The Companion of Enrico's Chart for Phase Noise and Two-Sample Variances

Enrico Rubiola, *Member, IEEE*, and François Vernotte

***Abstract*— Phase noise and frequency (in)stability both describe the fluctuation of stable periodic signals from somewhat different standpoints. Frequency is unique compared to other domains of metrology, in which its fluctuations of interest span at least 14 orders of magnitude, from 10^{-4} in a mechanical watch to 10^{-18} in atomic clocks. The frequency span of interest is some 12–15 orders of magnitude from μHz to GHz Fourier frequency for phase noise, while the time span over which the fluctuations occur ranges from sub-μs to years integration time for variances. Because this domain is ubiquitous in science and technology, a common language and tools suitable to the variety mentioned are a challenge. This article is at once: 1) a tutorial; 2) a review covering the most important facts about phase noise, frequency noise, and two-sample (Allan and Allan-like) variances; and 3) a user guide to "Enrico's Chart of Phase Noise and Two-Sample Variances." In turn, the Chart is a reference card collecting the most useful concepts, formulas, and plots in a single A4/A-size sheet, intended to be a staple on the desk of whoever works with these topics. The Chart is available under Creative Commons 4.0 CC-BY-NC-ND license from Zenodo, DOI: 10.5281/zenodo.4399218. A wealth of auxiliary material is available for free on Enrico's home page: http://rubiola.org.**

***Index Terms*— Frequency control and conversion, instrumentation and measurement, low phase noise oscillators, oscillators, phase noise.**

I. Introduction

PHASE noise and frequency instability are equivalent concepts, to the extent that frequency is the derivative of the instantaneous phase. However, the choice of terminology, mathematical tools, and experimental methods really depends on what aspect of the noise our system is sensitive to or which aspect we focus on. The PSD of the phase fluctuations of a periodic signal, as a function of the Fourier frequency, is the preferred tool to describe phase noise. Similarly, the frequency instability is described by the PSD of frequency fluctuations, or more often by the two-sample (Allan or Allan-like) variance of the fractional frequency, as a function of the measurement time.

A. Taste of Phase Noise and Frequency Stability

There are many reasons to be interested in phase noise and frequency instability, most of which are related to spectrum broadening, timing uncertainty, and reduced coherence time. The impact of phase noise and frequency instability is surprisingly ubiquitous, from everyday technology to fundamental science. Let us go through some examples.

In radio systems, sensitivity or dynamic range may be limited by phase noise sidebands appearing: 1) on the LO of a superheterodyne receiver, causing neighboring strong signals to mask the desired IF signal (reciprocal mixing) [1], [2, Sec. 2.5]; 2) on strong adjacent transmitted signals themselves in channelized systems (the near-far problem, in which sideband noise from local users can overpower distant cellular base stations); or 3) on strong reflected radar signals, such as ground clutter, limiting the ranging accuracy and the detection of small objects in the clutter [3, Chapter 6], [4], [5]. In these systems, amplitude fluctuation (AM) is limited in oscillators and nonlinear power amplifiers, so phase noise (PM) is generally the dominant source of noise sidebands. Similar problems are found in LiDAR, RADAR's optical counterpart. Phase- and frequency-modulated systems are, of course, directly subject to noise in the frequency domain; examples include analog frequency-division FM multiplex systems and digital QAM modulation systems.

Manuscript received 12 November 2022; revised 17 December 2022; accepted 27 December 2022. This work was supported in part by the Agence Nationale de Recherche (ANR) Programme d'Investissement d'Avenir under Grant ANR-11-EQPX-0033-OSC-IMP (Oscillator IMP project), Grant ANR-10-LABX-48-01 (FIRST-TF network), and Grant ANR-17-EURE-00002 (EIPHI); and in part by grants from the Région Bourgogne Franche Comté. *(Corresponding author: Enrico Rubiola.)*

Enrico Rubiola is with the Université de Franche Comté (uFC), 25000 Besançon, France, also with the Centre National de la Recherche Scientifique (CNRS), 75016 Paris, France, also with Supmicrotech-ENSMM, 25030 Besançon, France, also with the Institut FEMTO-ST, Department of Time and Frequency, 25030 Besançon, France, and also with the Division of Quantum Metrology and Nanotechnology, Istituto Nazionale di Ricerca Metrologica (INRiM), 10135 Turin, Italy (e-mail: enrico.rubiola@femto-st.fr; e.rubiola@inrim.it).

François Vernotte is with the Université de Franche Comté (UFC), 25000 Besançon, France, also with the Centre National de la Recherche Scientifique (CNRS), 75016 Paris, France, also with Supmicrotech-ENSMM, 25030 Besançon, France, and also with the Institut FEMTO-ST, Department of Time and Frequency, 25030 Besançon, France (e-mail: francois.vernotte@femto-st.fr).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TMTT.2023.3238267.

Digital Object Identifier 10.1109/TMTT.2023.3238267

Download Enrico's Chart from Zenodo

Download this article

Fig. 1. QR codes to download Enrico's Chart and this article.

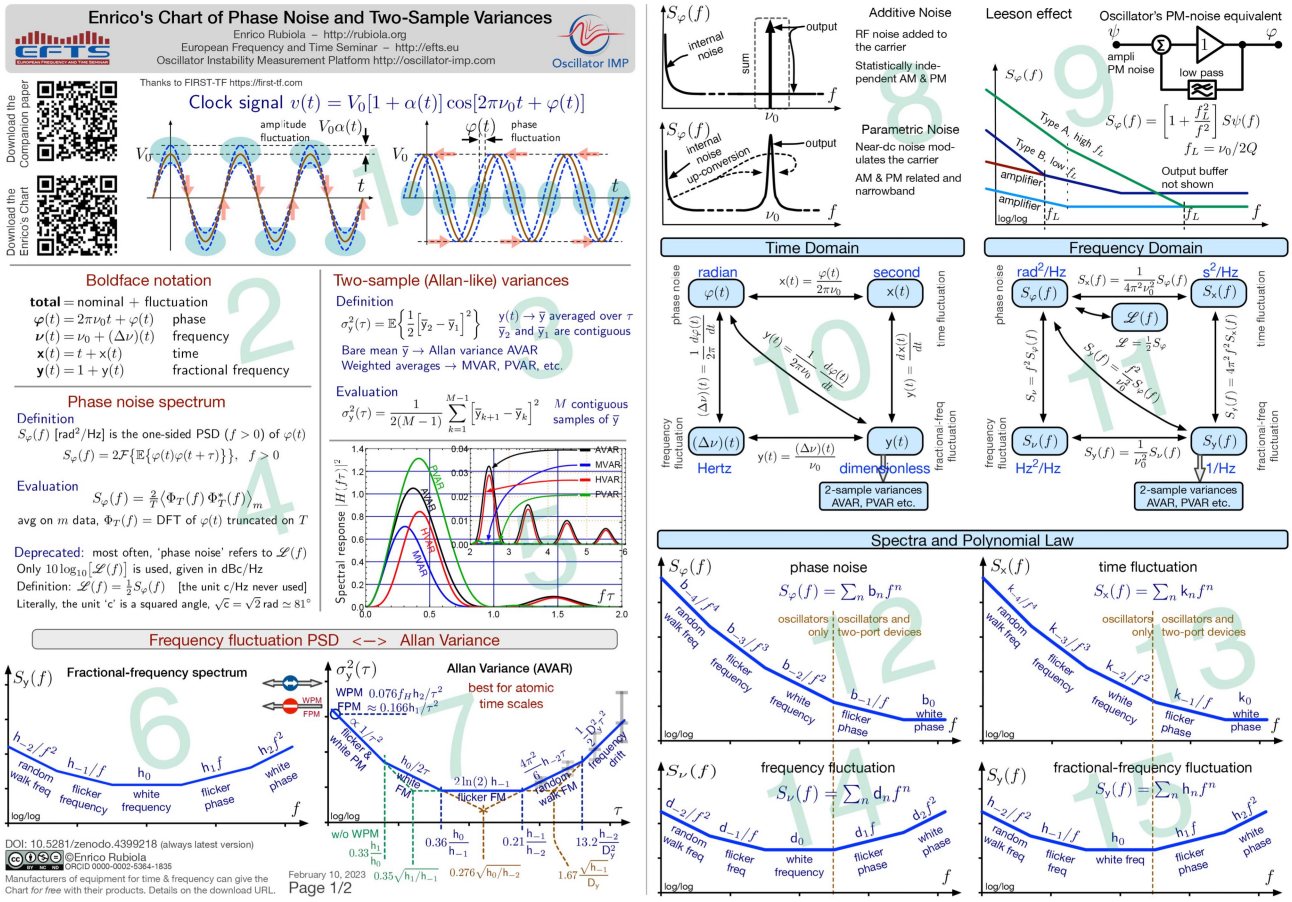

Fig. 2. Enrico's Chart of phase noise and two-sample variances, front side.

In digital systems and communications, and in microelectronics as well, the term "jitter" is often used for phase noise integrated over the appropriate bandwidth and converted into time fluctuations. In turn, time fluctuations increase the BER (see [6] and [7]). In such systems, the "quality" of the digital clock is often shown as the eye pattern (eye diagram), which is corrupted by the jitter (see [6, Sec 1.2] and [7, Sec. 4.1] and "Eye pattern" entry on Wikipedia). Accurate clock synchronization is a major leap forward in 5G/6G wireless systems [8], [9].

The community of power grids is looking at frequency stability for future power grids [10], [11], [12] because sustainability requires a spread of generator technologies with different inertia (latency between demand for increased/reduced power and actual power delivery). We believe that, at some point, the "generators' inertia" will be identified with the integration time τ of the AVARs. Insufficient synchronization is blamed as a cofactor of the 2003 Northeast Blackout in the USA [13].

Unique requirements arose with the development of the deep space network. The DSN required a blend of long-term stability for the initial acquisition of weak signals and short-term stability for maintaining the lock. This new combined need led to the 1964 IEEE-NASA Symposium on Short-term Frequency Stability [14], to a Special Issue of Frequency Stability published in 1966 in the PROCEEDINGS OF THE IEEE [15], and to the definitive 1971 Barnes et al.'s article [16]. The latter is a precursor of the IEEE Standard 1139 [17]. That effort to update this standard continues today.

The AVAR tools also find applications in geodesy and astrometry [18]. The short-term fluctuations of oscillators used in very-large baseline interferometry (VLBI) cannot be compensated for numerically and limit the detection sensitivity. Time fluctuations are critical in gravitational wave experiments (LIGO) because of the tiny space–time warp to be detected and in the RF cavities of particle accelerators [19, Chapters by F. Tecker and H. Damereau], where they cause intensity loss or energy loss in the beam.

In quantum computing, the correlation time limits the lifetime of a qubit [20].

Optics supports a major trend in high-purity signals [21] and in atomic oscillators and clocks as well [22] because the time interval associated with a phase angle is $\sim 10^{-4}$ smaller than at microwave frequencies. The FS laser, which enabled the first direct synthesis from RF to optics, played a major role in this trend. The Nobel Prize in Physics was awarded to John L. Hall and Theodor W. Hänsch in 2005 for the FS laser (together with Roy J. Glauber) (see [23], [24], and [25]).

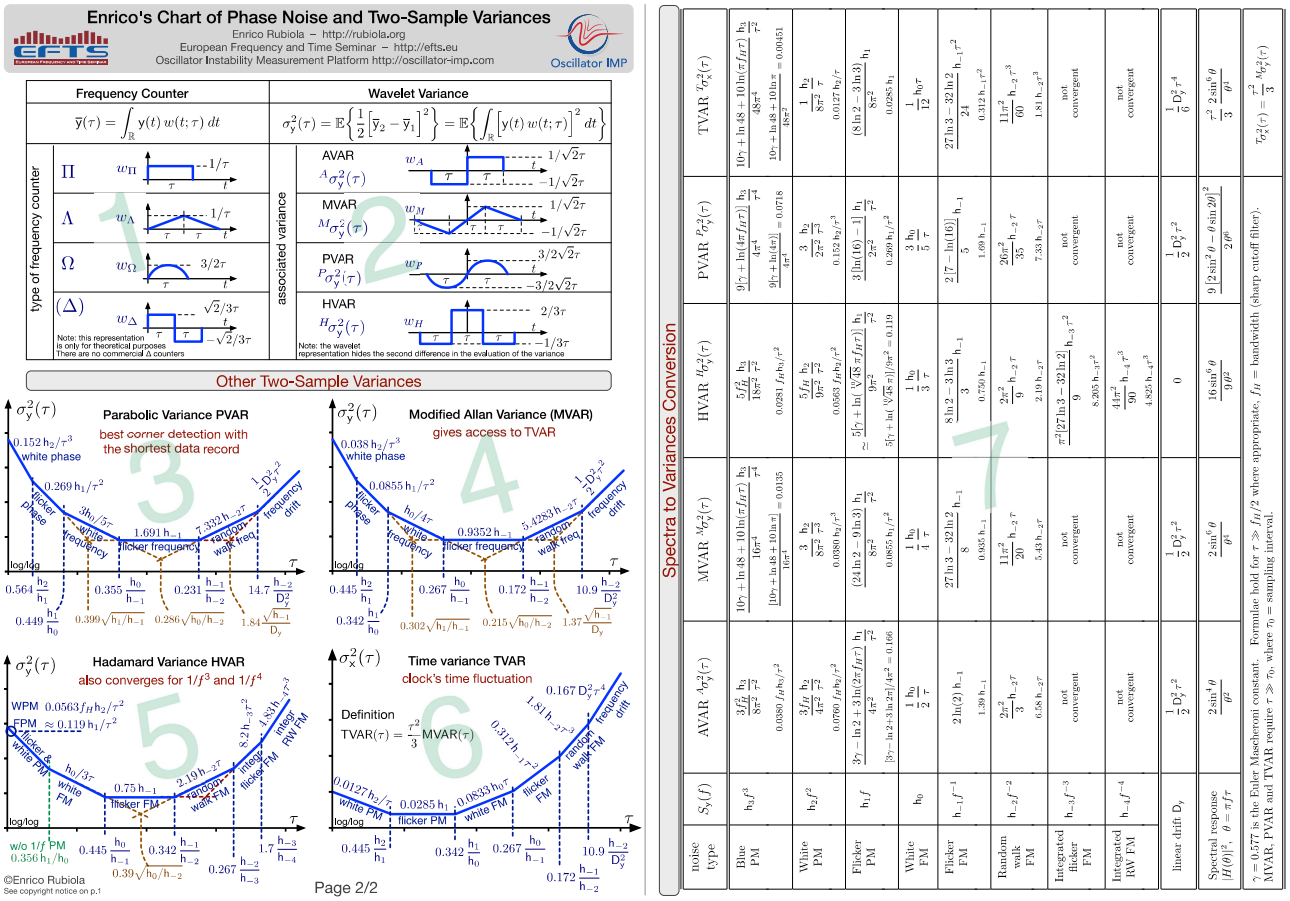

Fig. 3. Enrico's Chart of phase noise and two-sample variances, back side.

Phase noise challenges the networks intended for fundamental science and clock comparison between metrological labs using optical fibers shared Internet data traffic [26], [27], [28], [29].

In the new version of the International System of Units SI [30], virtually all measurement units rely on time. Thus, time fluctuations, and equivalently frequency fluctuations, impact all branches of metrology. For example, the ampere, the unit of electric current, relies on counting the electrons flowing in the unit of time. The Josephson voltage standard is based on the conversion from photon energy to voltage. The Nobel Prize in Physics was awarded to Brian D. Josephson in 1973 for the theoretical prediction of this effect.

We now discuss Enrico's Chart, the subject of this article, which brings the efforts to define and standardize phase noise and time-domain frequency instabilities to a concise two-page working summary.

B. Enrico's Chart

Enrico's Chart of phase noise and two-sample variances is a reference card collecting the most useful definitions, formulas, and plots in the domain of phase noise and frequency stability. Starting from the first draft in 2011, it has been improved from a loose single page on the back of the program of the European Frequency and Time Seminar[1] to quite a dense front/back format. At least 1000 plasticized prints have been distributed as learning material at the EFTS and at invited conferences, courses, and seminars. The current form results from the scientific contribution of the two authors and from feedback and amendments by the users. The name on the Chart is mainly historical, but it reflects Enrico's tenacity in maintaining, updating, improving the graphic design, and distributing the Chart.

C. License and Distribution

Enrico's Chart is a digital object (PDF file) available from Zenodo as DOI: 10.5281/zenodo.4399218 and released under Creative Commons 4.0 CC-BY-NC-ND license. This DOI is called *concept DOI* because it always resolves to the latest version. Albeit Zenodo delivers a separate DOI for each version, version DOIs should only be used in special cases. The QR codes of Fig. 1 point to the Chart and this article. Redistribution is encouraged as the Internet link or as the QR code, not as a file.

The "NC" and "ND" copyright attributions deserve a comment. The "NC" restriction is mainly intended to prevent

[1]The EFTS http://efts.eu is a full-week crash course with lectures and lab sessions funded in 2013 and held in Besançon every summer since.

selling the Chart for profit. In contrast, if companies distribute the Chart *for free* with their products, that is *not considered commercial use*, we *encourage* this practice. The equipment employed for measurements in time and frequency is a notable example. The "ND" restriction preserves the academic independence of the Chart, preventing companies from modifying it, or from building advertisements upon it.

For any other use, Enrico himself is the preferred contact.

D. About this Article

This article is intended: 1) to accompany the Chart; 2) as a short tutorial; 3) as learning material for lectures; and, besides, 4) it can be cited as a summary of modern notation, provided that the reader agrees with our choices.

The reader having an already good understanding of time and frequency may go straight to Enrico's Chart and put this article aside for later reading or refer to students and younger colleagues. Less experienced people may appreciate this article as a tutorial or as a review. To them, if they have been around for long enough, phase noise and frequency stability may relate to a well-identified problem.

A reduced copy of the Chart is included (see Figs. 2 and 3); the indication "*Region 1.n*" in the text means that we refer to the region *n* defined by the watermarks on page 1 (see Fig. 2) of the Chart, while "*Region 2.n*" refers to the region *n* on page 2 (see Fig. 2). In any event, it is a good idea to have on hand a separate copy of the Chart.

E. Suggested Introductory Readings

We advise starting from the latest version of the IEEE Standard 1139 [17]. The appendices of this standard cover topics similar to ours, but presentation and standpoint are surprisingly different. This article is our personal view, whereas the IEEE Standard is the outcome of a rather large committee. The recommendations of the International Telecommunication Union (ITU) [31], [32] are must too. The European Telecommunications Standards Institute (ETSI) provides a series of standards related to clock synchronization in networks [33, Parts 1–7].

Reference [16] is one of the very first articles that all readers should study. It defines the language and the notation still in use, it introduces the spectra, the variances, and their relations, and it explains the early experimental methods. The quantities φ, x, $\Delta\nu$, and y were first defined there. People interested in the birth of the scientific ideas that originated most concepts of phase noise and frequency stability should review [14]. This is the Proceedings book of a one-time IEEE-NASA conference that took impetus from the DSN. Reference [34] is a review article about the progress on the concepts of phase noise and frequency stability after the early ideas (see [35]).

Turning our attention to books and booklets, the Riley handbook [36] is freely available, sponsored by NIST. Emphasis is on Allan and Allan-like variances, rather than noise spectra, making extensive use of the Stable32 software package.[2]

[2]This and other software packages are discussed in Section XI.

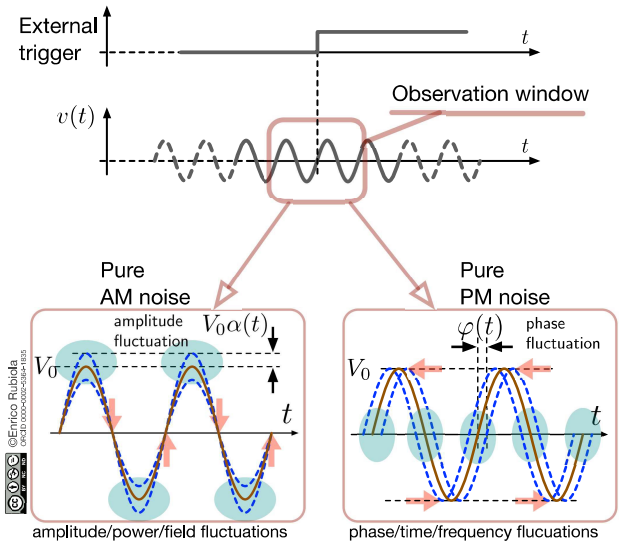

Fig. 4. Clock signal (1), observed with an oscilloscope triggered by an external noise-free reference. The shadowed (pink) arrows indicate where the signal is stationary, unaffected by noise. The elliptical shadowed (cyan) regions emphasize where noise shows up most.

Reference [37] is a rather extensive practical guide about phase noise, albeit elderly. Reference [38] is a collection of classical articles about phase noise and frequency stability, a few of which are cited elsewhere in this article. The same author also published a monograph [39]. Reference [40] is another edited book collecting classical articles from NIST. Reference [41, Chapter 2 (136 pages)] is a recent text entirely about phase noise and frequency stability.

Additional material and large slideshows are available on Enrico's home page: http://rubiola.org. Upon request, the original PPTX files may be released to qualified users. We turn now to the Chart and Fig. 2.

II. Phase Noise and Amplitude Noise

A. Clock Signal (Region 1.1)

A pure sinusoidal signal affected by AM and PM noises (see Fig. 4) can be written as

$$v(t) = V_0\big[1+\alpha(t)\big]\cos\big[2\pi\nu_0 t + \varphi(t)\big] \qquad (1)$$

where V_0 is the amplitude, ν_0 is the frequency, $\alpha(t)$ is the random fractional amplitude, and $\varphi(t)$ is the random phase. This representation is general, not limited to electrical signals as the symbol $v(t)$ suggests. In the IEEE Standard 1139, the amplitude is written as $V_0+\epsilon(t)$, and the fractional amplitude is defined as $\alpha(t) = \epsilon(t)/V_0$, which is equivalent to our notation. However, (1) is a simplification of (24).

The bandwidth B of $\varphi(t)$ and $\alpha(t)$ deserves some attention. Because modulation needs sidebands with appropriate symmetry around ν_0, the theoretical maximum B is equal to ν_0. That said, in numerous cases of interest, the quantity B/ν_0 is rather small, likely of $10^{-4},\dots,10^{-2}$, and not necessarily the same for $\varphi(t)$ and $\alpha(t)$. The autocorrelation function of noise spans over a time of the order of $1/B$ or equivalently ν_0/B oscillations. Thus, cycle-to-cycle noise is too small to

be visible. This is why we use the external trigger in Fig. 4. Otherwise, triggering on the signal itself would require a dual time base with an ideally stable delay $> 1/B$.

Naively, one may believe that $\overline{\alpha(t)} \approx 0$ and $\overline{\varphi(t)} \approx 0$ hold in general, where the "overline" means average or to take these conditions as necessary. In systems of practical interest, we observe quite small amplitude and frequency fluctuations, and thus, we can assume that $|\alpha| \ll 1$. We have measured $|\alpha|$ of the order of $10^{-7}, \ldots, 10^{-4}$ in quartz oscillators and synthesizers [42]. The hypothesis that $|\varphi(t)| \ll 1$ rad makes sense only for two-port systems, where $\varphi(t)/2\pi\nu_0$ is the fluctuation of the input-to-output delay. Conversely, the phase of oscillators always contains divergent processes, so we switch our attention to $|\dot{\varphi}|/2\pi\nu_0$. Practical values span from 10^{-4} for cheap watches to 10^{-16} (accuracy) and 10^{-18} (fluctuations) for the frequency-standard prototypes found in metrology labs.

B. Phase Noise Spectrum (Region 1.4)

The *variance*[3] (generalized power) of a quantity q, denoted by σ^2, is the mean square of q. Formally, $\sigma^2 = \mathbb{E}\{|q - \mu|^2\}$, where $\mathbb{E}\{\,\}$ is the mathematical expectation and $\mu = \mathbb{E}\{q\}$ is the average. The *PSD*, denoted with $S(f)$, tells us how σ^2 is distributed in frequency. The variable f is called "Fourier frequency" to differentiate it from the carrier frequency (constant). The single-sided PSD ($f > 0$) is generally preferred to the two-sided PSD with no need of saying so. Referring to the quantity q, the subscript q is optionally added, as in σ_q^2 and $S_q(f)$.

