Frequency Stability Measurement of Cryogenic Sapphire Oscillators With a Multichannel Tracking DDS and the Two-Sample Covariance

Claudio E. Calosso, François Vernotte, Vincent Giordano [ID],
Christophe Fluhr, Benoît Dubois, and Enrico Rubiola [ID]

Abstract—This paper shows the first measurement of three 100-MHz signals exhibiting fluctuations from 2×10^{-16} to parts in 10^{-15} for an integration time τ between 1 s and 1 day. Such stable signals are provided by three cryogenic sapphire oscillators (CSOs) operating at about 10 GHz, also delivering the 100-MHz output via a dedicated synthesizer. The measurement is made possible by a six-channel tracking direct digital synthesizer (TDDS) and the two-sample covariance tool, used to estimate the Allan variance. The use of two TDDS channels per CSO enables high rejection of the instrument background noise. The covariance outperforms the three-cornered hat (TCH) method in that the background converges to zero "out of the box," with no need of the hypothesis that the instrument channels are equally noisy, nor of more sophisticated techniques to estimate the background noise of each channel. Thanks to correlation and averaging, the instrument background (AVAR) rolls off with a slope $1/\sqrt{m}$, the number of measurements, down to 10^{-18} at $\tau = 10^4$ s. For consistency check, we compare the results to the traditional TCH method beating the 10-GHz outputs down to the megahertz region. Given the flexibility of the TDDS, our methods find an immediate application to the measurement of the 250-MHz output of the femtosecond combs.

Index Terms—Allan variance, digital instrumentation, frequency measurement, microwave oscillator, three-cornered hat (TCH).

I. INTRODUCTION

THIS paper is made possible by the simultaneous availability in the same place of three cryogenic sapphire six-channel tracking direct digital synthesizer (TDDS) for the

measurement of time fluctuations, and the know-how of clock statistics.

We demonstrate the first frequency-stability measurement of the three CSOs, taken simultaneously at the 10-GHz frequency of the oscillator loop and at the 100-MHz output of the dedicated synthesizer. The CSOs exhibit short-term fluctuations from 2×10^{-16} to parts in 10^{-15} (Allan deviation, ADEV) for measurement time τ between 1 s and 1 day. The synthesizer introduces a very small degradation to the purity of the microwave signal, and only for shortest τ, 1 min or less. Thus, the CSO exceeds by two orders of magnitude the short-term stability of H masers and other commercial atomic standards. The measurement of such CSOs is a challenging task because the target is significantly lower than the background noise of commercial instruments and is possible only by comparing three similar units. For this reason, the measurement was done, until now, by beating the microwave outputs, with no synthesizer [1], [2]. The use of $\approx$1–10-MHz beat notes relaxes the noise requirement for the instrument by 60–80 dB. In this paper, we demonstrate the direct ADEV measurement of the individual CSO at 100 MHz, with no need for the beat note method.

A reliable and stable signal is of paramount importance in strategic facilities, where the short-term stability (up to a few hours) of the Hydrogen maser is not sufficient. For example, the accuracy of very long baseline interferometry [3]–[5] would be improved with the use of the CSO. The CSOs have been developed at JPL for the Cassini mission [6] and used by ESA in the space station in Malargüe, Argentina [7]. The CSO proved to be the best flywheel for the Cesium fountains in a time scale [8]–[10], and the benefit for the ground stations of the Global Navigation Satellite System is just obvious. A 100-MHz reference is easy to distribute with regular low-temperature-coefficient Heliax cables (1 ppm/°C). Attenuation (2.7 dB/100 m at 100 MHz for 1/2" cables) limits the range to a few hundred meters. By contrast, microwave and optical signals are complex to distribute and, at the present time, do not fit the general requirements for a continuously running facility. The femtosecond laser locked to a Fabry–Pérot cavity can provide a VHF signal with stability in competition with the CSO. Albeit the laser technology will probably win in the long run thanks to the optical clocks, the reliability is still far from the requirements mentioned.

