at first approximation of the phase type, defined by the phasefluctuation spectrum being independent of frequency. This behavior is similar to that of amplifiers and other analog components. In conclusion, the phase noise follows the $(N/D)^2$ law at high N (high output frequency), and hits the limit set by the output stage at low N. The critical N, where output stage limits, depends on the noise type.

The above concepts need to be analyzed more in detail. The phase noise of a real DDS includes

- The quantization noise of the output DAC
- The analog noise at the DAC output, and of the stages which follow,
- The time fluctuation of the DAC sampling, with respect to the clock input.

We have already seen that the signal power is $P = RI_{FSR}^2/8$, and that the noise power is

$$\sigma_q^2 = \frac{1}{12} \frac{R I_{FSR}^2}{2^{2n}} \tag{110}$$

The noise bandwidth *B* is half the sampling frequency, thus $B = f_{ck}/2$. Starting with the provisional assumption that the spectrum is white, the PSD is equal to σ^2/B . The phase noise PSD results from σ_a^2/P

$$S_{\varphi q}(f) = \frac{4}{3} \frac{1}{2^{2n} f_{ck}}$$
(111)

The hypothesis of white noise is generally untrue in frequency synthesis because of truncation and nonlinearity. In turn, the spurs sink power from noise, or from some portion of, and the noise floor can be lower than in (111). Figure 44 shows an example.



Figure 44 – A simulation shows the effect of spurs on the noise floor. Adapted version, reproduced with permission from C. E. Calosso, Y. Gruson, and E. Rubiola, Phase noise and amplitude noise in DDS, Proc. 2012 International Frequency Control Symposium pp. 777-782 (Calosso, Gruson, & Rubiola, Phase Noise in DDS, 2012), © 2012 IEEE.

Flicker noise is a separate issue. It originates from the DAC output and from the analog electronics which follows. This type of noise is generally of the phase type, i.e., $S_{\varphi}(f)$ is rather constant versus the output frequency, and described by the experimental parameter b_{-1}

$$S_{\varphi}(f) = \frac{\mathsf{b}_{-1}}{f} \qquad \text{(flicker)} \tag{112}$$

The jitter of the output sampling must be added to the model. It originates in the DAC switching mechanism and in the DDS internal clock distribution from the input to the DAC. This noise is of the time type, thus the time fluctuation $S_x(f)$ is independent of f_0 . Such fluctuation appears on the phase noise PSD as

$$S_{\varphi \,\text{jit}}(f) = 4\pi^2 f_0^2 S_{\text{x jit}}(f) \tag{113}$$

Having on hand the phase noise spectrum taken at different values of f_o , we can identify the parameters of the phase-type and of the time-type noise processes, and use them to predict the phase noise in the general case.

4.4.5 Examples

The phase noise measured on commercial DDSs confirms our theoretical models, and also shows some unpredicted facts, inevitable in practical implementations.



Figure 45 – Phase noise of the AD9854 DDS measured at different output frequencies. Reproduced with permission from Claudio E. Calosso, Yannick Gruson, and Enrico Rubiola, Phase noise and amplitude noise in DDS, Proc. 2012 International Frequency Control Symposium pp. 777-782 (Calosso, Gruson, & Rubiola, Phase Noise in DDS, 2012), © 2012 IEEE.

Figure 45 shows the phase noise of an AD9854 DDS driven with 180 MHz clock frequency, measured at various output frequencies. The phase noise results from a differential measurement, where the noise of the 180 MHz source cancels because the same source is also used as the reference for the phase meter.

First, we observe that the flicker b_{-1} scales down as $1/f_o^2$, that is, 6 dB per factor of two. This is the signature of the time-type noise, where $S_x(f) = S_{\varphi}(f)/(2\pi f_o)^2$ is independent of f_0 , as shown on Table 7. The phase noise leaves the $1/f_o$ law only at the lowest values of f_o , where some phase-type noise shows up. The same happens with white noise, yet the phase-type noise starts to be visible at higher f_o .

<i>f</i> ₀ ,	Scale	b_1,	$\sqrt{b_{-1}}$,	k_1,	$\sqrt{k_{-1}}$,
MHz	factor, dB	dBrad ²	rad	dBs ²	S
22.5	0 (ref)	-114.5	1.9×10^{-6}	-277.5	1.33×10^{-14}
11.25	-6	-120.5	9.4×10^{-7}	-277.5	1.34×10^{-14}
5.63	-12	-126	5.0×10^{-7}	-277.0	1.42×10^{-14}
2.81	-18	-130	3.2×10^{-7}	-274.9	1.79×10^{-14}
1.41	-24	-133.5	2.1×10^{-7}	-272.4	2.39×10^{-14}
	$(f_o/f_{\rm ref})^2$	from the		$b_{-1}/4\pi^2 f_0^2$	
		plot			

Table 7 – Flicker noise of the AD9854 DDS. Values are taken from Figure 45.

In the region between 300 Hz and 30 kHz, the phase noise leaves the polynomial law, being higher than the 1/f asymptote. Other experiments show that this is due to the residual noise in power-supply lines. So, this behavior may be specific to our measurement rather than a property of the device. At very low Fourier frequencies, below 1 Hz, the phase noise follows an unexpected $1/f^3$ slope (-30 dB/decade). This is due to thermal effects. The evidence is that stabilizing the chip temperature with different heat sinks and thermal masses clamped onto the chip surface shifts the $1/f^3$ corner to lower frequencies. In more recent DDSs, a similar thermal effect appears at lower Fourier frequencies, or has not be detected. This makes us think that the thermal design in the newest components has been significantly improved.

In our measurements, we combined the two outputs of the DAC using a balun (transformer) to get the highest output power, and we used a low-noise RF amplifier at the balun output. Out of experience, we have identified two weak designs that result in unnecessarily higher white PM noise. The first is the use of single output of the DAC, and the second is the use of a high-speed operational amplifier at the DDS output, instead of the RF amplifier. Trite calculations using the noise parameters available in the data sheet give full account.

Most DDSs feature a digital amplitude control. In noise-critical application, the DDS should be operated close to the full-scale range, using this control only for fine tuning. The reason is that the amplitude control generally scales down the values at the LUT output, thus the DAC uses only a fraction of its dynamic range. At half amplitude one bit is lost, and the phase noise gets 6 dB worse. A practical example is seen on Figure 46. Of course, amplitude control acting on the reference of the DAC mitigates this problem.



Figure 46 – Phase noise of the AD9854 DDS measured at different output frequencies. Reproduced with permission from Claudio E. Calosso, Yannick Gruson, and Enrico Rubiola, Phase noise and amplitude noise in DDS, Proc. 2012 International Frequency Control Symposium pp. 777-782 (Calosso, Gruson, & Rubiola, Phase Noise in DDS, 2012), © 2012 IEEE.

Example 10 – Hacking the AD9912 DDS. The phase noise spectrum of the AD9912 is shown on Figure 47. We try to understand the noise parameters which describe this device. From this Figure, we calculate the flicker coefficients shown on Table 8. In this table, we use the coefficients of the polynomial law, b_i for $S_{\varphi}(f)$, and k_i for $S_x(f)$. We take 150 MHz as the reference. From 150 MHz to 50 MHz, $S_{\varphi}(f)$ scales down in exact agreement to the $(f_o/f_{ref})^2$ law, and k_{-1} is the same. This is the signature of the *time-type noise*. Conversely, at 10 MHz the flicker is 7 dB higher than the $(f_o/f_{ref})^2$ law. This happens because the scaled-down noise hits the noise of the output stage, which is *phase-type noise*. At 10 MHz, we calculate

 $[b_{-1}]_{\text{phase type}} = [b_{-1}]_{\text{total}} - [b_{-1}]_{\text{time type}}$

The time-type noise scaled down to 10 MHz is $[b_{-1}]_{\text{time type}} 1.78 \times 10^{-12} \text{ rad}^2$. Thus,

$$[b_{-1}]_{\text{phase type}} = 8.91 \times 10^{-12} - 1.78 \times 10^{-12} = 7.13 \times 10^{-12} \text{ rad}^2$$

In conclusion, the flicker of the DDS is given by

$$b_{-1} = 4.5 \times 10^{-28} (2\pi f_o)^2 + 7.13 \times 10^{-12} \text{ rad}^2$$

The evaluation of white PM noise is not trusted because we have only two values, thus we have no evidence of time-type noise at f_o . With this reservation, we can solve the system

$$[b_0]_{\text{phase type}} + (2\pi f_1)^2 [k_0]_{\text{time type}} = 10^{-14.5} \qquad f_1 = 50 \text{ MHz}$$
$$[b_0]_{\text{phase type}} + (2\pi f_2)^2 [k_0]_{\text{time type}} = 10^{-15.4} \qquad f_2 = 10 \text{ MHz}$$

which gives $[b_0]_{\text{phase type}} = 2.83 \times 10^{-16}$ and $[k_0]_{\text{time type}} = 2.83 \times 10^{-16}$, and an overall phase noise $b_0 = 2.92 \times 10^{-32} (2\pi f_o)^2 + 2.83 \times 10^{-16}$ rad²/Hz.



Figure 47 – Example of phase noise of a DDS (the units rad² and rad²/Hz in the polynomial-law coefficients are omitted). The PM noise spectrum is from the AD9912 data sheet, © Analog Devices, reproduced with permission. Graphical adaptation and comments are ours.

<i>f</i> ₀ ,	Scale	b_1,	$\sqrt{b_{-1}}$,	k_1,	$\sqrt{k_{-1}}$,
MHz	factor, dB	dBrad ²	rad	dBs ²	S
150	0 (ref)	-94	2×10^{-5}	-273.5	2.11×10^{-14}
50	-9.5	-103.5	6.68×10^{-6}	-273.4	2.14×10^{-14}
10	-23.5	-110.5	9.44×10^{-6}	-266.5	4.73×10^{-14}
	$(f/f_{c})^{2}$	from the		h $/4\pi^2 f^2$	
	() o/ Jref	plot		$D_{-1}/4\pi J_0$	

Table 8 – Flicker noise of the AD9912 DDS. Values are taken from Figure 47 (data sheet).

A small number of samples were measured at the Italian Institute of Metrology INRiM. The flicker noise of one of these samples is shown on Table 9. The value measured is some 12 dB lower than the spectrum found in the datasheet. Using the data of Table 9, the flicker phase noise is

$$b_{-1} = 2.5 \times 10^{-29} (2\pi f_o)^2 + 4.4 \times 10^{-14} \text{ rad}^2$$

We have little doubt about the measurement made in a laboratory of primary metrology, repeated several times in well controlled conditions, and reproducible over a small number of samples. Conversely, the datasheet says very little about how the spectrum is measured. Whether an unfortunate error was made in the datasheet, or the difference is due to samples from different batches, we cannot know. Nonetheless, the datasheet reports a conservative value, and a lower noise measured on a sample came as a good surprise. The general practitioner relies on data sheets, and has seldom the time and equipment for independent measurements.

As we have seen with amplifiers, the documentation about flicker noise is usually rather poor, as compared to our wishes. Occasional difficulties and frustrations, like in this example, are a part of the message we address to the reader.

<i>f</i> ₀ ,	Scale	b_1,	$\sqrt{b_{-1}}$,	k_1,	$\sqrt{k_{-1}}$,
MHz	factor, dB	dBrad ²	rad	dBs ²	S
100	0 (ref)	-110	-110 1.9×10^{-6} -286.0		5.03×10^{-15}
50	-6	-116	9.4×10^{-7}	-285.9	5.04×10^{-15}
25	-12	-122	5.0×10^{-7}	-285.9	5.06×10^{-15}
12.5	-18	-128	3.2×10^{-7}	-285.9	5.07×10^{-15}
6.25	-24	-131.5	2.1×10^{-7}	-283.4	6.78×10^{-15}
3.125	-30	-131.5	2.1×10^{-7}	-277.4	1.36×10^{-14}
	$(f_{\rm c}/f_{\rm ref})^2$	from the		$b_{-1}/4\pi^2 f_0^2$	
	(Joj Jref)	plot			

Table 9 – Flicker noise of the AD9912 DDS, measured at the Italian Institute of Metrology INRiM. Courtesy of Claudio E. Calosso, INRiM.

Example 11 – An application of the AD9912 DDS. We synthesize a high resolution 1.04–1.06 GHz signal using the scheme of Figure 48. The DDS has m = 48 bits, thus the resolution is $\Delta f_o = f_{ck}/2^{48} = 7.1 \,\mu$ Hz. Considering the oscillator as an external source, the noise of this synthesizer originates in the mixer and in the DDS. We see on Figure 47 that at 50 MHz the phase noise of the DDS is significantly higher than that of a mixer, and from the previous example we expect that the change in the DDS phase noise is a matter of 2 dB in a ±10 MHz region around 50 MHz. The expected phase noise is therefore

$$S_{\varphi}(f) = \frac{4.5 \times 10^{-11}}{f} + 3.2 \times 10^{-15} \text{ rad}^2/\text{Hz}$$

An output frequency in the same range can be obtained multiplying the DDS output by M = 21. In this case, the benefit of beating is lost. The resolution is $\Delta f_o = M f_{ck}/2^{48} = 149 \,\mu\text{Hz}$, and the phase noise by M^2 , thus $S_{\varphi}(f) = 2 \times 10^{-8}/f + 1.4 \times 10^{-12} \,\text{rad}^2/\text{Hz}$.



Figure 48 – Example of application of a DDS in a dedicated synthesizer.

4.5 Phase detectors

The designer has different options for the phase detector, the Double Balanced Mixer (DBM), transistor or FET mixers, the XOR gate, the SR flip-flop, and the Phase-Frequency Detector (PFD), to mention the most known types. Relevant parameters are frequency range, residual noise, and noise immunity. The detector noise is the most important parameter for low-noise PLLs because phase comparison generally occurs at a suitable low frequency, and the detector noise is scaled up according to the N^2 law.

The SR flip-flop and the PFD have internal memory which stores the last phase transition. This makes these detectors unsuitable to noisy signals, where multiple bounces occur at the lock point. Conversely, mixers and XOR suffer very little from bouncing because the bounces are averaged out by the loop. The consequence is just a reduction in the phase-to-voltage gain K_d . This discussion is useful for the detection of small signal in noise, while all signals inside a synthesizer are clean enough for the noise immunity to have little or no importance, as compared to the residual noise of the detector.

The DBM is considered the lowest-noise detector, and for this reason it is widely used in the measurement of phase noise.

The PFD replaces the SR flip-flop in virtually all applications because it provides a valid output signal also in unlock conditions, and because it mitigates the issue of metastable behavior when the S and R signals are nearly simultaneous. The analysis of noise in the DBM is postponed to Section 2.6, where we study the measurement of PM noise. Here, we analyze the PFD restricting or attention to phase noise.

4.5.1 Noise in the phase-frequency detector

The PFD is a digital circuit which receives clean digital signals at its inputs, thus we expect that the *time-type* noise is dominant. Related to this, the parameter commonly

found in the data sheets is the Figure Of Merit (FOM), which is the noise contribution of the PFD normalized to 1 Hz, and most often expressed in dBc/Hz.

The usual definition of the FOM refers to white PM noise in a PLL with a $\div N$ divider in the feedback path. Denoting with f_c the comparator input frequency and with f_{VCO} the VCO frequency, N results from $N = f_{VCO}/f_c$. The formula found in application notes is

$$FOM_{dBc/Hz^2} = [L(f)]_{dBc/Hz} + 10\log_{10}(f_c) - 20\log_{10}(f_{VCO})$$
(114)

A value of – 220 to -230 dBc/Hz^2 can be taken as the order of magnitude, but the direct comparison between components is not easy because the FOM changes with technology and frequency range. Notice that some commercial PFDs include a prescaler, and such prescaler may also work in fractional-N mode.

Removing the dB notation from (114), and using $N = f_{VCO}/f_c$ in different ways, the white phase noise at the output can be written in the following equivalent forms

$$L(f) = \text{FOM } N^2 f_c \tag{115}$$

$$L(f) = \text{FOM } Nf_{VCO} \tag{116}$$

$$L(f) = \text{FOM } \frac{f_{VCO}^2}{f_c} \tag{117}$$

The first two forms focus on the two obvious critical points, comparator and VCO. The third form is subtler because it reveals the aliasing at the input. References (Thompson & Brennan, 2003) (Homayoun & Razavi, 2013) provide insight and useful details.

Gao et al (Gao & Klumperink, 2009) give an alternate definition of the FOM, which differs in that the power is added, expressed in dBm. So, the FOM of a component taking 2 mW is 3 dB worse than that of an otherwise equal component requiring only 1 mW. The Gao definition is not followed by the manufacturers.

Information about flicker noise in commercial PFDs is often absent or difficult to understand, if not confusing. This reflects a quite general lack of documentation on this topic. The following formula is sometimes found in the technical literature

$$[L(f)]_{\rm dBc/Hz} = \left[\rm{FOM}_{1/f}\right]_{\rm dBc/Hz^2} + 20\log_{10}(f_{VCO}) - 10\log_{10}(f)$$
(118)

where the term $-10 \log_{10}(f)$ expresses the fact that that L(f) is proportional to 1/f. Removing the dB notation, (118) becomes

$$L(f) = f_{VCO}^2 \text{ FOM}_{1/f} \frac{1}{f}$$
(119)

Example 12 – Noise of a PLL with a PFD. We lock a 1 GHz VCO to a 100 MHz reference using a PFD which has a FOM of – 220 dBc/Hz for white noise, and of – 260 dBc for flicker, including the \div 10 internal prescaler. Accordingly, phase detection takes place

at 10 MHz. Let us evaluate the phase noise at the 1 GHz output, and also refer it to the 100 MHz input.

We prefer to convert the dBc into SI units. For white noise, we use $b_0 = FOM N f_{VCO}$, with FOM = 2 × 10⁻²², N = 10, and $f_{VCO} = 10^9$ Hz. Thus $b_0 = 2 \times 10^{-12}$ rad²/Hz, i.e., -117 dBrad²/Hz. Finally, L(f) = -120 dBc/Hz at 1 GHz, and -140 dBc/Hz at 100 MHz, after scaling as $1/N^2$.

For flicker noise, we use $b_{-1} = \text{FOM}_{\text{flicker}} f_{VCO}^2$ with $\text{FOM}_{\text{flicker}} = 2 \times 10^{-26}$ and $f_{VCO} = 10^9$ Hz. Thus, $b_{-1} = 2 \times 10^{-8}$ rad², i.e., -77 dBrad². Finally, L(f) = -80 dBc/Hz referred to 1 Hz at the 1 GHz output, and -100 dBc/Hz referred to 1 Hz at the 100 MHz input

Example 13 – Comparison of some multiplication schemes. Let us consider some options to multiply a high stability 10 MHz reference to 640 MHz by combining multiplication and cleanup PLL, which may include a divider. The phase noise of the reference and of two possible auxiliary VCOs are summarized in Table 10, and we look at the configurations shown on Figure 49. The reference has a stability $\sigma_{\rm Y} = 8.3 \times 10^{-13}$ (flicker floor). This is easily seen with the formula $\sigma_{\rm Y}^2(\tau) = 2 \ln(2) h_{-1}$ (holds for flicker of frequency), and $h_{-1} = b_{-3}/f_0^2$ (converts FM flicker from S_{φ} into $S_{\rm Y}$) seen in Section 2.2. Thus $b_{-3} = 10^{-103/10} = 5 \times 10^{-11}$, $h_{-1} = 5 \times 10^{-11}/(10^7)^2$, and $\sigma_{\rm Y}^2 = 6.95 \times 10^{-25}$

Oscillator type	Frequency	Noise types				
		RW FM	Flicker FM	White FM	Flicker PM	White PM
OCXO (reference)	10 MHz		-103		-131	-162
ОСХО	128 MHz		-67			-172
SAW	640 MHz	-47	-57			-170
		dB rad ² Hz ³	dB rad ² Hz ²	dB rad ² Hz	dB rad ²	dB rad ² /Hz

Table 10 – Data for the example 13.



Figure 49 – Schemes for frequency multiplication from 10 MHz to 640 MHz.

In order to compare the options, we first scale the PM noise of our oscillators to 640, 80, 10 and 2 MHz carrier frequency. These spectra are plotted in Figure 50. We start from the schemes A, B, and C, which do not involve phase locking. The scheme A is rather generic, and we expect that the output spectrum is that of the 10 MHz OCXO scaled up to 640 MHz. The output bandpass filter cannot be narrow enough to clean the spectrum, so the phase noise will have a floor of the order of – 120 dBrad²/Hz which spans over a wide band, hundreds of kHz. The scheme B suffers from the same problem. The scheme C makes use of quartz filters to clean up the phase noise of the 10 MHz OCXO after multiplication. The bandwidth of such filters can be of a few kHz. Besides complexity and cost, the problem is that the filters turn mechanical vibration into PM noise spurs, and temperature drift into phase drift. In most synthesizer applications, the slow phase drift due to temperature is probably not a problem. However, for special applications and timekeeping, where the phase stability is important, the scheme B probably wins because of its potentially high thermal stability. Unlike the bandpass filter, the phase of the lowpass filters is little affected by the temperature fluctuations.

The schemes D, E and F make use of phase locking. We see on Figure 50 that the preferred cutoff frequency f_c is of approximately 100 Hz, almost the same for the two VHF oscillators. After locking, the frequency flicker and the long-term stability are determined by the reference, and the two VHF oscillators are nearly equivalent. However, the 640 MHz SAW oscillator has lower white noise floor. With appropriate
design, the PLL filter can mitigate or reject the noise of the divider in the feedback path, and of the detector, at offset frequencies beyond f_c . This helps in the choice of the divider and phase detector.

The scheme D uses a 128 MHz oscillator phase locked to the 10 MHz reference. The problem is that the highest frequency at the phase detector input (the greatest common divider of 10 and 128) is 2 MHz. After scaling to that low carrier frequency, the PM noise of the 128 MHz oscillator is too low for a frequency divider to preserve it. In conclusion, the output noise will be limited by the divider and by the phase detector.

In the scheme E, the 640 MHz SAW is phase locked at 80 MHz. Phase detector and frequency divider must have a PM noise not greater than $-130 \text{ dBrad}^2/\text{Hz}$ at 130 Hz offset, which is not challenging. Comparing the schemes E and F, E is simpler because it uses a digital divider \div 8 and a multiplier \times 8, instead of a multiplier \times 64. Otherwise, the scheme F relaxes the specs for the phase detector.

E. Rubiola



Figure 50 – Phase noise of the oscillators of Table 5, scaled to different values of the carrier frequency.

4.6 Noise contribution from power supplies

We have mentioned the effect of line frequency pickup several times so far, the most direct being ripple on the dc supply voltage.

Power supplies can generally be built in one of two ways:

• Using a monolithic regulator.

• Using discrete components.

The safe approach is generally to use two cascaded regulator systems, starting with a monolithic regulator, followed by a discrete post-regulation stage.

In traditional synthesizers, it is typical to find the following voltage requirements: +5 V, \pm 12 V, +9 V, and +24 V. When using a power supply fed from a 117 or 230 V power line, generating these auxiliary voltages is fairly easy. As the 5 V probably has the highest current draw, this will be kept totally separate from the other voltages. The current consumption on the \pm 12 V is on the order of several hundred mA, and the 9 V is probably an auxiliary voltage that can be generated in a post-regulator from the +12 V.

Modern VLSI digital circuits usually require lower voltages, typically 1.8 V and 3.3 V, with rather high current. These voltages can be produced locally by a switching power supply, from a 12 V or 24 V line. To prevent spurs, the switching power supply must be carefully shielded, and powered by a separate 12-24 V line, not shared with analog electronics.

The +24 V requirement is generally of low power consumption and is required for the phase/frequency detector stages and the tuning diodes. If a dc amplifier translation stage is used following the phase/frequency comparators to drive the tuning diode, such a high voltage is necessary.

The dynamic regulation found in a monolithic regulator is typically 60 or sometimes 70 dB, which reduces the input ripple voltage to about 1 mV. This is insufficient for sensitive lines and a post-regulator of at least the same amount must be added. Here, a discrete circuit is the proper choice.

There are numerous regulators on the market, but the one with the lowest noise is probably the old National LM723. The typical output noise of this regulator is in the vicinity of a few microvolts. Figure 51 shows a regulator for extremely low noise output. It is based on the fact that the current generating PNP transistor produces much less noise than its emitter-follower equivalent.



Figure 51 – Schematic diagram of an extremely low noise output regulator based the Linear Technology LT3045 voltage regulator. Low noise is achieved by using a PNP transistor as the series regulator, with the output taken from on the collector. This solution produces much less noise than its emitter-follower equivalent. Also, this type of circuit has a much smaller voltage drop than the source follower. It operates quite well with voltage differences as low as 0.7 V. Reproduced with permission from the LT3045 data sheet.

In battery-operated synthesizers, especially if they operate from 12 V dc, it is somewhat difficult to generate the higher voltage for the tuning diodes. One of the best approaches is to use a switching dc/dc converter stage that is being driven from the reference oscillator at a rate of 10 kHz to 1 MHz. As the power consumption on the tuning line is very small, no special power transistors are required, and regulators take care of interference suppression. As these stages are being driven from a square wave generated from a regulated power supply, extremely high values of regulation can be obtained, and the tuning voltage is therefore very clean and noise-free. Attempts to generate the auxiliary voltage from asynchronous dc/dc converters have generally resulted in poor performance, and this approach is not recommended.

Switching circuits are an appealing choice because of high efficiency and potentially low ripple. The reason is the use of LC filters, made simple by the high operating frequency (generally a few hundreds of kHz). Isolation is another unique feature of switching circuits, thanks to ferrite transformers. This makes easy to break ground loops and to block the interferences conducted along power lines. The general performances of commercial modules, in terms of ripple and stability, can be improved by adding an external LC low-pass filter and a linear post regulator. Shielding is the major difficulty for lowest ripple and noise, and also to prevent interferences to other circuits. Low-noise power modules are now available, consisting of a flyback switching circuit followed by a linear regulator. The LTM8068 is an interesting example. It features 2.8-40 V input range, 1.2-18 V output range after the linear regulator, with 300 mA available current and 20 μ V rms noise and 2-kV isolation in a 1-cm² BGA package.

5 Phase noise in oscillators

In electrical engineering, the oscillator is a circuit that delivers a periodic signal, sinusoidal or square wave, with suitable purity and stability, powered by a source of energy. Different jargon terms may be encountered, like "self-oscillator" or "autonomous oscillator." By contrast, physicists often use the term "oscillator" for the "damped oscillator," which is the resonator in our terminology.

The simplest form of oscillator consists of a resonator and a sustaining amplifier in closed loop, so that the resonator sets the oscillation frequency f_0 , and the sustaining amplifier compensates for the loss of the resonator. A buffer is generally necessary, to amplify the output power and to isolate the oscillator from load perturbations. Denoting the amplifier gain with A, and the resonator transfer function with B, stationary oscillation requires that

$$AB = 1 \tag{120}$$

that is, |AB| = 1 and $\arg(AB) = 0$. This is known as the Barkhausen condition. A gain compression mechanism is necessary, which stabilizes the oscillation amplitude to a given level. Otherwise, even the smallest discrepancy from unity gain results in exponentially increasing oscillation (AB > 1) or in exponentially decaying oscillation (AB < 1).