For our purposes, $S_\varphi(f)$ is the quantity that should be used to describe the phase noise. In practical measurements, $S_\varphi(f)$ is best *evaluated* as

$$\widehat{S_\varphi}(f) = \frac{2}{T}\Big\langle \Phi_T(f)\, \Phi_T^*(f) \Big\rangle_m \quad [\mathrm{rad}^2/\mathrm{Hz}] \tag{2}$$

where the hat accent means "estimator" of S_φ (the reader may ignore it at first reading), $\Phi_T(f)$ is the FFT of $\varphi(t)$ sampled and truncated on an appropriate duration T, the superscript "$*$" means the complex conjugate, the $\langle\,\rangle_m$ operator is the average on m acquisitions, and the factor of 2 is needed for energy conservation after deleting the negative frequencies.[4] Equation (2) is used in the Welch algorithm for the estimation of power spectra [43]. The optional data overlapping used in the Welch algorithm and the optional window function are not explicit in (2). The most popular window functions are known under the names Bartlett, Blackman-Harris, flat-top, Hamming, Hann, Parzen, and, of course, Welch. The quantity $S(f) = \ldots$ defined by (2), with optional windowing but with no hat and no averaging ($m = 1$), is called a periodogram.

[3]The two-sample variance that we find later is a specialized instance of this general concept.

[4]Because $\varphi(t)$ is a real function, its Fourier transform $\Phi(f)$ is Hermitian function, that is, $\Phi(f) = \Phi^*(-f)$. Accordingly, $\Re\{\Phi(f)\}$ is the even function of f, and $\Im\{\Phi(f)\}$ is the odd function of f. Thus, all the information is contained in the $f > 0$ half-plane, and the $f < 0$ half-plane is redundant and can be deleted. In the case of the FFT, such redundant region is generally mapped to $f_s/2 > f > f_s$, where f_s is the sampling frequency.

C. Deeper Thoughts About the PSD

In statistics, the PSD $S(f)$ is formally defined as the Fourier transform of the autocovariance of a *random process*. In turn, the random process is a set of sample functions or distributions, called realizations, each of which is indexed by one outcome of a random experiment. Time statistics and ensemble statistics are different concepts, and they are interchangeable only in the case of ergodic processes.

In contrast, some authors take (2) as the definition of $S_\varphi(f)$, with no "hat." The problem with this choice is that it avoids some key concepts of statistics, making the uncertainty difficult to understand.

Using the mathematical concept of process requires one to define the random experiment. Since mathematics does not provide general rules for this, we need to make a choice. For example, we may identify the random experiment with a large abstract class or more pragmatically with the action of picking one oscillator from preproduction samples, from a batch, or from a larger part of a production. In turn, a realization may be identified with the waveform $\varphi(t)$ obtained by comparing such oscillator with a noise-free reference. This opens deeper questions, like the meaning and the scientific legitimacy of "typical" spectra (also stability, aging, drift, and other parameters) found in data sheets.

As a matter of fact, the right-hand side of (2) can always be calculated from experimental data. The question arises, whether or not the estimator converges to the PSD

$$\widehat{S_\varphi}(f) \to S_\varphi(f) \quad \text{for large } m. \tag{3}$$

This is true for stationary processes (the statistical properties are independent of the origin of time). That said, relevant processes often found in oscillators, such as flicker and random walk of frequency, are not stationary in a strict sense. Fortunately, evaluating (2) such processes can be treated as stationary.

Finally, we notice some analogies between *stationarity* and *repeatability*, and also between *ergodicity* and *reproducibility*, with the caveat that stationarity and ergodicity are mathematical concepts, while repeatability and reproducibility are defined by the International Vocabulary of Metrology VIM [44, Entries 2.21 and 2.25] and have quite a technical meaning related to experimental outcomes.

D. Quantity $\mathscr{L}(f)$ and the Related Measurement Units

Most often in the technical literature and data sheets, the term *phase noise* refers to the quantity $\mathscr{L}(f)$, defined since the IEEE Standard 1139-1988 [45] as

$$\text{definition: } \mathscr{L}(f) = \frac{1}{2} S_\varphi(f). \tag{4}$$

Plots and numerical values are always given as

$$10 \log_{10} \mathscr{L}(f) \quad [\mathrm{dBc/Hz}]. \tag{5}$$

After over 30 years in the field, we believe that $\mathscr{L}(f)$ is misleading, and if the community started from scratch, $S_\varphi(f)$ would be used instead. The roots are found in a one-time symposium coorganized by IEEE and NASA in 1964 [14].

Because $\varphi(t)$ is an angle, the dimension of $S_\varphi(f)$ is a square angle multiplied by time. Accordingly, the appropriate unit is rad for $\varphi(t)$ and $\mathrm{rad}^2\mathrm{s}$ for $S_\varphi(f)$, but the equivalent unit $\mathrm{rad}^2/\mathrm{Hz}$ is generally preferred. In logarithmic units, we use

$$10\log_{10} S_\varphi(f) \quad [\mathrm{dBrad}^2/\mathrm{Hz}]. \tag{6}$$

It follows from (4) that $\mathscr{L}(f)$ has the same dimension as $S_\varphi(f)$ but different units, such as a mass in kg or in lb. Accordingly, the unit associated with $\mathscr{L}(f)$ should be $\mathfrak{A}^2/\mathrm{Hz}$ where $\mathfrak{A}$ in a never-used unit of angle that equals $\sqrt{2}\,\mathrm{rad} \simeq 81°$. It is worth mentioning that there is no reason to change the symbol after switching unit, such as in $M = 1.5\,\mathrm{kg}$ and $M = 3.3\,\mathrm{lb}$. Thus, why should we change the symbol from S_φ to $\mathscr{L}$ because of the unit of angle?

The unit $\mathrm{dBc/Hz}$ is even more confusing. It is obvious that "c" cannot be read "referred to the carrier," as most people have in mind. Instead, taking (4) and (5) literally, "c" is a square unit of angle, $\mathfrak{A}^2 = 2\ \mathrm{rad}^2$. Notwithstanding this, "c" and "c/Hz" alone, with no "dB," are never seen in the literature and no unit of angle associated with $\mathscr{L}(f)$ either.

We advocate abandoning $\mathscr{L}(f)$ in favor of $S_\varphi(f)$ because $S_\varphi(f)$ is consistent with the International System of Units SI [30]—with the minor caveat that a decibel is a non-SI unit accepted for use with the SI units; in contrast, $\mathscr{L}(f)$ is not consistent. In order to prevent discontinuity in notation and units, $\mathscr{L}(f)$ may still be kept as an auxiliary scale.

E. Former Definition of $\mathscr{L}(f)$ and the Deprecated Terms "SSB Noise" and "Offset Frequency"

In the early time, $\mathscr{L}(f)$ was defined as

$$\mathscr{L}(f) = \frac{\text{noise power in 1 Hz bandwidth}}{\text{carrier power}} \quad \text{(wrong!)} \tag{7}$$

and always given as $10\log_{10}\mathscr{L}(f)$ in $\mathrm{dBc/Hz}$. The historical meaning of the symbol "c" is "referred to the carrier." For example, $-120\ \mathrm{dBc/Hz}$ means that "the noise sideband in 1-Hz bandwidth is 120 dB below the carrier power." Sadly, more than 30 years after the first version of the IEEE Standard 1139 [45], the old definition is still in the mind of numerous engineers and physicists.

The major problem is that PM and AM as well need sidebands with the appropriate symmetry with respect to the carrier. The relationships between LSB and USB define the type of modulation, AM, PM, or any combination of these. Thus, the noise power in one sideband does not say what fraction goes to phase noise and to amplitude noise. Consequently, the terms "SSB noise" and "offset frequency" are unsuitable to describe the phase noise and should be avoided. The variable f in $S_\varphi(f)$ should be referred to as the "modulation frequency" or as the "Fourier frequency" in formal mathematical language.

We point out that the use of (7) for phase noise is given as follows:

1) *conceptually incorrect*, because it does not divide PM noise from AM noise;
2) *experimentally incorrect*, because $\mathscr{L}(f)$ is always measured with a phase detector using $\mathscr{L}(f) = (1/2)S_\varphi(f)$, instead of the noise-to-carrier ratio;
3) *deprecated*, as a result of a major effort to adhere to the SI as the global system of units;
4) *unsuitable* to describe phase noise exceeding a small fraction of a radian.

Large phase swings are common at low Fourier frequencies (long measurement time) and in optics, where the carrier frequency is high. When $\int S_\varphi(f)\,df$ integrated over the full spectrum approaches or exceeds $1\ \mathrm{rad}^2/\mathrm{Hz}$, (7) breaks down because the noise sidebands in angular modulations come at the expense of the carrier power. Conversely, there is no reason to question the validity of (2) and (4), even in the case of a huge number of cycles. In the presence of large angles, the correct measurement of $S_\varphi(f)$ is only a matter of hardware design.

Finally, the power ratio as defined in the right-hand side of (7) is definitely not evil—provided that we do not call it "phase noise," and we do not use the same symbol $\mathscr{L}(f)$ defined by (4). In some domains, the noise per unit of bandwidth matters, regardless of the physical origin. This is the case of the reciprocal mixing in the superheterodyne receiver [46, Sec. 8.7.2], where a strong emission in neighbor channels leaks into the IF because of the local-oscillator noise sidebands. Similarly, Doppler radar visibility of subclutter targets is typically determined by noise in the oscillator's sideband power density. See [1], [3], [4], [5], and [47] as general references on radars. As noted previously, due to amplitude limiting, sideband noise in the cited examples is generally dominated by phase noise, leading to the persistence of spectrum sideband usage in these applications.

F. Suggested Readings

Digging into the old literature, the earliest use of the phase noise PSD is arguably Victor 1956 [48]. In the NASA Symposium [14], both $S_\varphi(f)$ and the sideband-to-carrier ratio were used, the latter mainly for radar and receivers. However, the symbol $\mathscr{L}(f)$ had still to come. Its early occurrences in archival documents are Blair [49] and Shoaf [50, pp. 167, 179, and 180] both in 1974. Interestingly, the experimental techniques described in [14] rely on phase detectors, which points to $S_\varphi(f)$. Reference [51] is the one and only remarkable counterexample that we found, where the sideband-to-carrier ratio is measured with practically useful dynamic range. Note that the history of phase noise is beyond our scope and our skills; we recommend Leeson's review article [52, Secs. III–V] and a blog page on Engineering and Technology History Wiki[5] submitted by Leeson, which refers to additional material of historical interest.

Most of the references about phase noise relate to experimental methods; thus, they are moved to Section X. The early ideas found in [14] were better formalized in [16] (see [34]).

There is a Special Issue of the PROCEEDINGS OF THE IEEE [15], three useful books, [53], [6], and [7], and a free booklet

[5]D. B. Leeson, *First-Hand: Phase Noise*, https://ethw.org/First-Hand:Phase_Noise

from the NPL [37]. There is very little about AM noise. The reader should refer to [54].

III. Useful Quantities (Region 1.10–1.11)

A. Phase Time Fluctuation (or Phase Time) $\mathsf{x}(t)$

The *phase time* is the phase fluctuation converted into time

$$\text{definition: } \mathsf{x}(t) = \frac{\varphi(t)}{2\pi\nu_0} \quad [\mathrm{s}]. \tag{8}$$

Its spectrum $S_{\mathsf{x}}(f)$ follows from the definition of $\mathsf{x}(t)$ using the property of the Fourier transform that the derivative operator d/dt maps into multiplication by $j2\pi f$, where $j^2 = -1$, thus, into multiplication by $4\pi^2 f^2$ in the PSD because of (2). Accordingly,

$$S_{\mathsf{x}}(f) = \frac{1}{4\pi^2\nu_0^2} S_\varphi(f) \quad [\mathrm{s}^2/\mathrm{Hz} \equiv \mathrm{s}^3]. \tag{9}$$

B. Frequency Fluctuation $(\Delta\nu)(t)$

The instantaneous *frequency fluctuation* is

$$\text{definition: } (\Delta\nu)(t) = \frac{1}{2\pi}\frac{d\varphi(t)}{dt} \quad [\mathrm{Hz}]. \tag{10}$$

Enclosing $\Delta\nu$ in parentheses emphasizes the fact that $(\Delta\nu)$ is an unbreakable quantity. The PSD is found using the property that $d/dt \to \times 4\pi^2 f^2$; thus,

$$S_\nu(f) = f^2 S_\varphi(f) \quad [\mathrm{Hz}^2/\mathrm{Hz} \equiv \ \mathrm{Hz}]. \tag{11}$$

Albeit the notation $S_{\Delta\nu}$ is more common than S_ν, the latter should be preferred because the PSD is already insensitive to the constant ν_0, by definition; thus, it shows only the fluctuation $\nu - \nu_0$.

C. Fractional Frequency Fluctuation $\mathsf{y}(t)$

The *fractional frequency fluctuation* is

$$\text{definition: } \mathsf{y}(t) = \frac{(\Delta\nu)(t)}{\nu_0} \quad [\text{dimensionless}]. \tag{12}$$

It follows from (8) to (12) that

$$\mathsf{y}(t) = \frac{d\mathsf{x}(t)}{dt} = \frac{1}{2\pi\nu_0}\frac{d\varphi(t)}{dt}. \tag{13}$$

Again, the PSD is found using $d/dt \to \times 4\pi^2 f^2$; thus,

$$S_{\mathsf{y}}(f) = \frac{f^2}{\nu_0^2} S_\varphi(f) \quad [\mathrm{Hz}^{-1} \equiv \mathrm{s}]. \tag{14}$$

D. Polynomial Law or Power Law (Region 1.12–1.15)

The polynomial fit, known as the *power law* or *polynomial law*, is widely used to model phase noise and related quantities. It is often written as

$$S_\varphi(f) = \sum_{n=-4}^{0} \mathsf{b}_n f^n \tag{15}$$

$$S_{\mathsf{x}}(f) = \sum_{n=-4}^{0} \mathsf{k}_n f^n \tag{16}$$

TABLE I
Polynomial Law and Basic Types of Noise

Noise type	$S_\varphi(f)$	$S_{\mathsf{x}}(f)$	$S_\nu(f)$	$S_{\mathsf{y}}(f)$
Blue phase noise*	b_1	k_1	$\mathsf{d}_3 f^3$	$\mathsf{h}_3 f^3$
White phase noise	b_0	k_0	$\mathsf{d}_2 f^2$	$\mathsf{h}_2 f^2$
Flicker phase noise	$\frac{\mathsf{b}_{-1}}{f}$	$\frac{\mathsf{k}_{-1}}{f}$	$\mathsf{d}_1 f$	$\mathsf{h}_1 f$
White frequency noise	$\frac{\mathsf{b}_{-2}}{f^2}$	$\frac{\mathsf{k}_{-2}}{f^2}$	d_0	h_0
Flicker frequency noise	$\frac{\mathsf{b}_{-3}}{f^3}$	$\frac{\mathsf{k}_{-3}}{f^3}$	$\frac{\mathsf{d}_{-1}}{f}$	$\frac{\mathsf{h}_{-1}}{f}$
Frequency random walk	$\frac{\mathsf{b}_{-4}}{f^4}$	$\frac{\mathsf{k}_{-4}}{f^4}$	$\frac{\mathsf{d}_{-2}}{f^2}$	$\frac{\mathsf{h}_{-2}}{f^2}$
Integrated flicker FM noise*	$\frac{\mathsf{b}_{-5}}{f^5}$	$\frac{\mathsf{k}_{-5}}{f^5}$	$\frac{\mathsf{d}_{-1}}{f^3}$	$\frac{\mathsf{h}_{-1}}{f^3}$
Integrated RW frequency*	$\frac{\mathsf{b}_{-6}}{f^6}$	$\frac{\mathsf{k}_{-6}}{f^6}$	$\frac{\mathsf{d}_{-2}}{f^4}$	$\frac{\mathsf{h}_{-2}}{f^4}$

$\star$ Process not shown in the plots of $S(f)$ and $\sigma_{\mathsf{y}}^2(\tau)$

where the values of n correspond to the noise types listed in Table I. These noise types are found in oscillators, with additional negative-exponent terms sometimes needed, $n < -4$. Limitations apply to two-port devices because the input-to-output delay is not allowed to diverge (cf. Sections IV and V).

Transposing (15) and (16) to frequency noise, the polynomial law is written as

$$S_\nu(f) = \sum_{n=-2}^{2} \mathsf{d}_n f^n \tag{17}$$

$$S_{\mathsf{y}}(f) = \sum_{n=-2}^{2} \mathsf{h}_n f^n. \tag{18}$$

Notice that, for a given process the exponents of f differ by 2 from (15) and (16) to (17) and (18), in agreement with the bounds of the sum.

In proper mathematical terms, (15)–(18) are Laurent polynomials, which is the extension of the regular polynomials to negative powers of the variable.

E. Quantities φ, x, $\Delta\nu$, and y in Frequency Synthesis

The ideal, noise-free synthesizer is the electrical analog of a play-free gearbox. It delivers an output frequency $\nu_o = (\mathcal{N}/\mathcal{D})\nu_r$, where $\mathcal{N}/\mathcal{D}$ is the rational number, which defines the synthesis ratio, and ν_r is the reference frequency. Thus, the synthesizer transfers the quantities $\mathsf{x}(t)$ and $\mathsf{y}(t)$ from the reference input to the output, unchanged. For example, shifting the reference by $+1.2$ ppm, the output frequency will be $+1.2$ ppm off the nominal value, thus 150 Hz higher if the output is set to 125 MHz. Similarly, introducing a 100-ps delay with a line stretcher at the input results in the output shifted by 100 ps. This value is the same at 5- and 125-MHz output frequencies.

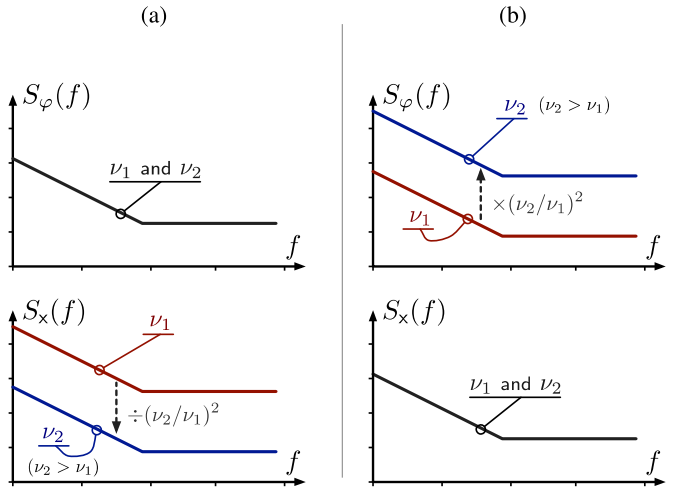

Fig. 5. Basic phase noise behavior of components and systems: (a) phase-type and (b) time-type. We show how the carrier frequency affects or does not affect $S_\varphi(f)$ and $S_\mathsf{x}(f)$. The effect of aliasing, not shown here, may change the white-noise scaling from $(\nu_2/\nu_1)^2$ into ν_2/ν_1. In contrast, flicker and other colored noise types are immune from aliasing.

By contrast, the synthesis is ruled by $\varphi = (\mathcal{N}/\mathcal{D})\psi$ when we express the phase shift as an angle, ψ at the input, and φ at the output. For example, a +1-mrad shift of the 10-MHz reference results in a +12.5-mrad shift if the output is set to 125 MHz. Accordingly, the phase noise spectrum is ruled by

$$S_\varphi(f) = \left(\frac{\mathcal{N}}{\mathcal{D}}\right)^2 S_\psi(f). \tag{19}$$

The above statements are simplistic, to the extent that we have not included the dynamic behavior, the noise bandwidth, and other phenomena. The following limitations apply.

1) Digital Dividers: In digital signals, phase noise exists only on rising and falling edges; thus, it is sampled at $2\nu_0$, twice its own frequency. Frequency division $\div\mathcal{D}$ results in lower sampling frequency, which originates aliasing. Aliasing impacts only white noise because folding multiple aliases proportional to $1/f$ has negligible effect. Thus, the white PM noise of the divided signal is ruled by $S_\varphi(f) = (1/\mathcal{D})S_\psi(f)$ instead of $(1/\mathcal{D}^2)S_\psi(f)$.

2) Output Stage: For $\mathcal{N}/\mathcal{D} \ll 1$, the scaled-down phase noise may hit the phase noise of the output stage. When this happens, $S_\varphi(f)$ is limited by the output stage.

3) High-Order Multiplication: Angular modulations are ruled by the property that the total power is constant; thus, the sideband power comes at expense of the carrier power. Carrier and nth sideband amplitudes are described by the Bessel functions $J_n(m)$. For $\mathcal{N}/\mathcal{D} \gg 1$, the random phase in the full bandwidth may approach or exceed 2.4 rad, where $J_0(m) = 0$. When this happens, all the power goes into the sidebands, and the carrier disappears. This is known as the "carrier collapse" in multiplication.

F. Phase-Type (φ-Type) and Time-Type (x-Type) PM Noise

A relevant aspect of the noise behavior of components and systems in the presence of large frequency changes is well described by the terms *phase-type* and *time-type* phase noise (see Fig. 5). We first introduced this idea in [55, Sec. 3] to describe noise in the clock distribution of digital electronics. The same terminology applies to the following:

1) amplifiers, logic gates, and so on, where the output frequency is the same as the input;
2) frequency multiplication, division, and synthesis, where the input frequency is scaled up/down by a rational number;
3) frequency synthesis, where the output frequency is set to different values.

1) Phase-Type (φ-Type) Phase Noise [see Fig. 5(a)]: The quantity $S_\varphi(f)$ is a parameter of the device or system, not affected by the carrier frequency. The dominant noise processes impacts $S_\varphi(f)$ independently of ν_0, and $S_\mathsf{x}(f)$ follows from (9). Additive noise in amplifiers (see Section IV), originating from the noise-to-carrier ratio (25), is like this. Often flicker PM in amplifiers is also this type, described by (26). Phase-type noise is also observed in frequency dividers, DDSs, and DACs at low ν_o (output), where the scaled-down input noise is lower than the limit set by the output stage.

For example, the white PM noise of an RF amplifier in a reference condition is $S_\varphi(f) = \mathsf{b}_0 = 10^{-16}$ rad^2/Hz, limited by noise figure and carrier power. Section IV-C explains how this can be calculated. Such noise is of the phase type, as shown in Fig. 5(a). Using (9), we get $S_\mathsf{x} = 2.53 \times 10^{-32}$ s^2/Hz when the carrier frequency takes the value $\nu_1 = 10$ MHz, and $S_\mathsf{x}(f) = 1.62 \times 10^{-34}$ s^2/Hz at $\nu_2 = 125$ MHz. At ν_2, the time fluctuation $\mathsf{x}(\mathsf{t})$ is a factor $\nu_1/\nu_2 = 0.08$ smaller than at ν_1, and its PSD is a factor $(\nu_1/\nu_2)^2 = 0.0064$ smaller.

2) Time-Type (x-Type) Phase Noise [see Fig. 5(b)]: The quantity $S_\mathsf{x}(f)$ is a parameter of the device or system, not affected by the carrier frequency. The dominant noise processes impacts the time fluctuation $S_\mathsf{x}(f)$ independently of ν_0, and $S_\varphi(f)$ follows from (9). This behavior is observed in ADCs, DACs, and DDSs at high frequencies, where the PM noise is dominated by the jitter of the internal clock distribution. Acoustic/seismic noise and other environment effects on coaxial cables and optical fibers are obviously like this.

For example, the clock distribution of a Cyclone III FPGA has a flicker $S_\mathsf{x} = \mathsf{k}_{-1} = 4 \times 10^{-28}$ s/Hz. Other way stated, $\sqrt{\mathsf{k}_{-1}} = 20$ fs. Such noise originates from the time fluctuation in the pipeline of logic gates, which is of the time type [see Fig. 5(b)]. Using (9), we expect a flicker $\mathsf{b}_{-1} = 1.6 \times 10^{-10}$ rad^2 (-98 dBrad^2) at 100-MHz clock, 6.3×10^{-10} rad^2 (-92 dBrad^2) at 200 MHz, and 2.5×10^{-9} rad^2 (-86 dBrad^2) at 400 MHz. The real experiment behind this example [55, Fig. 6] shows a rather small discrepancy from the $\times(\nu_2/\nu_1)^2$ law, ≈ 0.5 dB for b_{-1} at the three clock frequencies.