The TDDS [11] is a radically new concept in frequency metrology. In short, six direct digital synthesizers (DDSs)

are each phase-locked to one input signal, extracting the phase information from the phase-control word. At once, this eliminates the complexity of the dual-mixer system [12]–[14], mitigates the thermal instability by using only wideband components and filtering numerically at the output, and enables the simultaneous measurement of the six inputs at quite different frequencies scattered in a wide range (presently, 5–400 MHz). Our TDDS [15] exhibits a background noise of $1.5 \times 10^{-14}/\tau$ (ADEV) per channel at 100 MHz.

The most common approach for phase measurement with digital methods starts from sampling and digitizing the input signal [16]–[20]. The TDDS is superior to this approach because the phase noise of analog-to-digital converters (ADCs) is higher than that of a DDS [21]–[23] and because the digital signal processing is done at the speed of the phase fluctuations instead of at the carrier frequency.

It is worth mentioning that the general literature on the measurement of the Allan variance, and on the comparison of multiple clocks, is surprisingly old. The digital methods are recent because of the availability of fast ADCs and still limited by the flicker noise of the ADCs, which is of the order of -110 dBrad^2 (power spectral density at 1 Hz) in the best cases [22], [23]. This is 20–30 dB higher than a double-balanced mixer. Given a very small number of labs that have the technology of the CSO, the measurement at the 100-MHz output is a rather new problem.

We use the two-sample covariance [24] to reject the TDDS noise averaging on a large number of measures. The covariance is superior to the three-cornered hat (TCH) method in that the background noise converges to zero by theorem, and thus there is no need for sophisticated analysis to estimate and compensate the background noise. Finally, we compare the results to the traditional TCH with the beat note method.

In Sections II–VII, we go through the three tools, CSO, TDDS, and covariance, we describe the experiment, and we discuss the results.

II. CRYOGENIC SAPPHIRE OSCILLATOR

The CSO is a long-term project started in Besançon 25 years ago. A few laboratory prototypes have been built, demonstrating a stability (overlapped ADEV, drift removed) of parts in 10^{-16} floor, parts in 10^{-15} for τ up to 1 day, and a sufficient reliability for 1–2 years of unattended operation. The complete machine consists of the oscillator (in strict sense), the refrigerator, a dedicated synthesizer, and control equipment [25]. See also [26] for a general review. The CSO is based on the following ideas.

The sapphire (Al_2O_3) monocrystal is an ideal material for dielectric resonators because of its low loss, and good mechanical and chemical properties. It is hard (9 Mohs, by definition), stiff and stable, and suitable to precision machining. Our 10-GHz whispering-gallery-mode resonators achieve routinely Q of the order of two billions at liquid-He temperature. The value depends on the electromagnetic mode, crystal size, and specimen.

Certain modes exhibit a natural turning point of the resonant frequency at 5–8 K, just above the He boiling point [27]. This is due to the presence of paramagnetic Cr^{3+}, Fe^{3+}, and Mo^{3+}
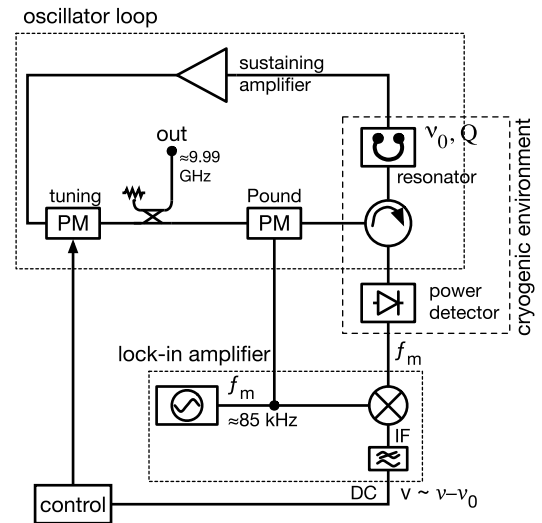

Fig. 1. Block diagram of the Pound–Galani cryogenic oscillator.

impurities. The growth process provides the right amount of such impurities, with very similar results over at least two growth technologies and manufacturers [28]. The WGH$_{15,0,0}$ mode of a cylinder of 54-mm diameter and 30-mm height resonates close to 10 GHz. The electromagnetic energy is confined in $\approx 1/10$ of the volume in the outer perimeter, indeed in a volume large enough to keep the resonator in a highly linear regime and to provide the high stability.