The oscillator's internal components introduce noise and fluctuations. Other types of fluctuation originate from power supply, temperature, and other quantities. Understanding and modeling the oscillator is a complex issue because of the multi-scale time range, from the period to the long time related to aging.

We focus on the Leeson model analyzed from the modern standpoint. The original article (Leeson, 1966) proposed a quasi-linear analysis, inherently limited to additive white noise. Adding very little complexity, we introduce the perturbation and modulation analysis, which is perfectly suitable to parametric noise (Rubiola E. , Phase Noise and Frequency Stability in Oscillators, 2010). An extension to AM noise is available in (Rubiola & Brendel, A Generalization of the Leeson Effect, 2010). Our approach is probably the simplest and the most suitable to understand real oscillators because we analyze the oscillator as a system, as opposed to a detailed schematic. Of course, simplicity and generality come at expenses of accuracy. A few alternate models are available in the literature. These models are either specialized (for example, the ring oscillator), or extremely complex (for example, the Fokker-Planck equation). We provide a list in the Suggested Readings, at the end of this Chapter. For a perspective on different approaches and oscillator models, we suggest starting from (Pankratz & Sanchez-Sinencio, 2013) and (Poddar & Rohde, 2013).

5.1 Modern view of the Leeson model

The oscillator is seen as a system consisting of an amplifier and a resonator in closed loop, shown on Figure 52 (top). The amplifier is operated at the compression point, where it stabilizes the amplitude. The phase noise is modeled as a phase modulator at the amplifier input, driven by a generator which introduces white noise, flicker noise, spurs, etc. The resonator is a bandpass filter exhibiting a sharp response which sets the oscillation frequency. The loop has an input for the virtual signal that starts the oscillation. Actual oscillation starts from noise or from the switch-on transient.

It is convenient to write the equations in terms of angular frequency ω , and to use the regular frequency f to represent the results. So, ω and f are used interchangeably implying that $\omega = 2\pi f$, and a quantity is uniquely identified by a subscript. For example, the natural frequency of a *LC* oscillator can be written as $\omega_n = 1/\sqrt{LC}$ or as $f_n = 1/2\pi\sqrt{LC}$, interchangeably. Without subscript, ω and f refer to the running variable in spectral analysis, denoted with f_m in other Sections.

For the sake of simplicity, we assume that the gain *A* is constant vs frequency, at least in the region around the oscillation frequency. If the gain flatness defect is not negligible, we move it from *A* to *B*, so that *A* is constant. Second, we assume that the resonator transfer function *B* is linear. Some resonators are nonlinear, and the resonant frequency depends on the amplitude. In quartz resonators, this is called "isochronism defect." Frequency may depend on amplitude, and bi-stability is observed at high amplitudes (Gufflet, Bourquin, & Boy, 2002) (Nayfeh & Mook, 2004). However, for the small amplitude fluctuations found in real oscillators, the assumption that *B* is linear at the operating point is fully satisfactory.



Figure 52 – General scheme of the oscillator loop (top), and (bottom) its companion circuit for phase noise. Reprinted from E. Rubiola, *Frequency and Amplitude Stability in Oscillators*, CC BY, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Modeling parametric noise and spurs requires that a phase modulator is introduced in the loop, as shown in Figure 52 (top) because these types of noise cannot be represented as additive processes. The presence of such phase modulator breaks the simplicity of the original Leeson model. We solve this difficulty by using the companion circuit for phase noise shown on Figure 52 (bottom). The companion circuit relies on the following ideas and simplifications:

- The startup transient is ended, and the oscillator is in its stationary condition.
- The gain compression, needed to stabilize the amplitude, has no effect on the phase.
- The *phase amplifier* has a gain equal to one, exact. This is consistent with the fact that time cannot be stretched or compressed, thus the noise-free amplifier delivers an exact copy of the input phase with no error.
- The amplifier's random phase is an additive process in this representation.
- The transfer function B_{φ} of the resonator is linear, and independent of the small fluctuations of amplitude.

The most important virtue of this companion scheme is that it is *inherently linear* because all the elements in the loop are linear.

Our approach differs from the original Leeson model in that we use a modulation method or a perturbation method, interchangeably and equivalently. More precisely, we introduce a phase perturbation ψ in the companion loop. The quantity ψ is the phase fluctuation of either the amplifier or the resonator. Thus, the phase noise transfer function is

$$H_{\varphi}(s) = \frac{\Phi(s)}{\Psi(s)} \tag{121}$$

The uppercase stands for the Laplace transform ($\Phi(s)$ is the Laplace transform of $\varphi(t)$), and the quantity $s = \sigma + j\omega$ is the complex frequency. Replacing $s \to 2\pi f$ and taking the square absolute value, we get

$$S_{\varphi}(f) = \left| H_{\varphi}(f) \right|^2 S_{\psi}(f) \tag{122}$$

This approach is surprisingly similar to the analysis of the response of a PLL, already familiar to the reader.

5.1.1 The resonator and its impulse response

Close to the resonance, the resonator can be approximated with a second-order linear differential equation

$$\ddot{v}_o + \frac{\omega_n}{Q}\dot{v}_o + \omega_n^2 v_o = \frac{\omega_n}{Q}\dot{v}_i$$
(123)

where ω_n is the natural angular frequency, Q is the quality factor in actual load conditions, and the term $(\omega_n/Q)\dot{v}_i$ is the driving force. We have chosen this type of driving force because it is homogeneous with the dissipative term $(\omega_n/Q)\dot{v}_o$, as it occurs in relevant cases like the series (parallel) RLC resonator driven by a voltage (current) source. Using the Laplace transforms, we find the resonator response $B(s) = V_o(s)/V_i(s)$

$$B(s) = \frac{\omega_n}{Q} \frac{s}{s^2 + (\omega_n/Q)s + \omega_n^2}$$
(124)

The square modulus $|B(f)|^2$, for $Q \gg 1$, describes a Lorentzian line shape of width f_n/Q centered at f_n .

The homogeneous equation, which is (123) with the force \dot{v}_i set to zero, describes the free decaying oscillation

$$v(t) = \cos(\omega_p t + \phi) e^{-t/\tau}$$
(125)

where ϕ is an arbitrary phase which results from the initial conditions,

$$\omega_p = \omega_n \sqrt{1 - \frac{1}{4Q^2}} \tag{126}$$

is the free decay angular pseudo frequency³, and

$$\tau = \frac{2Q}{\omega_n} \tag{127}$$

is the relaxation time.

In virtually all cases of interest for us, the resonator has large *Q*, and the oscillator oscillates close to the exact peak of resonance. Therefore, the following approximation holds

$$\omega_0 = \omega_n = \omega_p \tag{128}$$

Our phase-noise equivalent circuit relies on the knowledge of the resonator's response to the Dirac $\delta(t)$ impulse of phase, when the resonator is driven by a sinusoidal signal. This concept is illustrated in Figure 53. When the resonator is driven with the input signal $\cos[\omega_0 t + \delta(t)]$, it responds with $\cos[\omega_0 t + b_{\varphi}(t)]$. The function $b_{\varphi}(t)$ is the impulse response we need.

For most people, it is hard to figure out the meaning of a Dirac $\delta(t)$ in the argument of a sinusoid. However, the difficulty can be solved using a simple property of linear systems. The impulse response $b_{\varphi}(t)$ is related to the response $k_{\varphi}(t)$ to the Heaviside (step) function u(t) by

$$k_{\varphi}(t) = \int b_{\varphi}(t)dt \tag{129}$$

because

$$u(t) = \int \delta(t) \, dt \tag{130}$$

³ The quantity ω_p is not a valid angular frequency because v(t) is not periodic in a strict sense, being progressively attenuated by the term $e^{-t/\tau}$.



Figure 53 – The concept of phase-impulse response of a resonator (top), and (bottom) its derivation from the response to a small Heaviside (step) of phase. Simulation is also straightforward. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Let us apply to the resonator a small phase step $\varepsilon u(t)$, $\varepsilon \ll 1$, at the time t = 0 (Figure 53). The small value makes the derivation of $k_{\varphi}(t)$ simpler, but the result is general. The complete input signal is

$$v_i(t) = \cos(\omega_0 t) u(-t) + \cos(\omega_0 t + \varepsilon) u(t)$$
(131)

where u(t) is the Heaviside function

$$u(t) = \begin{cases} 0 & t < 0\\ 1 & t > 0 \end{cases}$$
(132)

Thus, at t = 0 the signal $\cos(\omega_0 t)$ is switched off by u(-t), and the signal $\cos(\omega_0 t + \varepsilon)$ is switched on by u(t). The resonator response results from

$$v_o(t) = v_{\rm off}(t) + v_{\rm on}(t)$$

where

$$v_{\rm off}(t) = \cos(\omega_0 t) \, e^{-t/\tau} \tag{133}$$

is the exponentially decaying response to $\cos(\omega_0 t)$, switched off at t = 0; and $v_{on}(t)$ is the growing response to the phase-shifted signal $\cos(\omega_0 t + \varepsilon) u(t)$, switched on at t = 0. Calculating the term $v_{on}(t)$ requires some manipulations

$$\begin{aligned} v_{\rm on}(t) &= \cos(\omega_0 t + \varepsilon) \left[1 - e^{-t/\tau} \right] & t > 0 \\ &= \left[\cos(\omega_0 t) \cos(\varepsilon) - \sin(\omega_0 t) \sin(\varepsilon) \right] \left[1 - e^{-t/\tau} \right] \\ &\simeq \left[\cos(\omega_0 t) - \varepsilon \sin(\omega_0 t) \right] \left[1 - e^{-t/\tau} \right] & \varepsilon \ll 1 \end{aligned}$$

Combining $v_{off}(t)$ and $v_{on}(t)$, we get

$$v_o(t) = \cos(\omega_0 t) - \varepsilon \sin(\omega_0 t) \left[1 - e^{-t/\tau}\right] \qquad \varepsilon \ll 1 \tag{134}$$

This is a sinusoid of phase $\varphi(t) = \varepsilon [1 - e^{-t/\tau}]$. The step response is obtained by deleting ε

$$k_{\varphi}(t) = \left[1 - e^{-t/\tau}\right] \tag{135}$$

Finally, the impulse response is obtained by differentiating the Heaviside response

$$b_{\varphi}(t) = \frac{1}{\tau} e^{-t/\tau} \tag{136}$$

The result is plotted in Figure 54. The Laplace transform of $b_{\varphi}(t)$ is

$$B_{\varphi}(s) = \frac{1/\tau}{s+1/\tau} \tag{137}$$

This is a single-pole low pass filter, like the *RC* lowpass filter. The cutoff frequency $f_L = 1/2\pi\tau = f_0/2Q$ is called Leeson frequency, and is equal to the half the resonator bandwidth.



Figure 54 – Impulse-of-phase response of a resonator at the exact resonant frequency. The response is equivalent to that of a first-order lowpass filter, like the *RC* network shown. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

5.1.2 The oscillator's phase-noise transfer function

We have all the pieces we need to calculate the phase-noise transfer function $H_{\theta}(s)$ of the complete oscillator

$$H_{\varphi}(s) = \frac{\Phi(s)}{\Psi(s)} \tag{138}$$

The oscillator is described in Figure 52 (bottom) as a classical feedback system, where

$$H_{\varphi}(s) = \frac{1}{1+B(s)}$$
 (139)

Using (137) in the above, we get immediately

$$H_{\varphi}(s) = \frac{s+1/\tau}{s} \tag{140}$$

Using $f_L = 1/2\pi\tau$, we get

$$\left|H_{\varphi}(f)\right|^{2} = \frac{f^{2} + f_{L}^{2}}{f^{2}} = 1 + \frac{f_{L}^{2}}{f^{2}}$$
(141)

and finally

$$\left|H_{\varphi}(f)\right|^{2} = 1 + \frac{f_{0}^{2}}{4Q^{2}} \frac{1}{f^{2}}$$
(142)

This is the simple function plotted in Figure 55. The physical meaning, related to the phase-noise scheme of Figure 52 (bottom), is surprisingly simple. At low offset frequency, $f \ll f_L$, the phase fluctuations are fed back to the input of the phase amplifier, and integrated. Because time cannot be compressed or stretched, the phase amplifier has a gain exactly equal to one. Thus, the loop is a perfect loss-free integrator, and $H_{\varphi}(s)$ has a pole in the origin. By contrast, at high offset frequencies, $f \gg f_L$, the resonator filters out the phase noise, preventing the fluctuations to be fed back to the input. Thus, the random phase φ at the output is equal to the random phase ψ introduced in the loop, and asymptotically $|H_{\varphi}(f)|^2 = 1$.



Figure 55 – Phase-noise transfer function of the complete oscillator. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola,

2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

5.1.3 The phase noise of the complete oscillator

We have seen in Section 2.3 that the amplifier phase noise is described by the polynomial law restricted to two terms, white and flicker

$$S_{\varphi}(f) = b_0 + \frac{b_{-1}}{f}$$
 amplifier (143)

In turn, b_0 is usually expressed as $b_0 = FkT_0/P_0$, where F is the noise factor of the sustaining amplifier, kT_0 is the thermal energy, and P_o is the power at the amplifier input; and b_{-1} is a parameter of the amplifier. Combining the amplifier noise with $|H(f)|^2$, we get

$$S_{\varphi}(f) = \left[\frac{FkT_0}{P_0} + \frac{b_{-1}}{f}\right] \left[1 + \frac{f_0^2}{4Q^2} \frac{1}{f^2}\right]$$
(144)



Figure 56 – The complete oscillator, including the frequency fluctuations of the resonator, and the phase noise of the output buffer. The tuning diode can be in series, in parallel, or take other configurations, depending on the type of oscillator and resonator. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Additional noise contributions are still to be included, namely, the frequency fluctuations of the resonator's natural frequency, the FM noise brought in by the tuning diode, and the phase noise of the output buffer. These noise perturbations are shown in Figure 56.

The resonator introduces flicker-of-frequency noise, originating a term

$$S_{\varphi}(f) = \frac{\mathsf{b}_{-3}}{f^3} \qquad \text{resonator} \tag{145}$$

in the phase noise plot. The coefficient b_{-3} can be derived from the floor of the resonator's Allan variance $\sigma_y^2(\tau)$, or from the floor of the modified Allan variance $mod\sigma_y^2(\tau)$, if known,

$$b_{-3} = \frac{f_0^2}{2\ln(2)}\sigma_y^2(\tau)$$
 AVAR floor (146)

$$b_{-3} = f_0^2 \frac{8}{27\ln(3) - 32\ln(2)} \mod \sigma_y^2(\tau) \qquad \text{MVAR floor} \qquad (147)$$

The modified Allan variance is preferred to the Allan variance because its superior capability to identify the fast noise processes, white and flicker PM, as we have seen in Section 3.3.1. However, the (regular) Allan variance is most often found in the technical documentation of oscillators. In practical cases, the b_{-3}/f^3 term due to the resonator fluctuations can be higher than the similar term from the sustaining amplifier (144). This the case, for example, of the high-stability 10 MHz OCXOs used as the frequency reference in electronic instruments.

Higher order fluctuations of the resonator, like the frequency random walk b_{-4}/f^4 , and the drift are more difficult to model in the phase noise spectrum, mainly because of the lack of data. Moreover, these perturbations depend on the environment, on the temperature control, etc.

If the oscillator is electrically tunable, it turns the voltage noise at the VCO input into FM noise. In the case of white noise across the tuning diode, the oscillator phase noise is

$$S_{\varphi}(f) = \frac{b_{-2}}{f^2}$$
(148)

with

$$\mathsf{b}_{-2} = e_n^2 K_o^2 \tag{149}$$

where e_n is the RMS voltage noise in 1 Hz bandwidth, and K_o is the VCO gain in (rad/s)/V. When the VCO input is connected to a resistor R, the noise cannot be lower than the thermal noise $\sqrt{4kT_0R}$. In this condition, the oscillator PM noise is

$$S_{\varphi}(f) = \frac{4kT_0RK_o^2}{f^2}$$
(150)

as explained in (Rohde, Poddar, & Boeck, Modern Microwave Oscillators for Wireless Applications: Theory and Optimization, 2005).

The contribution of the output buffer is an additional term like (143), yet with different noise factor, power, and flicker parameter b_{-1} . In most practical cases, the white phase noise of the buffer is lower than that of the sustaining amplifier because the carrier power is higher. By contrast, the phase flicker of the buffer is generally higher

than that of the sustaining amplifier. This happens for two reasons. First, the sustaining amplifier is in the loop, where the 1/f noise is turned into $1/f^3$ noise below f_L , thus a wise engineer spends a larger budget in a low-noise sustaining amplifier. Second, for proper isolation, the buffer consists of 2–3 cascaded stages, each of which contributes its own flicker.

In synthesis, the oscillator phase noise is given by

$$S_{\theta}(f) = \left[\frac{FkT_0}{P_0} + \frac{b_{-1}}{f}\right] \left[1 + \frac{f_0^2}{4Q^2} \frac{1}{f^2}\right]$$
sustaining amplifier
+ $\frac{f_0^2}{2\ln(2)}\sigma_y^2(\tau)\frac{1}{f^3}$ resonator flicker
+ $\frac{4kT_0RK_o^2}{f^2}$ tuning diode white
+ $\frac{FkT_0}{P_{\text{buf}}} + \frac{b_{-1}}{f}$ output buffer

+ FM, RW and higher order terms, and spurs

It is understood that symbols take their meaning from the context indicated on the right hand of the equation. For example, F in the first line refers to the sustaining amplifier, while the same symbol F in the fourth line refers to the output buffer.

Example 14 – 10.24 GHz DRO. We consider a 10.24 GHz DRO where the resonator has a loaded quality factor Q = 1000. The sustaining amplifier has a noise factor of 4 dB, and a flicker noise of -106 dBc/Hz extrapolated to 1 Hz. The power at the input of the sustaining amplifier is of -20 dBm. Let us calculate the phase noise spectrum, accounting only for the oscillator loop.

From the statement of the problem, we calculate the Leeson frequency

$$f_L = \frac{f_0}{2Q} = \frac{1.024 \times 10^{10}}{2 \times 1000} = 5.12 \text{ MHz}$$

and the amplifier noise parameters

$$b_0 = \frac{FkT_0}{P_0} = \frac{10^{4/10} \times 1.38 \times 10^{-23} \times 290}{10^{-20/10} \times 10^{-3}} = 10^{-15}$$

$$b_{-1} = 2 \times 10^{-106/10} = 5 \times 10^{-11}$$

Using (144), we calculate the oscillator phase noise PSD shown on Figure 57. The phase-noise coefficients are

$$b_{0} = 10^{-15} \qquad (-153 \text{ dBc/Hz})$$

$$b_{-1} = 5 \times 10^{-11} \qquad (-106 \text{ dBc/Hz at 1 Hz})$$

$$b_{-2} = \frac{f_{0}^{2}}{4Q^{2}} b_{0} = \frac{(1.024 \times 10^{10})^{2}}{4 \times 1000^{2}} \times 10^{-15} = 2.62 \times 10^{-2} \qquad (-18.8 \text{ dBc/Hz at 1 Hz})$$





Figure 57 – Phase noise PSD of the 10.24 GHz DRO discussed in the example.

Example 15 – 10 MHz OCXO. We consider a 10 MHz OCXO where the resonator has a loaded quality factor $Q = 10^6$. The sustaining amplifier has a noise factor of 1 dB, and a flicker noise of -140 dBc/Hz extrapolated to 1 Hz. The power at the input of the sustaining amplifier is -16 dBm. The buffer has a noise factor of 1 dB, and a flicker noise of -135.2 dBc/Hz extrapolated to 1 Hz. The power at the input of the buffer is of -7 dBm. The resonator has a stability of 3.2×10^{-13} (flicker term in the Allan deviation). Let us calculate the phase noise spectrum.

We calculate the oscillator loop, the buffer and the fluctuation of the resonator separately, and we add the results. The single contributions and the full spectrum are shown in Figure 58.

From the statement of the problem, we calculate the Leeson frequency

$$f_L = \frac{f_0}{2Q} = \frac{10^7}{2 \times 10^6} = 5 \text{ Hz}$$

and the noise parameters of the sustaining amplifier

$$b_0 = \frac{FkT_0}{P_0} = \frac{10^{1/10} \times 1.38 \times 10^{-23} \times 290}{10^{-16/10} \times 10^{-3}} = 2 \times 10^{-16}$$
$$b_{-1} = 2 \times 10^{-140/10} = 2 \times 10^{-14}$$

Using (144), we calculate the phase noise coefficients of the oscillator loop

$$b_{0} = 2 \times 10^{-16}$$

$$b_{-1} = 2 \times 10^{-14}$$

$$b_{-2} = \frac{f_{0}^{2}}{4Q^{2}}b_{0} = \left[\frac{10^{7}}{2 \times 10^{6}}\right]^{2} \times 2 \times 10^{-16} = 5 \times 10^{-15}$$

$$b_{-3} = \frac{f_{0}^{2}}{4Q^{2}}b_{-1} = \left[\frac{10^{7}}{2 \times 10^{6}}\right]^{2} \times 2 \times 10^{-14} = 5 \times 10^{-13}$$

The contribution of the output buffer is

$$b_0 = \frac{FkT_0}{P_0} = \frac{10^{1/10} \times 1.38 \times 10^{-23} \times 290}{10^{-7/10} \times 10^{-3}} = 2.5 \times 10^{-17}$$

$$b_{-1} = 2 \times 10^{-135.2/10} = 6 \times 10^{-14}$$

The phase noise due to the fluctuation of the resonator's natural frequency is

$$b_{-3} = \frac{f_0^2}{2\ln(2)}\sigma_y^2(\tau) = \frac{(10^7)^2}{1.386} \times (3.2 \times 10^{-13})^2 = 7.4 \times 10^{-12}$$

Adding all these terms, the phase noise PSD of the complete oscillator is

$$S_{\theta}(f) = 2.3 \times 10^{-16} + \frac{8 \times 10^{-14}}{f} + \frac{5 \times 10^{-15}}{f^2} + \frac{7.9 \times 10^{-12}}{f^3}$$



Figure 58 – Phase noise PSD of the 10 MHz OCXO discussed in the example.

5.1.4 Lessons from the examples

We discuss the phase noise spectra of the two above examples with respect to the scheme of Figure 56, extending some consideration to other types of oscillator.

First, we notice that in Figure 57 and Figure 58 there is either 1/f or $1/f^2$ noise, not both. This relates to the fact that we have two distinct corners, where the exponent of f changes by 1 crossing f_c , and by 2 crossing f_L .

If $f_c < f_L$, we get the $1/f^2$ slope term in the region between f_c and f_L , as in the DRO. This type of behavior is typical of microwave oscillators, where f_0 is of tens of GHz and Q generally not higher than a few thousands, thus we expect f_L of the order of a few MHz. With microwave oscillators, we often find f_c between 10 kHz and 100 kHz. A similar behavior is found in the internal VCOs of integrated circuits like FPGAs and DDSs. Such oscillators have an integrated LC tank having $Q \approx 10$ at f_0 of a few hundred of MHz, thus f_L of a few tens of MHz.

Oppositely, if $f_c > f_L$, we get the 1/f term in the region between f_L and f_c , as in the OCXO. This is typical of high-stability HF quartz oscillators, where technology suggests that $Qf_0 \approx 10^{13}$. Accordingly, we encounter typical $Q \approx 10^6$ at 10 MHz, thus $f_L \approx 5$ Hz. The flicker noise of HF amplifiers is rather low, in some cases even lower than -140

dBc/Hz extrapolated at 1 Hz. With a typical power of 10...100 μ W, f_c is of a few hundreds of Hz.

The noise pattern of VHF quartz oscillators is different. For example, using the thumb rule $Qf_0 \approx 10^{13}$, we expect $Q \approx 10^5$ at 100 MHz, thus $f_L \approx 500$ Hz. Since these resonators can work at higher power than the HF resonators, f_c is proportionally lower. So, f_L and f_c may more or less overlap, or even give $f_L > f_c$.

With a good design, the white noise of the buffer should not degrade the oscillator noise. We expect this because the Friis formula applies to white noise, and the power level is generally higher at the input of the buffer, than at the input of the sustaining amplifier.

By contrast, cascading several amplifiers, the b_{-1} coefficients add up in a way that is independent, or almost independent of the carrier power. When $f_L < f_c$, the oscillator loop has a 1/f region clearly visible, due to the sustaining amplifier (Figure 58). We expect that the 1/f noise at the oscillator output results from the contribution of the sustaining amplifier and of the buffer, and that the latter is generally dominant. A first reason is that the loop turns the 1/f noise of the sustaining amplifier into $1/f^3$ noise, while the 1/f noise of the buffer remains of the 1/f type at the output. In this respect, it is wise to put larger budget and design care in the sustaining amplifier. The second reason is that the output buffer has larger number of stages, because it has to isolate the loop from the output. Observing a 1/f region in an oscillator, in the absence of specific information, we may guess that 1/4 comes from the sustaining amplifier, and 3/4 from the buffer.

A further consequence of the buffer 1/f noise is that the corner between 1/f and $1/f^3$ noise is no longer at f_L . Interpreting the spectra gets more challenging.

The $1/f^3$ region results from the two contributions, the phase feedback in the oscillator (144), and the fluctuation of the resonator's natural frequency (146). These contributions are equal when

$$[b_{-1}]_{\text{sustaining amplifier}} = \frac{4Q^2}{2\ln(2)} \left[\sigma_y^2\right]_{\text{resonator FM flicker}}$$
(151)

or

$$[b_{-1}]_{\text{sustaining amplifier}} = 4Q^2 \frac{8}{27\ln(3) - 32\ln(2)} \left[\text{mod}\sigma_y^2 \right]_{\text{resonator FM flicker}}$$
(152)

The phase feedback is usually dominant in microwaves and in *LC* oscillators, which have moderate *Q*. Oppositely, the fluctuation of the resonator's natural frequency in the case quartz resonators (Figure 58) and other extremely high *Q* resonators.

5.2 Circumventing the resonator's thermal noise

The noise factor F and the Friis formula describe the noise of an amplifier impedancematched to the resistive input load at temperature T_0 . The equivalent noise PSD at the amplifier input is FkT_0 . This quantity is the sum of the available thermal energy kT_0 of the resistor, plus the contribution $(F - 1)kT_0$ of the amplifier. If the amplifier input is left open, is shorted to ground, or is connected to a noise-free load, the equivalent input noise is $(F - 1)kT_0$.

The Rohde oscillator (Figure 59 right) provides a means to circumvent the thermal noise of the resonator and of the sustaining amplifier by using the resonator also as an output filter. The series resonator consists of an inductance L, a capacitance C, and a resistance R_s . The latter represents the mechanical loss of the quartz resonator. At the resonance, the reactance ωL and $1/\omega C$ cancel one another, and the resonator is equivalent to the resistance R_s . Out of the resonator bandwidth the reactance is dominant, either $1/\omega C \gg R_s$ or $\omega L \gg R_s$, and the thermal noise is no longer coupled to the surrounding electrical circuit. Likewise, the noise of the sustaining amplifier falls in the stopband.