A PLL frequency multiplier is another example, where the output is generated by a VCO phase-locked to a 10-MHz reference after frequency division. The division modulo sets the output frequency. The PLL's own noise—not accounting for the external reference—is typically determined by the 10-MHz phase comparator. Assuming that the comparator's

flicker is $\mathsf{b}_{-1} = 3.2 \times 10^{-10}$ rad^2 (-95 dBrad^2), the output b_{-1} is of -95 dBrad^2 at 10-MHz output, -89 dBrad^2 at 20 MHz, -83 dBrad^2 at 40 MHz, and so on, in 6-dB steps per factor-of-2. This is exactly what we observe on the internal PLL of a Cyclone III FPGA, up to -59 dBrad^2 at 640 MHz, the maximum output frequency [55, Fig. 6].

3) Aliasing: The two noise types described may be modified by aliasing, which changes the scaling rule from $(\nu_2/\nu_1)^2$ to ν_2/ν_1. Aliasing only affects white noise, not flicker, and other reddish noise types because the tails of the PSD are too low for the folded-down aliases to be significant. Aliasing is typically seen in digital circuits, and in clock distribution of FPGAs and other complex devices as well, at low frequency [55, Sec. 3].

G. Notation (Region 1.2)

We use the sans serif font in $\mathsf{x}(t)$ and $\mathsf{y}(t)$, instead of the regular math font commonly found in the literature, to emphasize that x and y are special quantities defined by (8) and (12). This choice sets the regular x and y free for general use. The same applies to the coefficients b_i, d_i, h_i, and k_i.

Working with digital systems, we are regularly faced with phase exceeding $\pm\pi$ because IQ detectors and digital dividers keep a record of multiple cycles of the carrier. For this purpose, we find it useful to describe the clock signal with the quantity written in boldface, which is the sum of the deterministic (or nominal) quantity plus the fluctuation

$$\text{phase: } \boldsymbol{\varphi}(t) = 2\pi\nu_0 t + \varphi(t) \tag{20}$$

$$\text{frequency: } \boldsymbol{\nu}(t) = \nu_0 + (\Delta\nu)(t) \tag{21}$$

$$\text{time: } \mathbf{x}(t) = t + \mathsf{x}(t) \tag{22}$$

$$\text{fractional frequency: } \mathbf{y}(t) = 1 + \mathsf{y}(t). \tag{23}$$

The quantity $\mathbf{x}(t)$ is the most obvious. To the layman, $\mathbf{x}(t)$ is the readout of a clock, which is the sum of the "exact time" t plus the "error" $\mathsf{x}(t)$. A true layman would not consider relativity here and would have no idea about the technical meaning of words, such as "error" and "uncertainty." The quantity $\boldsymbol{\nu}(t)$ is the instantaneous frequency, measured in a sufficiently short time, $\boldsymbol{\varphi}(t)$ is the total phase accumulated after IQ detection, and $\mathbf{y}(t)$ differs from 1 by the small fractional fluctuation $\mathsf{y}(t)$.

IV. Two-Port Components (Region 1.8)

A. Additive and Parametric Noise

Most people find these concepts either quite simple, or rather confusing. The point is that (1) describes the clock signal as it is observed, hiding the physics of noise.

To understand, we start with the example of a noise-free radio broadcasting $v_e(t) = V_e\big[1+\alpha_m(t)\big]\cos\big[2\pi\nu_0 t+\varphi_m(t)\big]$, where $\alpha_m(t)$ is AM and $\varphi_m(t)$ is PM. AM and PM may be present simultaneously, as in the old analog television[6] and in QAM modulations. The *received* signal

$$v(t) = V_0\big[1 + \alpha_m(t)\big]\cos\big[2\pi\nu_0 t + \varphi_m(t)\big] + n(t) \tag{24}$$

includes the noise $n(t)$, which is the receiver's own noise, atmospheric noise, and other forms of noise collected by the antenna. However, the signal is *detected* as (1). In fact, $\varphi(t)$ is the sum of $\varphi_m(t)$ plus the image of $n(t)$ after phase detection. Similarly, $\alpha(t)$ is the sum of $\alpha_m(t)$ plus the image of $n(t)$ after amplitude detection. Random $n(t)$ is called additive noise, while random $\alpha_m(t)$ and $\varphi_m(t)$ are called parametric noise. The relevant difference is that *parametric noise* originates from the low-frequency signals $\varphi_m(t)$ and $\alpha_m(t)$ upconverted to sidebands, while *additive noise* originates around ν_0. If $\varphi_m(t)$ and $\alpha_m(t)$ are small[7] and have bandwidth $B_m \ll \nu_0$, their image in the RF spectrum is a pedestal $\pm B_m$ wide centered at ν_0.

As a relevant fact, white AM and PM noises result from both white $n(t)$ and white $\alpha_m(t)$ and $\varphi_m(t)$. The former is additive, the latter is parametric. Hence, the common belief that white AM/PM noise is only additive is incorrect. Conversely, "colored" noise ($1/f$, $1/f^2$, and so on) is generally of parametric origin. A narrowband $n(t)$ well centered at ν_0 appearing as $1/f$ noise is totally unrealistic, albeit conceptually possible.

B. Added Noise

The term "added noise" is seen in commercial phase noise analyzers to denote the phase noise *added* by a two-port component under test, usually an amplifier, when the component is inserted in the signal path. This is an unfortunate choice because the term "added" is easily mistaken for "additive." Of course, the "added" noise consists of white and flicker PM noise, thermal drift, aging, and so on, while the "additive" noise is white or quasi-white.

In proper terminology, added noise refers to a *differential measurement* because the instrument measures the quantity $\varphi - \psi$, i.e., output minus input, and its statistical properties (spectra and variances, at will). The oscillator noise ψ is a common-mode signal; thus, it is rejected.

C. White Noise

Given a white noise $n(t)$ of PSD N [W/Hz] added to a carrier of power P, the white PM noise's PSD is

$$\mathsf{b}_0 = \frac{N}{P} \quad [\text{rad}^2/\text{Hz}]. \tag{25}$$

Radio engineers often express the noise in terms of noise factor F. Accordingly, N is expanded as $N = FkT_0$, where $kT_0 = 4 \times 10^{-21}$ W/Hz is the thermal energy at the standard temperature $T_0 = 290$ K (17 °C). The noise figure is defined by NF $= 10\log_{10}(F)$. For reference, the phase noise of a noise-free device ($F = 1$ or NF $= 0$ dB) in the presence of a 1-mW carrier (0 dBm) is 4×10^{-18} rad^2/Hz, thus -174 dBrad^2/Hz or -177 dBc/Hz. Besides thermal energy, $n(t)$ may originate from shot noise. In our experience, experimentalists in optics are led to confusion by the fact that, because noise is represented as a temperature, in optics, the thermal energy is generally negligible.

[6]More precisely, TV audio is FM for compatibility with audio broadcasting, which is equivalent to PM.

[7]More precisely, small modulation angle $|\varphi_m(t)| \ll 1$ is necessary for higher order sidebands to negligible and the PM to be accurately described by the first USB and LSB. Conversely, only $|\alpha_m(t)| < 1$ is formally necessary, but the $|\alpha_m(t)| \ll 1$ condition makes sense in real systems.

Additional white PM noise, not accounted for in (25), can result from parametric effects, as explained in Section IV-A.

D. Flicker Noise

Flicker noise has PSD proportional to $1/f$. It has been observed that the flicker PM of RF and microwave amplifiers is rather constant versus P and ν_0

$$\mathsf{b}_{-1} = C \quad \text{(constant versus } P \text{ and } \nu_0\text{)}. \tag{26}$$

It follows from (25) and (26) that the corner frequency $f_c = \mathsf{b}_{-1}/\mathsf{b}_0$, where flicker equals white noise, depends on P.

Albeit the common integral $\int_a^b (1/f)\, df = \ln(b/a)$ diverges for $a \to 0$ or $b \to \infty$, the practical result is surprisingly small. To convince the reader, we evaluate $\ln(b/a)$ for the largest conceivable bounds, from the reciprocal of the age of the universe ($a = 2.3 \times 10^{-18}$ Hz) to the reciprocal of the Planck time ($b = 1.9 \times 10^{43}$ Hz). The results is $\ln(b/a) = 140.3$, i.e., 21.5 dB. Thus, if the flicker coefficient is $\mathsf{k}_{-1} = 10^{-24}$ s^2 ($\sqrt{\mathsf{k}_{-1}} = 1$ ps), the total $1{-}\sigma$ fluctuation is $(140.3 \times (10^{-24}\ \text{s}^2))^{1/2} = 11.8$ ps.

The general literature suggests that the spectrum of flicker is $1/f^\eta$ with η close to one. It turns out that the input-to-output delay never grows too large even for $\eta > 1$. Try this yourself with $\eta = 1.1$ and the integration bounds from 10^{-9} Hz (the reciprocal of 30 years, the supposed device's lifetime) to 10^9-Hz noise bandwidth (above the highest Fourier frequency found in any commercial noise analyzer).

E. Input-to-Output Delay

The fluctuation ($1{-}\sigma$) of the input-to-output delay of a two-port device is given by

$$\delta T = \sqrt{\int_a^b S_{\mathsf{x}}(f)\, df} \tag{27}$$

which follows from the time-fluctuation PSD integrated on the appropriate bandwidth $[a, b]$. Common sense suggests that δT does not diverge nor does it grow disproportionately during the device's life. This limits the phase noise to finite-bandwidth white PM noise and to flicker PM noise.

Environmental parameters, such as humidity and thermal drift, have only localized effects in time or are periodic. Random walk, aging, and other ever-growing phenomena are possible, but their amount is generally small enough not to significantly affect the delay over the lifetime of the device. Electrical engineers may be familiar with similar effects in voltage references or in the offset of analog components.

F. Example of Noise in a Two-Port Component

We consider an amplifier having NF = 2 dB and flicker coefficient $\mathsf{b}_{-1} = 2 \times 10^{-11}$ rad^2 (-107 dBrad^2), processing a 10-GHz carrier of power $P = 62.5\ \mu$W (-12 dBm). Let us calculate the noise spectrum and the rms delay fluctuation assuming a low cutoff $f_1 = 10^{-8}$ Hz (reciprocal of three years, rounded) and a bandwidth $f_2 = 50$ MHz.

1) PM Noise PSD: Using $\mathsf{b}_0 = FkT/P$, in dB, we get $10 \log_{10}(\mathsf{b}_0) = +2 - 174 + 12 = -160$ dBrad^2/Hz; hence, $\mathsf{b}_0 = 10^{-16}$ rad^2/Hz. The phase noise PSD is

$$S_\varphi(f) = \mathsf{b}_0 + \frac{\mathsf{b}_{-1}}{f} = 10^{-16} + \frac{2 \times 10^{-11}}{f} \quad \text{rad}^2/\text{Hz}.$$

Using $\mathsf{b}_0 = \mathsf{b}_{-1}/f$, the corner frequency is $f_c = 2 \times 10^5$ Hz.

2) Phase–Time PSD: Using (9) with $\nu_0 = 10$ GHz, we get

$$S_{\mathsf{x}}(f) = \mathsf{k}_0 + \frac{\mathsf{k}_{-1}}{f} = 2.6 \times 10^{-38} + \frac{5.1 \times 10^{-33}}{f} \quad \text{s}^2/\text{Hz}.$$

3) RMS Delay Fluctuation: This is evaluated using (27).
White noise $(\delta T)_0 = \sqrt{\mathsf{k}_0 (f_2 - f_1)} = 1.1 \times 10^{-15}$ s.
Flicker noise $(\delta T)_{-1} = \sqrt{\mathsf{k}_{-1} \ln(f_2/f_1)} = 4.3 \times 10^{-16}$ s.
Total $\delta T = \sqrt{(\delta T)_0^2 + (\delta T)_{-1}^2} = 1.2 \times 10^{-15}$ s.

4) Comment: Common sense suggests that the delay of an amplifier can be of a fraction of a nanosecond (several periods of the 10 GHz carrier) for a wideband device, with a thermal coefficient up to 10^{-3}/K, if no special design care is taken. As a result, we expect that the thermal effects exceed the random noise, and the random noise is visible only beyond 10 Hz, where the temperature may be stabilized by the thermal capacitance (inertia).

G. Suggested Readings About Two-Port Components

Reference [56] is an extensive treatise of phase noise in amplifiers. Arguably, [57] is the first article suggesting that the flicker noise (the parameter b_{-1}) in RF and microwave amplifiers is independent of power and frequency over a rather broad range.

The double-balanced mixer is a tool of paramount importance in PM noise and frequency stability. Reference [58] is a useful tutorial, and [59] provides experimental data about the noise of commercial double-balanced mixers for phase noise measurements.

We recommend the book [60] to understand analog–digital conversion. Reference [55] is a tutorial on phase noise digital systems. The origin and the propagation of phase noise in frequency division are explained in [61] and [62].

Phase locking and frequency synthesis are totally different kinds of two-port systems. We suggest [63], [64], and [65] for phase locking, [66] and [67] for digital frequency synthesis, [68] for the phase noise in digital synthesis, and [41] for a general treatise about modern synthesizes.

V. Oscillators and the Leeson Effect (Region 1.9)

The oscillator [see Fig. 6(a)] consists of a loop where the resonator sets the oscillation frequency ν_0 and the sustaining amplifier compensates for the resonator loss. Gain clipping (nonlinearity) is necessary to stabilize the amplitude, hence the typical dominance of PM over AM. The buffer isolates the loop from the load.

Everyday experience suggests that, unlike the two-port components, the time fluctuation $\mathsf{x}(t)$ of an oscillator can be quite large. The reason is that the oscillator accumulates the "error" of each cycle, however small it may be.

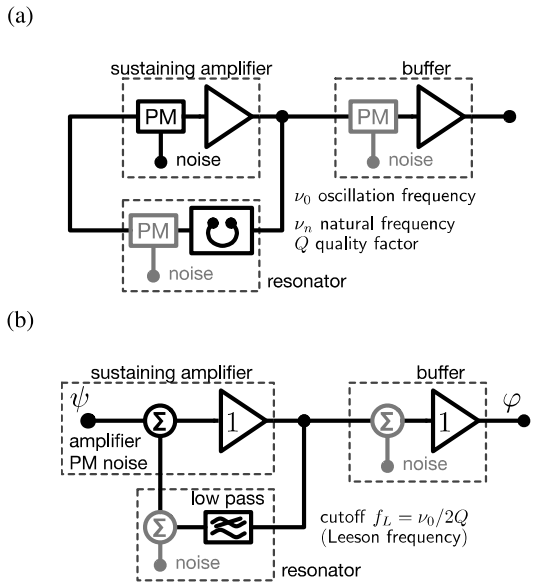

Fig. 6. Phase noise mechanisms inside an oscillator. (a) Conceptual schematic of the oscillator, where the signal is the microwave/RF sinusoidal oscillation at the frequency ν_0. (b) Phase noise equivalent of the oscillator, where the signal is the random phase of the microwave/RF circuit shown above. In this representation, phase noise is always additive, regardless of its physical origin. The grayed blocks in (a) and (b) represent the noise of the resonator and the buffer, not included in the equations.

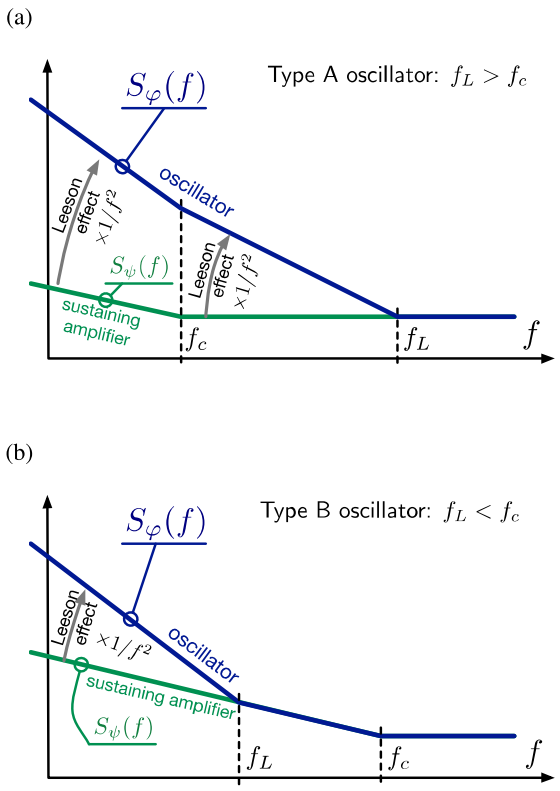

Fig. 7. Typical noise spectra occurring in oscillators. (a) Type A phase noise spectrum, often found in microwave oscillators. (b) Type B phase noise spectrum, often found in RF and quartz oscillators. (a) and (b) Noise $S_\psi(f)$ of the sustaining amplifier, the flicker corner at frequency f_c, the Leeson corner at frequency f_L, and the oscillator noise $S_\varphi(f)$. The fluctuation of the resonator's natural frequency and the noise of the output buffer are not included.

Resonator, sustaining amplifier, and buffer all introduce phase noise (amplitude noise too, but this is a more specialized topic, not considered here). The phase modulators in Fig. 6(a) are necessary to model the phase noise of the various part of the oscillator in the general case, including white and colored processes. However, these modulators make the mathematical treatise difficult. The *phase equivalent* model, introduced in [69, Chapter 4] and shown in Fig. 6(b), avoids such difficulty. The signal circulating in Fig. 6(b) is the phase fluctuation of Fig. 6(a). In this representation, the following holds.

1) Sustaining amplifier and buffer map into "phase amplifiers" of gain exactly equal one because they "copy" the input phase to the output, unchanged. Otherwise stated, the amplifier's delay cannot be stretched/shrunk.
2) Random PM maps into noise *added* in the signal path, regardless of the nature of the process.
3) The phase model is linear because gain clipping has no effect on phase—at least not first-order effects.
4) The resonator (narrowband second-order filter) maps into a single-pole LPF.

This representation holds for a large quality factor, say, $Q \gtrsim 50$, which ensures that signals are sinusoidal. Accepting this limitation, the approximation in Fig. 6(b) adequately models the more general oscillator in Fig. 6(a). Breaking the high-Q hypothesis, the simple mathematical treatise provided below is no longer valid because distortion introduces coupling between amplitude and phase.

We give a short summary of the original derivation [69, Secs. 4.4 and 4.5]. The low-pass impulse response[8] of the resonator is $b(t) = (1/\tau)\, e^{-t/\tau}$, where $\tau = Q/\pi\nu_n$ is the relaxation time, Q is the quality factor in actual load conditions, and ν_n is the natural frequency. In practice, the oscillation frequency ν_0 is so close to ν_n that they are interchangeable. The Laplace transform of $b(t)$ is $B(s) = (1/\tau)/(s + 1/\tau)$. The cutoff frequency of the low-pass equivalent is called *Leeson frequency*, equal to half the resonator's RF bandwidth ν_0/Q

$$f_L = \frac{\nu_0}{2Q}. \tag{28}$$

Letting aside the buffer noise, we introduce the phase noise transfer function of the oscillator

$$\text{definition:} \quad H(s) = \frac{\Phi(s)}{\Psi(s)} \tag{29}$$

where $\Phi(s)$ and $\Psi(s)$ are the Laplace transforms of the output $\varphi(t)$ and of the input $\psi(t)$, respectively. Inspecting the loop on the left-hand side of Fig. 6(b), we find that $H(s) = 1/[1-B(s)]$. The latter expression is derived from the classical rules of feedback, the same used in textbooks to analyze a simple control loop, of to derive the gain of the operational amplifier in a noninverting configuration. Expanding $H(s)$ and replacing $s \to j2\pi f$, we get $|H(f)|^2 = 1 + f_L^2/f^2$. The multiplication by $1/f^2$, which occurs at $f < f_L$, is the "soul" of the Leeson effect: the oscillator integrates the phase

[8]The standard symbol of the impulse response is $h(t)$, and $H(s)$ is its Laplace transform. Here, we use $b(t) \leftrightarrow B(s)$ to save $h(t) \leftrightarrow H(f)$ for the oscillator.

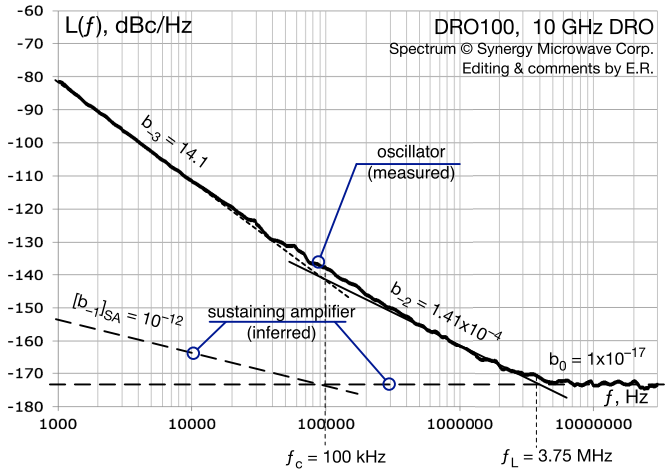

Fig. 8. Example of phase noise of an oscillator. The original plot is courtesy of Ulrich L. Rohde, Synergy Microwave Corporation, used with permission. We added the polynomial approximation (the units of b_i coefficients are omitted), the estimated amplifier noise, and other comments.

fluctuation $\psi(t)$ occurring in the loop. The oscillator phase noise is $S_\varphi(f) = |H(f)|^2 S_\psi(f)$; thus,

$$S_\varphi(f) = \left[1 + \frac{f_L^2}{f^2}\right] S_\psi(f). \tag{30}$$

Plugging the amplifier noise $S_\psi(f) = \mathsf{b}_{-1}/f + \mathsf{b}_0$ in (30), we find the two typical patterns of Fig. 7. Interestingly, the oscillator spectrum contains either $1/f^2$ or $1/f$ noise, not both. The $1/f^2$ noise is common in microwave oscillators, characterized by high ν_0 and low Q; thus, $f_L \gg f_c$. The $1/f$ noise is typical of RF and quartz oscillators, characterized by low ν_0 and high Q.

The noise of the resonator and the buffer, not included in (30) and not shown in Fig. 7, adds. In most quartz oscillators, the $1/f^3$ noise due to the $1/f$ fluctuation of the resonator natural frequency exceeds the $1/f$ contribution of the electronics turned into $1/f^3$ via the Leeson effect [70]. In such cases, the Leeson effect is hidden below the resonator's fluctuations.

A. Example of Noise in a Microwave Oscillator

Fig. 8 shows the phase noise of the DRO-100, a commercial 10-GHz oscillator based on a dielectric resonator. By comparison with Fig. 7, this oscillator is clearly of type A, characterized by $f_c < f_L$. The coefficients b_0, b_{-2}, and b_{-3} are obtained from the graphical approximation of the measured spectrum with the theoretical straight line having slope 0, -2, and -3. The frequencies f_c and f_L are obtained graphically from the straight-line approximation. We start our interpretation from the right-hand side ($f = 10$ MHz) to the left.

The white PM noise $\mathsf{b}_0 = 10^{-17}$ rad^2/Hz suggests that the power at the input of the sustaining amplifier (the location where the power is the lowest) is $P \approx 0.5$ mW. The power in the resonator is substantially the same. This result is obtained using (25), guessing that the noise figure of such an amplifier is 1 dB, quite a plausible value for good amplifiers.

The white frequency noise b_{-2}/f crosses the white phase noise b_0 at 3.75 MHz. This is clearly the Leeson frequency f_L. Using (28), we estimate the resonator's quality factor $Q = 1330$ in actual load conditions.

Finally, the flicker FM noise $\mathsf{b}_{-3} = 14.1$ rad^{2}Hz^2 crosses the white FM noise at $f_c = 100$ kHz. Equation (30) with $f \ll f_c$ suggests that the flicker PM of the sustaining amplifier is of 10^{-12} rad^2 (-120 dBrad^2).

B. Suggested Readings About Oscillators

Our presentation is based on [69], which includes the theoretical proof of the Leeson effect, the analysis of noise sources in the loop and at the output, and a chapter about the reverse engineering of oscillators from phase noise. Rather than focusing on the schematic, Rubiola [69] analyzes the oscillator as a system. A white paper [71] extends the theory to amplitude noise. The limitation of this approach is that it is suitable only to a high-Q resonator (say, $Q \gtrsim 50$), where all relevant signals are sinusoidal.

Numerous RF and quartz oscillator schemes make use of a negative-resistance amplifier in parallel to the resonator. We propose two classical references about the noise of such amplifiers: [72] and [73].

Other approaches deserve attention, chiefly: 1) models based on the limit cycle, derived from Einstein's diffusion theory, i.e., Loh et al. [74], [75] and Demir et al. [76], or from the Langevin equations, Kärtner et al. [77] and 2) the impulse sensitivity function proposed by Lee and Hajimiri [78], [79], which is good at describing low-Q oscillators, like those commonly found in microelectronics. In addition, Pankratz and Sánchez-Sinencio [80] provide a large survey specific to oscillators in integrated circuits.