The resonator is cooled by a two-stage pulse-tube refrigerator [29], and temperature is stabilized by heating to the turning point within 100 μK. The refrigerator is a special design exhibiting low vibes (<1 μm$_{pp}$).

The oscillator is a Pound–Galani, shown in Fig. 1. The Pound scheme [30] makes use of an auxiliary oscillator frequency-stabilized to the main resonator. The resonator is used in the reflection mode, with phase modulation sidebands out of the resonator bandwidth. The main virtues are the reduction of flicker and drift noises thanks to the ac modulation, and the inherent rejection of the fluctuations in the electrical path from the resonator to the detector. The Galani version [31] implements the auxiliary oscillator using the same resonator, in the transmission mode. This solution provides higher Q and a better resonator stability, as compared with an external voltage-controlled oscillator. In addition, the Galani version is suitable to a simple and effective design of the frequency control thanks to a pole-zero cancelation in the loop function, which results from using the same resonator in both oscillator and control. A power control, not shown in Fig. 1, keeps the power constant within 3 ppm, preventing the fluctuations from degrading the stability via radiation pressure [32] and self-heating.

The choice of the oscillation frequency is a key point of the design. Modeling and machining the sapphire limits the initial accuracy to ≈ 1 MHz, and we prefer not to go through measurement-and-trimming iterations to achieve a more accurate value. Such tolerances are useful in that they prevent electromagnetic interference between the oscillators. Therefore, we machine the sapphire for $\nu_0 = 10$ GHz $- \Delta$, where $\Delta = 10 \pm 5$ MHz, and the synthesizer adds Δ. However,
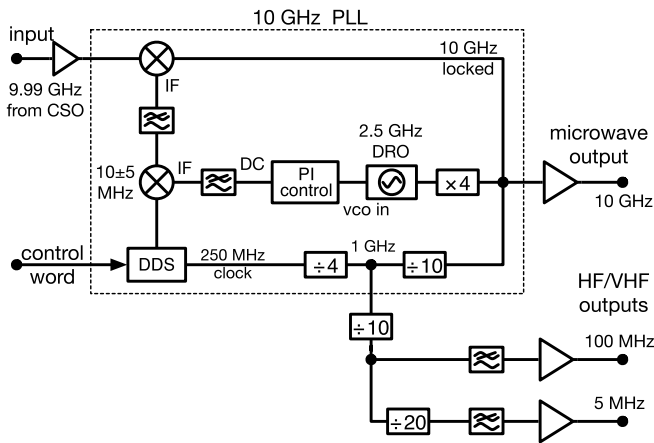
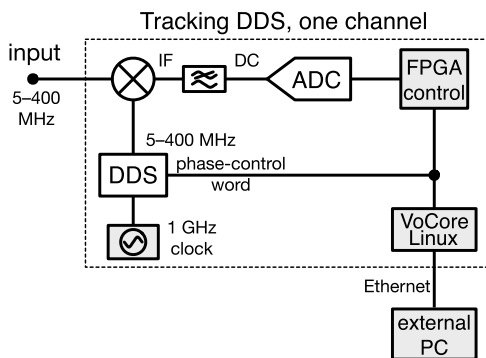Fig. 2.   Block diagram of the dedicated synthesizer.

Fig. 3.   Block diagram of the TDDS. The complete machine consists of six equal channels. The gray blocks are shared by the six channels.


Fig. 4.   Photograph of the instrument.

frequencies close to 5/10/15 MHz must be avoided for interference immunity. It is wise to keep a margin of at least 10 kHz. The design turns out to be quite simple (see Fig. 2), yet achieving μHz resolution at 10 GHz and 10^{-4} tuning range. The specs for the time fluctuation of the DDS are relaxed by the ratio $\nu_0/\Delta \approx 10^3$. Thanks to this leverage factor, the DDS contributes $\approx 10^{-17}$ to the stability.

III. MULTICHANNEL TRACKING DDS

Our instrument is a multichannel real-time phasemeter based on the TDDS technique [11], [15]. Fig. 3 shows the scheme of one channel, and Fig. 4 shows the complete machine. A proportional–integral control implemented in


Fig. 5.   Thermal image of the TDDS.