The original scheme (Rohde, Crystal oscillator provides low noise, 1975) derives from the Colpitts circuit, as shown on Figure 59. In the Colpitts oscillator, the white noise floor is determined by the transistor. In the Colpitts-Rohde oscillator, the resonator has the double role of the frequency reference and of the output filter. At the resonant frequency, the white phase noise is determined by R_s , by R, and by the noise of the transistor. Out of the resonator bandwidth, asymptotically, the resonator is open circuit, and the output noise is the thermal noise of the resistor R. Thus, the oscillator white phase noise PSD is given by

$$S_{\varphi}(f) = kT_0/P_0 \tag{153}$$

where P_0 is the carrier power dissipated by the resistor R.



Figure 59 – The Colpitts oscillator and the concept of the Rohde oscillator. The latter circumvents the thermal noise of the quartz resonator by using the resonator as the output filter.

The scheme of Figure 59 should be regarded as a concept, rather than a working oscillator. A problem is that R in series to the resonator reduces the resonator's Q. Another problem is that any perturbation to the load impacts on the frequency stability. Porting this idea to the complete circuit, we get the Rohde oscillator, shown in Figure 60. The entire circuit works in current mode. The sustaining amplifier is a feedback circuit implementing a negative resistance equal to $-(R_s + R_E)$. It could be the Colpitts scheme, or any other configuration suitable to oscillate with the resonator connected to ground. Instead of being grounded, the quartz resonator is connected to a grounded-base amplifier, which has low input impedance R_E thanks to local negative feedback. Neglecting the base current, the collector current I_C is equal to the current I_R flowing in the resonator. The conclusion is that the white phase noise is determined by the thermal noise of the collector resistance referred to the output carrier power

$$S_{\varphi}(f) = kT_0/P_{\text{out}} \tag{154}$$

From the ideal scheme to practical implementation, we notice that the condition $R_E \ll R_s$, necessary to preserve the resonator Q, can be obtained rather easily because R_E can be of the order of 1 Ω . The buffer contributes some white and flicker phase noise.

The main feature of the Rohde scheme is that it circumvents the thermal noise associated to the quartz internal dissipation, represented as the resistance R_s . However, it cannot remove the phase-to-frequency noise conversion mechanism inherent in the oscillator loop. Thus, the Rohde oscillator is an excellent solution for VHF quartz oscillators, typically 100 or 125 MHz, where the lowest phase noise floor is the most desirable feature, and the frequency stability is comparatively less important. When the highest frequency stability is of paramount importance, the Rohde oscillator may not be the best choice. The problem is that fluctuation of the buffer input impedance, however small, is converted into frequency noise.



Figure 60 – Principle of the Rohde oscillator.

5.3 Oscillator hacking

Inspecting on phase noise spectra provides information on the oscillator inside. Of course, our conclusions are only approximate, and mistakes are around the corner. Nonetheless, hacking oscillators from the phase noise spectra turns out to be surprisingly useful.

We provide guidelines based on the Leeson model, mainly addressed to the readers already familiar with the technology of oscillators and resonators.

The interplay between the Leeson frequency f_L and the corner frequency f_c of the sustaining amplifier defines the two main types of PM noise, shown on Figure 61 A and B. The type A, defined by $f_L < f_c$, is found with high Q resonators at low carrier frequency (HF). The type B, defined by $f_L > f_c$, is generally found in microwave oscillators and in low Q VHF and UHF oscillators. As we have seen, the spectrum contains either 1/f or $1/f^2$ phase noise types, not both. The Rohde oscillator (Section 5.2), presents additional difficulty in the interpretation, and for this reason it is not included in this analysis. The resonator is considered ideally stable in Figure 61 A and B, while Figure 61 C and D show the same spectra with the resonator's 1/f frequency fluctuation added, which is of the $1/f^3$ type in the phase noise plot.

The analysis starts from the identification of the coefficients b_i of the polynomial law. This is best done *by hand* sliding old-fashion *set squares* on the usual log-log plot of L(f), or on a comfortably large computer display by shifting a line. Don't forget the factor of two to convert L(f) into $S_{\varphi}(f)$. One of us (ER) is often seen at conferences doing this exercise on the data sheets found at the exhibitor boots, sliding two credit cards on the L(f) plots. With little training, the human eye does a good approximation close to the least square fit, which is exactly what we need. Rather than searching for the exact slope for the specific oscillator, we fit the spectrum with the canonical slopes f^0 , 1/f, $1/f^2$, etc. Frequency random walk, of the $1/f^4$ type and not shown in Figure 61, is almost always found at low f.

We proceed from the right-hand side of the spectrum to the left, thus from high f to low f.



Figure 61 – Basic types of oscillator PM noise spectra.

5.3.1 Moderate/low-Q (type A/C) oscillators

The spectrum of the type A/C is identified by the presence of $1/f^2$ PM noise, and the absence of 1/f.

First, the white phase noise tells us about the power P_0 at the input of the sustaining amplifier because the white noise of the buffer is negligible. This approximation makes sense because the power at the buffer input is generally higher than at the input of the sustaining amplifier. For the sake of simplification, we discard the effect of impedance mismatching out of the resonator bandwidth, which would give a white RF noise between $(F - 1)kT_0$ and FkT_0 , and we take F = 1 dB as an approximation. Using the formula $b_0 = FkT_0/P_0$, we calculate

$$P_0 = \frac{FkT_0}{b_0} \tag{155}$$

The power dissipated by the resonator is probably a little higher, yet of the same order of magnitude.

Second, we evaluate the Leeson frequency f_L as the intersection between the oscillator white phase noise b_0 , and the oscillator white frequency noise b_{-2}/f^2 . There is no need to account for the buffer because both white and 1/f noise types are negligible. Thus

$$f_L^2 = \frac{b_0}{b_{-2}}$$
(156)

and consequently

$$Q = \frac{f_0}{2f_L} \tag{157}$$

Third, we estimate the corner frequency $[f_c]_{SA}$ of the sustaining amplifier, which occurs when b_{-3}/f^3 equals b_{-2}/f^2

$$[f_c]_{SA} = \frac{b_{-3}}{b_{-2}} \tag{158}$$

These oscillators are rather simple to understand because the high value of f_L ends up in high value of b_{-3} (frequency flicker), which in practice exceeds the fluctuations of the resonator. Additionally, with this type of oscillator we spend comparatively little attention to low Fourier frequencies, say 100 Hz and below, where higher-slope phenomena show up.

Example 16 – Synergy Microwave DRO100, 10 GHz DRO. By inspection on the phase noise plot shown in Figure 62, we estimate

$$\begin{split} &b_0 = 10^{-17} \text{ rad}^2/\text{Hz}, \\ &b_{-1} \approx 0 \text{ (hidden below other noise processes),} \\ &b_{-2} = 1.41 \times 10^{-4} \text{ rad}^2\text{Hz}, \text{ and} \\ &b_{-3} = 14.1 \text{ rad}^2\text{Hz}^2. \end{split}$$

February 6, 2025

First, we calculate P_0 using F = 1 dB, thus $FkT = 6.2 \times 10^{-21}$ W/Hz at room temperature (300 K)

$$P_0 = \frac{FkT}{b_0} = \frac{5.2 \times 10^{-21}}{10^{-17}} = 520 \,\mu\text{W} \quad (-2.8 \,\text{dBm})$$

This is a reasonable value for a microwave DRO, optimized for low phase noise floor. The oscillator inside is not known, and we cannot know if the trick of Section 5.2 is implemented or not. In the absence of insider information, we are inclined to believe that the answer is "not" because the design of a low-impedance common-base amplifier is rather difficult at these frequencies.

Assuming that the floor gives the power at the amplifier input, the corner at 3.75 MHz is the Leeson frequency. Thus, we calculate

$$Q = \frac{f_0}{2f_L} = \frac{10^{10}}{2 \times 3.75 \times 10^6} = 1330$$

This is rather a typical value for a microwave dielectric resonator.

The frequency flicker b_{-3}/f^3 crosses the white frequency b_{-2}/f^2 at the corner frequency

$$f_c = \frac{\mathsf{b}_{-3}}{\mathsf{b}_{-2}}$$

This enables to calculate the flicker of phase of the sustaining amplifier

 $[b_{-1}]_{SA} = b_0 f_c = 10^{-17} \times 10^5 = 10^{-12} \text{ rad} (-120 \text{ dBrad}^2)$

This is quite a good value for a microwave amplifier, indeed well in the range of hightech devices.

Using the formulas of Table 6, the modified Allan deviation is

mod
$$\sigma_{\gamma}(\tau) = \frac{5.94 \times 10^{-13}}{\sqrt{\tau}} + 3.63 \times 10^{-10}$$

It is important to interpret correctly this result. First, the $1/\sqrt{\tau}$ term equals the flicker floor at $\tau = 2.7 \,\mu$ s. This short time is probably shorter than the sampling interval τ_0 of actual instruments (Section 3.2.1). The consequence is that the $1/\sqrt{\tau}$ term cannot be measured directly. Second, the estimation of mod $\sigma_{\rm V}(\tau)$ from $S_{\varphi}(f)$ makes sense only for very short measurement time, more or less up to $\tau = 1 \, \text{ms}$, which is the reciprocal of the lowest frequency (1 kHz) on the left hand side of Figure 62. Beyond, other terms may show up, like frequency random walk, temperature fluctuations, and aging.

After this digression, it is clear that the white PM and the flicker PM terms are too low and out of range for any practical measurement of mod $\sigma_{\rm y}(\tau)$, thus it have been discarded. For the sake of exercise, let us do the math. The white PM noise gives mod $\sigma_{\rm y}(\tau) = 6.16 \times 10^{-20}/(\tau \sqrt{\tau})$. The flicker PM is not directly accessible, but it can be guessed by adding the noise of three buffer stages similar to the sustaining amplifier, thus $b_{-1} = 4[b_{-1}]_{SA} = 4 \times 10^{-12} \text{ rad}^2$, and finally mod $\sigma_y(\tau) = 5.94 \times 10^{-13}/\tau$. The corner where the flicker PM noise equals the white FM noise occurs at $\tau = 9.7$ ns.



Figure 62 – Phase noise of the DRO100, 10 GHz DRO. The spectrum is from the DRO100 data sheet, © Synergy Microwave Corp., reproduced with permission. Graphical adaptation and comments are ours.

What happens if our guess is wrong, and the resonator is used to reduce the white noise as in Section 5.2? Let us say that the unfiltered white noise is

$$S_{\varphi}(f) = b'_0 = \lambda b_0$$

and take $\lambda = 4$ (6 dB) as an example. Equivalently, we allow that the unfiltered phase noise is a floor a factor of λ higher than the floor shown in Figure 62. Under this new hypothesis, the power at the amplifier input is

$$P_0 = \frac{fkT}{\lambda b_0} = 260 \,\mu W \quad (-5.8 \,\mathrm{dBm})$$

the "true" Leeson frequency is

$$f_L' = \lambda f_L = 1.88$$
 MHz,

and the resonator quality factor is

$$Q = \frac{f_0}{2f_L'} = \frac{f_0}{2\lambda f_L} = 665$$

However unsatisfied with this spread of values, we observe that the hacking process still provides useful information on the oscillator inside. We hope that the reader will dig in the literature and find his/her own way to improve on our estimates.

Example 17 – Synergy Microwave DCMO 1027, 100–270 MHz VCO. We proceed exactly as before, from the right-hand side of Figure 63 to the left-hand side, thus

 $b_0 = 2.82 \times 10^{-17} \text{ rad}^2/\text{Hz}$,

 $b_{-1} \approx 0$ (actually, hidden below other noise processes),

 $b_{-2} = 2.82 \times 10^{-4} \text{ rad}^2 \text{Hz}$, and

 $b_{-3} = 1.42 \text{ rad}^2 \text{Hz}^2$.

We calculate P_0 using F = 1 dB, at room temperature (300 K)

$$P_0 = \frac{FkT}{b_0} = \frac{5.2 \times 10^{-21}}{2.82 \times 10^{-17}} = 185 \,\mu\text{W} \quad (-7.3 \,\text{dBm})$$

With the same reservations, we assume that the floor indicates the power at the amplifier input, hence the corner at 3.16 MHz is the Leeson frequency. The phase noise is measured at $f_0 = 136.9$ MHz. Thus, we calculate

$$Q = \frac{f_0}{2f_L} = \frac{1.369 \times 10^8}{2 \times 3.16 \times 10^6} = 21.6$$

This is more or less what we expect in a *LC* oscillator with a wide tuning range.

The frequency flicker b_{-3}/f^3 crosses the white frequency b_{-2}/f^2 at the corner frequency

 $f_c = b_{-3}/b_{-2} = 4.5$ kHz. This enables to calculate the flicker of phase of the sustaining amplifier

 $[b_{-1}]_{SA} = b_0 f_c = 2.82 \times 10^{-17} \times 4.5 \times 10^3 = 1.26 \times 10^{-13} \text{ rad}^2$ (-129 dBrad²) The flicker of a good VHF amplifier should be lower, but this value probably includes the flicker of the tuning diodes. Notice that a tuning range of a factor of 2.7 requires a capacitance range of a factor of $2.7^2 = 7.3$. The burden of the tuning varactors is certainly higher because of the residual capacitance in the circuit.

Using the phase noise data we have, the modified Allan deviation is

mod
$$\sigma_{\gamma}(\tau) = \frac{6.1 \times 10^{-11}}{\sqrt{\tau}} + 7.9 \times 10^{-10}$$

We have discarded white and flicker PM for the same reasons of the previous example, and the same consideration about the meaning of the low value of mod $\sigma_y(\tau)$ apply here.



Figure 63 – Phase noise of the Synergy Microwave DCMO 1027, 100–270 MHz VCO. The spectrum is from the DCMO 1027 data sheet, © Synergy Microwave Corp., reproduced with permission. Graphical adaptation and comments are ours.

5.3.2 High-Q (type B/D) oscillators

With reference to Figure 61, the PM noise spectrum of the type B/D is identified by the presence of 1/f noise, and by the absence of $1/f^2$ noise.

First, we use the white phase noise to infer the power P_0 at the input of the sustaining amplifier. As with the A/C-type spectra, we assume that the white noise of the buffer is negligible because the power at the buffer input is generally higher than at the input of the sustaining amplifier. Likewise, we discard the effect of impedance mismatching out of the resonator bandwidth, and we take F = 1 dB as an approximation. Using the formula $b_0 = FkT_0/P_0$, we calculate

$$P_0 = \frac{FkT_0}{b_0} \tag{159}$$

We expect that the power dissipated by the resonator is close to this value, maybe a little higher.

Let provisionally neglect the resonator's frequency fluctuations, as in Figure 61-B. Highlighted with a circle, we see the "apparent corner" where the spectrum changes from 1/f to $1/f^3$. However related, this is not the Leeson frequency. The reason is that the oscillator 1/f phase noise is due to both the sustaining amplifier and the output buffer. We guess that 1/4 of such flicker is due to the sustaining amplifier, and 3/4 is due to the buffer

$$[b_{-1}]_{SA} = \frac{1}{4} [b_{-1}]_{osc} \qquad (-6 \text{ dB})$$
 (160)

This formula is used to estimate the 1/f noise of the sustaining amplifier. In terms of electrical circuit, we guess that the buffer consists of three stages similar in noise to the sustaining amplifier. This is sound because the buffer has to provide high isolation from the load circuit.

Next, still assuming that the resonator is free from fluctuations, the $1/f^3$ noise is given by $[b_{-3}]_{osc} = [b_{-1}]_{SA} f_L^2/f^3$. Hence, we estimate the Leeson frequency f_L as the intersection between the sustaining amplifier $[b_{-1}]_{SA}/f$ and the oscillator frequency flicker $[b_{-3}]_{osc}/f^3$

$$f_L^2 = \frac{[b_{-1}]_{SA}}{[b_{-3}]_{osc}}$$
(161)

The result is marked as f_L on Figure 61-A. The corner, highlighted with a circle, occurs *below* the oscillator noise, thus it is not visible. Accordingly, we find a first estimate of resonator's quality factor using $f_L = f_0/2Q$

$$Q = \frac{f_0}{2f_L} \qquad \text{(first estimate)} \tag{162}$$

At this point, we have to introduce the frequency flicker of the resonator, $[b_{-3}]_{res}/f^3$ on the phase noise plot. This is shown as the thick dashed lines on Figure 61-D. Understanding whether $[b_{-3}]_{res}$ is the dominant noise, or $[b_{-1}]_{SA} f_L^2/f^3$ prevails, requires experience and skill. We may start to collect information from the datasheet and from the literature, figuring out the stability (flicker of frequency floor on the Allan deviation plot) and the possible Q.

If the oscillator stability is limited by the Leeson effect, (162) is consistent with the technology, and the oscillator is fully described by Figure 61-B.

Oppositely, HF oscillators (5–10 MHz) exhibiting ultimate stability are generally limited by the fluctuations of the resonator. In this case, the Leeson effect is completely hidden. We can only do the academic exercise of guessing Q from the technology, or from other sources of information, calculating f_L , and identifying the $1/f^3$ part of the "oscillator loop" plot of Figure 61-D.

Example 18 – Rakon HSO 14, 5 MHz OCXO. This oscillator is intended for space and scientific applications which require ultimate stability, for example the VCO to be locked to the 1.42 GHz atomic transition in a Hydrogen maser. The phase noise spectrum is shown on Figure 64. By inspection on the plot, we estimate

 $b_0 = 1.6 \times 10^{-16} \text{ rad}^2/\text{Hz}$,

 $\mathsf{b}_{-1} = 8 \times 10^{-15} \ \mathrm{rad}^2$ upper bound,

 $b_{-2}\approx 0\,$ (actually, hidden below other noise processes), and

 $b_{-3} = 6.3 \times 10^{-13} \text{ rad}^2 \text{Hz}^2$.

The thick and irregular spectrum between 3 and 10 Hz may indicate that the correlation instrument has still not reached the final value, thus the true b_{-1} may be lower than indicated.

First, we calculate P_0 using F = 1 dB, thus $FkT = 6.2 \times 10^{-21}$ W/Hz at the oven temperature of 350 K (75–80 °C). Thus

$$P_0 = \frac{FkT}{b_0} = \frac{6.2 \times 10^{-21}}{16 \times 10^{-16}} = 33 \,\mu\text{W}.$$

This is quite a plausible value for this type of oscillator, which is optimized for stability rather than for low phase noise floor. The oscillator inside is not known. However, we believe that the trick of the Rohde oscillator is not implemented, first because it is not necessary for the target applications of this oscillator, and second because even the lowest instability introduced by the virtual ground would be detrimental to the stability at the ultimate level required.



Figure 64 – Phase noise spectrum of the Rakon HSO 14 OCXO measured with a Microsemi 5120A test set. The spectrum is © 2019 Rakon France SAS, courtesy of Patrice Canzian and Vincent Candelier. Graphical adaptation and comments are ours.

Assuming that the sustaining amplifier contributes 1/4 of the oscillator flicker $b_{-1} = 8 \times 10^{-15} \text{ rad}^2$, we get $[b_{-1}]_{SA} = 2 \times 10^{-15} \text{ rad}^2$.

By inspection on the plot, the flicker of frequency b_{-3}/f^3 crosses the sustaining amplifier flicker $[b_{-1}]_{SA}/f$ at f' = 6 Hz. If we interpret this as the Leeson frequency, we find a quality factor $Q' = f_0/2f' = 8.9 \times 10^5$. This seems too low for this class of oscillator. So, let us stick on the thumb rule $f_0Q \simeq 10^{13}$, thus $Q = 2 \times 10^6$.

Using the formulas of Table 6, the modified Allan deviation is

mod
$$\sigma_{\gamma}(\tau) = \frac{4.9 \times 10^{-16}}{\tau \sqrt{\tau}} + \frac{5.2 \times 10^{-15}}{\tau} + 4.9 \times 10^{-14}.$$

Unlike in the previous examples, white and flicker PM (the $1/\tau\sqrt{\tau}$ and $1/\tau$ terms) provide useful information because the corners where white PM crosses flicker PM, and where flicker PM crosses flicker FM, are of 8.9 ms and 106 ms, which is still in the range of practical measurements. Random walk, temperature fluctuations, and frequency drift are not visible on Figure 64. Such phenomena will inevitably show up, however only for $\tau > 10$ s, which is the reciprocal of the lowest frequency (0.1 Hz) available on the phase noise plot.

Given the applications this oscillator is intended for, rather specialized in the longterm performances, the Allan deviation is preferred to the modified Allan deviation. Setting $f_H = 5$ Hz, we find

$$\sigma_{\rm y}(\tau) = \frac{9.6 \times 10^{-15}}{\tau} + 5.9 \times 10^{-14}$$

with a corner at $\tau = 162$ ms.

It is interesting to compare mod $\sigma_y(\tau)$ to $\sigma_y(\tau)$. As we have seen, mod $\sigma_y(\tau)$ provides separate values for white PM and flicker PM, with no need of a lowpass filter. The lowpass, however, is implied in the sampling interval τ_0 . By contrast, $\sigma_y(\tau)$ provides a single value for both, proportional to $1/\tau$, with a strong effect of the lowpass filter on the contribution of white PM noise. Additionally, mod $\sigma_y(\tau)$ always gives values lower than those of $\sigma_y(\tau)$.

6 The measurement of phase noise

We have already seen that the measurement of SSB noise referred to the carrier power has been abandoned long time ago, replaced with the direct measurement of the phase fluctuations vs an appropriate reference. Some general-purpose spectrum analyzers include the dedicated hardware that enables the measurement of the phase noise associated to an input signal. However, these instruments are limited by the stability and by the noise of their internal oscillator and synthesizer, and they are usable only for the measurement of some rather noisy oscillators. Instead, dedicated instruments are the right choice.

Three basic ingredients are needed for the measurement of phase noise,

- Phase reference,
- Phase detector,
- Signal processing unit based on FFT and averaging.

The phase reference is an oscillator or a synthesizer which provides a suitably pure signal. The phase detector converts the phase difference, input vs reference, into a voltage or other signal. The double balanced mixer – or diode ring – saturated at both inputs is in most cases the preferred phase detector because of its low background noise. Digital detectors, like the XOR gate and the PFD, are not suitable to general test equipment, mainly because of their background noise. All these detectors require that the phase reference is at the same frequency of the input signal. The double balanced mixer is not the only option for the phase detector. Other types of instruments are found, based on direct digitization of the input signal, and on Software Defined Radio (SDR) techniques. These digital techniques are more flexible, overcome some of the problems of the saturated mixer, and enable to compare the phase of two signals that are not at the same frequency. However, the noise of the ADCs is the major problem of such instruments.

Most modern instruments make use of two separate and equal channels that measure simultaneously the input signal. The background noise is rejected thanks to an appropriate correlation-and-averaging algorithm which relies on the hypothesis that the two channels are statistically independent. The use of correlation relaxes the noise specifications for the reference oscillator and for the phase detector, at the cost of longer measurement time. The measurement of noise below the background noise of a single channel is possible. As a consequence, correlation and averaging allows the use of a synthesizer as the reference in each channel, which is generally noisier than a dedicated low-noise oscillator. Without synthesizers, a specific low-noise reference oscillator is necessary for each frequency of interest. In this section, we will learn about this type of equipment, principles, background noise and other limitations, and some tricks useful to extend the range of application.

6.1 Double-balanced-mixer instruments

The basic measurement scheme, shown on Figure 65, is straightforward. Unlike the regular use of the mixer, the LO and RF signals are synchronous ($f_{LO} = f_{RF} = f_0$), close to the quadrature, and large enough to saturate the input. In this condition, the difference $f_{LO} - f_{RF}$ degenerates to a DC signal sensitive to the phase φ

$$V = K_{\varphi}\varphi \tag{163}$$

The sum $f_{LO} + f_{RF}$ falls close to $2f_0$, which is filtered out. The value of K_{φ} can be up to 0.3–0.7 V/rad in favorable conditions (Figure 66). A low-noise amplifier (LNA) is needed at the mixer output to rise the small signal to a level suitable to the FFT analyzer.

For an introduction the double balanced mixer, the reader can refer to the author's earlier work (Rubiola E., Tutorial on the Double-Balanced Mixer, 2006), to an old but good white paper from Watkins Johnson (Kurtz, 2001), and to the classic Maas book (Maas, 1993).



Figure 65 – Basic phase noise measurement. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).


Figure 66 – Phase-to-voltage conversion of a double balanced mixer saturated at 15–20 dBm power at each input, plotted for different values of the load resistance. Reproduced from E. Rubiola, Tutorial on the Double Balanced Mixer, arXiv:physics/0608211v1 [physics.ins-det], August 2006 (Rubiola E. , Tutorial on the Double-Balanced Mixer, 2006).



Figure 67 – Double balanced mixer. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

The mixer is implemented with a diode ring and baluns, similar to the circuit shown on Figure 67. However, the actual implementation may be more complex. Baluns are present at both RF and LO input to match the unbalanced-mode input to the diode ring, which is balanced. Two types of balun are often found. At microwaves frequencies, multi-section microstrip lines are preferred, providing a typical bandwidth is of 1–3 octaves. Wider bandwidth comes at the cost of larger physical size and higher loss. The HF-UHF implementation is based on Iron powder transformers, which exhibit a bandwidth of up to 3 decades. A wider bandwidth, up to 4 decades, is achieved with a smart type of transformer, where twisted pairs are wound on an Iron powder core shaped in toroidal form or in binocular form. Primary and secondary windings are coupled magnetically at lower frequencies, while capacitive or electromagnetic-line coupling takes over at high frequencies. The boundary between RF-type and microwave-type implementation is in the region of 3 GHz, depending on design and manufacturing choices.

Schottky diodes are preferred because of low threshold and fast switching time. High-level mixers, up to 200 mW (+23 dBm) or more input power, are convenient because of the higher value of K_{φ} , and in turn the lower background noise. These mixers differ from Figure 67 in the use of 2–3 diodes in series in each arm of the ring. In each arm, an appropriate network distributes power and reverse bias equally between the diodes. By contrast, the double-double-balanced mixer (sometimes called triple balanced mixer) cannot be used as a phase detector because the IF output cannot be dc coupled. Special mixers intended as phase detectors achieve higher gain by increasing the IF impedance to 500 Ω typical.

Proper switching operation requires that the IF output current can circulate at both DC and $2f_0$. The problem arises in wideband mixers, where the circulation of the $2f_0$ current has to be ensured by the load at the IF output, the reason being that the upper IF frequency falls in the LO/RF frequency range. In such cases, the low-pass filter must have resistive or capacitive input impedance, not inductive impedance.

It has been reported that a series resonator at the IF output, tuned at $2f_0$, is useful in that it maximizes the IF current at $2f_0$, and increases K_θ . Of course, this trick is reserved to special cases, where the experimentalist is interested in a single value of f_0 , or at most in a small set of frequencies, and has full access to the system inside.