Digging into the origins, the widely cited van der Pol oscillator [81] is about chaos in oscillations, rather than phase noise. Reference [82] is arguably the first article that analyzes the phase noise in electronic oscillators, and [83] is the article that introduces the phase noise mechanism in feedback oscillators, later known as the Leeson effect. A review article by Leeson [52] is available.

VI. Allan Variance

The classical variance[9]

$$\sigma^2 = \frac{1}{n-1} \sum_{i=1}^{n} [x_i - \mu]^2 \tag{31}$$

where $\mu = (1/n) \sum_{i=1}^{n} x_i$ is the average, fails at describing time divergent processes because: 1) it depends on the averaging time used to take the samples x_i and 2) it depends on the number n of samples. Try this yourself, feeding $x_i = 1.0001, 1.0002, 1.0003, \ldots$, with $n = 2, 4, 8, \ldots$ in (31). A solution consists of introducing the averaging time as a parameter, denoted with τ, and to set $n = 2$. This is the minimum n, which gives a valid σ^2. Welcome to the AVAR.

[9]This is a simplified notation. More precisely, (31) describes an estimator; thus, it should be written as $\widehat{\sigma^2} = \ldots [x_i - \hat{\mu}]^2$, where $\hat{\mu} = \ldots$

A. Definition and Evaluation (Region 1.3)

The two-sample (Allan) variance AVAR of the quantity y is defined as

$$\text{2-sample variance: } \sigma_{\mathsf{y}}^2(\tau) = \mathbb{E}\left\{ \frac{1}{2}\left[\overline{\mathsf{y}}_2 - \overline{\mathsf{y}}_1\right]^2 \right\} \tag{32}$$

where $\mathbb{E}\{\ \}$ is the mathematical expectation, and the averages $\overline{\mathsf{y}}_1$ and $\overline{\mathsf{y}}_2$ are taken over contiguous time slots of duration τ. The quantity $(1/2)\left[\overline{\mathsf{y}}_2 - \overline{\mathsf{y}}_1\right]^2$ is the classical variance evaluated with two samples. This is immediately seen by replacing $n = 2$ in (31), and then, $\mu = (x_1 + x_2)/2$; finally, we identify the generic x_i with the average fractional frequency fluctuation $\overline{\mathsf{y}}_i$. The quantity ADEV, a deviation, the square root of AVAR—and similarly MDEV, PDEV, and so on defined later—can be seen as an estimator of the uncertainty of the quantity y, accumulated in the time τ after reset or calibration.

In experiments, $\mathbb{E}\{\ \}$ is replaced with the estimator

$$\widehat{\sigma_{\mathsf{y}}^2}(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} \left[\overline{\mathsf{y}}_{k+1} - \overline{\mathsf{y}}_k\right]^2 \tag{33}$$

which is the average on $M-1$ realizations of $\overline{\mathsf{y}}_2 - \overline{\mathsf{y}}_1$; thus, it requires M contiguous measures of $\overline{\mathsf{y}}$.

At first reading, one can take (33) as the formula to evaluate AVAR, ignoring the "hat." Some authors use (33) as the definition of AVAR. That said, keeping a clear difference between $\sigma_{\mathsf{y}}^2(\tau)$ and its estimate $\widehat{\sigma_{\mathsf{y}}^2}(\tau)$ is important in theoretical analysis and in the evaluation of the uncertainty.

It is often convenient to rely on time measurements x_k, using $\mathsf{y}_k = (\mathsf{x}_{k+1} - \mathsf{x}_k)/\tau$. Accordingly, (33) rewrites as

$$\widehat{\sigma_{\mathsf{y}}^2}(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} \left[\frac{\mathsf{x}_{k+2} - 2\mathsf{x}_{k+1} + \mathsf{x}_k}{\tau}\right]^2 \tag{34}$$

which requires $M+1$ measures of $\mathsf{x}(t)$ spaced by τ.

B. Spectral Response (Region 1.5)

The AVAR can be calculated from the spectrum using

$$\sigma_{\mathsf{y}}^2(\tau) = \int_0^\infty |H_A(f;\tau)|^2 \, S_{\mathsf{y}}(f)\, df \tag{35}$$

where the transfer function

$$|H_A(f;\tau)|^2 = 2\,\frac{\sin^4(\pi\tau f)}{(\pi\tau f)} \tag{36}$$

is similar to an octave BPF centered at $f \simeq 0.45/\tau$. Unfortunately, such a filter suffers from significant side lobes (inset in the plot of Region 1.5). Notice that $\sigma_{\mathsf{y}}^2(\tau)$ does not converge for white PM noise and flicker PM noise, unless an LPF is introduced to limit the bandwidth, whose cutoff frequency is denoted with f_H.

C. Overlapped Allan Variance

An efficient way to measure the AVAR is to sample $\mathsf{x}(t)$ at the rate $1/\tau_0$, taking $\tau = m\tau_0$, integer m. The kth value of the fractional frequency is evaluated as

$$\overline{\mathsf{y}}_k = \frac{\mathsf{x}_{k+m} - \mathsf{x}_k}{m\tau_0}. \tag{37}$$

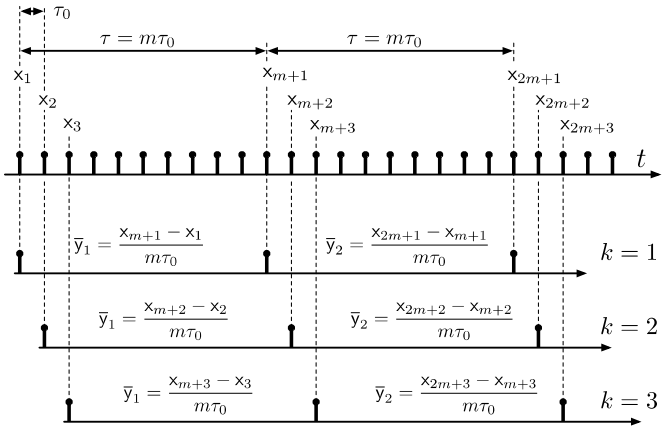

Fig. 9. Evaluation of the overlapped AVAR.

The overlapped AVAR consists of using partially overlapped realizations of $\overline{\mathsf{y}}_2 - \overline{\mathsf{y}}_1$ in (32), separated by the minimum amount τ_0. This concept is illustrated in Fig. 9. Accordingly, (34) becomes

$$\tau = m\tau_0$$
$$\widehat{\sigma_{\mathsf{y}}^2}(m\tau_0) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} \left[\frac{\mathsf{x}_{2m+k} - 2\mathsf{x}_{m+k} + \mathsf{x}_k}{m\tau_0}\right]^2 \tag{38}$$

which takes $2m + M - 1$ samples; thus, a measurement time $\mathcal{T} = (2m + M - 2)\tau_0$.

The first advantage of overlapping is smaller uncertainty. In fact, the confidence interval $\Delta\sigma_{\mathsf{y}}/\sigma_{\mathsf{y}} = \sqrt{2/\mathfrak{D}}$ is related to the number $\mathfrak{D}$ of degrees of freedom. In turn, $\mathfrak{D}$ is equal to the number of samples of $\overline{\mathsf{y}}_{k+1} - \overline{\mathsf{y}}_k$ in the case of white PM noise (uncorrelated samples), it gets progressively smaller for slower noise phenomena, and it degenerates to 2 in the case of pure drift. A simple example deserves attention. Suppose that we have a record of 3×10^5 samples spaced by $\tau_0 = 100$ ms; thus, $\mathcal{T} = 3 \times 10^4$ s $= 8.33$-h acquisition time. At $\tau = 10^4$ s, we have only two realizations of $\overline{\mathsf{y}}_{k+1} - \overline{\mathsf{y}}_k$ using (34), while, in the same conditions, we have 10^5 realizations if we opt for (38). At a deeper sight, such realizations are highly correlated, but the larger number is still beneficial to the reduction of uncertainty.

A further advantage of overlapping is that it solves the erratic response of AVAR in the presence of cyclic disturbances of period $T \approx \tau/0.45$ (2.2τ), such as the diurnal temperature. This happens because multiple realizations of AVAR are averaged based on measures of $\overline{\mathsf{y}}_2 - \overline{\mathsf{y}}_1$ progressively shifted by $m\tau_0$, as shown in Fig. 9.

The AVAR with no overlap is seldom used, if ever. Generally, the term "AVAR" *refers to the overlapping algorithm* with no need of saying so.

VII. Other Options for the Two-Sample Variance

A. Frequency Counters and Weighted Averages (Region 2.1)

Introducing the classical AVAR, we have defined $\overline{\mathsf{y}}$ as the uniform average of $\mathsf{y}(t)$ over the time τ. Other options make

sense based on the redefinition of the average $\overline{\mathrm{y}}$ as

$$\overline{\mathrm{y}}(\tau) = \int_0^\infty \mathrm{y}(t)\, w(t;\tau)\, dt. \tag{39}$$

The weight function $w(t;\tau)$ takes different forms, among which the following deserve attention.

1) $w_\Pi(t;\tau)$ gives the uniform average, the same used in Section VI (there is still no change). The classical reciprocal counter, also called Π counter after [84], uses w_Π. The symbol Π recalls the rectangular shape of $w_\Pi(t;\tau)$.
2) $w_\Lambda(t;\tau)$ gives the triangular average, which is calculated averaging on a sequence of highly overlapped rectangular averages. The corresponding instrument is the Λ *counter* after [84]. Of course, the Greek letter Λ is chosen because of its triangular shape. The benefit of such a counter is a high rejection of the wideband white PM noise of the trigger at the counter input. Some commercial instruments implement the Λ averaging, often without saying. They can be identified from the "precision" (response to the trigger noise) proportional to $1/\tau\sqrt{\tau}$ instead of $1/\tau$.
3) $w_\Omega(t;\tau)$ relates to a frequency measurement implemented as a linear regression on phase–time data. The corresponding instrument is called Ω counter after [85]. The benefit of the Ω counter is the highest rejection of the white PM noise, by theorem. The Greek letter Ω is the graphically closest to the parabolic shape of the frequency response. Besides, the letter Ω indicates that this is the ultimate counter, to the extent that no other counter performs better rejection of white PM noise.
   The "precision" (response to the trigger noise) is proportional to $1/\tau\sqrt{\tau}$ as in the Λ counter, just with 1.25 dB lower background noise (a factor of $3/4$). Very few commercial instruments implement this algorithm.
4) $w_\Delta(t;\tau)$ is equivalent to the difference between two contiguous measures taken with a Π frequency counter. This option is listed only for completeness because, to the best of our knowledge, it is not implemented in commercial counters.

Like the overlapped AVAR, $\mathrm{y}(t)$ is sampled at a suitable frequency $1/\tau_0$; thus, it holds that $\tau = m\tau_0$. As a consequence, the patterns of the Λ and Ω counters shown in Region 2.1 hold for $m \gg 1$, or equivalently $\tau \gg \tau_0$, so that the continuous approximation holds.

B. Generalized Two-Sample Variances (Region 2.2)

The definition (32) is more general than the classical AVAR. In fact, feeding the weighted averages (39) into (32) results in different types of variance, which can all be described by the same formula

$$\sigma_{\mathrm{y}}^2(\tau) = \mathbb{E}\left\{ \int_0^\infty \Big[\mathrm{y}(t)\, w(t;\tau)\Big]^2\, dt \right\}. \tag{40}$$

Notice that the difference $\overline{\mathrm{y}}_2 - \overline{\mathrm{y}}_1$ is now included in the wavelet-like function $w(t;\tau)$. The latter is similar to a wavelet but for the normalization of finite-power signals (power-type signals) instead of finite-energy signals (energy-type signals).

We have the following options:

$$\begin{aligned}
w_A(t;\tau) &= w_\Pi(t-\tau;\tau) - w_\Pi(t;\tau) \rightarrow \text{AVAR } {}^{A}\sigma_{\mathrm{y}}^2(\tau)\\
w_M(t;\tau) &= w_\Lambda(t-\tau;\tau) - w_\Lambda(t;\tau) \rightarrow \text{MVAR } {}^{M}\sigma_{\mathrm{y}}^2(\tau)\\
w_P(t;\tau) &= w_\Omega(t-\tau;\tau) - w_\Omega(t;\tau) \rightarrow \text{PVAR } {}^{P}\sigma_{\mathrm{y}}^2(\tau)\\
w_H(t;\tau) &= w_\Delta(t-\tau;\tau) - w_\Delta(t;\tau) \rightarrow \text{HVAR } {}^{H}\sigma_{\mathrm{y}}^2(\tau).
\end{aligned}$$

As with the frequency counters, the patterns of $w_M(t;\tau)$ and $w_P(t;\tau)$ shown in Region 2.2 are the limit for $\tau \gg \tau_0$, or equivalently for $m \gg 1$, where the continuous approximations hold. The formulas of Region 2.7 hold under this assumption.

Using commercial frequency counters to get a data stream $\overline{\mathrm{y}}_k$, we recommend attention to the contiguity of averages. Such contiguity is implied in the waveforms of Region 2.2. The case of MVAR is subtle because "contiguous" triangular averages overlap by one side, instead of touching one another by one edge. Something like /\/\/\... instead of /_/_/\... Thus, if the measurement of $\overline{\mathrm{y}}$ takes 2τ, the measurement of $\overline{\mathrm{y}}_2 - \overline{\mathrm{y}}_1$ takes 3τ. Dawkins et al. [86] point out that some Keysight counters provide a data stream of the latter type, where continuous averages touch one another by one edge.

Our notations ${}^{A}\sigma_{\mathrm{y}}^2(\tau)$, ${}^{M}\sigma_{\mathrm{y}}^2(\tau)$, ${}^{P}\sigma_{\mathrm{y}}^2(\tau)$, and ${}^{H}\sigma_{\mathrm{y}}^2(\tau)$ seem more elegant than AVAR, MVAR, PVAR, and HVAR, but it is not used in the literature.

We draw the reader's attention to the fact that (40) may be misleading in the case of HVAR because it hides the fact that HVAR is a second-difference variance. HVAR is rather different from the other variances, chiefly in the fact that it converges for integrated flicker FM and random run FM ($1/f^3$ and $1/f^4$ FM), and it is blind to frequency drift.

C. Time Variance (Region 2.6)

The time variance TVAR is defined as

$$\sigma_{\mathrm{x}}^2(\tau) = \frac{1}{3}\,\tau^2\, {}^{M}\sigma_{\mathrm{y}}(\tau). \tag{41}$$

TDEV, the square root of TVAR, is an estimator of the uncertainty of the time elapsed after the duration τ. TDEV is often used in telecom, e.g., to assess the TIE, defined in [32], and the holdover performance, i.e., the clock error accumulated between synchronizations [33, Part 5.1: Timing characteristics of slave clocks…].

D. Additional Options

The following options are highly specialized instances of the AVAR, beyond the scope of this article. They are mentioned for completeness and left to experts.

1) Dynamic Allan Variance: A time series of N data is sliced into m subseries of $n = N/m$ data. Computing the AVAR for each subseries, we end up with a 3-D plot, which shows the changes in AVAR versus time, most useful for diagnostic purposes. This variance and its properties are found in a series of articles by the same team [87], [88], [89], [90], [91], but it does not seem to have been followed by other authors and research teams.

2) Total Variance: The time series is circularized by joining a copy with time reversed, as often done in the domain of spectral analysis. Circularization cannot increase the number $\mathfrak{D}$ of degrees of freedom inherent in the experimental outcomes, but it makes their exploitation more efficient for the detection of certain phenomena. This concept, as described in [92] and [93], can be applied to all wavelet variances.

3) Thêo1, ThêoH, and ThêoBR: These estimators are based on the idea that two measures of duration $\tau' \ll \tau$ whose centers are spaced by τ provide a precise estimation of the slow processes occurring at τ, under the condition that τ' is long enough to average out the fast processes. These ideas are found in [94], [95], and [96]. The benefit is to extend the maximum τ beyond $\mathcal{T}/2$, where $\mathcal{T}$ is the duration of the data record. For example, applying ThêoH to a long data record of an atomic time scale, it is possible to extend the plot up to $\tau = 0.8\,\mathcal{T}$. Thus, these variances are even more efficient than AVAR at estimating the fluctuations occurring at large τ. This benefit comes at the cost of a higher computing burden and additional interpretation difficulty.

E. Choosing the Most Appropriate Variance

For general use, the variances described are broadly similar to one another, and none is really "the best" or just "bad." Each one has its own "personality," which makes it more suitable in some specific cases and weaker in other cases. Such personality follows from the wavelet-like patterns shown in Region 2.2. A summary of the main options is given in the following.

1) Normalization: In signal processing, normalization for white noise is the most common option. A different choice is made here because the two-sample variances are issued from timekeeping. The normalization is chosen for all the variances to have the same response $\sigma_{\mathsf{y}}^2(\tau) = (1/2)\mathsf{D}_{\mathsf{y}}^2\tau^2$ to the linear frequency drift D_{y} (the "linear drift" row in the lower part of Region 2.7). A consequence is that different variances have different responses to the other noise processes because the spectral response $|H(f\tau)|^2$ differs (Region 1.5). In addition, the corners separating the noise processes are not the same (Regions 1.7 and 2.3–2.5). Comparing the plots requires a small effort of interpretation. For example, the white FM noise $S_\varphi = \mathsf{b}_{-2}/f^2 \rightarrow S_{\mathsf{y}} = \mathsf{h}_0$ shows up as

$$
\begin{aligned}
{}^{A}\sigma_{\mathsf{y}}^2(\tau) &= \frac{1}{2}\,\frac{\mathsf{h}_0}{\tau} = 0.50\,\frac{\mathsf{h}_0}{\tau} \quad \text{(Allan)}\\
{}^{M}\sigma_{\mathsf{y}}^2(\tau) &= \frac{1}{4}\,\frac{\mathsf{h}_0}{\tau} = 0.25\,\frac{\mathsf{h}_0}{\tau} \quad \text{(Modified Allan)}\\
{}^{H}\sigma_{\mathsf{y}}^2(\tau) &= \frac{1}{3}\,\frac{\mathsf{h}_0}{\tau} \simeq 0.33\,\frac{\mathsf{h}_0}{\tau} \quad \text{(Hadamard)}\\
{}^{P}\sigma_{\mathsf{y}}^2(\tau) &= \frac{3}{5}\,\frac{\mathsf{h}_0}{\tau} = 0.60\,\frac{\mathsf{h}_0}{\tau} \quad \text{(Parabolic)}
\end{aligned}
$$

with a maximum difference of a factor of 2.4 (3.8 dB).

2) AVAR: This the best choice when we want to evaluate $\sigma_{\mathsf{y}}^2(\tau)$ up to the largest τ for a given data record of duration $\mathcal{T}$, that is, $\tau = \mathcal{T}/2$. This is the typical case of atomic time scales, where the oscillators are continuously monitored, and we focus on the slow processes.

Because $w_A(t;\tau)$ takes 2τ for one realization of $\sigma_{\mathsf{y}}^2(\tau)$, averaging on $M \gg 1$ realizations is made possible by overlapping the measures with $\tau_0 \ll \tau$ (see Section VI-C).

The uniform weight of $w_A(t;\tau)$ features the highest efficiency in picking up the energy of y. By contrast, AVAR is unsuitable for the measurement of white PM noise because ${}^{A}\sigma_{\mathsf{y}}^2(\tau) \simeq 0.076\,f_H\mathsf{h}_2/\tau^2$ (Region 2.7); thus, the result is highly dependent on the bandwidth f_H. This is not a problem for the slow phenomena that we mentioned. Finally, it is worth mentioning that, in the presence of white FM noise only, AVAR gives the same result as the classical variance.

3) MVAR: This is a good choice in the presence of wideband noise typical of fast processes. MVAR originates from optics, where precise and efficient measurement of white PM noise is a desired feature.

By contrast, MVAR is inferior to AVAR in the efficient use of $\mathcal{T}$ because the support of $w_M(t;\tau)$ is 3τ wide instead of 2τ. This may not be a problem when the physical phenomena that we are interested in occur at small or moderate τ, say hours.

4) HVAR: This variance is useful for the measurement of strong slow phenomena occurring in some circumstances, for example, in the absence of temperature stabilization. In fact, unlike the other variances described here, it converges for integrated flicker FM noise and for integrated random walk FM noise, also called "random run." Such processes are the h_{-3}/f^3 and h_{-4}/f^4 terms of $S_{\mathsf{y}}(f)$ or equivalently the b_{-5}/f^5 and h_{-6}/f^6 terms of $S_\varphi(f)$. By contrast, HVAR is blind to linear drift. This makes HVAR a specialized tool, particularly useful when high drift makes it difficult to estimate the other noise parameters.

Finally, HVAR, such as AVAR, gives ambiguous responses to white PM noise, as it depends on the instrument bandwidth f_H. Like MVAR, $w_H(t;\tau)$ takes a time equals to 3τ.

5) PVAR: The computation of $\sigma_{\mathsf{y}}^2(\tau)$ at $\tau = m\tau_0$ requires a data record of duration $\mathcal{T} = k\tau$, where $k > 2$ for AVAR and PVAR, and $k > 3$ for MVAR and HVAR. The minimum $\mathcal{T}$ depends on the noise process, the confidence level required, and the variance we choose. Running the measurement, we start seeing white PM noise at short $\mathcal{T}$, then flicker PM, white FM, and so on as $\mathcal{T}$ increases. Now, we take a different standpoint, asking which is the minimum $\mathcal{T}$ to detect a "new" noise phenomenon, out of the "previous," faster one. For example, which is the shortest $\mathcal{T}$ to see that flicker FM is above the white FM, with 95 % probability? Among the four variances considered here, PVAR is the best at doing this from white PM to RW FM. Yet, the corner between RW FM and frequency drift is still better detected by AVAR.

F. Some Pieces of Advice

Beginners should restrict their attention to ADEV and MDEV, the square root of AVAR and MVAR. ADEV and MDEV are both available in commercial instruments, and both benefit from the size effect of a wide community.

For historical reasons, ADEV is definitely the manufacturers' preferred option. Reading technical documentation, we recommend attention to possible confusion between ADEV and MDEV under the term "Allan deviation," with a possible

"modified" omitted or implied. The ITU-T Recommendation G.8260 [32] compares ADEV, MDEV, and TDEV from the standpoint of telecommunications. MDEV seems the favorite tool in telecommunications [97] and enables the direct calculation of TDEV using (41). However, MDEV takes 50% longer acquisition time $\mathcal{T}$.

Doing one's own measurements, MDEV is, in most cases, the best compromise. In fact, MDEV improves on ADEV in the detection of fast noise processes (white and flicker PM) and is as suitable as ADEV to detect all the other noise processes. AVAR is still the best option for the measurement of the atomic clocks[10] intended for time scales, where increasing $\mathcal{T}$ is costly or impossible. Finally, the reader should remember that normalization makes MVAR always "optimistic" compared to AVAR, as seen in the related columns in Region 2.7.

Looking at the future, PDEV may replace MDEV because it outperforms it in all parameters at no cost but computing power. Likewise, Thêo may replace ADEV because of the more efficient use of $\mathcal{T}$ at representing longer values of τ.

G. Example of MDEV

We measure the stability of a miniature Cs oscillator. The oscillator is a laboratory prototype based on the coherent population trapping (CPT) principle in a Cs microcell. The actual experiment is described in [98] and the fabrication of the microcell in [99]. For technical reasons specific to the experiment, the frequency is sampled at 150-ms intervals and preprocessed to provide a stream of values of y_k uniformly averaged on contiguous intervals of 1 s. This sets $\tau_0 = 1$ s. The total number of samples is 389 998, for a total duration $\mathcal{T} \approx 4.5$ days.

Fig. 10 shows the MDEV of the experiment described, processed with the SigmaTheta software tool. Starting from uniformly averaged data, a practical minimum of eight to ten samples is needed to approximate the triangular average (Regions 2.1 and 2.2.). Thus, MDEV is plotted for $\tau = 2^n\tau_0$ starting from $\tau = 8\tau_0 = 8$ s. At the scale of this experiment, the H maser used as the reference can be considered ideally stable, and the noise of the instrument is negligible as well.

The SigmaTheta software package provides the following pieces of information (see Sections IX and XI):

1) the bare values of ${}^M\sigma_y(\tau)$, shown as black crosses;
2) the Bayesian estimates of ${}^M\sigma_y(\tau)$, shown as green donuts;
3) the uncertainty bars;
4) the identification of the most relevant noise processes, shown as the colored straight lines of slope $1/\sqrt{\tau}$, constant versus τ, and $\sqrt{\tau}$.