FPGA phase-locks the DDS to the input acting on the phase-control word. Controlling the phase, instead of the frequency, requires a phase accumulator that counts the multiple cycles. The discriminator is a double-balanced mixer Mini Circuits ADE-1. The acquisition and lock range is 5–400 MHz. The lower limit is set by the mixer, and the upper limit by the sampling frequency of the DDS (1 GHz).

The phase error is digitized on 16 bits at 500 kS/s. This is also the sampling frequency of the digital control. The loop bandwidth is of 2–20 kHz, depending on internal parameters. Anyway, the value is not critical. Within the feedback-loop bandwidth, the phase-control word is equal to the phase difference between the input and the local clock. The FPGA guarantees that all the measures are simultaneous. This is necessary to cancel the fluctuation of the internal 1-GHz clock, common to all the channels.

The phase error is converted into a stream of phase-time data $\mathsf{x}(t)$, low-passed at the cutoff frequency $f_H = 5$ Hz, sampled at 10 S/s, and transferred to the external PC by a VoCore2 Linux computer [33].

For the best stability, we care about a low dissipated power and geometrical symmetry. The DDS is an Analog Devices AD9912, which has a typical dissipation of 650 mW, smaller than that of other high-frequency DDSs. The FPGA is an Altera CYCLONE III with 25 000 logic elements, 30% of which are actually used. For low dissipation, the 125-MHz clock is downconverted to 10 MHz using an internal PLL. The VoCore2 dissipates 1 W. The complete instrument takes 10-W power from a single +12-V supply. Fig. 5 shows the thermal image of the six channels. The temperature sensitivity is of about 1 ps/K (B-type uncertainty) on each channel, limited by the mixer.

Fig. 6 shows the noise budget of the TDDS at $\nu_{\mathrm{DDS}} = 100$ MHz carrier. At low Fourier frequency, the $1/f$ phase noise of the DDS ($\mathsf{b}_{-1} = -110$ dBrad^2) dominates, being 20 dB higher than the noise of the clock distribution and ≈ 30 dB higher than the noise of the mixer and the amplifier. The DDS flicker noise is of the time type. The $1/f$ term of $S_{\mathsf{x}}(f)$ is $\mathsf{k}_{-1} = \mathsf{b}_{-1}/(2\pi\nu_{\mathrm{DDS}})^2 = 2.5 \times 10^{-29}$ s^2, independent of ν_{DDS}, while b_{-1} scales proportionally to ν_{DDS}. Conversely, the $1/f$ noise of the mixer and the amplifier is of the phase type, with b_{-1} independent of ν_{DDS} in a
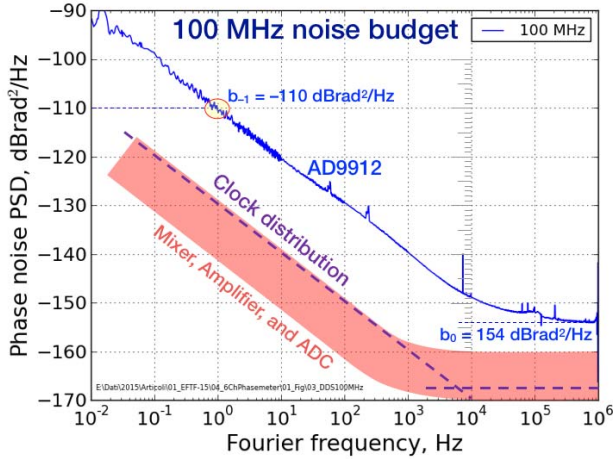
Fig. 6. Noise budget of the TDDS at 100 MHz.

Fig. 7. Block diagram of the experiment.

wide range, and $k_{-1} \propto 1/\nu_{DDS}$. The DDS $1/f$ noise is dominant from $\nu_{DDS} = 10$ MHz to the maximum carrier frequency. At $\nu_{DDS} < 10$ MHz, the $1/f$ noise of the mixer and the amplifier is no longer negligible, and the background noise starts degrading. The quantity $\sqrt{k_{-1}} = 5$ fs is the flicker fluctuation of the DDS. Converting the phase noise into frequency stability, with $f_H = 5$ Hz, we get $\sigma = 1.5 \times 10^{-14}$ at $\tau = 1$ s and 100-MHz carrier, with the slope close to $1/\tau$. This is the background noise of the instrument, one channel.