The double balanced mixer is an appealing choice for a phase detector because of the low background noise, the wide range of operating frequency, and the overall simplicity of the system. Most general-purpose double balanced mixers are suitable as phase detectors. In the absence of specific information, one can assume that best power is 3 dB above the nominal LO power, and that the same power should be used for the LO input and for the RF input. One can also assume that phase-detector bandwidth is ³/₄ of the nominal bandwidth.

The narrow power range, typical of the double-balanced mixer, can be annoying. The problem is that the input power must be sufficient to saturate the mixer, but smaller than the absolute maximum rating level, with with a safe margin. Unfortunately, the gap between nominal and maximum power is not comfortable, and leaves approximately ± 5 dB around a nominal power of 10–15 dBm. At lower power, the background noise increases. Further decreasing the power, K_{φ} drops abruptly and the mixer is no longer usable.

The mixer inputs, strongly saturated, have highly nonlinear behavior and the input impedance changes with frequency. Strong odd harmonics of the carrier frequency are reflected back, combining in rather unpredictable way depending on cable length. For this reason, it is a good practice to introduce a 3-dB attenuator as close as possible to the mixer inputs.

6.1.1 The measurement of oscillators

The basic scheme for the measurement of the PM noise of oscillators is shown on Figure 68. Taking the error voltage V as the output of the mixer, the PLL is used as a high-pass filter. So, beyond a cutoff frequency f_{HP} , the error signal is asymptotically equal to $V = K_{\varphi}(\varphi_{DUT} - \varphi_{REF})$. Below the cutoff, the error signal is small, but the phase noise can still be calculated using the equation of the PLL. It is useful to bring the reference oscillator as close as possible to the DUT frequency by adjusting the DC offset, so that the detector and the control work close to 0 V. In this condition K_{φ} is the highest, and the measurement starts with the instrument in the middle of the dynamic range.



Figure 68 - Phase noise measurement of an oscillator using the error signal of a PLL.

In production and in industrial applications, it is generally possible to rely a reference oscillator whose PM noise can be neglected, being $L_{REF}(f) \ll L_{DUT}(f)$ with a sufficient margin.

In rare cases, we may have to test special low-noise oscillators, where no lowernoise reference is available. As a first approximation, we can measure two equal oscillators, so that the noise of each is half (-3 dB) of the result displayed by the test set. Dropping the hypothesis that the two oscillators are equal, the reliable measurement of a single oscillator with the scheme of Figure 68 is a complex and time-consuming task because we need to compare all the possible pairs in a set of at least three similar oscillators, and to solve for the noise of each. However, the cross-spectrum method provides a simple and practical solution, discussed later in this Chapter.

Notice that in Figure 68 we have kept the reference oscillator outside the test set. This is often necessary in a general-purpose instrument because a low-noise reference is needed, at the same frequency of the oscillator under test. Neither a wideband VCO nor a synthesizer would feature the low noise needed to measure high purity oscillators. Introducing a synthesizer for flexible operation requires the dual-channel scheme, which we will study later.

The PLL error function $E(s) = \Phi_e(s)/\Phi_i(s)$ is given by

$$E(s) = 1 - B(s) = \frac{1}{1 + G(s)H(s)}$$
(164)

Including the gain of the LNA in K_{φ} , the closed-loop error voltage $K_{\varphi}E(s)$ is described by the transfer function

$$T(s) = \frac{V(s)}{\Phi(s)} = K_{\varphi}[1 - B(s)]$$
(165)

For the simplest loop, where $G(s)H(s) = K_{\varphi}K_o/s$, the function T(s) is a first order (single pole) high pass filter

$$T(s) = \frac{K_{\varphi}s}{s + K_{\varphi}K_o} \tag{166}$$

or equivalently

$$|T(f)|^2 = K_{\varphi}^2 \frac{f^2}{f^2 + f_{HP}^2}$$
(167)

where f_{HP} is the cutoff frequency

$$f_{HP} = \frac{1}{2\pi} K_{\varphi} K_o \tag{168}$$

In a commercial test sets, f_{HP} is generally chosen by an internal algorithm, and only advanced users can take control on it. However, the implications of f_{HP} deserve attention. Naively, one may be inclined to set f_{HP} at a value lower than the lowest analysis frequency. For example, being interested in L(f) from 10 Hz to 100 kHz, we would choose $f_{HP} = 1 \dots 2$ Hz, so that $|T(f)|^2 = K_{\varphi}^2$ (constant) in the full span. However, a tighter loop is a better choice, with f_{HP} set approximately at the corner between the $1/f^2$ noise, or the $1/f^3$ noise, and the white region (Figure 69). The instrument measures $S_v(f)$, that is, the PSD of v, and calculates L(f) as

$$L(f) = \frac{1}{2} \frac{S_{\nu}(f)}{|T(f)|^2}$$

Of course, this relies on the accurate measurement of $|T(f)|^2$ in actual conditions, which can be accomplished by modulating the VCO signal. A first advantage of this approach is the reduced burden for the FFT's dynamic range. This is quite obvious from Figure 69. The oscillator $S_{\varphi}(f)$ has a wide dynamic range (plot A) because of the $1/f^3$ and $1/f^4$ behavior. By contrast, V requires a comparatively smaller dynamic range because its spectrum (plot B) contains at most $1/f^2$ components at low frequency. A second and more subtle advantage is that the tighter lock overrides some uncontrolled effects of electromagnetic interferences, and in turn provide more reliable results.





Electromagnetic interference is sometimes a source of erratic or wrong results, difficult to identify and fix. RF/microwave leakage is to some extent inevitable, due to connectors, coaxial cables, power lines, grounding, insufficient shielding, etc. The problem arises from the fact that the reference oscillator and the oscillator under test are at the same frequency. The power leaking from one oscillator builds up as a significant energy in the second oscillator, after integration over the relaxation time of the internal resonator. The resonator's relaxation time may be unexpectedly long, up to hundreds of milliseconds in the case of high stability 5-10 MHz OCXOs. Of course,

reciprocity makes the stray coupling bidirectional. Leakage may injection-lock the two oscillators to one another, or corrupt the transfer function T(s) if coupling is insufficient for locking. Interestingly, injection locking is a phase sensitive phenomenon. With the same amount of power leakage, the oscillator may lock or not, depending on the electrical length of the path. Should a phase noise spectrum be suspected of being corrupted by leakage, the following tests are recommended.

- Opening the loop when the two oscillators are set as close as possible to the same frequency, they phase lock to one another.
- In open loop condition as above, the beat note is not sinusoidal. Instead, it slows down or almost stops when certain phase relationships are met. In this case, it is likely that the two oscillators *try* periodically to *lock* to one another when the phase relationships are favorable, but coupling is insufficient and injection locking fails.
- In the $1/f^3$ or $1/f^4$ region of $S_{\varphi}(f)$, the slope tends to decrease or get flat towards low frequencies instead of getting steeper.
- Changing the length of critical cables affects the low-frequency region $S_{\varphi}(f)$. The critical cables are those connecting the oscillator under test to the mixer, or the reference oscillator to the mixer.
- The shape of $S_{\varphi}(f)$ changes after introducing a common mode filter a ferrite ring or clamp along a cable, RF output, VCO input or power supply.

If any of the above symptoms show up, the experimentalist should be aware that the phase noise measurement is unreliable. Investigating on the transfer function T(s) is recommended.

6.1.2 Background noise, spurs, and other experimental issues

Recalling the scheme of Figure 68, we identify the following contributions to the instrument background noise

- Mixer
- Low-noise amplifier (LNA) between mixer and FFT analyzer
- Reference oscillator
- The DC reference, providing the tuning voltage
- Pollution from AM noise (CF Section 6.4.4)



Figure 70 – Typical background noise of a mixer, including the low-noise amplifier that follows. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Let us start with the mixer and the LNA. They must be analyzed together. The typical background noise is shown on Figure 70, and discussed below. Lowest noise operation requires high driving power, at least +10 dBm. The mixer adds little white PM noise because its noise factor is of some dB (Rotholz, 1984), but it has a low gain K_{φ} . The consequence is that the white noise floor is set by the LNA after the mixer. The reader can refer to (Rubiola & Lardet-Vieudrin, Low flicker-noise amplifier for 50 Ω sources, 2004) for the design of low-noise amplifiers specifically intended for the lowest background noise in this type of applications. Unfortunately, we do not have analytical expressions for K_{φ} and for the mixer noise factor. The flicker noise is an experimental parameter, for both the mixer and the amplifier. We see on Figure 70 that the total noise is significantly higher than the noise of a good amplifier divided by the mixer gain. The following examples show typical values of white and flicker PM noise, and their origin.

Example 19 – Mixer and LNA white noise. Let us calculate the white PM noise background assuming that the mixer has $K_{\varphi} = 500 \text{ mV/rad}$ driven at 40 mW (+16 dBm), and that the white noise of the LNA is $e_n = 1.25 \text{ nV/}\sqrt{\text{Hz}}$, including the 50 Ω input load.

We first convert e_n into $S_v(f) = e_n^2 = 1.56 \times 10^{-18} \text{ V}^2/\text{Hz}$, i.e., $-178 \text{ dBV}^2/\text{Hz}$. Going backwards to the input, the background noise is

$$S_{\varphi}(f) = S_{\nu}(f)/K_{\varphi}^2 = 6.25 \times 10^{-18} \text{ rad}^2/\text{Hz}$$
 (-172 dBrad²/Hz)

For comparison, the thermal noise at the mixer input is

$$S_{\varphi}(f) = kT/P = 4 \times 10^{-21}/4 \times 10^{-2} = 10^{-19} \text{ rad}^2/\text{Hz}$$
 (-190 dBrad²/Hz)

In this example, it takes a noise factor of 18 dB for the mixer noise to match the noise of the amplifier. Choosing different components and parameters, the result does not change significantly. For example, the highest K_{φ} found in a commercial instrument is of 1 V/rad, which requires +20 dBm input power.

Notice that the value of e_n of this example is quite optimistic because it includes the thermal noise $\sqrt{4kTR}$ of the input resistor, 0.9 nV/ $\sqrt{\text{Hz}}$ with $R = 50 \Omega$ at room temperature

Example 20 – Mixer and LNA flicker noise. We use the mixer of the previous example, $K_{\varphi} = 500 \text{ mV/rad}$, and th low-noise amplifier designed for PM noise applications (Rubiola & Lardet-Vieudrin, Low flicker-noise amplifier for 50 Ω sources, 2004), which exhibits 1.6 nV flicker (-176 dBV²). Referring this value to the input, we find

 $S_{\varphi}(f) = S_{\nu}(f)/K_{\varphi}^2 = (2.5 \times 10^{-18}/f)/0.5^2 = 10^{-17}/f \text{ rad}^2 (-170 \text{ dBrad}^2)$ This is 30 dB lower than the overall noise shown on Figure 70, which refers to the HF-UHF mixers. Such margin may be reduced by 10 dB with better mixers, if any, and with a not-as-good amplifier. Anyway, the result yield safely to the conclusion that the background noise is chiefly originated in the mixer.

At low frequencies, the dominant phase noise in oscillators is $1/f^2$, $1/f^3$ and higher slope types, while mixer and low-noise DC amplifier have only white and 1/f noise. In this region, the phase noise of the reference oscillator is generally the most severe limitation to the measurement.

It is often necessary to provide a DC voltage at the VCO input to bring the oscillator at the nominal frequency, as in Figure 68. The voltage noise of this source turns into FM noise at the oscillator output. Thus, white and flicker noise show up as white and flicker FM noise, whose slope is $1/f^2$ and $1/f^3$ on the phase noise spectrum. In principle, the contribution of the DC source should be made smaller than the oscillator noise. This is not always possible, chiefly in the case of oscillators having low phase noise and high voltage-to-frequency gain. By contrast, the control provides only the small correction needed to keep the oscillator locked during the measurement. For this reason, in a good design the fluctuations coming from the control fall below other noise contributions. Having seen unexpected experimental mistakes, we strongly recommend at least a quick check on the noise sent to the VCO input.

The FFT analyzer is preceded by a low-noise amplifier. Thus, an appropriate choice of the amplifier and of its gain makes the noise of the analyzer negligible. That said, the noise of the analyzer deserves more attention in earlier FFT analyzers, where the low-frequency decades were obtained by reducing the sampling frequency of the converter. The problem comes from the quantization noise, whose variance is $\sigma^2 = V_{LSB}^2/12$. The Parseval identity states that $\sigma^2 = S_v(f)f_s/2$, where $S_v(f)$ is the white noise floor, and $f_s/2$ is the bandwidth of the quantization noise, equal to half the sampling frequency f_s . Thus, the quantization noise is $S_v(f) = V_{LSB}^2/6f_s$. On the FFT analyzer, this is seen as a staircase-shaped noise floor, increasing steadily towards the low-frequency decades, where the sampling frequency is progressively lower. The problem is solved in modern analyzers. The input ADC runs always at full speed, and the lower sampling rate is obtained by data decimation after digital low pass filtering.

6.1.3 Asymmetric driving for low-power Signals

The mixer is unsuitable to low power signals because K_{θ} decreases. This impacts strongly on the white noise floor, and flicker PM noise tends to increase at low power. Below a threshold power, K_{θ} drops suddenly, and the mixer is no longer usable. However, asymmetric power driving is possible, with the LO input saturated, and the RF input in the linear regime, say at a power 10 dB lower than the LO nominal power, or even less. This may be convenient for the measurement of oscillators and of two-port components, when only the reference signal has a power sufficient to saturate the mixer.

In the asymmetric power driving, the mixer works as a synchronous detector. This mode is broadly similar to the regular "superheterodyne receiver," differing in that LO and RF frequency is the same, thus $|f_{LO} - f_{RF}|$ degenerates to dc. The LO signal is

$$V_{L0}(t) = V_{sat}\sin(\omega_0 t) \tag{169}$$

where the peak voltage V_{sat} results from saturation. Using the approximation $\cos \theta \simeq 1$ and $\sin \varphi \simeq \varphi$ for small φ , the RF signal $V_{RF}(t) = V_0 \cos(\omega_0 t + \varphi)$ becomes

$$V_{RF}(t) = V_0 \cos(\omega_0 t) - \varphi V_0 \sin(\omega_0 t)$$
(170)

Dropping the $2\omega_0$ term, the detected signal at the IF port is

$$V = V_0 A \varphi \tag{171}$$

where *A* is the mixer loss written as a "gain." For example, a loss of 6 dB translates into A = 0.5 because $10^{-6/20} = 0.5$. From the definition of the phase-to-voltage gain $K_{\theta} = V/\varphi$, we find

$$K_{\varphi} = V_0 A \tag{172}$$

Example 2-21. We use a mixer that has a loss of 8 dB (A = 0.4) when the LO port is saturated at +18 dBm ($V_{sat} = 2.5$ V across 50 Ω load). Sending a -12 dBm signal ($V_0 = 80$ mV across 50 Ω load) to the RF port, in quadrature with the LO signal, the phase-to-voltage gain is $K_{\varphi} = 32$ mV/rad.

6.1.4 Heterodyne measurement of oscillators

The heterodyne method (Figure 71) is a good option to extend the range of a phase noise test set to higher frequencies by exploiting a low-frequency beat at $f_b = |f_i - f_r|$, with $f_b \ll f_i$, and also $f_b \ll f_r$. In a typical case, we compare two microwave oscillators by bringing the beat down to the HF region. Of course, a suitable reference oscillator must be available, and some auxiliary pieces of hardware. In open-loop conditions, the phase fluctuation of the beat note is

$$\varphi_b = |\varphi_i - \varphi_r| \tag{173}$$

thus

$$S_{\varphi,b}(f) = S_{\varphi,i}(f) + S_{\varphi,r}(f)$$
 (174)



Figure 71 – Heterodyne (beat) method for the phase noise measurement of oscillators. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC

BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

Interestingly, the scheme of Figure 71 takes benefit from a leverage effect, which relaxes the frequency-stability specification for the VCO and for the synthesizer by a factor of f_b/f_i . This leverage effect is a direct consequence of the fact that the beat mechanism stretches the time associated to a unit of phase (radian) by the factor f_i/f_b . This is particularly useful in the $1/f^2$, $1/f^3$ and steeper regions of L(f). For the purpose of extending the frequency range, the

hetherodine scheme is preferred to a frequency divider because of the lower background noise.

In the scheme of Figure 71 (top), the reference oscillator also drives the synthesizer. This may be impractical because commercial synthesizers accept only some round values of the reference frequency, typically 5–10–100 MHz. An alternate scheme is possible, shown in Figure 71 (bottom). In this case, the main reference oscillator is free running, with no control, and the auxiliary reference is phase-locked to the beat note $f_b = |f_i - f_r|$. The hardware is clearly simpler than on Figure 71 (top), and the benefit of the leverage effect is the same. Besides the microwave practice, Figure 71 (bottom) solves some difficult problems of PM noise measurements, beyond our scope. For example, the metal-semiconductor (Schottky) diode can be used to down convert from the THz region to HF or VHF. Similarly, the fast PIN InGaAs photodetector is routinely used in metrology labs to beat 1550-nm telecom lasers down to microwaves.

6.1.5 The measurement of amplifiers and other two-port components

In the case of two-port components, we opt for the differential measurement scheme shown in Figure 72. Once the quadrature condition is set, the mixer delivers a voltage proportional to the instantaneous phase fluctuation of the DUT. The oscillator PM noise is common mode, thus it is rejected. In spite of this, practical measurements are way more difficult than Figure 72 lets us believe. The PLL scheme (Figure 68) simple to use, to the extent that the quadrature condition is set automatically and precisely by the feedback. By contrast, in Figure 72 the quadrature condition relies on an adjustable phase shifter manually set by the operator. The reference arm is the preferred location for this phase shifter because it is independent of the DUT. A tuning range of 180° , with a comfortable margin, is sufficient because any of the two quadrature points at $\pm 90^\circ$ can be used, with equivalent results. Different types of phase shifters can be used, depending on frequency. Mechanical phase shifters (U-shaped line stretchers) are appealing for their low noise, fine tuning capability, and wide frequency range. One of us (ER) has used extensively the phase shifters manufactured by ARRA for research applications. A problem with the line stretchers is the small delay range, related to the physical size. For reference, a range of 1 ns is equivalent to 30 cm change in the electrical length, thus of approximately 25 cm physical excursion. Of course, the range can be extended with a set of electrical cables of known length, but the operation is tedious and time consuming. A 90^o directional coupler terminated to varactors at two ports is an excellent phase shifter, provided the noise of such varactors is low enough compared to the DUT. The frequency range is limited by the 90^o coupler. For lower noise, the varactors can be replaced with variable capacitors, but in this case the adjustment is difficult and time consuming. A classic solution suitable to HF-VHF (Figure 73) is found in an article by Phillips (Phillips, 1987).

In summary, experience suggests that it is almost impossible to combine the suitable range of phase with a wide range of frequency. The reason is that electronics gives wideband control on delay, or narrowband control on phase. The fact that delay and phase are related, does not really help to get the $> 180^\circ$ excursion we need.



Figure 72 – Phase noise measurement of a two-port component. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).



Figure 73 – Example of variable phase shifter.

The measurement of two-port components challenges the background noise of the instrument. This happens because these components often exhibit very low noise, and because the noise processes are of the same type of those of the instrument, that is, white and flicker PM. For example, the flicker PM noise of a RF amplifier may be of the same order of that of a double balanced mixer. The background noise is discussed in Section 6.1.2.

The measurement of amplifiers is always tricky because it is necessary to match both input power and output power to the instrument. A problem is that the mixer has a narrow power range. Another problem is that white phase noise increases at low input power. In practice, it is often necessary to introduce appropriate attenuators at both input and output of the amplifier, whose attenuation must be determined for each case.

Finally, frequency multipliers, dividers and synthesizers are a special case because they deliver an output frequency that is not equal to the input frequency. The scheme of Figure 74 solves the problem by using two equal DUTs, so that the mixer receives the same frequency at the two inputs. Of course, this method gives the total noise of the two DUTs, with no means to divide the noise contribution of each. We rely on the assumption that the phase noise of the two DUTs is the same, and we take away 3 dB for the phase noise of one. This method may also be useful in other cases, for example in low noise amplifiers, where the enhanced sensitivity due to the presence of two DUTs helps to get out of the background noise.



Figure 74 – Phase noise measurement of frequency dividers and multipliers, and other devices whose output frequency is not equal to the input frequency. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

6.1.6 The discriminator method

Figure 75 shows a method to measure the PM noise of an oscillator using a delay line as the frequency reference, so that the signal at the mixer output is $K_{\varphi}[\varphi(t) - \varphi(t - \tau)]$. The measurement is possible because the delay line de-correlates the phase noise, under some conditions. This method is useful for fast phenomena, not for random walk and drift. The dynamic range is limited by the amount of delay that can be introduced without excessive attenuation, and by the background noise of the mixer and of the following circuits. Lance et al (Lance, Seal, & Labaar) used coaxial cables, enhancing the sensitivity with the cross-spectrum method discussed later. For long delay, up to 10-20 µs, the optical fiber proved to be an efficient solution (Rubiola, Salik, Huang, Yu, & Maleki, 2005) because the attenuation of the extremely low attenuation, 0.2 dB/km, or 0.04 dB per microsecond delay. The delay at microwave frequency is obtained by modulating and detecting the intensity of a laser beam. Improved sensitivity is achieved with the cross-spectrum method, using two statistically independent instruments which measure simultaneously the same oscillator (Salik, Yu, Maleki, & Rubiola, 2004), (Volyanskiy, et al., 2008).



Figure 75 – Phase noise measurement of an oscillator using a delay line as the reference. The lower sceme is the phase-noise equivalent circuit. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

The response of the system is defined as

$$|T(f)|^2 = \left|\frac{V(f)}{\Phi(f)}\right|^2$$
 (175)

where V(f) and $\Phi(f)$ are the Fourier transform of the output voltage, and of the oscillator random phase. This transfer function is easy to derive analitycally using the phase step method that we have seen in Section 5.1.1 with the response of the resonator. We use the Laplace transforms V(s) and $\Phi(s)$, where $s = \sigma + j\omega$ is the complex variable. By inspection on Figure 75, the line delays the phase perturbation and the impulse of phase by the same amount τ . The output voltage is

$$V(s) = K_{\varphi} [1 - e^{-s\tau}] \Phi(s)$$
(176)

With simple manipulations, we find

$$|V(f)|^{2} = |K_{\varphi}[1 - e^{-j2\pi f\tau}]|^{2} |\Phi(f)|^{2}$$

= $K_{\varphi}^{2}[1 - e^{-j2\pi f\tau}][1 - e^{j2\pi f\tau}]|\Phi(f)|^{2}$
= $4K_{\varphi}^{2}\sin^{2}(2\pi f\tau) |\Phi(f)|^{2}$

and finally

$$|T(f)|^2 = 4K_{\varphi}^2 \sin^2(2\pi f\tau)$$
(177)

The above equation is exploited to calculate

$$L(f) = \frac{1}{2} \frac{1}{4K_{\varphi}^2 \sin^2(2\pi f\tau)} S_{\nu}(f)$$
(178)

from the PSD $S_{\nu}(f)$ of the voltage measured by the FFT analyzer.

Notice that (178) has singularities at $f = n/2\tau$, integer n, where L(f) cannot be calculated. At $f \to 0$, the transfer function is approximated with $|T(f)|^2 = 16K_{\varphi}^2\pi^2\tau^2f^2$, and the instrument has a poor sensitivity due to background noise. For $n \ge 1$, L(f) shows large and sharp peaks due to the background noise. In practice, the system is usable up to $f \approx 0.8/2\tau$.

The delay line can be replaced with a reference resonator, as shown on Figure 76. In the first scheme from the top, the resonator is used as the reference for the measurement of an oscillator. Of course, it is necessary that the resonator is more stable than the oscillator. In the second scheme the roles are interchanged, and the resonator is the device under test.

In this case, it is convenient to use two equal DUTs having the same resonant frquency and the same *Q* because in this case the noise of the reference oscillator is rejected.



Figure 76 – Phase noise measurement of an oscilator using a resonator as the the reference (top). The roles can be inverted (middle), using the oscillator as the reference for the measurement of a resonator. In this case, it is convenient to use two equal DUTs. On the bottom, we see the phase-noise equivalent circuit of the first scheme. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

The third scheme, on the bottom of Figure 76, is the phase-noise equivalent circuit used to measure the input oscillator versus the reference resonator. We use this scheme to derive the frequency transfer function $|T(f)|^2$ assuming that the resonator natural frequency f_n is equal to the carrier frequency f_0 . We follow the same methods just used for the delay line, but the resonator's phase response is $B(s) = (1/\tau)/(s + 1/\tau)$, as seen in Section 5.1.1. There follows that

$$T(s) = K_{\varphi} \left[1 - \frac{1/\tau}{s + 1/\tau} \right] = K_{\varphi} \frac{s}{s + 1/\tau}$$
(179)

thus

$$|T(f)|^{2} = K_{\varphi}^{2} \frac{f^{2}}{f^{2} + \frac{1}{4\pi^{2}\tau^{2}}}$$
(180)

Using the quality factor Q and the natural frequency $f_n = f_0$ of the resonator, the transfer function is better rewritten as

$$|T(f)|^{2} = K_{\varphi}^{2} \frac{f^{2}}{f^{2} + \frac{f_{n}^{2}}{4O^{2}}}$$
(181)

Finally, the oscillator phase noise is

$$S_{\varphi}(f) = \frac{1}{|T(f)|^2} S_{\nu}(f) = \frac{1}{K_{\varphi}^2} \frac{f^2 + f_n^2/(4Q^2)}{f^2} S_{\nu}(f)$$
(182)

At $f \to 0$, the response of the instrument is poor, and dominated by the background noise. However, the measurement does not suffer from the infinite series of singularities as in the case of the delay line. The resonator has its minimum loss at $f_n = f_0$. This enables to pump correctly the mixer at the two inputs.

6.2 The cross-spectrum method

The scheme of the dual-channel measurement, shown on Figure 77, consists of two equal branches which measure the same oscillator using the PLL method. The main point is that the noise of the reference oscillators, of the mixers, and of the LNAs can be rejected using correlation and averaging. This is possible because the devices are physically separate, thus we can assume that their noise processes are statistically independent. By contrast, the DUT is common to the two branches, thus it is fully correlated, and captured by the statistical process.

The cross-PSD relates to the Fourier transfom of the correlation function. Thus, averaging on m measures of $S_{\varphi}(f)$, the single-channel background noise is rejected by a factor of approximately $1/\sqrt{m}$. It is therefore possible to measure a phase noise $S_{\varphi}(f)$ lower than the background noise of a single branch.

Figure 78 shows what happens during the measurement process. The DUT noise (C) is lower than the background noise in single-channel mode (A), thus the measurement is possible only after rejecting the background. The instrument displays the cross PSD averaged no *m* acquisitions. The cross PSD starts from the single-channel background (A), and is progressively reduced proportionally to $1/\sqrt{m}$. With small *m*, the single-channel background is not sufficiently rejected, and the instrument displays the plot (B).

When *m* is large enough, the single-channel background is well rejected (D), and the instrument displays the DUT noise (C).