[10]In the International Coordination of Metrology under the guidance of BIPM, an atomic oscillator is a "clock" only after six months of uninterrupted contribution to UTC/TAI. Some commercial oscillators are called clocks by the manufacturer, implying that they are suitable to contribute to UTC/TAI. This is perfectly sound because hundreds of such Cesium clocks actually do this, as well as hydrogen masers. However, the term clock is also used as a pretentious replacement for a precision oscillator even if reliability and long-term stability are insufficient for the oscillator to be even considered as a contributor to UTC/TAI.

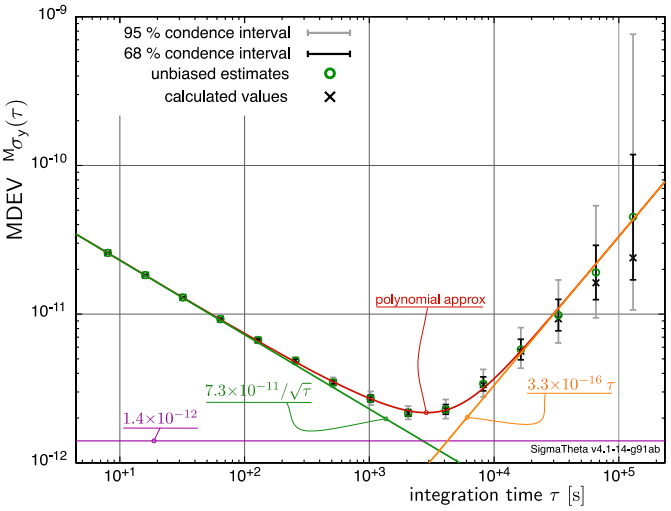

Fig. 10. Example of MDEV obtained with the SigmaTheta software tool. Notice the difference between the calculated values of MDEV (black crosses) and the Bayesian estimates (green donuts), and the asymmetry of the uncertainty bars. Data are courtesy of Moustafa "Mouss" Abdel Hafiz, FEMTO-ST Institute.

It is worth pointing out that the uncertainty bars systematically extend upward more than downward. This is a kind of "signature" of the inverse problem, as opposed to the bare simulation approach.

H. Suggested Readings About Variances

1) General References: The free booklet [36] is probably where most readers should start. Sponsored and distributed by NIST, it provides extensive coverage of most variances (AVAR, MVAR, HVAR, and so on) and the evaluation of the confidence intervals, with numerous examples and plots made with Stable32. Reference [100] is a review article about the AVAR, and [101] suggests that the AVAR can be used as a diagnostic tool. A wealth of information is available in a Special Issue of the IEEE TRANSACTIONS ON ULTRASONICS, FERROELECTRICS, AND FREQUENCY CONTROL celebrating the 50th anniversary of the AVAR [102]. A historical review is available [103], written by two of the most important contributors to the rise of this branch of knowledge. Reference [104] is the original article that introduces the sample variances, later called AVAR, and [105] introduces MVAR.

2) Π and Λ Counters, and the Related Statistics: Reference [84] is the first article that defines Π and Λ counters and the mathematical framework underneath, later extended to the case of nonoverlapping triangular averages [86]. However, the basic ideas were already in [106] and [107] yet without developing the statistical framework. A wealth of practical knowledge about the architecture of high-resolution counters is available in a review article [108].

3) Ω Counter and the Parabolic Variance: The linear regression of phase data is a rather obvious way to estimate a frequency. It was used in the HP5371A HP5372A time interval analyzers[11] in the late 1980s and at Pendulum [109]. The name

[11]Information provided by Magnus Danielsson, NetInsight, Sweden.

"Ω counter" comes from [85], which introduces the related mathematical framework for frequency metrology. Feeding the linear-regression estimates into (32), we get PVAR. This idea came independently of [110] and [111], and the name PVAR was decided together by the two teams. Reference [111] digresses the advanced statistical properties of PVAR, including the Bayesian statistics and the minimum duration of the data record to detect a noise process.

4) Aliasing: Vernotte et al. [112] provide theory and insight on spectral aliasing, and Calosso et al. [26] give an interesting perspective about aliasing in the AVAR and MVAR, covering the effect of spectral bumps and "blue noise," often found in long-range optical frequency-distribution systems. Finally, Bernier [113] provides useful insight into aliasing and cutoff frequency for MVAR.

VIII. Relations Between Phase Noise and Variance

Regions 1.6 and 1.7 show a plot of $S_y(f)$ besides AVAR $\sigma_y^2(\tau)$ for the noise processes from white PM to random-walk FM. The following facts deserve attention.

1) There is a kind of mirror symmetry between the plots of $S_y(f)$ and $\sigma_y^2(\tau)$. The fastest process, white PM, is on the right-hand side of $S_y(f)$ and on the left-hand side of $\sigma_y^2(\tau)$. Vice versa, the FM random walk is on the left-hand side of $S_y(f)$ and on the right-hand side of $\sigma_y^2(\tau)$.
2) The cutoff frequency f_H has a dramatic effect on white PM noise, only a weak effect on the flicker PM noise, and virtually no effect on slower processes.
3) The $1/\tau^2$ region of $\sigma_y^2(\tau)$ is ambiguous, in which it represents both white PM and flicker PM noise.
4) The conversion from $S_\varphi(f)$ to $\sigma_y^2(\tau)$ is always possible, while the opposite suffers from limitations. This is emphasized by the road signs between Regions 1.6 and 1.7.
5) The corner τ where the processes cross one another may occur rather far from the values one expects intuitively.

If we have both phase noise and variance measures, the conversion from $S_\varphi(f)$ to $\sigma_y^2(\tau)$ is a great way to check on consistency. However, this is possible only if $S_\varphi(f)$ extends to sufficiently low f to reveal the slow processes shown by $\sigma_y^2(\tau)$. Regardless, the conversion is an exercise of interpretation that we recommend.

A. Conversion from PM Noise to Allan Variance (Region 2.7)

The reader should first refer to Region 1.11 for the conversion $S_\varphi(f) \to S_y(f)$ and then to Region 2.7 for $S_y(f) \to \sigma_y^2(\tau)$, based on (35) and (36). A simple procedure is shown in Fig. 11 and detailed in the following.

0) Start with a log–log plot of phase noise. Since, in most cases, the phase noise is given as $10\log_{10}[\mathscr{L}(f)]$, you have to convert it into $S_\varphi(f)$ using $S_\varphi(f) = 2\mathscr{L}(f)$, i.e., add 3 dB (Region 1.11)
1) Approximate the true spectrum with the straight lines that represent the polynomial law. Proceed from the right-hand side (white PM noise) to the left, not vice versa. This is best done with a drawing app inserting a straight line of exact slope (0, −1, −2, and so on) and shifting it to fit the plot. The coefficients b_0, b_{-1}, b_{-2}, and so on are the value of the corresponding straight line found at $f = 1$ Hz. Even with little training, visual inspection is efficient at approximating the noise process and ignoring the artifacts. Sliding old-fashioned set drafting squares on a printed spectrum is a good

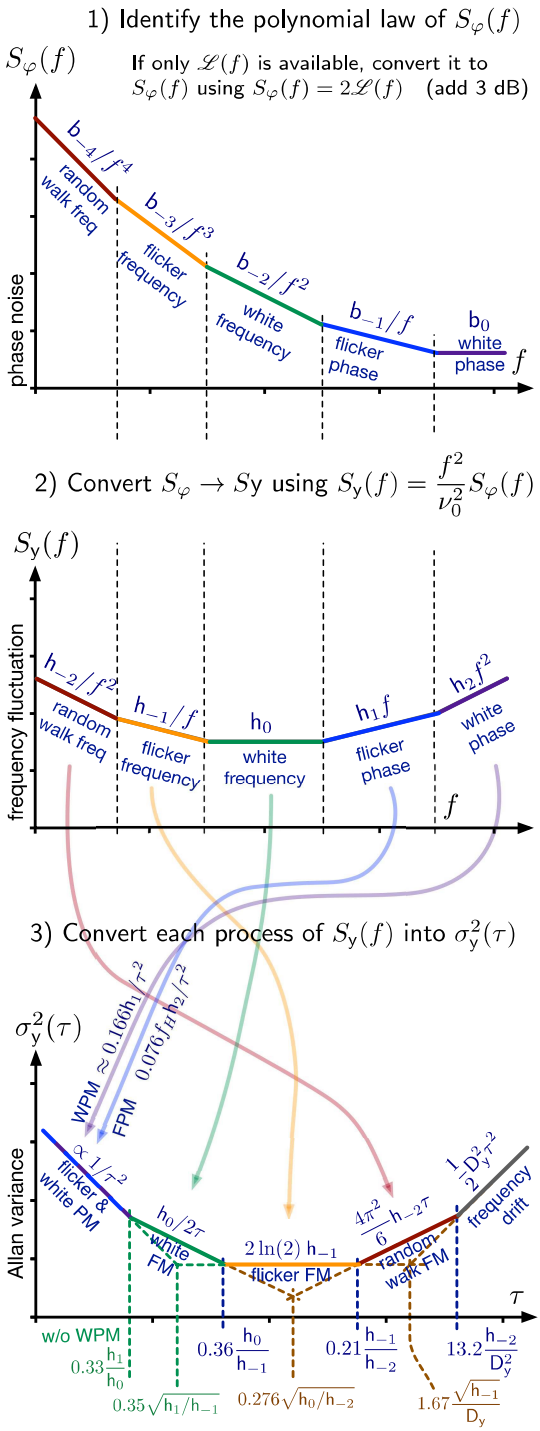

Fig. 11. Conversion from phase noise to AVAR. The colors of the straight line approximation recall the frequency, from reddish (low) to bluish (high). The conversion from phase noise to the other two-sample variances is an obvious extension.

alternative to a drawing app. Worked-out examples are available in [69, Chapter 6].

2) Convert $S_\varphi(f)$ into $S_\mathrm{y}(f)$ using

$$\mathsf{h}_n = \frac{1}{\nu_0^2}\,\mathsf{b}_{n-2}$$

which is equivalent to $S_\mathrm{y}(f) = (f^2/\nu_0^2)S_\varphi(f)$ (Region 1.11). Newcomers may be surprised by the small value of the h_n coefficients due to ν_0^2 in the denominator.

3) Sketch the AVAR using the pattern of Region 1.7 (right) and the formulas found in the "AVAR" column of the table in Region 2.7. Each process (white PM, flicker PM, and so on) requires its own formula, found in Region 2.7.

It goes without saying that the method described also applies to the other two-sample variances discussed in Section VII, just picking up the appropriate column in Region 2.7. The extension to noninteger slopes is found in [114] and [115].

B. Variance to Spectrum Conversion (Regions 1.6, 1.7, and 2.7)

The variance to spectrum conversion is not possible in the general case, but we can get useful information using the formulas given in Region 2.7, assuming that the spectrum is smooth and follows the polynomial law (see [116]).

The first limitation is that, in the case of AVAR and HVAR, both white and flicker PM show up as $\sigma_\mathrm{y}^2 \propto 1/\tau^2$; thus, they cannot be divided. In addition, AVAR and HVAR are almost unusable in the white PM region because $\sigma_\mathrm{y}^2 \propto f_H$. The cutoff f_H does not appear explicitly in MVAR, PVAR, and TVAR. However, the effect of f_H is hidden in the sampling process itself and then in the data. Since data are sampled at τ_0 interval, the Nyquist frequency is $(1/2\tau_0)$. Thus, if $f_H > (1/2\tau_0)$, there will be spectral aliasing, and the white noise level will be overestimated.

The second limitation is related to the resolution. In fact, the inherent resolution of $\sigma_\mathrm{y}^2(\tau)$ is one octave in τ, as it follows from the bandwidth of the main lobe in Region 1.5. Thus, $\sigma_\mathrm{y}(\tau)$ is usually plotted for τ in geometric series, such as $1, 2, 4, 8, \ldots$ Conversely, the frequency resolution of $S_\varphi(f)$ depends on the acquisition time, and plots usually represent $S_\varphi(f)$ with a resolution of 50–100 points/decade, that is, 15–30 points per octave.

C. Cutoff Frequency f_H and the Sampling Interval τ_0

The cutoff frequency f_H extensively used in the spectrum-to-variance conversions is often a source of confusion. Notice that an antialiasing filter is necessary; otherwise, the variance does not converge in the presence of white and flicker PM noise [113].

The phase time $\mathsf{x}(t)$ has finite bandwidth, which results from the measurement process and from the architecture of the instrument. This is generally described as the noise equivalent bandwidth, denoted with f_H.

The variance is evaluated after sampling $\mathsf{x}(t)$ at an appropriate frequency $1/\tau_0$, which requires that $f_H < 1/(2\tau_0)$. If this condition is not met, aliasing takes place. The white PM noise is folded $2f_H\tau_0$ times to the first Nyquist zone, and the observed white noise level is $2f_H\tau_0\mathsf{k}_0$ instead of k_0. Flicker noise too is subject to aliasing even though the impact on the results may be smaller. References [112], [113], and [117] detail the specific problems related to our domain.

(a)

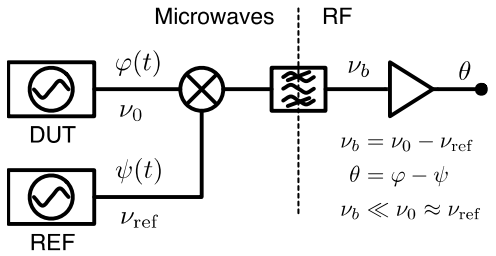

(b)

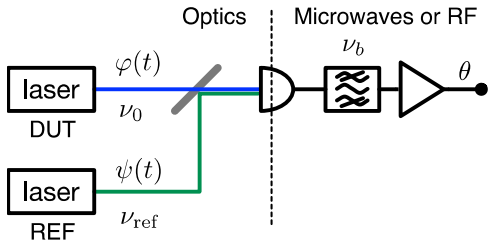

Fig. 12. Beat method, applied to (a) electrical signals and (b) optical signals. See Section X-F for the general description.

D. Example (f_H)

Let us take an example from optics, where we beat two 1550-nm lasers in a photodiode with the scheme of Fig. 12, getting an RF tone at $\nu_b = 60$ MHz. For technical reasons, we choose to filter a such signal with ± 0.5-MHz bandpass centered at ν_b. The filter halfwidth is 2.6×10^{-9} of the optical carrier; thus, the laser instability must be $< 5 \times 10^{-10}$ for ν_b to be decently centered in the filter passband. The noise bandwidth is equal to the filter half-width, i.e., $f_H = 500$ kHz.

1) First Option: We measure the beat note with a counter sampling at $\tau_0 = 1$-ms interval (for example, the old good K&K counters, or the more recent version made by Lange Electronic). In the conditions described, aliasing increases the white PM noise level by a factor $2f_H\tau_0 = 1000$.

2) Second Option: We measure the same beat note with an instrument based on direct digitization (see Section X-I and Fig. 16) or a TDDS (see Section X-J and Fig. 17). Inside the instrument, the ADCs are preceded by antialiasing filters, and a low pass reduces the bandwidth of $\mathsf{x}(t)$ after detection. Thus, the condition $f_H < 1/(2\tau_0)$ is met, and there is no aliasing. In some instruments, f_H can be set by the user.

IX. Confidence Intervals

Assessing the confidence intervals is all about understanding the interplay between the "true value" $\sigma_\mathrm{y}^2(\tau)$ and its estimate $\widehat{\sigma_\mathrm{y}^2}(\tau)$. The latter is evaluated with (33), or equivalently (34)

for AVAR, and with similar formulas for other variances. Since everything in this section applies separately to each value of τ, we let τ implied in $\widehat{\sigma_y^2}$ and σ_y^2.

A. Direct Problem

The *direct problem* is similar to computer simulations, where we add noise to a deterministic phenomenon. Without noise, the right-hand side of (33) gives σ_y^2. In this case, the term "true value" for σ_y^2 is fully legitimate because uncertainty is zero. Introducing the noise, the same equation gives the estimate $\widehat{\sigma_y^2}$, together with the associated PDF $p(\widehat{\sigma_y^2}|\sigma_y^2)$. The notation $p(|)$ stands for conditional probability, and the vertical bar "|" reads "given" or "knowing," as in $p(\widehat{\sigma_y^2})$ *given* $p(\sigma_y^2)$.

B. Inverse Problem

The simulation approach does not help to grab the hidden reality σ_y^2 from the outcome $\widehat{\sigma_y^2}$. The answer is the *inverse problem*, which consists of inferring the true σ_y^2 and the associated pdf $p(\sigma_y^2|\widehat{\sigma_y^2})$.

Otherwise stated, the direct problem targets the "cause $\rightarrow$ effect?" relationship, where "?" emphasizes the unknown. In contrast, the inverse problem targets "cause? $\leftarrow$ effect." Interestingly, the "true value" σ_y^2 in the direct problem and the "inferred true value" σ_y^2 in the inverse problem are different things, but they are generally denoted with the same symbol in the literature.

The evaluation of the confidence intervals in the inverse problem may be too complex for some readers. These readers may keep only the two results summarized in the following, skip Section IX-C, and go straight to the choice of the appropriate software package.

1) When a small number $M-1$ of realizations $\overline{y}_{k+1}-\overline{y}_k$ is available, the "error bars" are large and asymmetric. The half bar directed upward is wider than the downward half bar. This is seen in the example Fig. 10 for $\tau \geq 4000$ s and qualitatively illustrated in Region 1.7.
2) With large M, there is no practical difference between the direct problem and the inverse problem. The confidence intervals are small and symmetric, approaching $1/\sqrt{M-1}$, as seen in Fig. 10 for $\tau \leq 2000$ s.

C. Application of the Bayesian Statistics to Allan Variances

Everything starts with the Bayes theorem, which states that

$$p(\theta|\xi) = \frac{\pi(\theta)\, p(\xi|\theta)}{\pi(\xi)}, \quad \pi(\xi \neq 0) \tag{42}$$

where $p(\,)$ denotes the *posterior* pdf and $\pi(\,)$ denotes the *prior* pdf.[12] It is worth mentioning that some authors use the same symbol for both prior probability and posterior probability. In the measurement of frequency stability, the experimental value $\widehat{\sigma_y^2}$ is identified with ξ and the unknown "true" σ_y^2 with θ. Thus, we have to infer a confidence interval on σ_y^2.

[12]More specifically, $\pi(\theta)$ is the a priori knowledge before any measurement, and $\pi(\xi)$ is an unknown function that we do not care about because it is independent of θ. The prior $\pi(\xi)$ can be determined using the property that $\int p(\theta|\xi)\, d\theta = 1$.

The central limit theorem suggests that $\overline{y}_k$ are Gaussian, as they result from a lot of data. Thus, we assume that their differences are Gaussian centered. Thus, (33) indicates that $\widehat{\sigma_y^2}$ is described by a χ^2 distribution with $\mathfrak{D}$ degrees of freedom.[13] Such distribution is denoted with $\chi^2_{\mathfrak{D}}$. Of course, with M values of $\overline{y}_k$, it holds that $\mathfrak{D} \leq M-1$, where the equality indicates that all the terms $(\overline{y}_{k+1}-\overline{y}_k)$ of the sum are statistically independent. Greenhall and Riley [118] provide a very useful method to evaluate $\mathfrak{D}$.

Since the random variable ξ is $\chi^2_{\mathfrak{D}}$ distributed, the cumulative density function (CDF) of ξ knowing θ, denoted with $F(\xi|\theta)$, is also known in analytic form. The inverse CDF, available in the major mathematical libraries, enables the user to compute the confidence interval.

The above is for the direct problem. The inverse problem can be solved thanks to the relevant property that a $\chi^2_{\mathfrak{D}}$ distribution is defined by one and only parameter, $\mathfrak{D}$. It has been proven that such distributions are "fiducial" distributions [119], which means that the equality

$$F(\theta|\xi) = 1 - F(\xi|\theta) \quad \text{(fiducial)} \tag{43}$$

holds in both frequentist inference and Bayesian inference, provided that a $1/\theta$ prior (prior of total ignorance) is chosen [120], [121]. This implies that the confidence intervals given by frequentist or Bayesian methods are the same and are easy to compute.

Note that, because the $\chi^2_{\mathfrak{D}}(x)$ distribution is strongly asymmetric with a steep rise at small x and slow decay at high x, it results from (43) that $p(\theta|\xi)$ must have a steep side at some high x and a slow decay toward $x = 0$. The consequence is that the error bars on a log–log plot of AVAR are extended upward more than downward. This behavior is more remarkable at small $\mathfrak{D}$.

X. Measurement Techniques

A. Saturated-Mixer Phase Noise Analyzer

Fig. 13 shows the block diagram of a traditional phase noise analyzer. The instrument consists of two equal channels where a DBM compares the phase φ of the oscillator under test (DUT) to the phase θ and ψ of the references. The error signal sent to the LNA is equal to $k_d(\varphi-\theta)$, or to $k_d(\varphi-\psi)$, where k_d is the mixer's phase-to-voltage gain. The typical value of k_d is between 0.1 and 1 V/rad, depending on the signal level, technology, and frequency.

This operating mode requires that $\varphi - \theta = \pi/2$ within ± 0.1 rad, and likewise $\varphi - \psi$, which is ensured by the PLL. The error signal, corrected for the equation of the PLL, is proportional to the DUT phase fluctuation versus the references. A problem with mixers is that the power ranges are rather narrow, from a minimum of $\approx$5 dBm to a maximum of 17–20 dBm, depending on technology. The power range can be extended by adding an amplifier at each RF input.

Assuming that the two channels are independent (separate mixers and separate synthesizers, and no crosstalk), the average cross-spectrum is proportional to the DUT phase noise

[13]In the literature about statistics, the degrees of freedom are more often denoted with ν, but, in our notation, ν is used for the carrier frequency.

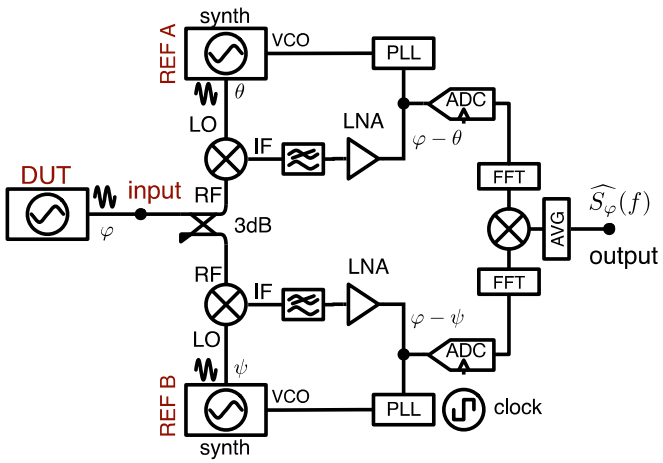

Fig. 13. Traditional phase noise analyzer. The mixers, saturated at both input, work as phase-to-voltage converters. Two equal channels are used to reject the background noise after averaging on multiple cross-spectra.

because the background noise of the individual channel is rejected by the averaging process. It is worth mentioning that averaging is essential to reject the noise of the synthesizers. The synthesizers are necessary for the instrument to be suitable to a wide range of frequencies, but they are generally noisier than oscillators.

Instruments of this type are commercially available from Anapico, Berkeley Nucleonics Corporation, Holzworth, Keysight Technologies, NoiseXT/Arcale, and Quantic Wenzel (a brand of Wenzel Associates). The frequency synthesizers may be inside the instrument or external modules. The input power splitter may differ from the 3-dB directional coupler shown in Fig. 13. We have seen Y resistive networks (Keysight) and V resistive networks (Holzworth).

B. Mathematical Tools for Spectral Analysis

The estimation of power spectra based on the FFT is common to all phase noise analyzers, analog or digital. The FFT is a fast algorithm for the evaluation of the discrete Fourier transform (DFT) and gives the same result. The algorithm used in most software packages and libraries is the "Fastest Fourier Transform of the West" (FFTW3) [122]. With a bit of humor, we note that the FFTW comes from the Eastern USA. The book [123] is Enrico's favorite reference about the FFT, while [124] is about the early use of the FFT in phase noise measurements.

The most commonly used algorithm for the estimation of the PSD, introduced in [43] and known as the Welch algorithm, stands on the FFT. In several software packages and libraries, the PSD function takes a name that recalls Welch. This is the case with MATLAB, Octave, and Python (SciPy). The classical article [125] discusses extensively the window (taper) functions used in spectral analysis.