IV. STATISTICS

The Allan variance can be written as

$$\sigma_y^2(\tau) = \mathbb{E}\left\{ \frac{(\overline{y}_2 - \overline{y}_1)^2}{2} \right\} = \mathbb{E}\left\{ \frac{(x_2 - 2x_1 + x_0)^2}{2\tau} \right\} \quad (1)$$

where $\mathbb{E}\{\}$ is the mathematical expectation, $\overline{y}$ is the fractional frequency fluctuation averaged over the measurement time τ, and the subscripts "1" and "2" refer to contiguous time slots. The alternate formulation in terms of the phase time x sampled at regular intervals τ relates to the "second difference" method [34], equivalent to the method used here. Making $x_2 - 2x_1 + x_0$ explicit emphasizes that there is no dead time in $\overline{y}_2 - \overline{y}_1$. Since all the variances in this paper are $\sigma_y^2(\tau)$, we will omit τ and the subscript y with no ambiguity. We will have numerous occurrences of $\overline{y}_2 - \overline{y}_1$ in the same formula, referring to different oscillators and instruments. It is, therefore, appropriate to define the *fractional frequency difference*

$$z = \overline{y}_2 - \overline{y}_1 = \frac{1}{\tau}(x_2 - 2x_1 + x_0). \quad (2)$$

Accordingly, the Allan variance is

$$\sigma^2 = \frac{1}{2}\mathbb{E}\{z^2\}. \quad (3)$$

The obvious extension to the two-sample covariance is

$$\sigma_{ij} = \frac{1}{2}\mathbb{E}\{z_i z_j\}. \quad (4)$$

We denote the oscillators with the subscripts A, B, and C (Capital), the instrument inputs with a, b, and c (lowercase), and the instrument readouts with α, β, and γ (Greek), as shown in Fig. 7. Following the path "A," the oscillator's z_A

is sent to the instrument, which contributes z_a, and delivers the readout:

$$z_\alpha = z_A + z_a. \quad (5)$$

This is similar to the Shannon channel, where the received signal is equal to the transmitted signal plus noise.

All our statistical measurements rely on the hypothesis that all oscillators and instrument channels have statistically independent noise processes, so that $\mathbb{E}\{z_i z_j\} = 0$ for $i \neq j$. Accordingly

$$\sigma_\alpha^2 = \sigma_A^2 + \sigma_a^2 \quad (6)$$

and likewise for the other channels.

A. Three-Cornered Hat Method

The TCH method is a well-established method to measure the variance of each oscillator by comparing three units [35]. The method is better understood by provisionally admitting that the counter noise is low enough to enable the direct measurement. Because the internal reference of the instrument does not have sufficient stability for absolute measurements, we rely on the differences $x_{\beta-\alpha} = x_\beta - x_\alpha$, $x_{\gamma-\beta} = x_\gamma - x_\beta$, and $x_{\alpha-\gamma} = x_\alpha - x_\gamma$ taken simultaneously. Using the corresponding fractional frequency differences, we calculate the variances

$$\sigma_{\beta-\alpha}^2 = \frac{1}{2}\mathbb{E}\{(z_\beta - z_\alpha)^2\} = \sigma_B^2 + \sigma_A^2 + \sigma_b^2 + \sigma_a^2 \quad (7)$$

$$\sigma_{\gamma-\beta}^2 = \frac{1}{2}\mathbb{E}\{(z_\gamma - z_\beta)^2\} = \sigma_C^2 + \sigma_B^2 + \sigma_c^2 + \sigma_b^2 \quad (8)$$

$$\sigma_{\alpha-\gamma}^2 = \frac{1}{2}\mathbb{E}\{(z_\alpha - z_\gamma)^2\} = \sigma_A^2 + \sigma_C^2 + \sigma_a^2 + \sigma_c^2. \quad (9)$$

This is immediately seen by expanding the terms inside the $\mathbb{E}\{\}$ operator. In (7), $(z_\beta - z_\alpha)^2 = (z_B + z_b - z_A - z_a)^2$. The four "$z$"s are statistically independent (separate oscillators and separate channels of the instrument), and thus $\mathbb{E}\{z_i z_j\} = 0$ for all the cross terms $2z_i z_j$, $i \neq j$. Likewise, (8) and (9).