Figure 77 – Dual-channel phase-noise measurement system. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019.



Figure 78 – Rejection of the background noise in the dual-channel phase noise measurement.

In this version of the PLL method we have introduced a synthesizer between the reference oscillator and the phase detector. The obvious benefit is that the system is flexible and suitable to a wide range of frequencies, without need of a separate reference oscillator for each frequency of interest. The higher noise of the synthesizer, as compared to an oscillator, can be tolerated thanks to the noise rejection of the dual-

channel scheme. Figure 79 shows a simpler version of the dual-channel system. It differs from the previous version in that there is only one reference oscillator, driving both the synthesizers. The trick is that each synthesizer has an internal frequency reference, locked to the main reference. The appropriate cutoff frequency inside such synthesizers may be of the order of 0.1-1 Hz, depending on the interplay between the stability and phase noise spectra of the internal and external references. Thus, for f beyond the cutoff, the two synthesizers are statistically independent, and their noise is rejected. Below the cutoff, the entire measurement relies on the stability and on the spectral purity of the main reference.



Figure 79 – Alternate dual-channel phase-noise measurement system. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E., Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

6.2.1 The rejection of background noise

After the heuristic reasoning, we explain the mathematics underneath the rejection of the single-channel noise. A more detailed treatise of the cross spectrum is available in (Rubiola & Vernotte, The cross-spectrum experimental method, 2010).

With reference to Figure 77, the two signals at the input of the FFT analyzer are

$$x(t) = K_{\varphi}[\varphi(t) - \psi(t)]$$
(183)

$$y(t) = K_{\varphi}[\varphi(t) - \theta(t)]$$
(184)

where the gain K_{φ} is the same for the two channels, and includes the trivial gain of the LNAs. The random phases $\psi(t)$ and $\theta(t)$ account for the noise of the references, the mixers and the LNAs. Accordingly, $\varphi(t)$ is the random phase of the oscillator under test,

with no additional terms. It is sound to assume that $\varphi(t)$, $\psi(t)$ and $\theta(t)$ are statistically independent because they come from separate hardware. Only $\theta(t)$ appears in both x(t) and y(t).

As usual, we denote the Fourier transform with the uppercase letters

$$X(\omega) = K_{\varphi}[\Phi(\omega) - \Psi(\omega)]$$
(185)

$$Y(\omega) = K_{\varphi}[\Phi(\omega) - \Theta(\omega)]$$
(186)

The cross PSD is given by

$$S_{yx}(f) = \frac{2}{T} Y(f) X^*(f)$$
(187)

where *T* is the measurement time, the superscript "*" stands for complex conjugate, and the factor "2" fixes the scale factor from the two-sided Fourier transform into to the one-sided PSD. Using a lighter notation where the frequency is implied, the above formula is expanded as

$$S_{yx} = K_{\varphi}^2 \frac{2}{T} \left(\Phi \Phi^* - \Phi \Psi^* - \Theta \Phi^* + \Theta \Psi^* \right)$$
(188)

The dual-channel FFT analyzer measures the cross spectrum $(S_{yx}(f))_m$ averaged on m data records of x(t) and y(t) acquired simultaneously

$$\langle S_{yx} \rangle_m = K_{\varphi}^2 \frac{2}{T} [\langle \Theta \Theta^* \rangle_m - \langle \Phi \Psi^* \rangle_m - \langle \Theta \Phi^* \rangle_m + \langle \Theta \Psi^* \rangle_m]$$
(189)

A rather intuitive theorem states that if two random variables are statistically independent in the time domain, their Fourier transforms are also statistically independent. Thus, we expect that $\langle \Phi \Psi^* \rangle_m \to 0$, $\langle \Theta \Phi^* \rangle_m \to 0$, and $\langle \Theta \Psi^* \rangle_m \to 0$ for large *m*. Consequently

$$\langle S_{yx} \rangle_m = K_{\varphi}^2 \frac{2}{T} \langle \Phi \Phi^* \rangle_m = K_{\varphi}^2 \langle S_{\varphi} \rangle_m \tag{190}$$

The process takes a time mT, not counting the computing time.

It is useful to write the instrument readout as the *estimation*, denoted with the "hat" accent

$$\widehat{S_{\varphi}}(f) = \frac{1}{K_{\varphi}^2} \left\langle S_{yx}(f) \right\rangle_m \tag{191}$$

The estimation is a powerful concept because the simple average is not the one and only option, and we can consider other estimators.

The above reasoning gives account for the most interesting feature of the crossspectrum method, which is the possibility to measure $S_{\varphi}(f)$ below the limit set by the single-channel background noise. However, it takes infinite averaging for (191) to fully eliminate the background. For finite m, the terms $\langle \Phi \Psi^* \rangle_m$, $\langle \Theta \Phi^* \rangle_m$, and $\langle \Theta \Psi^* \rangle_m$, are not completely averaged out, and set the measurement limit. It is worth mentioning that $S_{\varphi}(f)$ is a real and positive quantity because $\Theta\Theta^*$ is obviously real and positive. By contrast, the Fourier transform is a complex quantity, thus $\Theta\Phi^*$, $\Theta\Phi^*$ and $\Theta\Psi^*$ are complex, and consequently (191) is complex. In the mixed terms $\Theta\Phi^*$, $\Theta\Phi^*$ and $\Theta\Psi^*$, the background noise is equally split between real part and imaginary part. Therefore, (191) can be replaced with

$$\widehat{S_{\varphi}}(f) = \frac{1}{K_{\varphi}^2} \left\langle \Re \{ S_{yx}(f) \} \right\rangle_m \tag{192}$$

This results in a faster measurement because the unnecessary part of the background noise is removed, at no cost in terms of hardware and computation complexity. It can be proved that (192) is the optimum estimator for white noise, that is, the estimator that converges to $S_{\varphi}(f)$ with the lowest *m*, or equivalently in the shortest measurement time. Assuming that the background noise is the same for the two channels, $S_{\text{one-ch}}(f) = S_{\psi}(f) = S_{\theta}(f)$, the averaging limit is

$$S_{\varphi}(f) = \frac{S_{\text{one-ch}}(f)}{\sqrt{2m}}$$
 averaging limit (193)

A problem with the estimator (192) is that negative values are incompatible with the logarithmic scale (dBrad²/Hz). The problem is explained on Figure 80, which shows the Probability Density Function (PDF) of $\widehat{S_{\varphi}}(f)$ at a single frequency, i.e., one bin of the FFT, for different values of m. The PDF is quite large at low m, where the single-channel background noise is dominant. Being the background noise dominant at small m, a significant amount of negative outcomes occur. Increasing m, the PDF shrinks and converges to the final value of $S_{\varphi}(f)$, and the negative occurrences get progressively rare.



Figure 80 – Probability Density Function (PDF) of $\widehat{S_{\varphi}}(f) = (1/K_{\varphi}^2) \langle \Re \{S_{yx}(f)\} \rangle_m$ at a single frequency f, i.e., one bin of the FFT.

Most commercial instruments use the estimator

$$\widehat{S_{\varphi}}(f) = \frac{1}{K_{\varphi}^2} \left| \left\langle S_{yx}(f) \right\rangle_m \right|$$
(194)

instead of (192). This guarantees that all values are positive, and suitable to the logarithmic scale. In this case, the averaging limit is given by

$$S_{\varphi}(f) = \frac{S_{\text{one-ch}}(f)}{\sqrt{m}}$$
 averaging limit (195)

The consequence is that, for the same $S_{\text{one-ch}}(f)$ and for the same averaging-limit target, the value of m set by (194) is four times larger than that set by (192). Accordingly, the full measurement process takes four times longer time.

We believe that the choice of (194) is mainly due to historical reasons. If we want to preserve the logarithmic display, more efficient options are possible. A good choice consists of taking the real part as in (192), but replacing all the negative outcomes of with the smallest positive number. With this choice, the estimator discards all the noise associated to the imaginary part, and reduces the bias.

The above digression is for a given number m of averaged spectra. Focusing on the measurement time T, taken for the m acquisitions, provides a totally different perspective on the noise rejection. Given the time T for one acquisition, it holds that m = T/T. This can be rewritten as $m = T\Delta f$ because the resolution (distance between contiguous bins) of the FFT is $\Delta f = 1/T$. Measuring phase noise, we always represent the frequency on a log scale. Hence, we like a logarithmic frequency resolution with

 $\Delta f/f = C$, a constant. This gives a constant number μ of bins per decade, related to the resolution by

$$\frac{\Delta f}{f} = e^{\ln(10)/\mu} - 1 \tag{196}$$

Working in logarithmic resolution, we rewrite *m* as $m = Tf(\Delta f/f)$, and (195) becomes

$$S_{\theta}(f) = \frac{1}{\sqrt{\mathcal{T}(\Delta f/f)}} \frac{S_{\text{one-ch}}(f)}{\sqrt{f}} \qquad \text{averaging limit}$$
(197)

The rejection law is shown on Figure 81. The $1/\sqrt{f}$ term of (197) introduces a -5 dB/decade slope, which adds to the background noise of the instrument. Thus, the flicker region is seen as a slope of -15 dB/decade, and the white region is seen as a slope of -5 dB/decade. Of course, additional limitations apply, due to crosstalk and to other hardware problems which introduce a correlation between the two channels.

The logarithmic resolution cannot be obtained directly from the FFT algorithm. Other Specific algorithms exist (see for example (Barash & Ritov, 1993)). A popular solution is the FFT implemented in segments, more or less one decade wide. The resolution Δf is constant inside each segment, but proportionally narrower Δf is adopted in the lower-frequency segments after decimating the time series. The corresponding pattern is a step function which approximates the $1/\sqrt{f}$ term of (197).



Figure 81 – Rejection of the background noise in logarithmic resolution, with constant $\Delta f/f$ over the full span and fixed measurement time \mathcal{T} .

6.3 Digital instruments

The double balanced mixer has been the preferred phase detector since the 1970s. More recently a new generation of digital instruments appeared, when fast ADCs were available, capable of about 12 bit resolution at 50-100 MHz sampling rate. This enables the direct digitization of a RF signal, and the extraction of the instantaneous amplitude and phase.

Figure 82 shows a rather general scheme of the instrument, consisting of two equal branches which compare the input and the external reference to the internal clock. Each branch implements a classical I-Q detection in FPGA exploiting the stream of digitized data. In principle, the ADCs should operate close to the full speed because the lowest background is achieved in his condition.

The classical sampling theorem states that the input frequency must be lower than the Nyquist frequency, that is, $f_0 < f_N = f_{ck}/2$. However, the input frequency can be extended beyond f_N by under sampling the input signal. Numerous modern ADCs are intended for under sampling operation, and for this purpose have an input-frequency range significantly wider than f_N . The input bands, called *Nyquist zones*, are selected by introducing an appropriate antialiasing filter at the input, which is a lowpass for the first zone, and bandpass for the subsequent zones. Of course, under sampling comes at the cost of higher background PM noise. In practice, the maximum frequency of the first Nyquist zone is of 0.8 f_N because the antialiasing filter has a roll-off region before achieving the appropriate attenuation. A similar reasoning applies to the bandpass filter for the next Nyquist zone, which leaves a dark region between zones set by the roll-off region of the filters. This can be fixed by shifting the sampling frequency.

The NCO (Numerically Controlled Oscillator) provides two orthogonal phases of a sinusoid at the same frequency of the input, or of the external reference. The digital down conversion is free from the usual defects of analog I-Q detection, like orthogonality error and gain asymmetry. The digital low-pass filters are necessary to remove everything beyond f_N , and to reduce the sampling rate to a value suitable for further processing. The maximum baseband frequency is of 0.8 f_N , again limited by the filter roll off. However, more stringent limitations may apply, due to the architecture of the instrument, and to the processing speed.

The CORDIC algorithm (Volder, 1959) (Meher, Valls, Juang, Sridharan, & Maharatna, 2009) is most often used to calculate the phase. Interestingly, the digital technology enables the calculation of phase, and also of amplitude, with so high accuracy that it exceeds the general metrological performance of the instrument. The phase of the reference signal is scaled according to the frequency ratio f_i/f_r , so that it can be compared to the input phase. An alternate and elegant solution consists of converting the phase of both input and reference to phase-time. The two-branch configuration is necessary to bring the external reference (5-10-100 MHz, or arbitrary frequency) in the machine because the clock frequency takes fixed values determined by design considerations.



Figure 82 – Basic scheme of the direct-digitization phase detector.

The configuration of Figure 82 has three relevant features, advantageous versus the double-balanced mixer scheme

- It operates at arbitrary frequencies, with no need for the input and the external reference to be at the same frequency.
- The oscillator under test and the reference are free running, with no need of phase or frequency lock.
- Measuring a two-port device, there is no need for a line stretcher or for a variable phase shifter to set the quadrature condition.

These features enable the measurement of frequency dividers, multipliers, etc. in a straightforward way, without need of comparing two equal DUTs.

The main problem of the scheme shown is the background noise, generally limited by the noise of the ADCs. For reference, the noise of a selected 12-14 bit ADC at 100 MHz sampling frequency, operated at full range is

$$S_{\varphi}(f) = 10^{-11}/f + 10^{-15} \text{ rad}^2/\text{Hz}$$
 (198)

that is, -110 dBrad^2 flicker, and $-150 \text{ dBrad}^2/\text{Hz}$ white floor. The flicker PM noise is a technical parameter of the ADC, as we have seen with amplifiers. The white noise results from the quantization noise and from the clock jitter

$$S_{\varphi}(f) = S_{\varphi,q} + S_{\varphi,ck} \tag{199}$$

The quantization noise can be calculated as follows. At full range input $V_{pp} = V_{FSR}$, the carrier power on a 1-Ohm resistance is

$$P = \frac{V_{FSR}^2}{8} \tag{200}$$

With *n* bits (ENOB), the quantization noise power is

$$\sigma^2 = \frac{V_{LSB}^2}{12} = \frac{V_{FSR}^2}{12 \times 2^{2n}}$$
(201)

uniformly distributed from 0 to the bandwidth $B = f_s/2$. Thus

$$N = \frac{\sigma^2}{B} = \frac{V_{FSR}^2}{6 \times 2^{2n} f_s}$$
(202)

The phase noise is given by $S_{\varphi}(f) = N/P$, thus

$$S_{\varphi,q} = \frac{4}{3 \times 2^{2n} f_s}$$
(203)

For example, a 14-bit ADC with ENOB = 12 bits, and sampling at 32 MS/s, has a quantization noise

$$S_{\varphi,q} = \frac{4}{3 \times 2^{24} \times (32 \times 10^6)} = 2.5 \times 10^{-15}$$

that is, $-146 \text{ dBrad}^2/\text{Hz}$, or -149 dBc/Hz.

The quantity $S_{\varphi,ck}$ is the clock-distribution PM noise, which is a technical parameter of the ADC. It hits on PM noise only, not on AM noise. For this reason, there is an asymmetry between AM and PM noise floor, and $S_{\varphi,ck}$ can be measured as

$$S_{\varphi,ck} = S_{\varphi} - S_{\alpha}$$

Values of 0-3 dB and more are observed, depending on the operating conditions and on frequency. In fact, $S_{\varphi,ck}$ is of the time type, while $S_{\varphi,q}$ is of the phase type.

Spurs and artifacts are another problem of digital systems. A first type of spurs results from sampling and digital synthesis. The sampling process produces

$$f_{spur} = \mathcal{M}f_i - \mathcal{N}f_s \tag{204}$$

$$f_{spur} = \mathcal{M}f_r - \mathcal{N}f_s \tag{205}$$

Additionally, the NCO produces spurs at multiples of the grand repetition rate

$$f_{spur} = \frac{f_{ck}}{2^n} \tag{206}$$

where f_{ck} s the NCO clock frequency, and n is the number of bits of the NCO. However, the equivalent value of n to be used here can be smaller than the actual number of bits in the NCO register, depending on the frequency control word. Torosyan suggests that the equivalent n is the number of bits of the frequency control word from the MSB to the rightmost "1," which of course depends on the output frequency. For example, a 24bit NCO has n = 24 bits when the control word is 01110101 11010111 11000001, and n = 18 bits when the control word is 01110101 11010111 11000000. Details are found in (Torosyan & Wilson, Exact Analysis of DDS Spurs and SNR due to Phase Truncation and Arbitrary Phase-to-Amplitude Errors, 2005) and (Torosyan, DDS Complete Analysis and Design Guidelines, PhD thesis, 2003). Distortion produces spurs at high frequencies, and aliasing brings them down to baseband. The digression we have seen with the DDS applies. At the state of the knowledge, the spurs cannot be eliminated, so they are generally removed from the displayed data in order to give the best representation of the DUT noise.

Because of the high noise of the ADCs, a cross spectrum configuration is necessary to reduce the background noise of the instrument. Two equal blocks like Figure 82 are used instead of the double-balanced mixers, and measure simultaneously the quantity $\varphi = \varphi_i - \varphi_r$. As a result of design choices, and probably also of marketing choices, commercial instruments often use a single input for the external reference. Thus, the noise of the external reference cannot be rejected.

A small number of digital instruments is commercially available, listed in Table 11. Some of them will be briefly discussed in the following pages.

Type and Brand	Input frequency	Analysis frequency	Note
5125A	1 - 400 MHz	100 μHz -	Performs $L(f)$ and ADEV
Microsemi		1MHz	Discontinued, June 2018
5120A	1 - 30 MHz	100 μHz -	Performs $L(f)$ and ADEV
Microsemi		1MHz	
3120A	0.5 - 30 MHz	1 Hz - 100	Performs $L(f)$, AM noise and ADEV
Microsemi		kHz	Requires a host PC for all measurements
			Supersedes the Miles Design 5330A
PhaseStation 53100A	1 - 200	1 mHz - 1	Performs $L(f)$, AM noise and ADEV
Jackson Labs	MHz	MHz	Requires a host PC for all measurements
FSWP8	1 MHz -	10 mHz -	Performs $L(f)$ and AM noise
FSWP26	8/26.5/50 GHz	300 MHz	Additional functions (VCO test, baseband
FSWP50			and microwave spectrum analyzer, pulsed
Rohde Schwarz			signals, etc.), some optional

Table 11 – Digital phase-noise analyzers.

6.3.1 The Microchip family of phase noise and Allan deviation testers

Microchip (formerly Microsemi and Symmetricom) manufactures three instruments for the measurement of phase noise and Allan deviation. All these instruments are based on a scheme broadly similar to Figure 82. The background noise is rejected thanks to the cross-spectrum method. Input and reference are symmetrical, so they can be chosen independently in the range shown. Unfortunately, there is only one input for the external reference, thus it is impossible to reject the noise of the reference as we did in Figure 77.

Figure 83 shows the block diagram of the 5120A (Microsemi Corp.), and Table 12 shows the background noise. This instrument, probably the first implemented with fully digital architecture, has with no capability to work beyond the first Nyquist zone, thus the input frequency is limited to 30 MHz.



Figure 83 – Block diagram of the Microsemi 5120A PM noise test set. Based on the documentation available from the Microsemi web site.

Analysis frequency	L(ƒ), dBc/Hz		
	External ref	Internal ref	
1 Hz	-145	-120	
10 Hz	-155	_	
100 Hz	-165	_	
≥10 kHz	-175	-170	

Table 12 – Sensitivity of the Microsemi (Symmetricom) 5120A PM noise test set. Data are from the documentation of the instrument.

The Microsemi 5125 is a more modern instrument, which works up to 400 MHz input frequency by exploiting several Nyquist zones. Depending on the input frequency, it switches the antialiasing input filter and optimizes the sampling frequency by choosing the value between 104 MHz to 128 MHz. Of course, it uses the average cross-spectrum method to reduce the noise of the ADCs. The background noise is shown on Table 13. Unfortunately, the 5125A was discontinued in 2018.

Table 13 – Sensitivity of the Microsemi (Symmetricom) 5125A PM noise test set. Data are from the documentation of the instrument.

Offset frequency	Input frequency			
	10 MHz	100 MHz	400 MHz	
1 Hz	-140	-120	-110	
10 Hz	-150	-130	-120	
100 Hz	-157	-140	-130	
1 kHz	-162	-150	-140	
10 kHz	-165	-160	-150	
≥100 kHz	-165	-162	-155	

The 3120A (Microsemi Corp.) is a low-cost solution which derives from the TimePod designed by John Miles, and operates at the fixed sampling frequency of 78 MHz. Since the instrument is a sophisticated analog-to-digital interface which relies on an external computer for control, data analysis and display, it can have some additional features available as software update. The most interesting of them, according to our taste, is the measurement of AM noise.

Besides flexibility and sensitivity, the strength of the Microsemi family is the minimalist look of the front panel, however with a complete set of functions that focuses

strictly on PM noise and Allan variance analysis. A very small set of function keys enables to choose the quantity to display (L(f), ADEV, phase vs time, etc.) and to set the vertical scale, while most parameters are set automatically. The capability to work with arbitrary input and reference frequencies makes these instruments great for the measurement of DDSs, DACs, frequency dividers, and frequency multipliers. With these instruments, we could measure the PM noise of some DACs over 9 decades of Fourier frequency, observing flicker PM noise with a clean 1/f slope over 7.5 decades.

Three features we miss on the 5120A/5125A family. The first is the measurement of AM noise. All the building blocks are already there, thus we believe that this feature is only a matter of internal software and user interface. The second is the option to use two external references instead of sending the same reference to two input DACs. The obvious benefit is the rejection of the PM noise of the references. Based on the scheme of Figure 83, this feature is expected to cost only a minimum change in the internal software, and one additional connector on the front panel. The third feature is to allow some control on the software filters used to remove the spurs generated in the analog-to-digital conversions. In some cases, the user is interested in a noise component of the oscillator under test which falls in the narrow spectral regions where the filters remove the spurs. When this happens, the results are difficult to understand. This is an advanced topic, which was discussed in three workshops on the cross-spectrum method (European Cross Spectrum Phase Noise Measurement Workshop, 2014), (Cross Spectrum L(f) Measurement Workshop, 2015), (Cross Spectrum L(f) Measurement Workshop, 2017) organized by one of us (ER).

6.3.2 The Jackson Labs PhaseStation 53100A⁴

The PhaseStation is a new instrument, whose production started in August 2019. Having little first-hand experience, this Section is based on the material provided by the design team.

The architecture (Figure 84) derives from the TimePod. The core of the instrument an analog-to- digital interface with classical I-Q detection, which sends the baseband data to an external computer via USB interface for further processing. The bandwidth of the I-Q data is of the order of 1 MHz. Compared to the TimePod, there are significant differences. The passive power splitters (ferrite transformers) are replaced with active devices. The 1-200 MHz range is analyzed in four Nyquist zones, set by switching the antialiasing filters and the sampling frequency. The use of single-channel ADCs is an obvious choice for minimum crosstalk, because the crosstalk limits the noise rejection in the cross-spectrum analysis. The converters are now Analog Devices AD9265-125

⁴ The PhaseStation is now with Microchip.

because this component has a good thermal stability and reasonably low dissipation, in part related to the not-too-high clock frequency. These thermal characteristics help in achieving the thermal stability instrument, which is useful for the Allan variance beyond 1 s measurement time.

The external computer runs the TimeLab app. TimeLab performs five conceptual tasks.

- Evaluation of amplitude, phase and frequency from the I-Q data, with no loss of information. This feature is necessary for the data analysis down to low frequency.
- Decimation of the amplitude, phase and frequency data for multiresolution analysis. In fact, the frequency span of AM and PM spectral analysis is of multiple decades, which must be segmented for proper operation of the FFT algorithm. Therefore, the sampling frequency has to be progressively scaled down going towards the lower decades.
- Calculation of the cross spectrum, averaging on multiple acquisitions.
- Calculation of the ADEV, MDEV, Hadamard deviation (HDEV), and time deviation (TDEV).
- Plotting and storing the data



Figure 84 – Block diagram of the PhaseStation 53100A, as it is expected to appear on the User Manual. Courtesy of Jackson Labs Technologies, Inc. and Miles Design LLC, used with permission.

The hardware of the PhaseStation contains a radical piece of innovation, however obvious it may seem: the four NCOs and I-Q detectors are completely separated and independent machines. The four ADC inputs are all accessible and can be fed by four signals of different amplitude and frequency. Albeit some limitations set by the software may apply, let us see some fancy and useful examples of what the hardware can do.

The first example is the PM-noise measurement of a 500 MHz oscillator, which is out of the range of the instrument. A frequency divider by four brings the frequency in the range (500/4 = 125 < 200 MHz), but does not solve the problem because the result is corrupted by the noise of the divider. The use of two dividers is expected to solve the problem. Removing some jumpers, we send the two "125 MHz" to channels 3 and 4, so that the cross spectrum rejects the divider noise because the two dividers are statistically independent.

The second example is the stability measurement of a 5 MHz OCXO, where we expect an ADEV of less than 10^{-13} at $\tau = 1 \dots 10$ s (we have seen an example of such oscillator in Section 5.3.2). No OCXO is stable enough to be taken as a noise-free reference, and we have no access to exotic and expensive sources, like the Hydrogen maser or the cryogenic sapphire oscillator (150-300 k\$). The solution consists of using two reference OCXOs, sent to channel 1 and 2. The instability of such OCXOs can be a factor of 2-3 higher than that of the OCXO under test because their fluctuations are independent, and can be rejected.

A third example is the PM-noise measurement of a low-noise 100 MHz OCXO prototype. For technical reasons the reference has to be at 100 MHz (the PM noise of the 5-10 MHz OCXO is too high), and we suspect that the measure is corrupted by RF leakage. Let us proceed with two 100-MHz OCXOs used as the references as in the previous example, but we misalign them by random amounts (for example -170 Hz and +230 Hz, which is in the typical range of mechanical tuning). In this way the leakage is either eliminated owing to the high Q of the resonators, or its effect occurs at clearly identified frequencies in the PM-noise spectrum.



Figure 85 – Background phase noise of the PhaseStation 53100A. Courtesy of Jackson Labs Technologies, Inc. and Miles Design LLC, used with permission. Graphical editing and comments are ours.

The background noise of the PhaseStation is shown on Figure 85. This phase-noise spectrum is measured connecting one oscillator to the two inputs with the all jumpers inserted in the normal place, thus the four ADCs receive the same signal. Thus, the spectrum represents the PM-noise of the machine, not including the noise of the oscillator. The latter is common mode, and cancels. The FFT is calculated in segments broadly approximating the logarithmic resolution, which introduces a -5 dB/decade slope due to the reduction in the number of averages per unit of time in the segments at low frequency. Accordingly, the flicker region is seen as a slope of -15 dB/decade, as expected. Because the spectrum is quite irregular in the flicker region, the frequency segments cannot be identified, and a longer measurement time may result in further reduction of the background noise. At 5 MHz the white noise appears flat and regular. This makes us think that *m* is large enough for the average to hit the ultimate limit set by the hardware. The same is less clear for the other plots, obtained with shorter measurement time.