Bennett [126] introduces the spectral analysis of the quantized signal, and Widrow et al. [127] and Widrow and Kollár [128] provide an extensive treatise of quantization noise.

C. How the Background-Noise Rejection Works

The noise rejection is based on the *cross-spectrum method*, as described in the following. Referring to Fig. 13, we define

$$x = \varphi - \theta \leftrightarrow X = \Phi - \Theta \tag{44}$$

$$y = \varphi - \psi \leftrightarrow Y = \Phi - \Psi \tag{45}$$

where "↔" stands for Fourier transform inverse–transform pair, time and frequency are implied, and x and y should not be mistaken for x and y. For us, φ is the *signal*, and θ and ψ are the *noise* to be rejected. The quantities θ and ψ account for all the noise sources, i.e., references, synthesizers, mixers, and low-noise amplifiers.

The cross-spectrum averaged on m acquisitions is

$$S_{yx} = \frac{2}{T}\Big\langle YX^* \Big\rangle_m \tag{46}$$

$$= \frac{2}{T}\Big\langle (\Phi - \Psi)(\Phi - \Theta)^* \Big\rangle_m \tag{47}$$

$$= \frac{2}{T}\Big\langle \Phi\Phi^* - \Phi\Theta^* - \Psi\Phi^* + \Psi\Theta^* \Big\rangle_m . \tag{48}$$

It is seen in Fig. 13 that φ, θ, and ψ are statistically independent (separate and independent hardware), and likewise Φ, Θ, and Ψ. Thus, the mathematical expectation of $\Phi\Theta^*$, $\Psi\Phi^*$, and $\Psi\Theta^*$ is zero, and

$$S_{yx} \to \frac{2}{T}\Big\langle \Phi\Phi^* \Big\rangle_m \equiv S_\varphi \quad \text{for large } m \tag{49}$$

which is the same as (2). Averaging on m acquisition, the background noise of the instrument is rejected by a factor of $\sim 1/\sqrt{m}$, i.e., 5 dB per factor-of-10 in m. The actual rejection depends on the estimator.

Most phase noise analyzers use the estimator

$$\widehat{S_{yx}} = \frac{2}{T}\Big|\big\langle YX^* \big\rangle_m\Big|. \tag{50}$$

This estimator is biased and suboptimal [129]. The residual noise terms $\langle \Phi\Theta^* \rangle_m$, $\langle \Psi\Phi^* \rangle_m$, and $\langle \Psi\Theta^* \rangle_m$ are proportional to $1/\sqrt{m}$.

Reference [129] shows that the real-part estimator

$$\widehat{S_{yx}} = \frac{2}{T}\Re\Big\{\big\langle YX^* \big\rangle_m\Big\} \tag{51}$$

is the best option. First, it follows from (48) that $\Phi\Phi^* \in \mathbb{R}$, the other terms are complex, and the background noise is equally split between $\Re\{YX^*\}$ and $\Im\{YX^*\}$. Thus, (51) is obviously advantageous versus (50) because it keeps the entire signal and discards the unnecessary noise in $\Im\{YX^*\}$. The residual single channel is proportional to $1/\sqrt{2m}$, improving by a factor of 2. Otherwise stated, the measurement time for a given noise rejection is shorter by a factor of 4. A further advantage of (51) is that it is unbiased.

Sadly, (50) is often the one and only option, and when (51) is available, it is not the default.

The cross-spectrum method derives from radio astronomy [130] and from the early attempts to assess the frequency fluctuations of hydrogen masers [131]. Reference [132, Sec. IV] is arguably the first application to the measurement of phase noise. It goes without saying that the cross-spectrum method is of broader usefulness than just phase

noise. See [129] for a tutorial, [133] for the use of Bayesian statistics, needed to assess the statistical uncertainty when m is small, and [134] for the extension to multiple instruments simultaneously measuring the same physical quantity.

D. Limitations of the Cross-Spectrum Method

The noise limit $\sim 1/\sqrt{m}$ is often shown on the analyzers' screen under fancy names, such as "correlation factor," "number of correlations," or "Xcorr," together with $\mathscr{L}(f)$. However, such noise rejection is useful only if the correlated disturbances are sufficiently small. In fact, a disturbing term $\delta \leftrightarrow \Delta$ introduced in (44) and (45) changes the instrument readout from S_φ to $S_\varphi + \varsigma S_\delta$, where $\varsigma = \pm 1$ is a sign which depends on whether δ has the same sign or opposite sign in the two channels. This is a systematic error of unknown sign ς. Consequently, the common belief that the background noise of the instrument always results in the *over-estimation* of the DUT noise is *untrue*.

The obvious benefit of large m is high sensitivity,[14] i.e., high capability to measure small S_φ. The inevitable drawback is that a proportionally small S_δ can spoil the measurement. A gross error occurs when $S_\delta \approx S_\varphi$, but $S_\varphi + \varsigma S_\delta > 0$ still holds. Worse, if $S_\varphi + \varsigma S_\delta < 0$, the result is blatant nonsense.

Among the causes for the disturbances δ, we mention the thermal energy in the input power splitter, the effect of AM on the mixers, and the RF crosstalk inside the instrument.

Most power splitters are either four-port directional couplers terminated at one port or Y resistive network. Such splitters, inherently, add an amount of anticorrelated noise originating from the thermal energy associated with the internal dissipation. Thus, $S_\varphi < kT/P$ (thermal energy divided by the carrier power) at the splitter output is a physically legitimate outcome, but it goes with a measurement error. The experimental evidence and the full theoretical proof are found in [135]. The splitter's thermal noise may result in the collapse of the cross-spectrum in the measurement of oscillators [136]. Different options for the power splitter are discussed in [137]. Another perspective on the power splitters is proposed in [138], with original results. In [139], we propose a method for the measurement of the bias error due to both power splitter and internal crosstalk.

The correction for the splitter's thermal energy is trivial if the instrument measures the carrier power. NoiseXT/Arcale informed us that this correction now has been implemented in the DNA.[15]

E. Metrologist's Perspective of Phase Noise Uncertainty

The essential concepts that we need to understand uncertainty are defined by the International Vocabulary of Metrology [44], generally known as VIM. The VIM is published by the BIPM, the international organization in charge to ensure the worldwide unification of measurements.[16]

The statistical noise limit $\sim 1/\sqrt{m}$ falls into *Type A evaluation of measurement uncertainty* [44, Entry 2.28].

Other components of uncertainty (crosstalk, AM noise, thermal energy, and so on) cannot be calculated from the measurement outcomes. They fall into *Type B evaluation of measurement uncertainty* [44, Entry 2.29], which relies on the analysis of the system.

Literature of the past ten years [136], [137], [138], [139] and three workshops [140], [141], [142] point to substantial errors and discrepancies in the measurement of commercial oscillators exhibiting very-low phase noise. The concept of *null measurement uncertainty* applies [44, Entry 4.29], which is the uncertainty in the special case of signal approaching zero. Broadly speaking, this is the extrapolation of uncertainty to the measurement of a noise-free oscillator and tells us the minimum detectable amount of phase noise.

We encourage the reader to get the awareness of the *definitional uncertainty* [44, Entry 2.27], the component of uncertainty resulting from the lack of knowledge of details in the definition of the DUT, and of the *influence quantity*, a quantity affecting the relation between the indication and the measurement result. However, the direct application of these concepts to phase noise is not straightforward.

Reading the instruments' specs, we often find the *uncertainty* (or improperly, *error*) and the *sensitivity* parameters. For example, we may read that the uncertainty is ± 3 dB for $f < 1$ kHz and ± 2 dB for $f \geq 1$ kHz. Similar information may be more detailed, in the form of a table. Forgiving the "$\pm$" symbol, we interpret this as the type B uncertainty. The sensitivity takes the form of a table or a spectrum, indicating $\mathscr{L}(f)$ for relevant values of ν_0 and f, for example, -160 dBc/Hz at 1-GHz carrier and 100-kHz "offset" (actually, modulation frequency). We are inclined to interpret this parameter as the "null measurement uncertainty." Note that the term "sensitivity" found in this context is quite different from the VIM definition [44, Entry 4.12]. That said, we observed that users tend to trust a measured noise level below the "sensitivity" of the instrument, provided that it is still higher than the "$\sim 1/\sqrt{m}$" statistical limit seen on the screen. This practice is encouraged by ads showing a "typical" background noise, clearly lower than specs.

F. Beat Method

The method, as shown in Fig. 12, makes use of the leverage effect, which results from beating ν_0 down to $\nu_b = \nu_0 - \nu_{\text{ref}}$. Mathematically,

$$2\cos\big[2\pi\nu_0 t + \varphi(t)\big]\cos\big[2\pi\nu_{\text{ref}} t + \psi(t)\big] = \cos\big[2\pi\nu_b t + \varphi(t) - \psi(t)\big] \quad (52)$$

after deleting the $\nu_b + \nu_{\text{ref}}$ term with the LPF. The beat note preserves the frequency fluctuations $\nu_0 - \nu_{\text{ref}}$ and the phase

[14] Strictly speaking, this is an improper use of the term sensitivity.

[15] Informal discussion with L. Adrien at the IEEE EFTF IFCS Joint Meeting, Paris, April 2022.

[16] The international coordination of metrology is a complex topic even for specialists because it has both scientific and political implications under the *Metre Convention*. The reader interested can find all information on the BIPM site https://bipm.org.

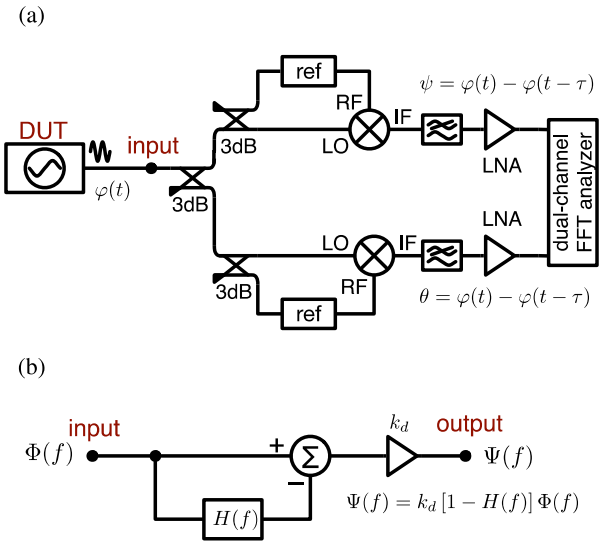

Fig. 14. Discriminator method for the measurement of phase noise in oscillators. (a) Block diagram of the two-channel instrument. (b) Phase space representation of one channel.

fluctuation $\varphi - \psi$, and stretches the time fluctuation by a factor $\kappa \simeq \nu_0/\nu_b$.

This method is often used to beat microwave signals down to the RF region, where we can use digital instruments. In optics, this is the preferred option to measure the fluctuations of a stabilized laser with electrical instruments.

G. Discriminator Method

The use of a reference discriminator, either a resonator or a delay line, is another way to measure the phase noise of an oscillator by comparing the oscillator output to a delayed version of the same signal.

Fig. 14 shows the principle and the equivalent scheme in phase space. The latter follows the same approach used with the oscillator (see Section V). Using the upper case for the Fourier transform of the lowercase function of time, the output is $\Psi(f) = k_d\,[1 - H(f)]\,\Phi(f)$, where the discriminator's phase response is $H(f) = (1/\tau)/(jf + 1/\tau)$ for a resonator characterized by the relaxation time τ and $H(f) = e^{-j2\pi\tau f}$ for a delay line characterized by the delay τ. Accordingly, $\Phi(f)$ is evaluated from the analyzer readout as

$$\Phi(f) = \frac{1}{k_d\,[1 - H(f)]}\,\Psi(f). \tag{53}$$

Finally, using the two-channel configuration, the oscillator phase noise is evaluated as

$$S_\varphi(f) = \frac{1}{k_d^2\,|1 - H(f)|^2}\,S_{\psi\theta}(f). \tag{54}$$

This idea is for sure old, albeit we could not track the origin. The scheme (see Fig. 14) is suitable for implementations based on most or all the phase noise analyzers described in this Section. The standard 1.55-μm optical fiber of appropriate length carrying a modulated beam is a good idea to delay a microwave signal. We implemented a single-channel version [143], and we further developed this concept by adding the cross-spectrum method [144], [145]. To the best of our knowledge, the OEwaves OE8000 is the one and only commercial phase noise analyzer based on photonic delay.

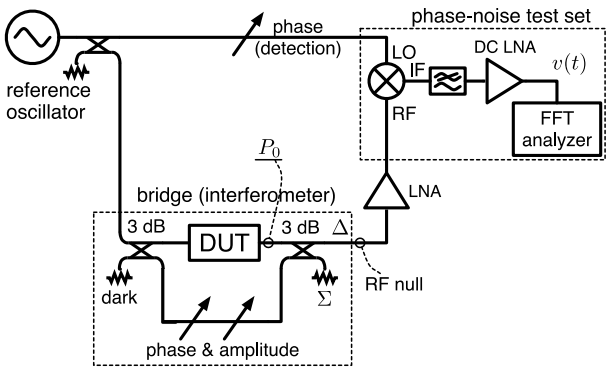

Fig. 15. Scheme of the bridge (interferometric) phase noise measurement.

H. Bridge (Interferometric) Method

The bridge (see Fig. 15) is an advanced method exhibiting extremely low background noise. The bridge must be at equilibrium, achieved by setting the bridge's phase and amplitude. At the equilibrium, the carrier is suppressed at the Δ port, and all the DUT power P_0 goes to the Σ port. The DUT noise sidebands, equally split between Σ and Δ, are amplified and synchronously detected by the mixer. The variable phase at the mixer LO port sets the detection of PM, AM, or a combination of them. Unlike the previous schemes, the mixer works in a linear regime, where the RF input is a small signal, and of course, the LO input is saturated.

Low background flicker b_{-1} is the main benefit. In fact, the passive components used in the bridge (directional couplers, line stretchers, attenuators, and so on) feature extremely low flicker compared to semiconductors and active devices, and the amplifier PM flicker is reduced proportionally to the carrier rejection because flicker comes from upconversion of the near-dc noise.

The bridge exhibits low white noise as well. Neglecting trivial losses, the background noise is $\mathsf{b}_0 = 2FkT_0/P_0$, where F is the amplifier noise factor, kT_0 is the thermal energy at room temperature, and P_0 is the power at the DUT output. The factor "2" means that only half the DUT noise goes to the amplifier, and the other half goes to the Σ port, but it can be reduced using an asymmetrical power combiner.

Finally, low sensitivity to ac magnetic fields from the power grid is easily achieved because RF amplification (20–40 dB) takes place before downconverting to dc.

No commercial solutions exist, but the scheme of Fig. 15 can be built on top of a commercial phase noise analyzer, including the modern dual-channel instruments.

The main ideas derive from Sann [146], but the RF amplification was introduced later by Labaar [147]. The method knew a sudden popularity in the late 1990s in Australia [148], [149] and France [150]. We used it for the measurement of frequency stability in piezoelectric quartz resonators [151] and of the phase noise of DACs and DDSs [152]. We demonstrated

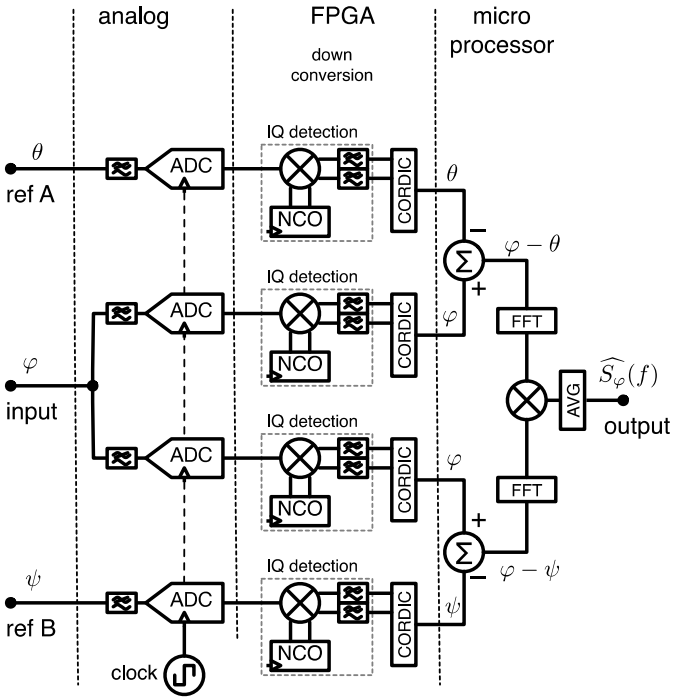

Fig. 16. Digital phase noise and AVAR analyzer.

ultimate sensitivity by adding the correlation for lowest white noise [135] and a multistage bridge for lowest flicker [153].

I. Digital Phase Noise and Allan-Variance Analyzer

Fig. 16 shows the block diagram of a digital phase noise analyzer. The architecture differs from Fig. 13 in the use of SDR techniques. The input signal is digitized and downconverted to an I-Q stream at zero or near-zero frequency by multiplying the input data with sinusoidal signals from an NCO. The CORDIC algorithm [154], [155], [156] is probably the best option to convert the IQ stream into polar coordinates, phase, and amplitude. The reference input signals A and B cannot be used as the sampling signal because the ADCs do not work well at arbitrary clock frequency. Consequently, the measurement of $\varphi - \theta$ requires two ADCs clocked by a common-mode oscillator, whose fluctuation is rejected. In turn, four ADCs are needed to reject the instrument noise.

The digital architecture has the following interesting features.

1) It is great at measuring large values of $\varphi(t)$, not limited to $\pm\pi$, and even millions of cycles are not a problem.
2) The use of a separate NCO at each input removes the requirement that the inputs are at the same frequency because all phases can be referred to the same ν_0 after a trivial numerical conversion.
3) Full control on f_H in the measurement of AVARs.

These properties make digital architecture great at measuring both phase noise and variances, and open new perspectives.

Reference [157] pioneered the measurement of phase noise and the AVAR with direct digitization of the RF signal. Later, Mochizuki et al. [158] provide a more detailed treatise, and Sherman and Jördens [159] focus on SDR techniques. The experimental methods for the noise characterization of ADCs for AM/PM noise measurements are shown in [160], and Calosso et al. [152] introduce a new method for the measurement of AM and PM noises in DACs and DDSs based on the amplification of the (random) modulation index, with optional AM/PM and PM/AM conversion.

Digital instruments are commercially available from Arcale (NoiseXT) and Microchip (Jackson Labs Technologies and Microsemi). The maximum frequency is limited to 30–400 MHz, depending on the instrument. Such instruments may have the reference oscillator(s) inside or rely entirely on external oscillators. If the inputs A and B are accessed through a single connector instead of being available separately, the instrument cannot reject the noise of such oscillator. A hobby project is available from Andrew Holme's home page,[17] based on commercial boards, which even includes the source code and the binaries.

The Rohde & Schwarz analyzers FSWP and FSPN are based on similar concepts, but they implement microwave-to-IF downconversion to extend the input range to 8, 26, or 50 GHz, depending on the model. They include two OCXOs and synthesizers as references and work with external references as well. The internal architecture of the FSWP is explained in [161].

J. Tracking DDS

The TDDS [162] is a PLL where the VCO is replaced with a DDS. The lock differs from a regular PLL in that the FPGA acts on the numerical phase of the DDS, instead of on the frequency of a VCO. Thus, an additional integrator is needed in the control loop. Interestingly, the TDDS locks in a wider frequency range than most PLLs, being limited by the frequency range of the mixer and the DDS. The background noise is limited by the DDS (see [68] for the DDS noise), noisier than double-balanced mixers.

Fig. 17 shows a multichannel TDDS. At the start of the operation, the FPGA sets the numerical frequency of all the DDSs to the frequency of the respective input. Then, the FPGA phase-locks each DDS to the input. At the same time, the FPGA provides the estimation of all the phase–time differences $\widehat{\mathrm{x}_i - \mathrm{x}_j}$ by combining the numerical phases and dc errors.

A multichannel TDDS was successfully used for the direct measurement of the frequency stability at the 100-MHz output of cryogenic sapphire oscillators [163]. The Italian institute of metrology INRiM has developed the *Time Processor*, a multichannel TDDS for the continuous monitoring and stability measurement of high-end oscillators and atomic frequency standards. Albeit some prototypes have already been transferred under contract to qualified users (European Space Agency and the FEMTO-ST Institute), up until now, there is only a conference publication [164]. The Principal Investigator informed us[18] that a startup company is being created, under the provisional name Kairos TeX. The PicoPak,[19] a single channel TDDS, was produced by

[17] http://www.aholme.co.uk/PhaseNoise/Main.htm

[18] Claudio E. Calosso, Private Communication, October 2022.

[19] Hamilton Technical Services, http://www.wriley.com/7510A.pdf

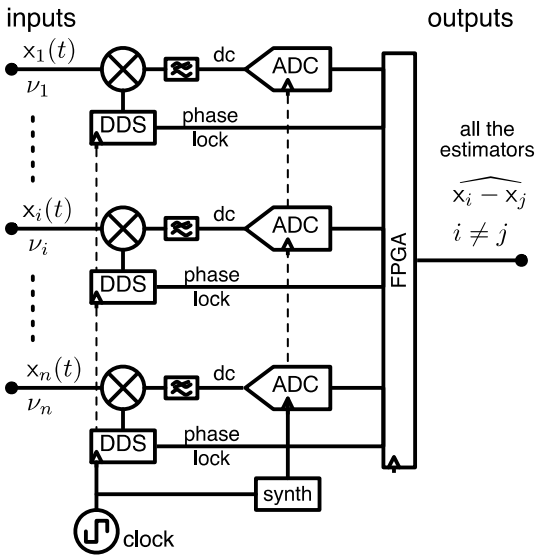

Fig. 17. TDDS.

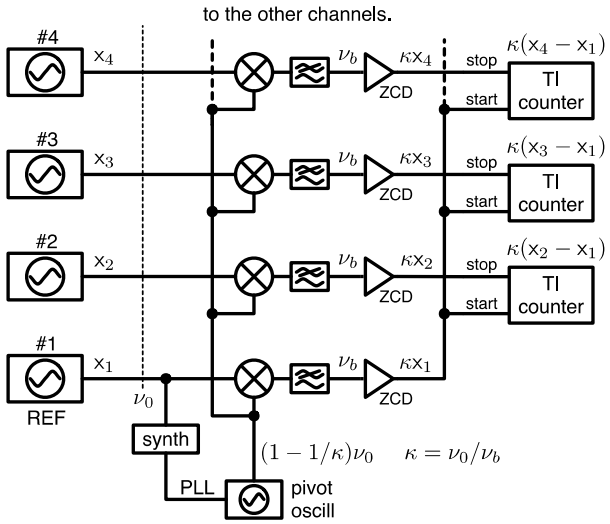

Fig. 18. Multichannel analog AVAR analyzer.

Precision Time and Frequency, LLC, but it is unsure whether it is still available.

K. Multichannel Analog Allan Variance Analyzer

This type of analyzer is a highly specialized instrument intended to simultaneously monitor multiple oscillators in a time scale (Cs beams, fountains, and H masers), usually comparing the 10- or 100-MHz outputs. Fig. 18 shows an example, easily extended to more than four inputs.

The machine exploits the leverage effect seen in Section X-F, but, in this case, the pivot frequency ν_p is just below ν_0 so that ν_b is in the subaudio or low-audio range. Accordingly, $x(t) = \varphi(t)/2\pi\nu$ is stretched by a factor $\kappa = \nu_0/\nu_b$. With $\kappa = 10^6$, a 100-MHz signal can be measured with 10-fs resolution using a simple counter that has a resolution of 10 ns. However, the actual resolution is limited by other factors.

The dual-mixer method, later extended to the multichannel system shown in Fig. 18, is introduced in [165]. Brida [166] discuss the design of the dual-mixer system, providing interesting experimental data. The best choice for the ZCD is a multistage amplifier where the first stage has narrow bandwidth for low noise, and the bandwidth increases progressively toward the output to allow a high slew rate in the saturated signal [167], [168].

Major metrology labs have been using instruments based on this principle for many years. However, those instruments are house-built prototypes. Commercial units are available from Quartzlock (U.K.), Time Tech (Germany), and VREMYA-CH (Russia).