Note that the subscripts form a group $A \to B \to C \to A \ldots$, and thus we can derive all equations for B from the

homologous equation for A by replacing $A \to B$, $a \to b$, and $\alpha \to \beta$ in the same equation for A. Likewise for C from B and for A from C. This is clearly seen in (7)–(9). Hereafter, we will write only the equations for the oscillator A, because the other two equations can be written with the above-mentioned rule.

Solving (7)–(9), the variance of the oscillator A is

$$\sigma_A^2 = \frac{1}{2}\left(\sigma_{\beta-\alpha}^2 - \sigma_{\gamma-\beta}^2 + \sigma_{\alpha-\gamma}^2\right) - \frac{1}{2}\sigma_a^2. \tag{10}$$

The terms σ_a^2, σ_b^2, and σ_c^2 in (7)–(9) are the white and flicker phase noise introduced by the instrument, and thus they roll off as $1/\tau^2$. Nonetheless, they are annoying because no instrument has noise low enough to measure the short-term stability of our cryogenic oscillators. For this reason, we have to combine the TCH with the beat method [see Fig. 7 (bottom)]. Beating the ν_0 output (≈ 10 GHz) down to the megahertz region relaxes the stability requirement by a factor ν_{beat}/ν_0, which is $10^{-4} \ldots 10^{-3}$.

B. Covariance Method

The two-sample covariance method is a different way to process the same data, $x_{\beta-\alpha}$, $x_{\gamma-\beta}$, and $x_{\alpha-\gamma}$, measured simultaneously [24], [36]. In this case, we exploit the product of fractional frequency differences

$$\frac{1}{2}\mathbb{E}\{(z_\beta - z_\alpha)(z_\gamma - z_\alpha)\} = \frac{1}{2}\mathbb{E}\{z_\alpha^2\} = \sigma_A^2 + \sigma_a^2 \tag{11}$$

for the oscillator A, and likewise for the oscillators B and C. The proof is immediate. Assuming that z_α, z_β, and z_γ are statistically independent (separate oscillators and instrument channels), all the cross terms in $\mathbb{E}\{(\ldots)(\ldots)\}$ are equal to zero, and only the square term remains. The result is the same of the TCH method and not suitable to our purposes for the same reasons.

The background noise can be rejected by using two channels per oscillator. Let us start with A, which splits into a "prime" channel and a "second" channel. Thus, $\mathbb{E}\{(z_\beta - z_\alpha)(z_\gamma - z_\alpha)\}$ of (11) becomes $\mathbb{E}\{(z_\beta - z_{\alpha'})(z_\gamma - z_{\alpha''})\}$. Assuming that $z_{\alpha'}$ and $z_{\alpha''}$ are statistically independent, it holds that $\mathbb{E}\{z_{\alpha'}z_{\alpha''}\} = \mathbb{E}\{z_A^2\}$. This solves the problem.

A more efficient use of the hardware is possible. Replacing $z_\beta \to (z_{\beta'} + z_{\beta''})/2$ and $z_\gamma \to (z_{\gamma'} + z_{\gamma''})/2$ results in lower background noise and, hence, in faster convergence

$$\frac{1}{2}\mathbb{E}\left\{\left(\frac{z_{\beta'} + z_{\beta''}}{2} - z_{\alpha'}\right)\left(\frac{z_{\gamma'} + z_{\gamma''}}{2} - z_{\alpha''}\right)\right\} = \sigma_A^2. \tag{12}$$

Averaging (12) with the same after interchanging α' with α'' results in even lower background noise. Accordingly, the final equation we use is

$$\frac{1}{4}\mathbb{E}\left\{\left(\frac{z_{\beta'} + z_{\beta''}}{2} - z_{\alpha'}\right)\left(\frac{z_{\gamma'} + z_{\gamma''}}{2} - z_{\alpha''}\right) \right. $$
$$\left. + \left(\frac{z_{\beta'} + z_{\beta''}}{2} - z_{\alpha''}\right)\left(\frac{z_{\gamma'} + z_{\gamma''}}{2} - z_{\alpha'}\right)\right\} = \sigma_A^2. \tag{13}$$