6.3.3 The NoiseXT DNA

6.3.4 The Rohde & Schwarz FSWP family of phase noise analyzers⁵

The FSWP is a recent and highly innovative, sophisticated and complex family of microwave phase noise analyzers working from 1 MHz up to 8/26.5/50 GHz, using the cross-spectrum technique. These instruments provide a variety of features, the AM and PM measurement of oscillators, the AM and PM measurement of two-port components via an internal synthesizer, the analysis of pulsed signals, the measurement of noise factor, and the test of VCOs, to mention the most important. Additionally, they can be used as a regular microwave spectrum analyzer, and also as a dual-channel FFT analyzer up to 10 MHz. This Section is based on the product documentation available online, on the article (Feldhaus & Roth, 2016), and on personal experience.

The scheme (Figure 86), however derived from the general principles stated earlier in this Chapter, looks rather different. The input signal is split into two channels, downconverted to an appropriate IF using two separate references, separate synthesizers, and separate mixers. The IF signal is digitized, at 100 MS/s on 16 bits, and processed by a sophisticated FPGA module (Figure 87).



Figure 86 – Basic scheme of the FSWP family of phase noise test sets. Based on the documentation available from the Rohde Schwarz web site, and on [98].

⁵ Also FSPN


Figure 87 – Detail of the FPGA processing inside the FSWP family of phase noise test sets. Based on the documentation available from the Rohde Schwarz web site, and on [98].

The LO signals are derived from two different reference oscillators, one of which is phase-locked to the other with a bandwidth of less than 0.1 Hz. Consequently, the PM noise in the two channels de-correlates progressively starting from 0.1 Hz, and the full benefit of the cross-spectrum method is achieved one decade beyond, thus at $f \ge 1$ Hz. Two high-level I-Q mixers are used to down convert the input signal. The I-Q conversion is more complex that a regular conversion. The advantage is that it is possible to fix the asymmetry error and the quadrature errors of the mixer after digitizing the outputs, and in turn to attenuate the residual AM. With this trick, the receiver achieves a typical AM rejection of 40 dB, instead of the 20 dB usually found in regular mixers. The I-Q down conversion does not increase the number of ADCs required with respect of the conceptual scheme of Figure 82. Thus, the cross-spectrum measurement is done with a total of four ADCs, two per channel.

The value of the IF frequency results from a technical tradeoff. Beyond 1 MHz Fourier frequency, the IF is set to zero (dc), where the mixer exhibits the highest gain and sensitivity. Beyond 1 MHz Fourier frequency, the IF is set to an appropriate value above 1 MHz. The reason for this choice is that the oscillator under test is free running, therefore a residual frequency offset Δf_0 is inevitable. At low Fourier frequencies, the harmonics of Δf_0 would fall in the analysis band and produce artifacts. Beyond 1 MHz Fourier frequency, the zero IF is allowed because the harmonics of Δf_0 fall in a region where they do not interfere with the measurement. The phase detection is based on the CORDIC algorithm. The aforementioned Δf_0 causes an increasing phase which wraps at the limits of $\pm \pi$, incompatible with the FFT processing. This is fixed with a feed-forward structure that converts the PM signal into a non-wrapping FM signal. The 20 dB per decade slope introduced by the PM to FM conversion is later removed by a digital filter. The noise of the two references and of the two channels is rejected as explained in Section 6.2, using (194) as the estimator

$$\widehat{S_{\varphi}}(f) = \frac{1}{K_{\varphi}^2} \left| \left\langle S_{yx}(f) \right\rangle_m \right|$$

Accordingly, the single-channel noise is rejected by 5 $\log_{10}(m)$ dB. The frequency range is divided into half-decade segments, so that the ratio RBW/*f* between the resolution bandwidth and the center frequency of each FFT bin is broadly constant. The sensitivity (background noise), vs carrier frequency f_0 and Fourier frequency *f*, is shown on Table 14 and Figure 88. The value shown include the phase noise and the instability of the two internal reference oscillators. Table 15 shows the improvement, by the number of averaged spectra. For offset frequencies below 1 Hz, such improvement impact of correlation is limited by the coupling between the two reference oscillators. The improvement achievable in this case ranges from 15 dB (nominal) at 0.1 Hz to 3 dB (nominal) at 30 mHz. The low phase noise of the two internal synthesizers (Figure 89) is one of the main virtues of the FSWP because this is the reference where the correlation engine starts improving.

Star	t offset 1 H	z, correlati	on factor =	1, freque	ncy referei	nce interna	I, signal le	vel ≥ 10 dE	3m	
RF input frequency	Offset frequency from the carrier									
	1 Hz	10 Hz	100 Hz	1 kHz	10 kHz	100 kHz	1 MHz	10 MHz	30 MHz	
1 MHz	-124	-142	-154	-172	-182	-182				
10 MHz	-121	-138	-148	-166	-176	-176	-176			
100 MHz	-101	-123	-146	-172	-176	-179	-181	-181	-181	
1 GHz	-81	-103	-126	-156	-172	-179	-179	-179	-179	
3 GHz	-71	-93	-116	-146	-162	-164	-169	-176	-176	
7 GHz	-64	-86	-109	-139	-158	-159	-163	-172	-172	
10 GHz	-61	-83	-106	-139	-158	-159	-163	-179	-181	
16 GHz	-57	-79	-102	-135	-154	-155	-159	-176	-177	
26 GHz	-53	-75	-98	-131	-150	-151	-155	-172	-173	
50 GHz	-47	-69	-92	-125	-144	-145	-149	-166	-167	

Table 14 – Typical Phase noise sensitivity with R&S®FSWP-B61 cross correlation option. Data are provided by A. Roth, R&S.



Figure 88 – Background noise of the FSWP phase noise test set measured at 100 MHz carrier with 1 hour averaging time. © Synergy Microwave Corp, reproduced with permission.

Table 15 – Improvement of phase sensitivity by number of correlations. Data are from the data sheet of the R&S®FSWP-B61 cross correlation option.

Improvement of phase noise sensitivity by number of correlations							
Offset frequencies ≥ 1 Hz							
Correlations	10	100	1000	10 000			
Improvement	5 dB	10 dB	15 dB	20 dB			



Figure 89 – Typical phase noise of the FSWP internal synthesizer. ©Rohde & Schwarz, reproduced with permission.

We have explained at the beginning of this Chapter that $S_{\varphi}(f)$, and equivalently L(f) defined correctly as $L(f) = (1/2)S_{\varphi}(f)$, give a valid measurement the phase noise even in the presence of large-angle and multiple-cycle swings, provided the phase detector work in this regime. This situation is often encountered at the highest microwave frequencies, where the phase noise gets inevitably large because of frequency-stability limitations. Figure 90 shows an example of measurement of 2, 20 and 40 GHz signals, where L(f) greatly exceeds 0 dBc/Hz in the left-hand part of the spectrum.



Figure 90 – Phase noise of some microwave signals measured with a FSWP phase noise analyzer.

6.4 Pitfalls and limitations of the cross-spectrum

measurements

The capability of averaging out the single-channel noise is the strength and the weakness of the cross-spectrum method. Nowadays digital electronics provides so large memory and computing power for cheap, that it is easy to average over millions of cross spectra. For example, the rejection of the single-channel noise given by (194) is of 30 dB with $m = 10^6$. This gives the *false impression* that it is sufficient to use larger m to increase the the sensitivity of the instrument. Common sense suggests that at some point other limitations apply.

Another common belief is that the background noise results always in a positive bias to the result. In other words, most people believe that the plot of $S_{\varphi}(f)$, or L(f) seen on the display is always higher than the true phase noise of the DUT. This is wrong in the case of cross spectrum measurements. The instrument may *underestimate* the DUT noise.

We first show how correlated effects hit on the measurement, then we go through the concept of uncertainty, and we discuss the physical phenomena. Some of the effects are so subtle that they escaped from the attention of Manufacturers and Government Labs for a long time. Recently, the limitations of the cross-spectrum method were addressed in three dedicated workshops (European Cross Spectrum Phase Noise Measurement Workshop, 2014), (Cross Spectrum L(f) Measurement Workshop, 2015), (Cross Spectrum L(f) Measurement Workshop, 2017) organized by one of us (ER). Unlike in regular conferences, the material circulated only among the participants, and no proceedings were published. However, some relevant material is published in the References (Gruson, Giordano, Rohde, Poddar, & Rubiola, 2017), (Nelson, Hati, & Howe, A collapse of the cross-spectral function in phase noise metrology, 2014), (Hati, Nelson, & Howe, 2016).

6.4.1 The effect of a disturbing signal

Let us introduce a disturbing signal represented as a phase $\xi(t)$ which affects the two channels with arbitrary coefficients ς_x , and ς_y . Accordingly, we replace x(t) and y(t) with

$$x(t) = K_{\varphi}[\varphi(t) - \psi(t) + \varsigma_x \xi(t)]$$
(207)

$$y(t) = K_{\varphi} \big[\varphi(t) - \theta(t) + \varsigma_{y} \xi(t) \big]$$
(208)

The quantities $\varsigma_x \xi(t)$ and $\varsigma_y \xi(t)$ can be any signal of unwanted origin impacting on the two channels. Since the amount of $\varsigma_x \xi(t)$ and $\varsigma_y \xi(t)$ is unknown, these signals fall in the category of uncertainties. The hypothesis that the two channels are equal has some practical limitations which allow ς_x and ς_y to be different, and to have different sign. For example, in a double balanced mixer the thermal coefficient comes from internal-diode mismatch. Thus, two nominally equal mixers will almost certainly have different thermal coefficient, and there is no reason for these coefficients to have the same sign. Likewise, the offset sensitivity to power.

Looking at the above definition of x(t) and y(t), all signals are *either* statistically independent or correlated, and *either* desired or unwanted. The following scheme accounts for all possible cases

Correlation	Desired	Unwanted
Fully correlated	$\varphi(t)$	$\xi(t)$
Fully independent	(none)	$\psi(t)$ and $ heta(t)$

The Fourier transforms of x(t) and y(t) are

$$X = K_{\varphi}(\Phi - \Psi + \varsigma_{\chi} \Xi)$$
(209)

$$Y = K_{\varphi} \left(\Phi - \Theta + \varsigma_{y} \Xi \right) \tag{210}$$

and the cross PSD $S_{yx} = \frac{2}{T} YX^*$ becomes

$$S_{yx} = K_{\varphi}^{2} \frac{2}{T} \left(\Phi \Phi^{*} - \Phi \Psi^{*} + \varsigma_{x} \Phi \Xi^{*} - \Theta \Phi^{*} + \Theta \Psi^{*} - \varsigma_{x} \Theta \Xi^{*} + \varsigma_{y} \Xi \Phi^{*} - \varsigma_{y} \Xi \Psi^{*} + \varsigma \Xi \Xi^{*} \right) (2-211)$$

where $\varsigma = \varsigma_{x} \varsigma_{y}$. When such signal is averaged on a large number of acquisitions, the

cross terms $\Phi \Psi^*$, $\Phi \Xi^*$, $\Theta \Phi^*$, $\Theta \Psi^*$, $\Theta \Xi^*$, $\Xi \Phi^*$, and $\Xi \Psi^*$ null. There results

$$S_{yx} = K_{\varphi}^{2} \frac{2}{T} (\Phi \Phi^{*} + \varsigma \Xi^{*})$$
(212)

The instrument readout is based on an estimator. For example, choosing the estimator (191), that is, $\widehat{S_{yx}} = \langle S_{yx} \rangle_m$, we get

$$\widehat{S_{\varphi}} = \frac{2}{T} \langle \Phi \Phi^* \rangle_m + \frac{2}{T} \langle \Xi \Xi^* \rangle_m \tag{213}$$

and therefore

$$\widehat{S_{\varphi}} = \langle S_{\varphi} \rangle_m + \varsigma \langle S_{\xi} \rangle_m \tag{214}$$

The measure is affected by the error term

$$\Delta S_{\varphi} = \varsigma S_{\xi} \tag{215}$$

Interestingly, $S_{\varphi}(f)$ and $S_{\xi}(f)$ are both positive because $\Phi\Phi^* = |\Phi|^2$ and $\Xi\Xi^* = |\Xi|^2$, but the sign of ς is arbitrary. Therefore, the error ςS_{ξ} can be positive or negative. Measuring a very-low-noise DUT, or in other extreme cases, a negative error term may prevail at some frequencies, where $\widehat{S_{\varphi}}(f) < 0$. Such outcome is a total nonsense. If we choose the absolute-value estimator (194), that is, $\widehat{S_{yx}} = |\langle S_{yx} \rangle_m|$, the result is always positive, with no warning.

6.4.2 Some concepts related to the measurement uncertainty

Whoever is faced to serious measurements of any physical quantity will at some point come across the concepts of uncertainty explained in the *International Vocabulary of Metrology* (VIM) (JCGM, 2012), the *Guide to the Expression of Uncertainty in Measurement* (GUM) (JCGM, 2008), and a bundle of related documents from the Joint Committee for Guides in Metrology (JCGM). All these documents are available free of charge from the BIPM web site https://bipm.org.

The approach described in the JCGM is formally correct, and also operational, and it is used routinely in the laboratories of primary metrology. Unfortunately, this culture is still absent in the practice of phase noise measurement, and learning it may require effort and patience. We will summarize the main points, and explain their implications on cross-spectrum measurements.

The components of measurement uncertainty are grouped into two categories, called Type A and Type B, depending on how the available data and information are combined to form the overall uncertainty.

The *Type A evaluation* of uncertainty is done by a statistical analysis of a time series of data obtained under defined measurement conditions.

The *Type B evaluation* of uncertainty is determined by means other than a Type A evaluation. The evaluation relies on authoritative published data, for example from metrological institutes, on calibration certificates, on environmental parameters like temperature and humidity, on available knowledge of aging, drift, etc., and also on personal experience. Ensemble statistics applies to this type of evaluation, but the uncertainty cannot be reduced by increasing the number of measurements or by increasing the measurement time.

We have seen in Section 6.2.1 that the single-channel background noise can be reduced by increasing *m*. This type of noise is suited to the *Type A* analysis. Conversely, the signals $\varsigma_x \xi(t)$ and $\varsigma_y \xi(t)$ yield $\Delta S_{\varphi} = \varsigma S_{\xi}$. This falls in the *Type B* evaluation, and the uncertainty cannot be reduced with statistical analysis on the time series.

A third concept found in the JCGM documents is the *null measurement uncertainty*. This concept applies to the measurement of a quantity that cannot be negative (in some cases, cannot be positive), when the outcome of the measurement is close to zero. The null measurement uncertainty is the smallest signal that the instrument can detect with a given probability. If the error bar crosses zero, the instrument gives only the upper bound.

The reader may be used to a format like a measured value S_{φ} with uncertainty ΔS_{φ} , usually meaning that the "true value" of S_{φ} falls in the interval $S_{\varphi} \pm \Delta S_{\varphi}$ with a probability of 95%. What happens if the lower error bar $S_{\varphi} - \Delta S_{\varphi}$ crosses zero in a region of f? Of course, a negative PM noise is not allowed for physical reasons. In this case, the *null measurement uncertainty* applies. In proper metrological terms, the outcome of the experiment is that we have *measured zero phase noise*, with *zero uncertainty* equal to ΔS_{φ} . The nonsense associated to the negative outcomes of S_{φ} is no longer a problem.

In the case of the null measurement uncertainty, expressing the uncertainty in dB is often the result of a wrong approach.

After learning about the null measurement uncertainty, whenever we come across measurements taking large values of *m*, long measurement time, or too low phase noise spectra, we should try to understand better the uncertainty. For reference, a measurement time in excess of a few minutes starting at f = 100 Hz is probably the maximum that can be accepted without warning.

6.4.3 Thermal energy in the input power divider

A power divider is necessary to split the DUT signal into the two channels. In most cases, the power divider is a directional coupler internally terminated at one port (Figure 91). This choice provides low loss, uniform loss in a wide range of frequency, good isolation between the two outputs, and the good impedance matching at all ports. The Y resistive power splitter is seldom used because of the inherent 6-dB loss and poor isolation. However, the resistive power splitter is superior to the directional coupler in impedance matching and in loss flatness over extreme frequency range.

The problem is that the thermal energy *kT* inherent in the splitter's internal dissipation introduces a systematic error, or bias, in the spectrum. The same type of error is found in the directional coupler and in the Y resistive coupler, just with a different value of the bias. We explain what happens with the directional coupler, addressing the reader to the References (Gruson, Giordano, Rohde, Poddar, & Rubiola, 2017) and (Hati, Nelson, & Howe, 2016) the details of both.



Figure 91 – A loss-free power splitter is a directional coupler terminated to one port. To the extent of our digression about the thermal energy, the 90^o version of the directional coupler is fully equivalent to that shown.

Let us recall that the random voltage e_n across a resistor R at temperature T has mean square value $\langle e_n^2 \rangle = 4kTR$ in 1 Hz bandwidth. We apply this signal to the inputs of a power splitter, as shown on Figure 91. The circuit is impedance matched to R_0 at all ports. We focus on the white PM noise region, where $S_{\theta}(f)$ is the lowest. The oscillator under test delivers a power P_0 , and its output resistor R_a has the equivalent temperature T_a that results from $S_{\theta} = kT_a/P_0$. The dark port is terminated to the resistor R_b at the temperature T_b , the room temperature or the temperature inside the instrument. The two signals at the power-splitter output are

$$c = \frac{V_g}{2\sqrt{2}} + \frac{e_a - e_b}{2\sqrt{2}}$$
(216)

$$d = \frac{V_g}{2\sqrt{2}} + \frac{e_a + e_b}{2\sqrt{2}}$$
(217)

The term "2" at the denominators is necessary because of impedance matching at the left-hand side of the coupler, and the term " $\sqrt{2}$ "is due to the energy conservation in the loss-free power splitter. The negative sign, represented as a 180° phase shift on Figure 91, is necessary for energy conservation in a loss-free device. The explanation is found in most microwave textbooks, for example (Pozar, 2012), in the section about the Scatter Matrix. Now we expand the cross PSD as we did earlier in this Chapter

$$S_{dc} = \frac{2}{\mathcal{T}} YX^* = \frac{2}{\mathcal{T}} \frac{E_a - E_b}{2\sqrt{2}} \frac{E_a^* + E_b^*}{2\sqrt{2}} = \frac{2}{\mathcal{T}} \frac{E_a E_a^* + E_b E_b^*}{8}$$
(218)

Notice that the measurement time is temporarily denoted with T because here T is the temperature. Using the PSD of the thermal voltage $S_e = (2/T) EE^* = 4kTR_0$, we get

$$S_{dc} = \frac{1}{2}k(T_a - T_b)R_0$$
(219)

Interestingly, this technique has been known long time ago as a noise comparator tool (Allred, 1962), and in the early time of radio astronomy (Hanbury Brown & Twiss, 1954). For phase noise measurements, the trivial factor $\frac{1}{2}$ cancels with half the input power going in each output of the coupler. The relevant fact about (219) is that S_{dc} is proportional to $k(T_a - T_b)$, instead of kT_a . Obviously, the same happens with S_{yx} after phase detection. If this error is not accounted for, the instrument readout is

$$S_{\varphi} = \frac{k(T_a - T_b)}{P_0} \qquad \text{instead of} \qquad S_{\varphi} = \frac{kT_a}{P_0} \tag{220}$$

with a systematic error

$$\Delta S_{\varphi} = -\frac{kT_b}{P_0} \tag{221}$$

Example 22. A 125-MHz OCXO has output power of 13 dBm, and the white PM noise floor displayed by the test set is of -186 dBc/Hz at $f \ge 10 \text{ kHz}$. We evaluate the error due to the thermal energy in the input coupler at the internal instrument temperature of 40 °C.

First, we convert the data into proper SI units. Thus, the dark port temperature is $T_C =$ 313 K, the carrier power is $P_0 = 20$ mW, and the white PM noise is $S_{\varphi} =$

 $5 \times 10^{-19} \text{ rad}^2/\text{Hz}$. Using $S_{\varphi} = kT_{read}/P_0$, we get the equivalent temperature $T_{read} = 724$ K associated to the noise floor, from the instrument readout. Since the measure is

biased by the thermal energy as explained, the correct value is $T_{eq} = T_{read} + T_c = 1038$ K, hence $S_{\varphi} = kT_{eq}/P_0 = 7.16 \times 10^{-19} \text{ rad}^2/\text{Hz}$. The bias error is $5 \times 10^{-19} - 7.16 \times 10^{-19} = -2.16 \times 10^{-19} \text{ rad}^2/\text{Hz}$, that is, -1.5 dB referred to the instrument readout.

Notice that the fractional error gets larger as the displayed floor approaches kT_C/P_0 . With a displayed noise of -190 dBc/Hz and 13 dBm power, the error exceeds 3 dB.

6.4.4 The effect of AM noise

The measurement of phase noise should be independent of amplitude noise. However, the saturated mixer suffers from a residual sensitivity to AM noise through the sensitivity of the DC offset to power. In simple terms, when RF and LO are in quadrature, a residual DC offset is present at the IF output. Such offset is due to the imperfect balance of baluns and diodes, and it is affected by the power at both inputs. Hence, a power fluctuation (AM noise) ends into a fluctuation of the DC offset, which adds to the regular signal $K_d \varphi$. This concept is described quantitatively by replacing $V = K_{\varphi} \varphi$ with

$$V = K_{\varphi}\varphi + K_{RF}\alpha_{RF} + K_{LO}\alpha_{LO}$$
(222)

where the AM noise of the RF and LO inputs appears explicitly. Of course, looking at the mixer output it is impossible to divide which part of *V* is due to phase noise and which is due to power fluctuations. Notice that K_{RF} and K_{LO} can be unpredictably positive or negative, and change with power and frequency.

Our treatise is limited to the double balanced mixer. The ADCs used in digital phase detectors suffer from AM noise leakage into the estimation of PM noise because of nonlinearity and other problems. This is an open issue, still not addressed in the technical literature.

When the mixer is saturated correctly, we expect a rejection of AM of the order of 20 dB. This means that the overall effect of K_{φ}/K_{LO} and K_{φ}/K_{RF} of the order of 10, a comfortable value in most cases, but poor or insufficient in other cases. For example, microwave-photonic devices often have larger 1/f AM noise than 1/f PM noise. Likewise, some quartz oscillators optimized for the lowest PM noise floor have larger AM noise.



Figure 92 – The effect of AM noise on a cross-spectrum PM noise test system. Reprinted from *Frequency and Amplitude Stability in Oscillators*, lecture slideshow, CC BY E. Rubiola, 2019 (Rubiola E. , Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists, Public material, Creative Commons 4.0 CC-BY, 2019).

In the cross-spectrum measurement, the effect of AM noise may lead to wrong or nonsensical results because the DUT AM noise is not rejected, and becomes comparatively large after averaging out the instrument background noise. This idea is sketched on Figure 92. We apply the simple concept expressed by (222) to the scheme of Figure 77. The two signals at the input of the FFT analyzer are

$$x = K_{\varphi}(\varphi - \psi) + K_A \alpha_A + K_X \alpha \tag{223}$$

$$y = K_{\varphi}(\varphi - \theta) + K_B \alpha_B + K_Y \alpha \tag{224}$$

where K_A and K_X are the offset sensitivity of the upper mixer to the amplitude α_A of the reference, and to the amplitude α of the DUT; K_B and K_Y are the same sensitivities for the lower mixer. Applying the statistical reasoning we are now familiar with, for large m we get

$$S_{yx}(f) = K_{\varphi}^2 S_{\varphi}(f) + K_Y K_X S_{\alpha}(f)$$
(225)

In fact, by inspection on the scheme, α , α_A and α_B are statistically independent, and only α is common to the two channels. As before, φ , ψ and θ are statistically independent, and only φ is common to the two channels. Consequently, the usual readout formula gives the wrong value

$$S_{\varphi}(f) = \frac{1}{K_{\varphi}^2} S_{yx}(f)$$

instead of

$$S_{\varphi}(f) = \frac{1}{K_{\varphi}^{2}} S_{yx}(f) - \frac{K_{Y}K_{X}}{K_{\varphi}^{2}} S_{\alpha}(f)$$
(226)

The "error" term

$$\Delta S_{\varphi}(f) = -\frac{K_Y K_X}{K_{\varphi}^2} S_{\alpha}(f)$$
(227)

is actually a Type-B uncertainty because $S_{\alpha}(f)$, K_X and K_Y are not known. One may object that the error can be easily fixed by measuring these quantities. At a closer sight, only $S_{\alpha}(f)$ can be measured in a simple way. Eventually, some commercial instruments already have the additional hardware for the measurement of AM noise. Conversely, K_X and K_Y depend on the carrier power and frequency, and on the experimental conditions. Thus, a factory characterization is not an option. The direct measurement is not simple because it requires an amplitude modulator in series to the DUT, and such modulator leaves a residual PM.

When the sensitivity to AM noise is annoying, it is sometimes possible to find a working point where the effect of AM is mitigated (Rubiola & Boudot, The Effect of AM Noise on Correlation Phase-Noise Measurements, 2007), (Cibiel, Regis, Tournier, & Llopis, 2002). The two main options are

- Setting LO and RF slightly off the quadrature
- Inject a weak DC current at the mixer IF output.

In all cases, the working point must be found experimentally, by measuring the output in the presence of a small and controlled AM in the same conditions of the PM noise measurement. Experience indicates that an optimum point is more easily found in microwave mixers (microstrip balun) than in RF mixers (transformer balun), and that it may not exist at all.

In the end, fixing $\Delta S_{\varphi}(f)$ makes the difference between an industrial measurement and a scientific experiment.

6.5 The bridge (interferometric) method

The bridge method, shown on Figure 93, enables the measurement of two-port components with the lowest background noise, and the highest rejection of spurs and interferences.

After adjusting the phase and the amplitude in the bridge, the carrier is nulled in the $\Delta(t)$ signal. Thus, all the carrier power goes to $\Sigma(t)$, and is dissipated in the termination. However, the signal $\Delta(t)$ contains the noise sidebands of the DUT, after the obvious loss of the directional coupler. The noise sidebands are amplified and down-converted to DC by the mixer. The latter detects the DUT phase or amplitude, or any combination of, depending on the phase of the LO signal.

The appealing features of this method rely on the following concepts.

- The bridge is implemented with passive components (directional couplers, attenuators, and phase shifters), which exhibit low PM noise as compared to active components.
- The amplifier works in small-signal regime, where it is fully linear. We have seen earlier in this Chapter that the 1/f PM noise originates by up-conversion of the near-DC flicker, thus the noise sidebands introduced by the amplifier are approximately proportional the input power. In turn, PM noise is proportional to the sideband-to-carrier ratio. Consequently, with high carrier suppression the 1/f PM noise is virtually absent.
- Microwave amplification rises the useful signal (the DUT noise) from the background before down conversion. This results in low interference from 50-60 Hz lines because the sensitivity of microwave circuits to low-frequency magnetic fields is very low. Additionally, amplification helps to achieve low residual white PM noise because the low noise figure of the amplifier is more favorable than the relatively high conversion loss of a saturated mixer.
- Unlike the traditional systems described earlier, the mixer operates in linear regime, with the LO input saturated and $P_{RF} \ll P_{LO}$. Thus, flicker comes only from the LO, and its impact is reduced proportionally to the microwave gain.



Figure 93 – Scheme of the bridge phase measurement.