XI. Software Tools

A quick Internet search reveals that there is a rather broad choice of software packages for the AVAR but almost none for phase noise. This relates to the fact that the variances are used to investigate slow phenomena; thus, they require rather low sampling rate and small digital storage space. A value of 1 kS/s is generally sufficient for all practical cases, but 1 or 10 S/s is most often used in atomic time scales. Conversely, it is quite common to plot the phase noise up to 1-MHz Fourier frequency. This requires a practical minimum of 2.5 MS/s, allowing a mere 250 kHz for the antialiasing filter to roll off. In turn, a transfer rate of 10 MB/s is necessary for a cross-spectrum system under the hypothesis that the phase is encoded on 16 bits. Thus, a measurement lasting 100 s takes 1-GB disk space. To the best of our knowledge, noise analyzers do not save or transfer raw data of this size and at this rate. TimeLab[20] is no exception, to the extent that continuous gap-free IQ data are transferred from the 3120A/5330A/53100A instruments to the computer at reduced bandwidth (100 kHz). Faster IQ data required to plot spectra up to the maximum frequency of 1 MHz are transferred in bursts with a 10% duty cycle. Only the spectra bins are saved in the .TIM files, together with gap-free data at a significantly lower rate (1 kS/s) for time measurements and AVARs.

We present a selection of software packages (see Table II), chosen for their scientific value or for their wide use. The grayed area highlights some interesting features.

A. Features

1) Main Purposes: AllanTools, SigmaTheta, and Stable32 are intended for data analysis. By contrast, TimeLab is a tool for data acquisition, with limited analysis capabilities. It supports the phase noise analyzers from Jackson Lab and Microchip (formerly Microsemi), and a few frequency counters.

2) Graphical Interface Versus Scripting: Albeit elderly, Stable32 has an efficient graphical interface, which makes it a great choice for occasional users. For this reason, it is by far the most widely used. AllanTools and SigmaTheta require

[20]Information provided by John Miles, September 2022.

TABLE II
SOFTWARE TOOLS FOR THE TWO-SAMPLE VARIANCES

Features[1]		Tool: AllanTools	SigmaTheta	Stable32	TimeLab
Fractional frequency	**ADEV**	**Y**[2]	**Y**	**Y**[2]	**Y**
	GCODEV	beta	Y	–	–
	GRADEV	Y	–	–	–
	HDEV	Y[2]	Y	Y[2]	Y
	MDEV	**Y**	**Y**	**Y**	**Y**
	PDEV	–	Y	–	–
	Thêo1	Y	–	Y	–
	ThêoH	–	–	Y	–
	TOTDEV	Y	–	Y	–
Time	**MTIE**	**Y**	**–**	**Y**	**Y**
	TDEV	Y	–	Y	Y
	TIErms	Y	–	Y	–
Other	Estimator	direct	Bayes	direct	direct
	Error bars	LA & GR[3,4,5]	GR[4]	GR[4,5]	LA[3]
	Curve fit	–	Y	–	–
General	Type	Python lib	CL[6], API[6], GUI[6]	GUI[6]	GUI[6]
	OS[7]	L/M/W	L/M[8]/W[8]	L[9]/W	W
	License	LGPLv3+ (open)	CeCILL (open)	©[10] (free)	© (free)

(1) Here, "DEV" implies that both DEV and VAR are available
(2) Also, non-overlapped algorithm
(3) LA = Lesage-Audoin algorithm, based on stdev(σ^2)
(4) GR = Greenhall-Riley algorithm, based on $\chi^2(\sigma^2)$
(5) See main text
(6) CL = Command line, API = Application Programming Interface, GUI = Graphical user interface
(7) L = Linux, M = macOS, W = Windows
(8) API and GUI on Mac and Windows require recompiling
(9) Wine environment, or Windows virtual machine
(10) The team is working at cleaning the code for public release

TABLE III
SOFTWARE AVAILABILITY AND DEVELOPERS

AllanTools	
Anders E. E. Wallin, Danny Price, Cantwell G. Carson, Frédéric Meynadier, Yan Xie, and Eric Benkler	
Anaconda	https://anaconda.org/conda-forge/allantools
GitHub	https://github.com/aewallin/allantools
PyPI	https://pypi.org/project/AllanTools/
SigmaTheta	
François Meyer, François Vernotte, and Attila Kinali, Benoit Dubois	
Cmd line	https://gitlab.com/fm-ltfb/SigmaTheta
API	https://gitlab.com/bendub/libsigmatheta Not tested on Mac/Win, some funct. still missing.
GUI	https://gitlab.com/bendub/pysigmatheta β-test on Linux, lives on API via Python wrapper
Stable32	
Magnus Danielson, Vivek D. Dwivedi & al. Originally, William J. Riley (retired)	
GitHub	https://github.com/IEEE-UFFC/stable32
IEEE	https://ieee-uffc.org/software
Stable32	http://www.stable32.com/162Stable32.exe
TimeLab	
John Miles, KE5FX	
Miles Design	http://www.miles.io/timelab/beta.htm

programming skills. On the other hand, scripting is great in that it enables the analysis of bulk of datasets at once.

3) Mathematical Functions: All the packages provide ADEV, MDEV, and the other most used functions. For simplicity, DEV in Table II stands for DEV or VAR. When it comes to more exotic functions, such as Thêo and the Groslambert covariance, the choice of packages is smaller.

4) Estimator: The unique feature of SigmaTheta is that it uses Bayesian statistics (inverse problem) to estimate the variances. The other packages evaluate the *average* value using the classical formulas such as (34) for AVAR.

5) Missing Data: The GRADEV, included only in AllanTools, enables the evaluation of ADEV in the case of missing data during the measurement.

6) Error Bars: The basic choice is between the classical Lesage–Audoin algorithm [169] and the more sophisticated Greenhall–Riley algorithm [118]. The LA algorithm uses the square root of the fourth moment, assuming that the error distribution is Gaussian centered around the estimate. By contrast, the GR algorithm estimates the degrees of freedom using the χ^2 distribution. For each value of τ, the number of samples and the dominant noise type are taken into account.

However, Wallin [170] points out that the identification of noise processes is also needed. In the same blog, it is said that AllanTools has no "intelligent" top-level algorithm that takes a time series as input and would choose between lag-1-ACF, $B1$, and $R(n)$, and automatically determine the power-law noise type. Wallin suggests that Stable32 may use a combination of lag-1-autocorrelation [171], $B1$-ratio [172], [173], and $R(n)$.

Finally, the principles used inside SigmaTheta are discussed in Section IX-C.

B. Software Availability

The availability is summarized in Table III. AllanTools and SigmaTheta are open codes released to the public domain under very similar licenses. Stable32 is IEEE. The developers informed us[21] that the code itself is covered by the MIT license, but compiling needs code from other sources. They are trying to clean the code for it to be released into the public domain, provided that the IEEE agrees, with the ultimate goal of having at least Linux, macOS, and Windows distributions available.

[21] Magnus Danielson, private email.

ACKNOWLEDGMENT

The authors express gratitude to the Nobel laureate John Lewis "Jan" Hall for his authoritative and friendly encouragement[22]

> *…But the best thing for me is the link to your new Chart, collecting the multiple information about how to analyze measured phase noise spectra to begin to grasp the underlying causes. I have been using a 10th-generation photocopy of an early precursor to your chart, which came many years ago from either Don Halford or Helmut Hellwig at the NBS/NIST…*

They thank the users of the earlier versions of the Chart for feedback and corrections.

APPENDIX

A. Notation and Symbols

dot ($\dot{a}$)	Derivative over time.
bar ($\bar{a}$)	Mean or weighted average.
hat ($\hat{a}$)	Estimation.
$\leftrightarrow$	Fourier or Laplace transform inverse–transform pair.
$\langle\,\rangle$	Mathematical expectation, same as $\mathbb{E}\{\,\}$.
$\langle\,\rangle_\tau$	Average over a specified time interval τ.
$\langle\,\rangle_n$	Average over an integer number n of samples.
B	Bandwidth of the baseband signal and AM/PM noise (the bandwidth of the RF signal is $2B$).
b_n	Coefficients of the polynomial law of $S_\varphi(f)$.
D_y	Fractional-frequency drift.
$\mathfrak{D}$	Degrees of freedom (often denoted with ν, dof or DOF in the literature about statistics).
d_n	Coefficients of the polynomial law of $S_\nu(f)$.
$\mathbb{E}\{\,\}$	Mathematical expectation.
F	Noise factor (see NF in the following).
f	Fourier frequency, in spectral analysis.
f_c	Corner frequency, where flicker noise equals white noise.
f_L	Resonator's Leeson frequency (baseband bandwidth), half of the RF bandwidth.
h_n	Coefficients of the polynomial law of $S_\mathsf{y}(f)$.
$H(f)$	Transfer function, typically in $\lvert H(f)\rvert^2$.
k_n	Coefficients of the polynomial law of $S_\mathsf{x}(f)$.
$\mathscr{L}(f)$	Phase noise, usually $10\log_{10}[\mathscr{L}(f)]$ [dBc/Hz].
N	White noise PSD [W/Hz].
NF	Noise figure, $\mathrm{NF} = 10\log_{10}(F)$.
P	Carrier power.
Q	Quality factor, in resonators.
$S_x(f)$	One-sided power spectral density of the random variable $x(t)$. For us, chiefly, α, φ, x, y, and ν.
T	Time interval, or period.
T	Acquisition time (data record for one FFT).
T	(Subscript) signal truncated over a duration T, as in $x_T(t)$ and $X_T(f)$.
$\mathcal{T}$	Total acquisition time (full data record used in the measurement).
t	Time.
V_0	Peak amplitude (of the clock signal).
$w(t;\tau)$	Weight function (counters) or wavelet-like function (variances).
$\mathsf{x}(t)$	Time fluctuation.
$\mathsf{y}(t)$	Fractional frequency fluctuation.
$\alpha(t)$	Fractional amplitude fluctuation.
$(\Delta\nu)(t)$	Frequency fluctuation.
$\epsilon(t)$	Amplitude fluctuation.
$\theta(t)$	Random phase, replacement for $\varphi(t)$, when needed.
Λ	Triangular average or a frequency counter implementing triangular average.
ν	Carrier frequency.
Π	Uniform average or a frequency counter implementing uniform average.
$\sigma^2_\mathsf{x}(\tau)$	Same as TVAR, used in formulas.
${}^{A}\sigma^2_\mathsf{y}(\tau)$	Same as AVAR, used in formulas.
${}^{M}\sigma^2_\mathsf{y}(\tau)$	Same as MVAR, used in formulas.
${}^{P}\sigma^2_\mathsf{y}(\tau)$	Same as PVAR, used in formulas.
${}^{H}\sigma^2_\mathsf{y}(\tau)$	Same as HVAR, used in formulas.
τ	Delay (delay line).
τ	Measurement (integration) time.
τ	Relaxation time, in resonators.
τ	Time shift (correlation, convolution, and so on).
$\varphi(t)$	Random phase.
$\psi(t)$	Random phase, replacement for $\varphi(t)$, when needed.
ω	Shorthand for $2\pi\nu$.
Ω	Shorthand for $2\pi f$.
Ω	Parabolic-weight average or a frequency counter implementing such average.

B. Acronyms

ADC	Analog-to-digital converter.
AM	Amplitude modulation (also AM noise).
ADEV	Square root of AVAR.
AVAR	Allan VARiance.
BER	Bit error rate (digital electronics and telecom).
BPF	Bandpass filter.
CORDIC	COordinate Rotation DIgital Computer or Volder's algorithm.
DAC	Digital-to-analog converter.
DBM	Double balanced mixer.
DDS	Direct digital synthesizer.
DUT	Device under test (oscillator or two-port component).

[22] Private email, February 5, 2020.

FFT	Fast Fourier transform.
FM	Frequency modulation (also FM noise).
FPGA	Field-programmable gate array.
FS	Femtosecond (laser).
GRADEV	Gap-resistant ADEV.
HDEV	Square root of HVAR.
HVAR	Hadamard VARiance.
IF	Intermediate frequency (superheterodyne receiver). Also IF output of a mixer.
IQ	In-phase and quadrature (detection, modulation).
LiDAR	Light detection and ranging.
LIGO	Laser Interferometer Gravitational-Wave Observatory.
LO	Local oscillator (superheterodyne). Also LO input of a mixer.
LPF	Low-pass filter.
LSB	Lower sideband, in modulated signals.
MDEV	Square root of MVAR.
MVAR	Modified [Allan] VARiance.
NCO	Numerically controlled oscillator.
NIST	National Institute of Standards and Technology (USA).
PLL	Phase-locked loop.
PM	Phase modulation (also PM noise).
PDF	Probability density function.
PSD	Power spectral density.
PDEV	Square root of PVAR.
PVAR	Parabolic VARiance.
RF	Radio frequency. Also RF input of a mixer.
SDR	Software-defined radio.
TDDS	Tracking direct digital synthesizer.
TIE	Time interval error [32].
TVAR	Time VARiance.
USB	Upper sideband, in modulated signals.
VCO	Voltage-controlled oscillator. Also voltage-control input of an oscillator.
ZCD	Zero-crossing detector.

REFERENCES

[1] C. J. Grebenkemper, "Local oscillator phase noise and its effect on receiver performance," Watkins Johnson, San Jose, CA, USA, Tech. Rep., 86, Nov. 1981. [Online]. Available: https://www.rfcafe.com/references/articles/wj-tech-notes/effect-LO-phase-noise-receiver-sensitivity-v8-6.pdf

[2] B. M. Sosin, "H.F. Communication receiver performance requirements and realization," *Radio Electron. Engineer*, vol. 41, no. 7, pp. 321–329, Jul. 1971.

[3] M. I. Skolnik, *Radar Handbook*, 3rd ed. New York, NY, USA: McGraw-Hill, 2008.

[4] D. B. Leeson and G. F. Johnson, "Short-term stability for a Doppler radar: Requirements, measurements, and techniques," *Proc. IEEE*, vol. 54, no. 2, pp. 244–248, Feb. 1966.

[5] D. B. Leeson and G. F. Johnson, "Short-term stability for a Doppler radar: Requirements, measurements and techniques," in *Proc. IEEE-NASA Symp. Definition Meas. Short-Term Freq. Stability*, A. R. Chi, Ed., Greenbelt, MD, USA, Nov. 1964, pp. 3–9. [Online]. Available: https://ntrs.nasa.gov/citations/19660001092

[6] N. D. Dalt and A. Sheikoleslami, *Understanding Jitter and Phase Noise*. Cambridge, U.K.: Cambridge Univ. Press, 2018.

[7] M. P. Li, *Jitter, Noise, and Signal Integrity at High-Speed*. Boston, MA, USA: Prentice-Hall, 2008.

[8] J.-C. Lin, "Synchronization requirements for 5G," *IEEE Veh. Technol. Mag.*, vol. 13, no. 3, pp. 91–99, Sep. 2018.

[9] H. Tataria, M. Shafi, A. F. Molisch, M. Dohler, H. Sjoland, and F. Tufvesson, "6G wireless systems: Vision, requirements, challenges, insights, and opportunities," *Proc. IEEE*, vol. 109, no. 7, pp. 1166–1199, Jul. 2021.

[10] F. Conte, S. Massucco, M. Paolone, G.-P. Schiapparelli, F. Silvestro, and Y. Zuo, "Frequency stability assessment of modern power systems: Models definition and parameters identification," 2021, *arXiv:2104.07330*.

[11] J. Kruse, B. Schäfer, and D. Witthaut, "Revealing drivers and risks for power grid frequency stability with explainable AI," *Patterns*, vol. 2, no. 11, Nov. 2021, Art. no. 100365.

[12] "Frequency stability evaluation criteria for the synchronous zone of continental Europe—Requirements and impacting factors," Eur. Netw. Transmiss. Syst. Oper. Electr. (ENTSOE), Brussels, Belgium, Tech. Rep., Mar. 2016. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/RGCE_SPD_frequency_stability_criteria_v10.pdf

[13] J. Amelot et al., "Towards timely intelligence in the power grid," in *Proc. Precis. Time Time Interval Conf.*, Reston, VA, USA, Nov. 2012, pp. 399–412.

[14] A. R. Chi, Ed., (Nov. 1964). *Short Term Frequency Stability*. NASA SP-80. Goddard Space Flight Center, Greenbelt, MD, USA. [Online]. Available: https://ntrs.nasa.gov/citations/19660001092

[15] A. R. Chi, Ed., *Proc. IEEE*, vol. 54, no. 2, Feb. 1966, Special Issue on Frequency Stability. [Online]. Available: https://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=31081&punumber=5

[16] J. A. Barnes et al., "Characterization of frequency stability," *IEEE Trans. Instrum. Meas.*, vol. IM-20, no. 2, pp. 105–120, May 1971.

[17] E. Donley, Chair, et al., *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*, IEEE Standard, Dec. 2022. [Online]. Available: https://standards.ieee.org/ieee/1139/7585/

[18] Z. Malkin, "Application of the Allan variance to time series analysis in astrometry and geodesy: A review," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 582–589, Apr. 2016.

[19] B. Holzer, Ed., "Beam injection, extraction and transfer," CERN, Geneva, Switzerland, Tech. Rep. CERN-2018-008-SP, 2018, doi: 10.23730/CYRSP-2018-005.

[20] H. Ball, W. D. Oliver, and M. J. Biercuk, "The role of master clock stability in quantum information processing," *NPJ Quantum Inf.*, vol. 2, no. 1, pp. 1–8, Nov. 2016.

[21] S. A. Diddams, K. Vahala, and T. Udem, "Optical frequency combs: Coherently uniting the electromagnetic spectrum," *Science*, vol. 369, no. 6501, Jul. 2020.

[22] T. Udem, R. Holzwarth, and T. W. Hänsch, "Optical frequency metrology," *Nature*, vol. 416, no. 6877, pp. 233–237, Mar. 2002.

[23] S. T. Cundiff and J. Ye, "Colloquium: Femtosecond optical frequency combs," *Rev. Mod. Phys.*, vol. 75, no. 1, pp. 325–342, Mar. 2003.

[24] N. Kuse and M. E. Fermann, "Electro-optic comb based real time ultra-high sensitivity phase noise measurement system for high frequency microwaves," *Sci. Rep.*, vol. 7, no. 1, p. 2847, Jun. 2017.

[25] X. Xie et al., "Photonic microwave signals with zeptosecond-level absolute timing noise," *Nature Photon.*, vol. 11, no. 1, pp. 44–47, Jan. 2017.

[26] C. E. Calosso, C. Clivati, and S. Micalizio, "Avoiding aliasing in Allan variance: An application to fiber link data analysis," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 646–655, Apr. 2016.

[27] C. Lisdat et al., "A clock network for geodesy and fundamental science," *Nature Commun.*, vol. 7, no. 1, p. 12443, Aug. 2016.

[28] P. Delva et al., "Test of special relativity using a fiber network of optical clocks," *Phys. Rev. Lett.*, vol. 118, pp. 1–6, Jun. 2017, doi: 10.1103/PhysRevLett.118.221102.

[29] D. Nicolodi, B. Argence, W. Zhang, R. Le Targat, G. Santarelli, and Y. Le Coq, "Spectral purity transfer between optical wavelengths at the 10^{-18} level," *Nature Photon.*, vol. 8, no. 3, pp. 219–223, Mar. 2014.

[30] *Le Système International d'unités/The International System of Units ('The SI Brochure')*, 9th ed. Sèvres, France: Bureau International des Poids et Mesures (BIPM), 2019. [Online]. Available: https://www.bipm.org/documents/20126/41483022/SI-Brochure-9.pdf

[31] *Measures for Random Instabilities in Frequency and Time (Phase)*, document TF.538-4, Recommendation ITU-R, Geneva, Switzerland, Jul. 2017.

[32] *Definitions and Terminology for Synchronization in Packet Networks*, document G.8260, ITU Working Group 15, Geneva, Switzerland, Mar. 2000.

[33] *Transmission and Multiplexing (TM); Generic Requirements for Synchronization Networks*, European Telecommunications Standards Institute (ETSI), Sofia Antipolis, France, May 2002. [Online]. Available: https://etsi.org

[34] J. Rutman, "Characterization of phase and frequency instabilities in precision frequency sources: Fifteen years of progress," *Proc. IEEE*, vol. 66, no. 9, pp. 1048–1075, Sep. 1978.

[35] J. Rutman and F. L. Walls, "Characterization of frequency stability in precision frequency sources," *Proc. IEEE*, vol. 79, no. 7, pp. 952–960, Jul. 1991.

[36] W. J. Riley. *Handbook of Frequency Stability Analysis*, NIST. Accessed: Feb. 1, 2023. [Online]. Available: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=50505

[37] D. Owen, "Good practice to phase noise measurement," National Physical Laboratory, Teddington, U.K., Tech. Rep., May 2004.

[38] V. F. Kroupa, *Frequency Stability: Fundamentals and Measurement*. New York, NY, USA: IEEE Press, 1983.

[39] V. F. Kroupa, *Frequency Stability*. Hoboken, NJ, USA: Wiley, 2012.

[40] D. B. Sullivan, D. W. Allan, and D. A. Howe. (Mar. 1990). *Characterization of Clock and Oscillators*. NIST. [Online]. Available: https://tf.nist.gov/general/pdf/868.pdf

[41] U. L. Rohde, E. Rubiola, and J. C. Whitaker, *Microwave and Wireless Synthesizers*. Hoboken, NJ, USA: Wiley, Apr. 2020.

[42] E. Rubiola and R. Boudot, "The effect of AM noise on correlation phase-noise measurements," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 5, pp. 926–932, May 2007.

[43] P. D. Welch, "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Trans. Audio Electroacoust.*, vol. AE-15, no. 2, pp. 70–73, Jun. 1967.

[44] *International Vocabulary of Metrology—Basic and General Concepts and Associated Terms (VIM)*, document JCGM 200:2012, Joint Committee for Guides in Metrology (JCGM), The VIM is available free of charge, 2012. [Online]. Available: https://www.bipm.org/en/publications/guides/

[45] H. Hellwig, *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology*, IEEE Standard 1139-1988, 1988.

[46] B. Razavi, *RF Microelectronics*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2012.

[47] D. K. Barton and S. A. Leonov, *Radar Technology Encyclopedia*. Nordwood, MA, USA: Artech House, 1998.

[48] W. K. Victor, "The evaluation of phase-stable oscillator for coherent communication systems," in *Proc. 10th Annu. Symp. Freq. Control*, Asbury Park, NJ, USA, May 1956, pp. 268–304.

[49] J. H. Shoaf, "Specification and measurement of frequency stability," NIST, Gaithersburg, MD, USA, NBS Tech. Rep., 74-396, Nov. 1974.

[50] B. E. Blair, *Time and Frequency: Theory and Fundamentals*. Washington, DC, USA: U.S. Dept. Commerce, May 1974.

[51] C. H. Horn, "A carrier suppression technique for measuring S/N and carrier/sideband ratios greater than 120 dB," in *Proc. 23rd Annu. Symp. Freq. Control*, Fort Monmouth, NJ, USA, May 1969, pp. 223–235.

[52] D. B. Leeson, "Oscillator phase noise: A 50-year review," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 8, pp. 1208–1225, Aug. 2016.

[53] W. P. Robins, *Phase Noise in Signal Sources*. Maricourt, France: Peter Peregrinus and IEE, 1984.

[54] E. Rubiola, "The measurement of AM noise of oscillators," 2005, *arXiv:physics/0512082*.

[55] C. E. Calosso and E. Rubiola, "Phase noise and jitter in digital electronics," 2016, *arXiv:1701.00094*.

[56] R. Boudot and E. Rubiola, "Phase noise in RF and microwave amplifiers," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 59, no. 12, pp. 2613–2624, Dec. 2012.

[57] D. Halford, A. E. Wainwright, and J. A. Barnes, "Flicker noise of phase in RF amplifiers and frequency multipliers: Characterization, cause, and cure," in *Proc. 22nd Annu. Symp. Freq. Control*, Apr. 1968, pp. 340–341.

[58] E. Rubiola, "Tutorial on the double-balanced mixer," 2006, *arXiv:physics/0608211*.

[59] C. A. Barnes, A. Hati, C. W. Nelson, and D. A. Howe, "Residual PM noise evaluation of radio frequency mixers," in *Proc. Joint Conf. IEEE Int. Freq. Control Eur. Freq. Time Forum (FCS)*, San Francisco, CA, USA, May 2011, pp. 1–5.

[60] W. Kester, *Analog–Digital Conversion*. Wilmington, MA, USA: Analog Devices, 2004.

[61] W. F. Egan, "Modeling phase noise in frequency dividers," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 37, no. 4, pp. 307–315, Jul. 1990.