The method was first proposed by Sann in 1968 for the measurement of microwave amplifiers (Sann, 1968). At that time, the bridge was followed immediately by the mixer, with no microwave amplification. The amplifier was introduced by Labaar in 1982, implementing a RF version of the method (Labaar, 1982). Further development had to wait until 1998, when this method was used extensively at the University of Western Australia and at the Laboratoire de Physique et Metrologie des Oscillateurs (now FEMTO-ST Institute) in France for rather extreme measurements (Ivanov, Tobar, & Woode, 1998), (Rubiola, Giordano, & Groslambert, Very high frequency and microwave interferometric phase and, 1999). For example, the flicker frequency fluctuations of high-stability HF quartz resonators ($\sigma_{\rm y}$ as low as a few parts in 10⁻¹⁴) can only be measured with this method (Rubiola, Groslambert, Brunet, & Giordano, 2000). A more sophisticated version uses the cross spectrum method at the output of the bridge (Rubiola & Giordano, Advanced interferometric phase and amplitude noise measurements, 2002). With this scheme, we could measure the 1/*f* phase noise of VHF power splitters (of the order of -170 dBc/Hz at f = 1 Hz, with 100 MHz carrier). The measurements reported are a challenge for scientists. By contrast, for less demanding applications, still out of reach for commercial instruments, the bridge method is not difficult for an experienced engineer or for a skilled amateur.

6.5.1 Phase-to-voltage gain and background noise

The phase-to-voltage gain K_{φ} is given by

$$K_{\varphi}^{2} = \frac{A^{2}R_{0}P_{0}}{\ell^{2}}$$
(2-228)

where A^2 is the power gain of the amplifier, R_0 is the characteristic impedance, P_0 is the DUT output power, and ℓ^2 is the SSB loss of the mixer. This is easily proved taking a signal of carrier power is P_0 with sinusoidal PM where each sideband has power P_s . Combining the two sidebands, the mean square phase modulation is $\langle \varphi^2 \rangle = 2P_s/P_0$. Now we follow the signal path on Figure 93, from the DUT to the mixer. Neglecting the loss of the second coupler for monitoring, one sideband has power $P_s/2$ at the amplifier input, $P_s A^2/2$ at the amplifier output, and $P_s A^2/2\ell^2$ at the mixer output. Combining the two sidebands at the mixer output, we get a power $2P_s A^2/\ell^2$, thus a mean square voltage $\langle V^2 \rangle = 2P_s A^2 R_0/\ell^2$. Finally, the gain results from $K_{\varphi}^2 = \langle V^2 \rangle/\langle \varphi^2 \rangle$.

The background noise is

$$S_{\varphi}(f) = 2 \frac{FkT}{P_0}$$
 background noise (2-229)

To prove this, we start from the random noise FkT at the amplifier input, hence A^2FkT at the amplifier output. The PSD at the mixer output is $2A^2FkT/\ell^2$, where the factor 2 comes from adding USB and LSB as statistically independent signals. Thus, the voltage PSD is $S_v = 2A^2FkTR/\ell^2$, and the background results from $S_{\varphi} = S_v/K_{\varphi}^2$.

Trivial losses in the bridge, and in the directional coupler used to monitor he carrier suppression, apply. They impact on K_{φ} and on the background noise.

6.5.2 Building your own system

This kind of system has to be assembled for specific DUT. The best approach consists of assembling the bridge first, and pre-adjusting it using a network analyzer. At microwave frequencies, the phase difference inside the bridge is difficult to predict, but setting the bridge is rather simple because commercially available line stretchers enable precise phase adjustment in a comfortable range. For reference, a range of 100 ps is equivalent to one period at 10 GHz. Conversely, in the HF and VHF region the wavelength is too long for line stretchers, but the phase can be predicted based on the components. In this case, phase matching can be achieved by try-and-error with cables, and fine-tuned with a narrow-range variable phase shifter. Semirigid cables are preferred. Inspecting on the bridge with a network analyzer, we observe a dip at an unpredictable frequency determined by the delay difference between the two arms. The dip can be deepened by adjusting the attenuation, and shifted to the frequency of interest by adjusting the delay. Because we need only the absolute value of the transfer function, the network analyzer can be replaced with a spectrum analyzer with tracking oscillator. Fairly high carrier suppression at the desired frequency is usually obtained after a small number of iterations. Fine tuning should be done with the complete system, with a spectrum analyzer at the monitor tap.

Best results are obtained with 20-45 dB RF/microwave gain. Higher gain results in higher K_{θ} and lower 50-60 Hz spurs, but it is more difficult to adjust the bridge for sufficiently low residual carrier at mixer input. For full linear operation, the total integrated noise should be kept at least 30 dB lower than the 1-dB compression point of the amplifier. Bandpass filtering at an intermediate stage of the amplifier may be useful.

The phase shifter at the LO input of the mixer must be set for the detection of PM noise. A low-index amplitude modulation in the DUT path is the best choice. Driving this modulator with an audio-frequency tone, we adjust the phase shifter to null the corresponding spectral line on the FFT analyzer. A lock-in amplifier, if available, provides the best result.

Interestingly, the component in the dashed rectangle on the right-hand side of Figure 93 (mixer, DC low-noise amplifier and FFT analyzer) are the basic ingredients of a traditional single-channel saturated-mixer phase noise analyzer. If available, a commercial instrument should be used as a replacement, keeping the power at the RF port of the mixer low as described. More than the hardware, the reader will appreciate the computer interface, the software and the ergonomics of the commercial instrument.



Figure 94 – Example of implementation intended for 8-10 GHz operation. Reprinted from E. Rubiola, V. Giordano and J. Groslambert, "Very high frequency and microwave interferometric phase and amplitude noise measurements," *Rev Sci Instrum*, vol. 70, no. 1, pp. 220-225, January 1999, with the permission of AIP Publishing. Comments are ours.

6.5.3 A practical example

Figure 94 shows an example intended for the measurement of phase noise in the 8-10 GHz band (Rubiola, Giordano, & Groslambert, Very high frequency and microwave interferometric phase and, 1999). This implementation relates to our early experiments. At that time, we focused on the lowest background noise, rather than seeking for a decent compromise between complexity, manpower and background noise. The bridge is implemented with a Wilkinson power splitter, a multi-turn variable attenuator, a micrometric line stretcher, and a 4-port hybrid junction. Unlike shown on Figure 93, the Σ port of the hybrid junction is re-used to pump the mixer, instead of dissipating the power in a termination. This choice was made before collecting extensive experience, and is not recommended. The problem is that adjusting the phase at the mixer LO input interacts with the bridge balance, and consequently multiple interactions are needed to null the carrier and to achieve proper detection of PM noise with minimum interference from AM noise. The amplifier has a gain of 42.5 dB in two stages, with a \approx 30 MHz bandpass between the first and the second stage, and a noise figure of 2 dB. Isolators here and there proved to be useful. Accounting for gain and losses, we measured $K_{\varphi} = 34.3 \text{ dBrad}^2/\text{Hz}$, fairly close to the calculated value of 33.9 dBrad²/Hz. The background noise (Figure 95), measured at +14 dBm DUT power with a null DUT in the bridge (short cable), is of $-154 \text{ dBrad}^2/\text{Hz}$ flicker extrapolated to 1 Hz, and $-182 \text{ dBrad}^2/\text{Hz}$ white.



Figure 95 – Background noise of the implementation shown on Figure 94, measured at +14 dBm DUT output power, and 9.13 GHz carrier frequency. Reprinted from E. Rubiola, V. Giordano and J. Groslambert, "Very high frequency and microwave interferometric phase and amplitude noise measurements," *Rev Sci Instrum*, vol. 70, no. 1, pp. 220-225, January 1999, with the permission of AIP Publishing. Graphical editing and comments are ours.

6.6 Artifacts and oddities often found in real world

We provide some typical examples of artifacts, spurs, and other problems often found in the measurement of phase noise.

Figure 96 shows an example of phase noise spectrum of a two port device, measured with a saturated mixer. Two nominally equal 2.3 GHz Al-N-Sapphire HBARs are inserted, one in each arm of the mixer, because this configuration rejects the noise of the master oscillator. The background noise consists of flicker and white noise, -132 dBrad² and -162 dBrad²/Hz respectively, thus

$$S_{\varphi}(f) = 6.3 \times 10^{-14}/f + 6.3 \times 10^{-17} \text{ rad}^2/\text{Hz}$$

This is in agreement with earlier discussions and examples in this Chapter. Let us look closer at this spectrum.

In the region A, we see that the fractional frequency resolution, i.e., the number of points per decade, is significantly smaller than in the other regions. This is rather usual because the frequency resolution in this region is limited by the measurement time. In the region B, we see spurs from the power grid at 50 Hz and multiples (these experiments are done in Europe) on both DUT and background noise spectra. In most cases, these spurs are picked up at the input of the LNA that follows the mixer. As usual, odd-order harmonics (50, 150, 250... Hz) are quite strong, while even-order harmonics (100, 200... Hz) are barely noticeable. These spurs seem to end at 1 kHz, leaving room to a bump in the small region C. At closer sight, this bump is the envelope of the power-grid spurs. These spurs cannot be seen separately because the analyzer has not sufficient resolution in this region. Other spurs show up in the region D, between 35 kHz and 80 kHz. This type of spurs is usually due to switching power circuits in the experiment, or around in the room, conducted through cables. Common-mode filters (ferrites) often help to reduce these spurs.



Figure 96 – Example of phase noise spectrum (2.3 GHz AlN-Sapphire HBAR resonator). Reprinted from R. Boudot, G. Martin, J.-M. Friedt and E. Rubiola, "Frequency flicker of 2.3 GHz AlN-sapphire high-overtone bulk acoustic resonators," *J Appl Physics*, vol. 120, no. 22, pp. 223103 1-8, 14 December 2016 (Boudot, Martin, Friedt, & Rubiola, 2016), with the permission of AIP Publishing. Comments are ours.

February 6, 2025

Figure 97 shows the phase noise of a 100 MHz OCXO, measured with the crossspectrum method using saturated double-balanced mixers. At $f \le 200$ Hz, the plot is quite smooth, and we identify clearly the $1/f^4$ and the $1/f^3$ regions, typical of quartz oscillators. The white PM noise is also well identified, in the region between 100 kHz and 10 MHz. By contrast, the spectrum is quite irregular between 200 Hz and 100 kHz. There is something wrong in the measurement, and we are unable to identify the 1/f or $1/f^2$ terms of the polynomial law we expect.

Starting from the left-hand side of the plot, we see a small discontinuity in A, at the transition between two decades. The most probable explanation is a change in the sampling frequency, affecting the white noise floor of the ADC. However small in this case, a signature like this is often seen in PM noise spectra. In B we see a bump, where the plot is quite irregular. The best interpretation we have is that the averaging process, necessary in the cross-spectrum method, is still not converging. This may also reveal Btype (systematic) uncertainty. The region C is weird because the plot is most irregular, and lower than the white PM floor. We have good reasons to believe that there is some anti-correlated process polluting the measure in this region, probably due to AM noise. The cross spectrum may even be negative, made positive by the absolute-value estimator. Of course, region C cannot be trusted. Reference (Gruson, Rus, Rohde, Roth, & Rubiola, 2019) addresses these problem in depth, yet it may take a real effort to get in because of the complexity of the problem underneath. The region D contains an artifact which appears as a notch. The narrowness recalls a spur, but spurs take the shape of a peak, not of a notch. The one and only credible explanation for this pattern is the presence of a spur with negative correlation in the two branches. However narrow and irrelevant it seems, this spur provides the evidence of anti-correlated artifacts, and reinforces the hypothesis that something similar happens in the region C. The region E is the signature of a low-pass filter. This pattern is the bandwidth limit of the system, which is in principle known, and should not be regarded as an artifact.



Figure 97 – Phase noise spectrum of a 100 MHz oscillator. The spectrum is from DDC Electronics Ltd., a subsidiary of Data Device Corporation (DDC). Graphical editing and comments are ours.

The last piece of our collection is the set of phase noise spectra shown on Figure 98. The device under test is a 1-GHz source having unusually high level of third harmonic distortion, -6.9 dBc. This source is measured with double balanced mixers and the cross-spectrum method. The different curves are obtained by shifting the third harmonic from 0° to 360° in 30° steps. This plot is intended to alert the reader that large harmonic distortion and impedance mismatch should always avoided in PM noise measurement. Why this large spread of values occurs, up to 12 dB at 10 kHz, and why the effect of the distortion is so irregular vs frequency, is not clear. The mixer is highly nonlinear, and for this reason it cannot be impedance matched at the inputs. Impedance mismatch and back reflections are not the same for first harmonic and for third harmonic. Furthermore, impedance mismatch affects the out-to-out isolation of the input power splitter, and in turn introduces coupling between the two channels of the instruments. Of course, none of these effects is under control, and erratic inconsistent results are around the corner. Experience suggests that a 3-dB attenuator inserted on each input connector of the mixer improves the impedance matching and helps significantly in avoiding inconsistencies. Introducing a low pass filter at the input of the instrument is also a good idea when the source under test has significant harmonic distortion. In the laboratory practice, even a small set of filters covers most cases. A

relatively new generation of filters is now commercially available, impedance matched also in the stopband. The use of these filters is highly recommended.



Figure 98 – Phase noise of a 1-GHz oscillator having -6.9 dBc third harmonic distortion. The third harmonics is shifted from 0° to 360° in 30° steps. © Synergy Microwave Corp., reproduced with permission.

7 References

- Adler, R. (1946, june). A Study of Locking Phenomena in Oscillators. *Proc IRE, 34*(6), 351-357.
- Allan, D. W., & Barnes, J. (1981). A Modified "Allan Variance" With Increased Oscillator Characterization Ability. *Proc Internat Frequency Control Symp*, (pp. 470-475).
- Allred, C. M. (1962, October-December). A Precision Noise Spectral Density Comparator. *J Research NBS C, 66C*(4), 323-330.
- Arthur, M. G. (1974). *The Measurement of Noise Performance Factors: A Metrology Guide* (Vol. NBS Monograph no 142). Boulder, CO, USA: US Dept of Commerce / NBS (now NIST).
- Ashburn, P. (2003). SiGe Heterojunction Bipolar Transistors. John Wiley & Sons.
- Baghdady, E. J., Lincoln, R. N., & Nelin, B. D. (1965). Short-Term Frequency Stability: Characterization Theory, and Measurement. In A. R. Chi, *Short-Term Frequency Stability* (pp. 65-87). NASA SP-80.
- Barash, S., & Ritov, Y. (1993, March). Logarithmic Pruning of the FFT Frequencies. *IEEE Transact Signal Processing*, *41*(3), 1398-1400.
- Barnes, J. A., & et al. (1971, May). Characterization of Frequency Stability. *IEEE Transact Instrum Meas*, *20*(2), 105-120.
- Baugh, R. A. (1972). Low-Noise Frequency Multiplication. *Proc Internat Frequency Control Symp*, (pp. 50-54).
- Benkler, E., Lisdat, C., & Sterr, U. (2015, July). On the relation between uncertainties of weighted frequency averages and the various types of Allan deviations. *Metrologia*, 52(4), 565-574.
- Bennett, W. R. (27, July). Spectra of Quantized Signals. Bell Sys Tech J, 27(4), 446-472.
- Boudot, R. (2006, December 7). *Oscillateur micro-ondes à haute pureté spectrale, PhD thesis.* PhD Thesis, University of Franche Comté, Besançon, France.
- Boudot, R., & Rubiola, E. (2012, December). Phase noise in RF and microwave amplifiers. *IEEE Transact Ultrason Ferroelec Frequency Control*, 59(12), 2613-2624.
- Boudot, R., Guerandel, S., & de Clercq, E. (2009, October). Simple-Design Low-Noise NLTL-Based Frequency Synthesizers for a CPT Cs Clock. *IEEE Transact Instrum Meas, 58*(10), 3659-3665.
- Boudot, R., Martin, G., Friedt, J.-M., & Rubiola, E. (2016, December 14). Frequency flicker of 2.3 GHz AlN-sapphire high-overtone bulk acoustic resonators. *J Appl Physics*, *120*(22), 223103 1-8.

- Bregni, S. (2016, April). Twenty-Five Years of Applications of the Modified Allan Variance in Telecommunications. *63*(4), 520-530.
- Brochard, P., Sudmeyer, T., & Schlit, S. (2017, November). Power Spectrum Computation for an Arbitrary Phase Noise Using Middleton's Convolution Series: Implementation Guideline and Experimental Illustration. *IEEE Transact Ultrason, Ferroelec, Frequency Control,* 64(11), 1766-1765.
- Calosso, C. E., & Rubiola, E. (2013). The Sampling Theorem in Pi and Lambda Digital Frequency Dividers. *Proc Internat Frequency Control Symp* (pp. 960-962). IEEE.
- Calosso, C. E., & Rubiola, E. (2016., December). Phase Noise and Jitter in Digital Electronics. arXiv:1701.00094 [physics.ins-det].
- Calosso, C. E., Clivati, C., & Micalizio, S. (2016, April). Avoiding Aliasing in Allan Variance: An Application to Fiber Link Analysis. *IEEE Transact Ultrason Ferroelec Frequency Control, 63*(4), 646-655.
- Calosso, C. E., Gruson, Y., & Rubiola, E. (2012). Phase Noise in DDS. *Proc Internat Frequency Control Symp* (pp. 777-782). IEEE.
- Chi, A. R. (1965). Short-Term Frequency Stability. NASA SP-80.
- Christiansen, C. J., & Pearson, G. L. (1937, April). Spontaneous Resistance Fluctuations in Carbon Microphones and Other Granular Resistances. *Bell Sys Tech J*, *15*(2), 197-223.
- Cibiel, G., Regis, M., Tournier, E., & Llopis, O. (2002, June). AM Noise Impact on Low Level Phase Noise Measurements. *IEEE Transact Ultrason Ferroelec Frequency Control, 49*(6), 784-788.
- Cross Spectrum L(f) Measurement Workshop. (2015). Side event of the European Frequency and Time Forum and Internat Frequency Control Symp Joint Meeting, organized by E. Rubiola, C. W. Nelson, A. Hati, D. A. Howe, Denver, CO, USA: The Proceedings were distributed only to the participants.
- Cross Spectrum L(f) Measurement Workshop. (2017, July 12). Side event of the European Frequency and Time Forum and Internat Frequency Control Symp Joint Meeting, organized by E. Rubiola and C. W. Nelson, Besançon, France: The Proceedings were distributed only to the participants.
- Donley, E., & et al. (2022). *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities. IEEE Standard 1139* (2022 ed.). IEEE. doi:10.1109/IEEESTD.2022.9973001
- Driscoll, M. M. (1990, April). Phase Noise Performance of Analog Frequency Dividers. *IEEE Transact Ultrason Ferroelec Frequency Control*, 37(4), 295-301.
- Egan, W. F. (1990). Modeling Phase Noise in Frequency Dividers. *IEEE Transact Ultrason Ferroelec Frequency Control, 37*(4), 307-315.

- Elliott, D. S., Roy, R., & Smith, R. J. (1982, July). Extracavity Laser Band-Shape and Bandwidth Modification. *Phys Rev A*, *26*(12), 12-18.
- European Cross Spectrum Phase Noise Measurement Workshop. (2014, December 18). organized by E. Rubiola, LNE Headquarters, Paris, France: The Proceedings are available only to the participants.
- Feldhaus, G., & Roth, A. (2016). A 1 MHz to 50 GHz Direct Down-Conversion Phase Noise Analyzer with Cross-Correlation. *Proc European Frequency and Time Forum*, (pp. 1-4).
- Ferre-Pikal, E. S., Walls, F. L., & Nelson, C. W. (1997, March). Guidelines for Designing BJT Amplifiers with Low 1/f AM and PM Noise. *IEEE Transact Ultrason Ferroelec Frequency Control*, 44(2), 335-343.
- Ferre-Pikal, E., & al., e. (2009). IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities. IEEE Standard 1139.
- Francois, B., Calosso, C. E., Danet, J. M., & Boudot, R. (2014, September). A low phase noise microwave frequency synthesis for a high-performance cesium vapor cell atomic clock. *Rev. Sci. Instrum.*, *85*(9), 094709 1-7.
- Friis, H. T. (1944, July). Noise Figures of Radio Receivers. Proc IRE, 32(7), 419-422.
- Gao, X., & Klumperink, A. M. (2009, February). Jitter Analysis and a Benchmarking Figure-of-Merit for Phase-Locked Loops. *IEEE Transact Circuits Systems II*, 56(2), 117-122.
- Gentile, K. (1999). Chapter 4: The Effect of DAC Resolution on Spurious Performance. In *A Technical Tutorial on Digital Signal Synthesis* (pp. 16-38). Analog Devices Inc.
- Godone, A., Micalizio, S., & Levi, F. (2008). RF Spectrum of a Carrier with a Random Phase Modulation of Arbitrary Slope. *Metrologia*, 45, 313-324.
- Gruson, Y., Giordano, V., Rohde, U., Poddar, A., & Rubiola, E. (2017, March). Cross-Spectrum, PM Noise Measurement, Thermal Energy, and Metamaterial Filters. *IEEE Transact Ultrason Ferroelec Frequency Control*, 64(3), 634-641.
- Gruson, Y., Rus, A., Rohde, U. L., Roth, A., & Rubiola, E. (2019, December). Artifacts and Errors in Cross-Spectrum Phase Noise Measurements. arXiv:1912.10449 [physics.ins-det].
- Gufflet, N., Bourquin, R., & Boy, J.-J. (2002, April). Isochronism Defect for Various Doubly Rotated Cut Quartz Resonators. *IEEE Transact Ultrason Ferroelec Frequency Control, 49*(4), 514-518.
- Halford, D. (1968, March). General Mechanical Model for f[^]α Spectral Density Random Noise with Special Reference to Flicker Noise 1/f. *Proc IEEE*, *53*(3), 251-258.

- Halford, D., Wainwraight, A. E., & Barnes, J. A. (1968). Flicker Noise of Phase in RF
 Amplifiers and Frequency Multipliers, Characterization, Cause, and Cure. *Proc Internat Frequency Control Symp* (pp. 340-341). IEEE.
- Hanbury Brown, R., & Twiss, R. Q. (1954, March). LXXIV. A new type of interferometer for use in radio astronomy. *Phil Mag Series 7, 45*(366), 663-682.
- Hati, A., Nelson, C. W., & Howe, D. A. (2016, March). Cross-spectrum measurement of thermal-noise limited oscillators. *Rev Sci Instrum*, *87*(3), 034708 1-8.
- Haus, H. A., & et al. (1960, January). IRE Standards on Methods of Measuring Noise in Linear Twoports, 1959. *Proc. IRE, 48*(1), 60-68.
- Hellwig, H., & et al. (1988). IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology. IEEE Standard 1139.
- Homayoun, A., & Razavi, B. (2013, March). Analysis of Phase Noise in Phase-Frequency Detectors. *IEEE Transact Circuits Systems I*, 60(3), 1-11.
- Howe, D. A., & Hati, A. (2005). Low-Noise X-Band Oscillator and Amplifier Technologies: Comparison and Status. *Proc Internat Frequency Control Symp and PTTI Joint Meeting*, (pp. 481-487).
- Ikezi, H., Woitowicz, R. E., Waltz, R. E., & Baker, D. R. (1988, December). Temporal Contraction of Solitons in a Nonuniform Transmission Line. J Appl Physics, 64(12), 6836-6838.
- Ivanov, E. N., Tobar, M. E., & Woode, R. A. (1998, November). Microwave Interferometry: Application to Precision Measurement and Noise Reduction Techniques. *IEEE Transact Ultrason Ferroelec Frequency Control*, 45(6), 1526-1536.
- JCGM. (2008). Evaluation of measurement data Guide to the expression of uncertainty in measurement.
- JCGM. (2012). International vocabulary of metrology Basic and general concepts and associated terms (VIM).
- Johansson, S. (2005). New frequency counting principle improves resolution. *Proc Internat Frequency Control Symp* (pp. 628-635). IEEE.
- Johnson, J. B. (1928, July). Thermal Agitation of Electricity in Conductors. *Phys Rev*, 32(1), 97-110.
- Kalisz, J. (2004). Review of methods for time interval measurements with picosecond resolution. *Metrologia*, *41*(1), 17-32.
- Keysight Technologies. (September 2017). Fundamentals of RF and Microwave Noise Figure Measurements.
- Kuleshov, V. N., & Boldyreva, T. I. (1997). 1/f AM and PM Noise in Bipolar Transistor Amplifiers: Sources, Ways of Influence, Techniques of Reduction. *Proc Internat Frequency Control Symp* (pp. 446-455). IEEE.

- Kurokawa, K. (1973, October). Injection Locking of Microwave Solid-State Circuits. *Proc IEEE*, *61*(10), 1386-1410.
- Kurtz, S. R. (2001). Mixers as Phase Detectors. Watkins Johnson Communications Inc.
- Labaar, F. (1982, March). New Discriminator Boosts Phase Noise Testing. *Microwaves,* 21(3), 65-69.
- Lance, A. L., Seal, W. D., & Labaar, F. (n.d.). Phase Noise and AM Noise Measurements in the Frequency Domain. In K. Button, & K. Button (Ed.), *Infrared and Millimeter Waves* (Vol. 11, pp. 239-289). Academic Press.
- Leeson, D. (1966). A Simple Model of Feedback Oscillator Noise Spectrum. *Proc. IEEE*, 54(2), 329-330.
- Levitin, D. J., Chordia, P., & Menon, V. (2012, February). Musical Rythm Spectra from Bach to Joplin Obey to 1/f Power Law. *Proc Nat Academy of Science*, *109*(10), 3716-3720.
- Maas, S. A. (1993). Microwave Mixers. Artech House.
- Maestrini, A., Thomas, B., Wamg, H., Jung, C., Treuttel, J., Jin, Y., . . . Beaudin, G. (2010, August-October). Schottky diode based terahertz frequency multipliers and mixers. *Comptes Rendus de l'Academie des Sciences, Physique, 11*, 480-485.
- Meher, P. K., Valls, J., Juang, T. B., Sridharan, K., & Maharatna, K. (2009, September). 50 Years of CORDIC: Algorithms, Architectures, and Applications. *IEEE Transact Circuits Systems I, 56*(9), 1893-1907.
- Microsemi Corp. (n.d.). 3120A Phase Noise and Allan Deviation Test Set Specifications. Retrieved February 2018, from https://www.microsemi.com/productdirectory/clocks-frequency-references/3834-test-measurement
- Microsemi Corp. (n.d.). 5120A Phase Noise and Allen Deviation Test Set Specifications. Retrieved February 2018, from https://www.microsemi.com/productdirectory/clocks-frequency-references/3834-test-measurement
- Microsemi Corp. (n.d.). 5125A Phase Noise and Allan Deviation Test Set Specifications. Retrieved February 2018, from https://www.microsemi.com/productdirectory/clocks-frequency-references/3834-test-measurement
- Miller, R. R. (1939, July). Fractional-Frequency Generators Utilizing Regenerative Modulation. *Proc IRE, 27*, 446-457.
- Milotti, E. (2002, April). 1/f Noise: A Pedagogical Review. arXiv:physics/0204033 [physics.class-ph].
- Nayfeh, A. H., & Mook, D. T. (2004). Nonlinear Oscillations. Wiley-VCH.
- Nelson, C. W., Hati, A., & Howe, D. A. (2014, February). A collapse of the cross-spectral function in phase noise metrology. *Rev Sci Instrum*, *85*(2), 024705 1-7.