[62] V. F. Kroupa, "Jitter and phase noise in frequency dividers," *IEEE Trans. Instrum. Meas.*, vol. 50, no. 5, pp. 1241–1243, Oct. 2001.

[63] W. F. Egan, *Phase-Lock Basics*, 2nd ed. Hoboken, NJ, USA: Wiley, 2008.

[64] D. Banerjee, *PLL Performance, Simulation, and Design*, 5th ed. Indianapolis, IN, USA: Dog Ear Publishing, May 2017.

[65] K. Shu and E. Sáncez-Sinencio, *CMOS PLL Synthesizers*. Boston, MA, USA: Springer, 2005.

[66] B.-G. Goldberg, *Digital Frequency Synthesis Demystified*. Eagle Rock, VA, USA: LLH Technology Publishing, 1999.

[67] V. F. Kroupa, *Direct Digital Frequency Synthesizers*. New York, NY, USA: IEEE Press, 1999.

[68] C. E. Calosso, Y. Gruson, and E. Rubiola, "Phase noise and amplitude noise in DDS," in *Proc. IEEE Int. Freq. Control Symp.*, Baltimore, MD, USA, May 2012, pp. 777–782.

[69] E. Rubiola, *Phase Noise and Frequency Stability in Oscillators*. Cambridge, U.K.: Cambridge Univ. Press, Nov. 2010.

[70] E. Rubiola and V. Giordano, "On the $1/f$ frequency noise in ultra-stable quartz oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 1, pp. 15–22, Jan. 2007.

[71] E. Rubiola and R. Brendel, "A generalization of the Leeson effect," 2010, *arXiv:1003.0113*.

[72] P. Penfield, "Noise in negative-resistance amplifiers," *IRE Trans. Circuit Theory*, vol. 7, no. 2, pp. 166–170, Jun. 1960.

[73] A. van der Ziel, "On the noise figure of negative-conductance amplifiers," *IRE Trans. Circuit Theory*, vol. 9, no. 1, pp. 83–84, Mar. 1962.

[74] W. Loh, S. Yegnanarayanan, R. J. Ram, and P. W. Juodawlkis, "Unified theory of oscillator phase noise I: White noise," *IEEE Trans. Microw. Theory Techn.*, vol. 61, no. 6, pp. 2371–2381, Jun. 2013.

[75] W. Loh, S. Yegnanarayanan, R. J. Ram, and P. W. Juodawlkis, "Unified theory of oscillator phase noise II: Flicker noise," *IEEE Trans. Microw. Theory Techn.*, vol. 61, no. 12, pp. 4130–4144, Dec. 2013.

[76] A. Demir, A. Mehrotra, and J. Roychowdhury, "Phase noise in oscillators: A unifying theory and numerical methods for characterization," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 47, no. 5, pp. 655–674, May 2000.

[77] F. X. Kärtner, "Analysis of white and $f^{-\alpha}$ noise in oscillators," *Int. J. Circuit Theory Appl.*, vol. 18, pp. 485–519, Jan. 1990.

[78] A. Hajimiri and T. H. Lee, "A general theory of phase noise in electrical oscillators," *IEEE J. Solid-State Circuits*, vol. 33, no. 2, pp. 179–194, Feb. 1998.

[79] T. H. Lee and A. Hajimiri, "Oscillator phase noise: A tutorial," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 326–336, Mar. 2000.

[80] E. Pankratz and E. Sánchez-Sinencio, "Survey of integrated-circuit-oscillator phase-noise analysis," *Int. J. Circuit Theory Appl.*, vol. 42, no. 9, pp. 871–938, Sep. 2014.

[81] B. van der Pol and J. van der Mark, "Frequency demultiplication," *Nature*, vol. 120, no. 3019, pp. 363–364, Sep. 1927.

[82] W. A. Edson, "Noise in oscillators," *Proc. IRE*, vol. 48, no. 8, pp. 1454–1466, Aug. 1960.

[83] D. B. Leeson, "A simple model of feed back oscillator noise spectrum," *Proc. IEEE*, vol. 54, pp. 329–330, Feb. 1966.

[84] E. Rubiola, "On the measurement of frequency and of its sample variance with high-resolution counters," *Rev. Sci. Instrum.*, vol. 76, no. 5, May 2005, Art. no. 054703.

[85] E. Rubiola, M. Lenczner, P.-Y. Bourgeois, and F. Vernotte, "The Ω counter, a frequency counter based on the linear regression," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 7, pp. 961–969, Jul. 2016.

[86] S. T. Dawkins, J. J. McFerran, and A. N. Luiten, "Considerations on the measurement of the stability of oscillators with frequency counters," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 5, pp. 918–925, May 2007.

[87] L. Galleani and P. Tavella, "The dynamic Allan variance," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 56, no. 3, pp. 450–464, Mar. 2009.

[88] L. Galleani, "The dynamic Allan variance II: A fast computational algorithm," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 1, pp. 182–188, Jan. 2010.

[89] L. Galleani, "The dynamic Allan variance III: Confidence and detection surfaces," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 58, no. 8, pp. 1550–1558, Aug. 2011.

[90] L. Galleani and P. Tavella, "The dynamic Allan variance IV: Characterization of atomic clock anomalies," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 62, no. 5, pp. 791–801, May 2015.

[91] L. Galleani and P. Tavella, "The dynamic Allan variance V: Recent advances in dynamic stability analysis," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 624–635, Apr. 2016.

[92] C. A. Greenhall, D. A. Howe, and D. B. Percival, "Total variance, an estimator of long-term frequency stability [standards]," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 46, no. 5, pp. 1183–1191, Sep. 1999.

[93] D. A. Howe, "The total deviation approach to long-term characterization of frequency stability," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 47, no. 5, pp. 1102–1110, Sep. 2000.

[94] D. A. Howe and T. K. Peppler, "Very long-term frequency stability: Estimation using a special-purpose statistic," in *Proc. IEEE Int. Freq. Control Symp. PDA Exhib. 17th Eur. Freq. Time Forum*, Tampa, FL, USA, May 2003, pp. 233–238.

[95] D. A. Howe, "ThêoH: A hybrid, high-confidence statistic that improves on the Allan deviation," *Metrologia*, vol. 43, no. 4, pp. S322–S331, Aug. 2006.

[96] J. A. Taylor and D. A. Howe, "Fast TheoBR: A method for long data set stability analysis," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 9, pp. 2091–2094, Sep. 2010.

[97] S. Bregni, "Twenty-five years of applications of the modified Allan variance in telecommunications," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 520–530, Apr. 2016.

[98] M. Abdel Hafiz, C. Carlé, N. Passilly, J. M. Danet, C. E. Calosso, and R. Boudot, "Light-shift mitigation in a microcell-based atomic clock with symmetric auto-balanced Ramsey spectroscopy," *Appl. Phys. Lett.*, vol. 120, no. 6, Feb. 2022, Art. no. 064101.

[99] M. Hasegawa et al., "Microfabrication of cesium vapor cells with buffer gas for MEMS atomic clocks," *Sens. Actuators A, Phys.*, vol. 67, pp. 594–601, Jan. 2011.

[100] S. R. Stein, "The Allan variance—Challenges and opportunities," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 3, pp. 540–547, Mar. 2010.

[101] D. G. Enzer, D. W. Murphy, and E. A. Burt, "Allan deviation of atomic clock frequency corrections: A new diagnostic tool for characterizing clock disturbances," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 68, no. 7, pp. 2590–2601, Jul. 2021.

[102] J. Levine, P. Tavella, and G. Santarelli, "Introduction to the special issue on celebrating the 50th anniversary of the Allan variance," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 511–512, Apr. 2016.

[103] D. W. Allan and J. Levine, "A historical perspective on the development of the Allan variances and their strengths and weaknesses," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 513–519, Apr. 2016.

[104] D. W. Allan, "Statistics of atomic frequency standards," *Proc. IEEE*, vol. 54, no. 2, pp. 221–230, Feb. 1966.

[105] D. W. Allan and J. A. Barnes, "A modified 'Allan variance' with increased oscillator characterization ability," in *Proc. Ann. Freq. Control Symp.*, Ft. Monmouth, NJ, USA, May 1981, pp. 470–474.

[106] G. Kramer and W. Klische, "Multi-channel synchronous digital phase recorder," in *Proc. IEEE Int. Frequncy Control Symp. PDA Exhib.*, Seattle, WA, USA, Jun. 2001, pp. 144–151.

[107] G. Krame and W. Klische, "Extra high precision digital phase recorder," in *Proc. 18th Eur. Freq. Time Forum (EFTF)*, Guildford, U.K., Apr. 2004, pp. 595–602.

[108] J. Kalisz, "Review of methods for time interval measurements with picosecond resolution," *Metrologia*, vol. 41, no. 1, pp. 17–32, 2004.

[109] S. Johansson, "New frequency counting principle improves resolution," in *Proc. IEEE Int. Freq. Control Symp. Expo.*, Vancouver, BC, Canada, Aug. 2005, pp. 628–635.

[110] E. Benkler, C. Lisdat, and U. Sterr, "On the relation between uncertainties of weighted frequency averages and the various types of Allan deviations," *Metrologia*, vol. 55, no. 4, pp. 565–574, Aug. 2015.

[111] F. Vernotte, M. Lenczner, P.-Y. Bourgeois, and E. Rubiola, "The parabolic variance (PVAR): A wavelet variance based on the least-square fit," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 611–623, Apr. 2016.

[112] F. Vernotte, G. Zalamansky, and E. Lantz, "Time stability characterization and spectral aliasing. Part II: A frequency-domain approach," *Metrologia*, vol. 35, no. 5, pp. 731–738, Oct. 1998.

[113] L. G. Bernier, "Theoretical analysis of the modified Allan variance," in *Proc. 41st Annu. Symp. Freq. Control*, Philadelphia, PA, USA, May 1987, pp. 116–121.

[114] T. Walter, "Characterizing frequency stability: A continuous power-law model with discrete sampling," *IEEE Trans. Instrum. Meas.*, vol. 43, no. 1, pp. 69–79, Feb. 1994.

[115] F. Vernotte, S. Chen, and E. Rubiola, "Response and uncertainty of the parabolic variance PVAR to noninteger exponents of power law," *IEEE Trans. Instrum. Meas.*, vol. 70, pp. 1–6, 2021.

[116] C. A. Greenhall, "Spectral ambiguity of Allan variance," *IEEE Trans. Instrum. Meas.*, vol. 47, no. 3, pp. 623–627, Jun. 1998.

[117] F. Vernotte, G. Zalamansky, and E. Lantz, "Time stability characterization and spectral aliasing. Part I: A time-domain approach," *Metrologia*, vol. 35, no. 5, pp. 723–730, Oct. 1998.

[118] C. Greenhall and W. Riley, "Uncertainty of stability variances based on finite differences," in *Proc. 35th Annu. Precise Time Time Interval Meeting*, San Diego, CA, USA, Dec. 2003, pp. 267–280.

[119] R. A. Fisher, "The fiducial argument in statistical inference," *Ann. Eugenics*, vol. 6, no. 4, pp. 391–398, Dec. 1935.

[120] D. V. Lindley, "Fiducial distributions and Bayes' theorem," *J. Roy. Stat. Soc., B, Methodol.*, vol. 20, no. 1, pp. 102–107, Jan. 1958.

[121] F. Vernotte and E. Lantz, "Statistical biases and very-long-term time stability analysis," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 59, no. 3, pp. 523–530, Mar. 2012.

[122] M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," *Proc. IEEE*, vol. 93, no. 2, pp. 216–231, Feb. 2005.

[123] O. E. Brigham, *The Fast Fourier Transform and Its Applications*. Upper Saddle River, NJ, USA: Prentice-Hall, 1988.

[124] D. Babitch and J. Oliverio, "Phase noise of oscillators at very low Fourier frequencies," in *Proc. Ann. Freq. Control Symp.*, Atlantic City, NJ, USA, May 1974, pp. 150–159.

[125] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proc. IEEE*, vol. 66, no. 1, pp. 172–204, Feb. 1978.

[126] W. R. Bennett, "Spectra of quantized signals," *Bell Syst. Tech. J.*, vol. 27, no. 3, pp. 446–472, Jul. 1948.

[127] B. Widrow, I. Kollár, and M.-C. Liu, "Statistical theory of quantization," *IEEE Trans. Instrum. Meas.*, vol. 45, no. 2, pp. 353–361, Apr. 1996.

[128] B. Widrow and I. Kollár, *Quantization Noise*. Cambridge, U.K.: Cambridge Univ. Press, 2008.

[129] E. Rubiola and F. Vernotte, "The cross-spectrum experimental method," 2010, *arXiv:1004.5539*.

[130] R. H. Brown, R. C. Jennison, and M. K. D. Gupta, "Apparent angular sizes of discrete radio sources," *Nature*, vol. 170, no. 4338, pp. 1061–1063, Dec. 1952.

[131] R. F. C. Vessot, R. F. Mueller, and J. Vanier, "A cross-correlation technique for measuring the short-term properties of stable oscillators," in *Proc. Symp. Definition Meas. Short-Term Freq. Stability*, 1964, pp. 231–235.

[132] F. L. Walls, S. R. Stein, J. E. Gray, and D. J. Glaze, "Design considerations in state-of-the-art signal processing and phase noise measurement systems," in *Proc. 30th Annu. Symp. Freq. Control*, Atlantic City, NJ, USA, 1976, pp. 269–274.

[133] A. Baudiquez, E. Lantz, E. Rubiola, and F. Vernotte, "Cross-spectrum measurement statistics: Uncertainties and detection limit," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 67, no. 11, pp. 2461–2470, Nov. 2020.

[134] A. Baudiquez, E. Lantz, E. Rubiola, and F. Vernotte, "The statistics of the cross-spectrum and the spectrum average: Generalization to multiple instruments," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 69, no. 8, pp. 2585–2594, Aug. 2022.

[135] E. Rubiola and V. Giordano, "Very high frequency and microwave interferometric PM and AM noise measurements," *Rev. Sci. Instrum.*, vol. 71, no. 8, pp. 3085–3091, Aug. 2000.

[136] C. W. Nelson, A. Hati, and D. A. Howe, "A collapse of the cross-spectral function in phase noise metrology," *Rev. Sci. Instrum.*, vol. 85, no. 3, pp. 1–7, Mar. 2014.

[137] A. Hati, C. W. Nelson, and D. A. Howe, "Cross-spectrum measurement of thermal-noise limited oscillators," *Rev. Sci. Instrum.*, vol. 87, no. 3, Mar. 2016, Art. no. 034708.

[138] Y. Gruson, V. Giordano, U. L. Rohde, A. K. Poddar, and E. Rubiola, "Cross-spectrum PM noise measurement, thermal energy, and metamaterial filters," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 64, no. 3, pp. 634–642, Mar. 2017.

[139] Y. Gruson, A. Rus, U. L. Rohde, A. Roth, and E. Rubiola, "Artifacts and errors in cross-spectrum phase noise measurements," *Metrologia*, vol. 57, no. 5, pp. 1–12, Oct. 2020, open access.

[140] *European Workshop on Cross-Spectrum Phase Noise Measurements*, Laboratoire National de métrologie et d'Essays (LNE), Paris, France, Dec. 2014.

[141] *Cross Spectrum $\mathscr{L}(f)$ Workshop*, document, Denver, CO, USA, Side Event at the 2015 European Frequency and Time Forum and Frequency Control Symposium, Joint Meeting, Apr. 2015.

[142] *Cross Spectrum $\mathscr{L}(f)$ Workshop*, document, Besançon, France, Side Event at the 2017 European Frequency and Time Forum and Frequency Control Symposium, Joint Meeting, Jul. 2017.
[143] E. Rubiola, E. Salik, S. Huang, and L. Maleki, "Photonic delay technique for phase noise measurement of microwave oscillators," *J. Opt. Soc. Amer., B, Opt. Phys.*, vol. 22, no. 5, pp. 987–997, May 2005.
[144] E. Salik, N. Yu, L. Maleki, and E. Rubiola, "Dual photonic-delay line cross correlation method for phase noise measurement," in *Proc. IEEE Int. Freq. Control Symp. Expo.*, Montreal, QC, Canada, Aug. 2004, pp. 303–306.
[145] K. Volyanskiy et al., "Applications of the optical fiber to the generation and to the measurement of low-phase-noise microwave signals," *J. Opt. Soc. Amer. B, Opt. Phys.*, vol. 25, no. 12, pp. 2140–2150, Dec. 2008.
[146] K. H. Sann, "The measurement of near-carrier noise in microwave amplifiers," *IEEE Trans. Microw. Theory Techn.*, vol. MTT-16, no. 9, pp. 761–766, Sep. 1968.
[147] F. Labaar, "New discriminator boosts phase noise testing," *Microwaves*, vol. 21, no. 3, pp. 65–69, Mar. 1982.
[148] E. N. Ivanov, M. E. Tobar, and R. A. Woode, "Applications of interferometric signal processing to phase-noise reduction in microwave oscillators," *IEEE Trans. Microw. Theory Techn.*, vol. 46, no. 10, pp. 1537–1545, Oct. 1998.
[149] E. N. Ivanov, M. E. Tobar, and R. A. Woode, "Microwave interferometry: Application to precision measurements and noise reduction techniques," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 45, no. 11, pp. 1526–1535, Nov. 1998.
[150] 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, Jan. 1999.
[151] E. Rubiola, J. Groslambert, M. Brunet, and V. Giordano, "Flicker noise measurement of HF quartz resonators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 47, no. 2, pp. 361–368, Mar. 2000.
[152] C. E. Calosso, A. C. C. Olaya, and E. Rubiola, "Phase-noise and amplitude-noise measurement of DACs and DDSs," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 67, no. 2, pp. 431–439, Feb. 2020.
[153] E. Rubiola and V. Giordano, "Advanced interferometric phase and amplitude noise measurements," *Rev. Sci. Instrum.*, vol. 73, no. 6, pp. 2445–2457, Jun. 2002.
[154] J. E. Volder, "The CORDIC trigonometric computing technique," *IRE Trans. Electron. Comput.*, vol. 8, no. 3, pp. 330–334, Sep. 1959.
[155] J. E. Volder, "The birth of CORDIC," *J. VLSI Signal Process.*, vol. 25, no. 2, pp. 101–105, 2000.
[156] P. K. Meher, J. Valls, T.-B. Juang, K. Sridharan, and K. Maharatna, "50 years of CORDIC: Algorithms, architectures, and applications," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 9, pp. 1893–1907, Sep. 2009.
[157] J. Grove et al., "Direct-digital phase-noise measurement," in *Proc. Int. Freq. Control Symp.*, Montreal, QC, Canada, Aug. 2004, pp. 287–291.
[158] K. Mochizuki, M. Uchino, and T. Morikawa, "Frequency-stability measurement system using high-speed ADCs and digital signal processing," *IEEE Trans. Instrum. Meas.*, vol. 56, no. 5, pp. 1887–1893, Oct. 2007.
[159] J. A. Sherman and R. Jördens, "Oscillator metrology with software defined radio," *Rev. Sci. Instrum.*, vol. 87, no. 5, May 2016, Art. no. 054711.
[160] A. C. Cárdenas-Olaya et al., "Noise characterization of analog to digital converters for amplitude and phase noise measurements," *Rev. Sci. Instrum.*, vol. 88, pp. 1–9, Jun. 2017.
[161] G. Feldhaus and A. Roth, "A 1 MHz to 50 GHz direct down-conversion phase noise analyzer with cross-correlation," in *Proc. Eur. Freq. Time Forum (EFTF)*, York, U.K., Apr. 2016, pp. 1–14.
[162] C. E. Calosso, "Tracking DDS in time and frequency metrology," in *Proc. Joint Eur. Freq. Time Forum Int. Freq. Control Symp.*, Prague, Czech Republic, Jul. 2013, pp. 747–749.
[163] C. E. Calosso, F. Vernotte, V. Giordano, C. Fluhr, B. Dubois, and E. Rubiola, "Frequency stability measurement of cryogenic sapphire oscillators with a multichannel tracking DDS and the two-sample covariance," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 66, no. 3, pp. 616–623, Mar. 2019.
[164] C. E. Calosso, "The time processor: A new platform for next generation timescales," in *Proc. Workshop Synchronization Timing Syst.*, Vancouver, BC, Canada, Mar. 2021, pp. 1–15. [Online]. Available: https://wsts.atis.org/wp-content/uploads/sites/9/2021/03/AddTalk_TimeProcessor_Calosso.pdf
[165] D. W. Allan and H. Daams, "Picosecond time difference measurement system," in *Proc. 29th Annu. Symp. Freq. Control*, Atlantic City, NJ, USA, May 1975, pp. 404–411.
[166] G. Brida, "High resolution frequency stability measurement system," *Rev. Sci. Instrum.*, vol. 73, no. 5, pp. 2171–2174, May 2002.
[167] G. J. Dick, P. F. Kuhnle, and R. L. Snydor, "Zero-crossing detector with sub-microsecond jitter and crosstalk," in *Proc. 22th Annu. Precise Time Time Interval Syst. Appl. Meeting*, Vienna, VI, USA, Dec. 1990, pp. 269–282.
[168] O. Collins, "The design of low jitter hard limiters," *IEEE Trans. Commun.*, vol. 44, no. 5, pp. 601–608, May 1996.
[169] P. Lesage and C. Audoin, "Characterization of frequency stability: Uncertainty due to the finite number of measurements," *IEEE Trans. Instrum. Meas.*, vol. IM-22, no. 2, pp. 157–161, Jun. 1973.
[170] A. E. E. Wallin. (Apr. 2018). *Power Law Noise Identification for Allantools*. [Online]. Available: https://ui.adsabs.harvard.edu/abs/2018ascl.soft04021W
[171] W. J. Riley and C. A. Greenhall, "Power law noise identification using the lag 1 autocorrelation," in *Proc. 18th Eur. Freq. Time Forum*, Guildford, U.K., Apr. 2004, pp. 576–580.
[172] D. A. Howe, R. Beard, C. Greenhall, F. Vernotte, and B. Riley, "The total estimator of the Hadamard function used for GPS operations," in *Proc. 32th Annu. Precise Time Time Interval Syst. Appl. Meeting*, 2000, pp. 255–268.
[173] J. A. Barnes, *Tables of Bias Functions, B1 and B2, for Variances Based on Finite Samples of Processes With Power Law Spectral Densities*, Standard 375, US Department of Commerce, National Bureau of Standards, 1969.

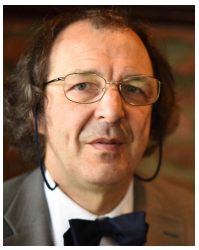

Enrico Rubiola (Member, IEEE) is a Professor of electronics with the Université de Franche Comté (uFC), Besançon, France, and a Researcher with the FEMTO-ST Institute, Besançon, in turn affiliated to the National Research Council (CNRS) and hosted by Supmicrotech-ENSMM. He is also an Associated Researcher with the Italian National Institute of Metrology (INRiM), Torino, Italy. He is the PI of the Oscillator IMP project, a platform for the measurement of short-term frequency stability and AM/PM noise he founded in 2012. He also founded the European Frequency and Time Seminar (http://efts.eu) in 2013, a no-profit crash course in Time and Frequency, and he has been chairing and running it since. His primary interests are high-purity oscillators, AM and PM, noise in digital systems, time and frequency metrology, frequency synthesis, Allan variances, microwave photonics, and precision electronics. He is known for AM/PM noise measurement instruments with ultimate sensitivity, for the theory of the "Leeson effect," for the statistics of Π, Λ, and Ω frequency counters, and for hacking oscillators from phase noise.

Dr. Rubiola received the IEEE W. G. Cady Award "for groundbreaking contributions" in the field in 2018. A wealth of free scientific and learning material is available on his home page: http://rubiola.org.

François Vernotte is a Professor of astronomy and astrophysics with the Université de Franche Comté (uFC), Besançon, France, and a Researcher with the FEMTO-ST Institute, Besançon, in turn affiliated to the National Research Council (CNRS) and hosted by Supmicrotech-ENSMM. After receiving the Ph.D. degree in engineering from the uFC in 1991, he joined the Observatory 'THETA' of Besançon, where he served as the Director in 2002 to 2012, then he moved to the FEMTO-ST Institute in 2019. His main research topic is the frequency stability of atomic frequency standards and millisecond pulsar signals, which he pursues using statistical data processing (time series and spectral analysis), parameter estimation (inverse problem and Bayesian statistics), and Monte-Carlo simulation.

Open Access funding provided by 'Istituto Nazionale di Ricerca Metrologica - INRiM' within the CRUI CARE Agreement