- Nelson, C. W., Walls, F. L., Siccardi, M., & De Marchi, A. (1994). A new 5 and 10 MHz High Isolation Distribution Amplifier. *Proc Internat Frequency Control Symp* (pp. 567-571). IEEE.
- North, D. O. (1942, January). The Absolute Sensitivity of Radio Receivers. *RCA Review*, 6(3), 332-343.
- Nyquist, H. (1928, July). Thermal Agitation of Electric Charges in Conductors. *Phys Rev*, *32*(1), 110-113.
- Pankratz, E., & Sanchez-Sinencio, E. (2013, September). Survey of Integrated-Circuit-Oscillator Phase-Noise Analysis. *Int. J. Circ. Theor. Appl., 42*(9), 871-938.
- Penfield, P., & Rafuse, R. P. (1962). Varactor Applications. MIT Press.
- Phillips, D. E. (1987). Random Noise in Digital Gates and Dividers. *Proc Internat Frequency Control Symp*, (pp. 507-511).
- Poddar, A., & Rohde, U. L. (2013, September-October). How Low Can They Go? *Microwave Mag.*, 14(6), 50-72.
- Posener, D. W. (1959, June). The Shape of Spectral Lines: Table of the Voigt Profile. *Australian J Phys*, *12*, 184-196.
- Pozar, D. M. (2012). *Microwave Engineering* (4th ed.). John Wiley & Sons.
- Rategh, H. R., & Lee, T. H. (1999, June). Superharmonic Injection-Locked Frequency Dividers. *IEEE J Solid State Circuits*, *34*(6), 813-821.
- Rategh, H. R., Samavati, H., & Lee, T. H. (2000, May). A CMOS Frequency Synthesizer with an Injection-Locked Frequency Divider for a 5-GHz Wireless LAN Receiver. *IEEE J Solid State Circuits*, *35*(5), 780-787.
- Rohde, U. L. (1975, October). Crystal oscillator provides low noise. *Electroncs Design, 22*, 11, 14.
- Rohde, U. L., Poddar, A. K., & Boeck, G. (2005). *Modern Microwave Oscillators for Wireless Applications: Theory and Optimization.* John Wiley & Sons.
- Rothe, H., & Dahlke, W. (1956, June). Theory of Noisy Fourpoles. *Proc. IRE, 44*(6), 811-818.
- Rotholz, E. (1984, September 13). Phase Noise of Mixers. *Electronics Lett, 20*(19), 786-787.
- Rubiola, E. (2005, May). On the measurement of frequency and of its sample variance with high-resolution counters. *Rev Sci Instrum*, *76*(5), 054703 1-6.
- Rubiola, E. (2005, December). The measurement of AM noise of oscillators. arXiv:physics/0512082v1 [physics.ins-det].
- Rubiola, E. (2006, August). Tutorial on the Double-Balanced Mixer. arXiv:physics/0608211v1 [physics.ins-det].

- Rubiola, E. (2010). *Phase Noise and Frequency Stability in Oscillators*. Cambridge University Press.
- Rubiola, E. (2016, July). The Ω Counter, a Frequency Counter Based on the Linear Regression. *IEEE Transact Ultrason Ferroelec Freq Control*, 63(7), 961-969.
- Rubiola, E. (2023). The Enrico's chart of phase noise and two-sample variances. doi:DOI: 10.5281/zenodo.4399218
- Rubiola, E. (Public material, Creative Commons 4.0 CC-BY, 2019). *Frequency and Amplitude Stability in Oscillators, slides of a lecture series for PhD students and young scientists.* University of Bourgogne Franche Comte, Besancon, France.
- Rubiola, E., & Boudot, R. (2007, May). The Effect of AM Noise on Correlation Phase-Noise Measurements. *IEEE Transact Ultrason Ferroelec Frequency Control, 54*(5), 926-932.
- Rubiola, E., & Brendel, R. (2010, April). A Generalization of the Leeson Effect. arXiv:1004.5539 [physics.ins-det].
- Rubiola, E., & Giordano, V. (2002, June). Advanced interferometric phase and amplitude noise measurements. *Rev Sci Instrum*, *73*(6), 2445-2457.
- Rubiola, E., & Lardet-Vieudrin, F. (2004, May). Low flicker-noise amplifier for 50 Ω sources. *Rev Sci Instrum*, *75*(5), 1323-1326.
- Rubiola, E., & Vernotte, F. (2010, February). The cross-spectrum experimental method. arXiv:1003.0113 [physics.ins-det].
- Rubiola, E., & Vernotte, F. (2023, July). The Companion of Enrico's Chart for Phase Noise and Two-Sample Variances. *IEEE Transact. Microwave Theory and Techniques*, 71(7), 2996-3025. doi:10.1109/TMTT.2023.3238267
- Rubiola, E., Giordano, V., & Groslambert, J. (1999, January). Very high frequency and microwave interferometric phase and. *Rev Sci Instrum*, *70*(1), 220-225.
- Rubiola, E., Groslambert, J., Brunet, M., & Giordano, V. (2000, March). Flicker Noise Measurement of HF Quartz Resonators. *IEEE Transact Ultrason Ferroelec Frequency Control*, 47(2), 361-368.
- Rubiola, E., Olivier, M., & Groslambert, J. (1992, June). Phase Noise in Regenerative Frequency Dividers. *IEEE Transact Instrum Meas*, *41*(3), 353-360.
- Rubiola, E., Salik, E., Huang, S., Yu, N., & Maleki, L. (2005, May). Photonic delay technique for phase noise measurement of microwave oscillators. *J Optical Soc America B*, *22*(5), 987-997.
- Rutman, J. (1978, September). Characterization of Phase and Frequency Instability in Precision Frequency Sources: Fifteen Years of Progress. *Proc IEEE*, 66(9), 1048-1078.

- Salik, E., Yu, N., Maleki, L., & Rubiola, E. (2004). Dual photonic-delay-line cross correlation method for the measurement of microwave oscillator phase noise. *Proc Internat Frequency Control Symp*, (pp. 303-306).
- Sann, K. H. (1968, September). The Measurement of Near-Carrier Noise in Microwave Amplifiers. *IEEE Transact Microwave Theory Techniques*, *16*(9), 761-766.
- Schottky, W. (1918). Über spontane Stromschwankungen in verschiedenen Elektrizitätsleitern. *Annalen der Physik, 362*(23).
- Snyder, J. J. (1981). An Ultra-High Resolution Frequency Meter. *Proc Internat Frequency Control Symp*, (pp. 470-475).
- Stewart, J. L. (1954, October). The Power Spectrum of a Carrier Frequency Modulated by Gaussian Noise. *Proc IRE, 42*, 1539-1543.
- Thompson, I., & Brennan, P. V. (2003, February). Phase Noise Contribution of the PhaseFrequency Detector in a Digital PLL Frequency Synthesiser. *Proc IEEE*, *150*(1), 15.
- Torosyan, A. (2003). *DDS Complete Analysis and Design Guidelines, PhD thesis.* Univ. California Los Angeles.
- Torosyan, A., & Wilson, A. (2005). Exact Analysis of DDS Spurs and SNR due to Phase Truncation and Arbitrary Phase-to-Amplitude Errors. *Proc Internat Frequency Control Symp*, (pp. 50-58).
- Vernotte, F., Lenczner, M., Bourgeois, P.-Y., & Rubiola, E. (2016, April). The Parabolic Variance (PVAR), a Wavelet Variance Based on the Least-Square Fit. *IEEE Transact Ultrason Ferroelec Frequency Control*, 63(4), 611-623.
- Vig, J., & et al. (1999). IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities. IEEE Standard 1139.
- Volder, J. E. (1959, September). The CORDIC Trigonometric Computing Technique. *IRE Transact Electron Comput, 8*(3), 330-334.
- Volyanskiy, K., Cussey, J., Tavernier, H., Salzentein, P., Sauvage, G., Larger, L., & Rubiola,
 E. (2008, December). Applications of the optical fiber to the generation and to
 the measurement of low-phase-noise microwave signals. *Journal Optical Society* of America B, 25(12), 2140-2150.
- Walls, F. L., & De Marchi, A. (1975). RF Spectrum of a Signal after Frequency Multiplication, Measurement and Comparison with a Simple Calculation. *IEEE Transact Instrum Meas*, 24(3), 210-217.
- Walls, F. L., Ferre-Pikal, E. S., & Jefferts, S. R. (1997, March). Origin of 1/f PM and AM Noise in Bipolar Transistor Amplifiers. *IEEE Transact Ultrason Ferroelec Frequency Control*, 44(2), 326-334.

Walls, F. L., Stein, S. R., Gray, J. E., & Glaze, D. J. (1976). Design Consideration in State-ofthe-Art Signal Processing and Phase Noise Measurement Systems. *Proc Internat Frequency Control Symp*, (pp. 269-274).

8 Suggested readings

8.1 Power spectra and Fourier transform

- O. E. Brigham, The Fast Fourier Transform and Its Applications, Prentice Hall 1988, ISBN 0-13-307505-2.
- R. B. Blackman, J. W. Tukey JW, The measurement of power spectra, Dover 1959, ISBN 0-486-60507-8.
- A. V. Oppenheim, R. W. Schafer, Discrete-time signal processing 2nd ed., Prentice Hall 1989, ISBN 0-13-754920-2.
- A. Papoulis, The Fourier Integral and Its Applications, McGraw Hill 1962.
- D. B. Percival, A. T. Walden, Spectral Analysis for Physical Applications, Cambridge 1993, ISBN 0-521-43541-2.

8.2 Electromagnetic Compatibility

- J. Goedbloed, Electromagnetic Compatibility, Prentice Hall 1990, ISBN 0-13-249293-8.
- H. W. Ott, Electromagnetic Compatibility Engineering, John Wiley & Sons 2009, ISBN 978-0-470-18930-6.
- C. R. Paul, Introduction to Electromagnetic Compatibility, John Wiley & Sons 2006, ISBN 978-0-471-75500-5.
- R. Perez, Handbook of Electromagnetic Compatibility, Academic Press 1995, ISBN 0-12-550710-0.

8.3 General Aspects of Noise

- M. G. Arthur, The Measurement of Noise Performance Factors: A Metrology Guide, NBS Monograph 140, US Department of Commerce and NBS (now NIST), June 1974.
- D. A. Bell, Electrical Noise, Van Nostrand 1960.
- M. J. Buckingham, Noise in Electronic Devices and Systems, Ellis Horwood, UK, 1983. ISBN 0-85312-218-0.
- W. B. Davenport, Jr., W. L. Root, An Introduction to the Theory of Random Signals and Noise, IEEE Press 1987, ISBN 0-87942-235-1. Reprint of the McGraw Hill edition, 1958.
- W. F. Egan, Practical RF System Design, John Wiley & Sons 2003. ISBN 0-471-20023-9. Chapter 3 focuses on noise.
- H. T. Friis, The Noise Figures of Radio Receivers, Proc. IRE Vol. 32 no. 7 pp. 419-422, July 1944.

- H. A. Haus, R. B. Adler, Circuit Theory of Linear Noisy Networks, John Wiley & Sons 1959.
- A. R. Kerr, J. Randa, Thermal Noise and Noise Mesurements—A 2010 Update, IEEE Microwave Mag. Vol. 11 no. 6, October 2010.
- Keysight Technologies, Fundamentals of RF and Microwave Noise Figure Measurements, September 2017.
- Sh. Kogan, Electronic Noise and Fluctuations in Solids, Cambridge University Press 1996. ISBN 0-521-46034-4.
- S. A. Maas, Noise in Linear and Nonlinear Circuits, Artech House 2005. ISBN 1-58053-849-5.
- H. Rothe, W. Dahlke, Theory of Noise Fourpoles, Proc. IRE Vol. 44 no. 6 pp. 811-818, June 1956.
- G. Vasilescu, Electronic Noise and Interfering Signals, Springer 2005, ISBN 3-540-40741-3.
- N. Wax (Ed.), Selected Papers on Noise and Stochastic Processes, Dover 1954.
- B. Widrow, I. Kollar, Quantization Noise, Cambridge 2008. ISBN 978-0-511-40990-5.

8.4 Phase Noise, Frequency Stability, and Measurements

- A. R. Chi (ed.), Short Term Frequency Stability, NASA SP-80, 1965.
- N. Da Dalt, A. Sheikoleslami, Understanding Jitter and Phase Noise, Cambridge 2018. ISBN 978-1-107-18857-0.
- C. A. Greenhall, A method for using a time interval counter to measure frequency stability, IEEE Transact. Ultrason., Ferroelec. Freq. Control Vol. 36 no 5 pp. 478-480, September 1989.
- S. Henzler, Time-to-Digital Converters, Springer 2010. ISBN 978-90-481-8628-0
- J. Kalisz, Review of methods for time interval measurements with picosecond resolution, Metrologia Vol 41 no 1 p. 17-32, 2004
- H. G. Kimball, Handbook of Selection and Use of Precise Frequency and Time Systems, International Telecommunications Union 1997.
- V. F. Kroupa, Frequency Stability: Introduction and Applications, IEEE Press 2012. ISBN 978-1-118-15912-5.
- M. P. Lee, Jitter, Noise, and Signal Integrity at High Speed, Prentice Hall 2008. ISBN 978-0-13-242961-0.
- W. J. Riley, Handbook of Frequency Stability Analysis, NIST SP 1065, NIST 2008.
- W. P. Robins, Phase Noise in Signal Sources, Peter Peregrinus 1984. ISBN 0-86341-026-X.

- E. Rubiola, On the measurement of frequency and of its sample variance with high-resolution counters, Rev. Sci. Instrum. Vol.76 no.5 article no. 054703, May 2005.
- E. Rubiola, Phase noise and frequency stability in oscillators, Cambridge University Press, 2008, 2010. ISBN 978-0-521-15328-7. Also available in Simplified Chinese, Cambridge University Press and Science Press, 2014. ISBN 978-7-03-041231-7.
- E. Rubiola, M. Lenczner, P.Y. Bourgeois, F. Vernotte, The Ω Counter, a Frequency Counter Based on the Linear Regression, IEEE Transact. Ultrason., Ferroelec. Freq. Control Vol 63 no 7 p.961-969, July 2016.
- S. R. Stein, The Allan variance—Challenges and opportunities, IEEE Transact. Ultrason., Ferroelec. Freq. Control Vol. 57 no 3 pp. 540-547, March 2010.
- IEEE Transact. Ultrason., Ferroelec. Freq. Control, Special Issue of the 50th anniversary of the Allan variance, Vol. 63 no. 4 p.611-623, April 2016.

8.5 Amplifiers

- E. S. Ferre-Pikal ES, F. L. Walls, C. W. Nelson, Guidelines for designing BJT amplifiers with low 1/f AM and PM noise, IEEE Transact. Ultrason., Ferroelec. Freq. Control Vol. 44 no 2, March 1997.
- M. C. Delgado Aramburo, E. S. Ferre-Pikal, F. L. Walls FL, H. D. Ascarrunz, Comparison of 1/f PM noise in commercial amplifiers, Proc Internat Frequency Control Symp pp. 470-477, 1997.
- P. Penfield, jr, Noise in negative-resistance amplifiers, IRE Transact Circuit Theory Vol. 7 no 2 pp. 166-170, June 1960.
- K. Theodoropoulos, J. Everard, Residual phase noise modeling of amplifiers using Silicon bipolar transistors, IEEE Transact Ultrason Ferroelec Frequency Control Vol. 57 no 3 pp. 562-573, March 2010.
- T. D. Tomlin, K. Finn, A. Cantoni, A model for phase noise generation in amplifiers, IEEE Transact Ultrason Ferroelec Frequency Control Vol. 48 no. 6 p.1547-1554, November 2001.
- D. A. Howe, A. Hati, Low-noise X-band oscillator and amplifier technologies, Comparison and status, Proc Internat Frequency Control Symp pp. 481-487, 2005.
- S. Moreschi, J. Skeen, Skyworks Inc., Ultra Low Noise Amplifiers, white paper, June 2014.

8.6 Frequency Dividers

- R. Adler R, A study of locking phenomena in oscillators, Proc IRE Vol. 34 no. 6 pp 351-357, June 1946
- C. E. Calosso, E. Rubiola, The Sampling Theorem in Pi and Lambda Digital Frequency Dividers, Proc Internat Frequency Control Symp pp. 60-962, 2013

- M. Driscoll, Phase noise performance of analog frequency dividers, IEEE Transact Ultrason Ferroelec Frequency Control Vol. 37 no 4, April 1990.
- W. F. Egan, Modeling Phase Noise in Frequency Dividers, IEEE Transact Ultrason Ferroelec Frequency Control Vol 37 no 4 pp. 307-315, July 1990.
- V. F. Kroupa, Jitter and Phase Noise in Frequency Dividers, IEEE Transact Instrum Meas Vol. 50 no. 5, October 2001.
- K. Kurokawa, Injection Locking of Microwave Solid-State Circuits Proc IEEE Vol. 61 no. 10 pp. 1386-1410, October 1973.
- R. L. Miller, Fractional-Frequency Generators Utilizing Regenerative Modulation, Proc IRE Vol. 27 no. 7 pp.446-457, July 1939.
- P. Penfield, jr, R. P. Rafuse, Varactor Applications, MIT Press, 1962.
- D. E. Phillips, Random Noise in Digital Gates and Dividers, Proc Internat Frequency Control Symp pp. 507-511, 1987.
- H. R. Rategh, T. H. Lee, Superharmonic Injection-Locked Frequency Dividers IEEE J Solid State Circuits Vol. 34 no. 6 pp. 813-821, June 1999.
- J. R. Rategh, H. Samavati, T. H. Lee, A CMOS Frequency Synthesizer with an Injection-Locked Frequency Divider for a 5-GHz Wireless LAN Receiver, IEEE J Solid State Circuits Vol. 35 no. 5 pp. 780-787, May 2000
- B. Razavi, A Study of Injection Locking and Pulling in Oscillators, IEEE J Solid State Circuits Vol. 39 no. 9 pp. 1415-1424, September 2004.
- E. Rubiola, M. Olivier, J. Groslambert, Phase Noise in Regenerative Frequency Dividers IEEE Transact Instrum Meas Vol. 41 no. 3 pp. 353-360, June 1992.

8.7 Frequency Multipliers

- E. Camargo, Design of FET frequency multipliers and harmonic oscillators, Artech House 1998. ISBN 0-89006-481-4.
- M. T. Faber, J. Chramiek, M. E. Adamski, Microwave and millimeter-wave diode frequency multipliers, Artech House 1995. ISBN 0-89006-611-6

8.8 DDS

Analog Devices Inc., A Technical Tutorial on Digital Signal Synthesis, 1999.

- C. E. Calosso, Y. Gruson and E. Rubiola, "Phase Noise in DDS," in Proc. Internat. Frequency Control Symp. pp 2012.
- B. G. Goldberg, Digital frequency synthesis demystified, LLH 1999, ISBN 1-878707-47-7.
- V. Kroupa (ed.), Direct digital frequency synthesis, IEEE 1999, ISBN 0-7803-3438-8
- E. Murphy, C. Slattery, All About Direct Digital Synthesis, Analog Dialogue 38-08 pp. 1-6, August 2004.
- H. T. Nicholas, III, H Samueli, An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase-Accumulator Truncation, *Proc Internat Frequency Control Symp* pp. 495-502, 1987
- H. T. Nicholas, III, H Samueli, The Optimization of Direct Digital Frequency Synthesizers Performance in the Presence of Finite Word Length Effects, *Proc Internat Frequency Control Symp* pp. 357-363, 1988.
- A Torosyan, A. N. Wilson, Jr, Exact Analysis of DDS Spurs and SNR due to Phase Truncation and Arbitrary Phase-to-Amplitude Errors, Proc Internat Frequency Control Symp pp. 50-58, 2005.
- B. Widrow, I. Kollar, Quantization Noise, Cambridge 2008. ISBN 978-0-511-40990-5.

8.9 Phase-Frequency Detectors

- X. Gao, A. M. Klumperink, P. F. J. Geraedts, B. Nauta, Jitter Analysis and a Benchmarking Figure-of-Merit for Phase-Locked Loops, *IEEE Transact Circuits Systems II* Vol. 56 no. 2, February 2009.
- F. M. Gardner, Charge-Pump Phase-Lock Loops, *IEEE Transact Communications* Vol. 28 no. 11, November 1980.
- A. Homayoun, B. Razavi, Analysis of Phase Noise in Phase/Frequency Detectors, IEEE Transact Circuits Systems I Vol. 60 no. 3 pp. 529-539, March 2013.

8.10 Oscillators

- A. M. Apte, A. K. Poddar, U. L. Rohde, E.Rubiola, Colpitts Oscillator, A New Criterion of Energy Saving for High Performance Signal Sources, *Proc Internat Frequency Control Symp* pp. 70-76, 2016.
- A. M. Apte, U. L. Rohde, A.K. Poddar, M. Rudolph, Optimizing Phase Noise Performance, *Microwave Mag*, June 2017 pp. 108-123.
- A. Demir, Fully nonlinear oscillator noise analysis, An oscillator with no asymptotic phase, *Internat J Circuit Theory and Applications* Vol. 35 pp. 175-203, October 2005.
- D. S. Greywall, B. Yurke, P. A. Busch, A.N. Pargellis, R.L.Willett, Evading Amplifier Noise in Nonlinear Oscillators, *Phys Rev Lett* Vol. 72 no. 19 pp. 2992-2995, May 1994.
- A. Hajimiri, , T. H. Lee, The design of low noise oscillators, Kluwer 1999. ISBN 0-7923-8455-5.
- A. Hajimiri, S. Limotyrakis, T. H. Lee, Jitter and Phase Noise in Ring Oscillators, *IEEE J Solid State Circuits* Vol. 34 no. 6 pp. 790-804, June 1999.
- T. H. Lee, A. Hajimiri, Oscillator Phase Noise, A Tutorial, IEEE J Solid State Circuits Vol. 35 no. 3 pp. 326-336, March 2000.

- D. Ham, A. Hajimiri, Virtual Damping and Einstein Relation in Oscillators, IEEE J Solid State Circuits Vol. 38 no. 3 pp. 407-418, March 2003.
- A. P. S. Khanna, Microwave Oscillators: The State of the Technology, *Microwave J*, April 2006
- D. B. Leeson, A Simple Model of Feedback Oscillator Noise Spectrum, *Proc IEEE* Vol. 54 no. 2 pp. 329-330, February 1966
- J. C. Nallatamby, M. Prigent, M. Camiad, J. Obregon, Phase Noise in Oscillators, Leeson Formula Revisited, *IEEE Transact Microwave Theory Techniques* Vol. 51 no. 4 pp. 1386-1394, April 2003
- E. Pankratz, E. Sanchez-Sinencio, Survey of integrated-circuit-oscillator phase-noise analysis, *Internat J Circuit Theory and Applications* Vol. 42 no. 9 pp. 871-938, September 2014
- A. K. Poddar, U. L. Rohde, A. M. Apte, How Low Can They Go?, *Microwave Mag*, September/October 2013 pp. 50-72.
- A. K. Poddar, U. L. Rohde, Latest Technology, Technological Challenges, and Market Trends for Frequency Generating and Timing Devices, *Microwave Mag*, September/October 2012 pp. 120-134.
- U. L. Rohde, A. M. Apte, Everything You Always Wanted to Know About Colpitts Oscillators, *Microwave Mag*, August 2016 pp. 59-76.
- U. L. Rohde, A. K. Poddar, G. Böck, The Design of Microwave Oscillators for Wireless Applications, John Wiley & Sons 2005. ISBN 0-471-72342-8.
- E. Rubiola, Phase noise and frequency stability in oscillators, Cambridge University Press, 2008, 2010. ISBN 978-0-521-15328-7. Also available in Simplified Chinese, Cambridge University Press and Science Press, 2014. ISBN 978-7-03-041231-7.
- F. L. Traversa, F. Bonani, Oscillator Noise, A Nonlinear Perturbative Theory Including Orbital Fluctuations and Phase-Orbital Correlation, *IEEE Transact Circuits Systems I* Vol. 58 no. 10 pp. 2485-2497, October 2011
- T. C. Weigandt, B. Kim, P. R. Gray, Analysis of Timing Jitter in CMOS Ring Oscillators, Proc. IEEE Internat Symp Circuits and Systems (ISCAS) pp. 191-194, 1994.

8.11 Resonators

- V. B. Braginsky, V. P. Mitrofanov, V. I. Panov, Systems with Small Dissipations, University of Chicago Press, 1985. ISBN 0-226-07073-5.
- E. A. Gerber, A. Ballato (editors), Precision Frequency Control Volume 1: Acoustic Resonators and Filters, Academic Press, 1985. ISBN 0-12-280601-8.
- N. Gufflet, R. Bourquin, J.-J. Boy, Isochronism Defect for Various Doubly Rotated Cut Quartz Resonators, *IEEE Transact Ultrason Ferroelec Frequency Control* Vol. 49 no. 4 pp. 514-518, April 2002.

- J. Helszajn, YIG Resonators and Filters, John Wiley & Sons 1985. ISBN 0-471-90516-X.
- D. Kayfez, P. Guillon (editors), Dielectric Resonators, 2nd edition, Noble Publishing, 1998. ISBN 1-884932-05-3.
- A. H. Nayfeh, D. T. Mook, Nonlinear Oscillators, John Wiley & Sons 2004. ISBN 0-471-03555-6.
- U. L. Rohde, A. Poddar, D. Sundararajan, Printed Resonators: Moebius Strip Theory and Applications, Microwave J. Vol. 56 no. 11 pp. 24-54, November 2013.
- P. Vizmuller, Filters with Helical Resonators, Artech House 1987. ISBN 0-89006-244-7.

8.12 Double Balanced Mixer

- S. R. Kurtz, Mixers as Phase Detectors, Watkins Johnson Communications Inc., 2001.
- S. A. Maas, Microwave Mixers, Artech House 1993, ISBN 0-89006-605-1.
- D. Owen, Good Practice Guide to Phase Noise Measurement, National Physical Laboratory, UK, Crown Copyright, May 2004.
- E. Rubiola, F. Lardet-Vieudrin, Low flicker-noise amplifier for 50 Ohm sources, *Rev. Sci. Instrum.* Vol.75 no.5 pp. 1323-1326, May 2004.
- E. Rotholz, Phase Noise of Mixers, *Electronics Lett*, Vol. 20, no. 19, pp. 786-787, 13 September 1984.
- E. Rubiola, Tutorial on the double-balanced mixer, arXiv:physics/0608211v1 [physics.ins-det], 2006
- J. Sevick, Transmission line transformers, 4th edition, Noble Publishing, 2001. ISBN 1-884932-18-5.