C. Averaging on a Finite Data Record

The mathematical expectation is replaced with the average on a finite time series of m samples. We use the overlapped Allan variance in all the cases. The ultimate limit to the detection of fractional frequency fluctuations is

$$\sigma = \frac{\sigma_0}{\tau}\frac{1}{m^{1/4}} \tag{14}$$

where σ_0 is the background at $\tau = 1$ s, m is the number of averages, and the ratio σ_0/τ is the usual "$1/\tau$" law for the white and flicker phase noise. In turn, m results from the duration $\mathcal{T}$ of the time series according to $m = \mathcal{T}/\tau$. Combining the latter with (14), we get

$$\sigma = \frac{\sigma_0}{\mathcal{T}^{1/4}}\frac{1}{\tau^{3/4}}. \tag{15}$$

The averaging process uses a large amount of samples at the shorter τ, where the background noise is higher, and progressively smaller amount of samples at longer τ. According to (14), it takes $m = 10^4$ (2 H 47 M) to reduce the background noise by a factor 10, from 2.1×10^{-14} to 2.1×10^{-15} at $\tau = 1$ s.

V. Experiment

The block diagram of the experiment follows Fig. 7. The 100-MHz output of the three CSOs is measured with the TDDS in the six-channel mode, using the covariance. The 10-GHz outputs are beaten down to HF and measured with a multichannel counter. The latter is a dedicated "Time and Frequency Monitor" made by K&K Messtechnik (now Lange-Electronic [37]), originally described in [38] and [39]. The data averaged on $\tau = 1$ s are directly available at the output of the K&K. With the TDDS, the average on $\tau = 1$ s is obtained by the decimation of the 10-S/s output stream. We collected all the data measured simultaneously for a duration of 4.7×10^5 s (5.5 days).

The cutoff frequency f_H is 5 Hz for the TDDS and 0.5 Hz for the K&K. This difference is irrelevant because there is no white phase noise in our results (see Section VI), and thus f_H does not get in the noise equations [40].

We had only two synthesizers, and thus we used a frequency divider (Hittite) instead. This divider, at the output of the oscillator B, does not have a thermal shield.

The oscillator B had the power control not operating properly. This problem was discovered when the experiment was already running.

The most common correlated phenomena, breaking the hypothesis of statistically independent noise processes, are the microwave leakage and temperature fluctuations of the environment. The microwave leakage is in principle absent in our CSOs because the resonator bandwidth is of the order of 10 Hz, a few orders of magnitude smaller than the frequency difference between the oscillators. Crosstalk in the instruments is a lesser problem because it impacts on the phase, instead of on the frequency. Nonetheless, we set the synthesizers at three different frequencies slightly off the nominal value of 100 MHz. We observed that spurs are also reduced in this way. The HF beat notes are substantially immune from leakage, being well separated.

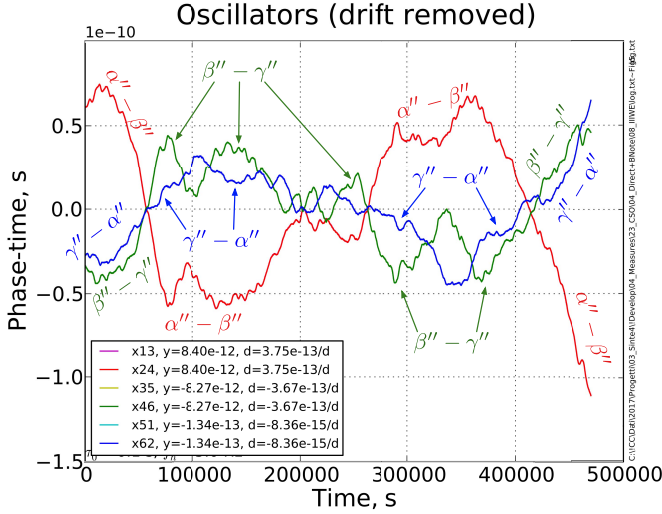Fig. 8. Experiment in its environment. The Helium pumps are in the next room.

Fig. 9. Time fluctuation of the differences $(x_{\alpha''} - x_{\beta''})$, $(x_{\beta''} - x_{\gamma''})$, and $(x_{\gamma''} - x_{\alpha''})$ taken from the TDDS.

Fig. 10. Allan deviation of the difference between oscillators.

The experimental setup is shown in Fig. 8. The He pumps are located in a nearby room. The thermal fluctuations are strongly reduced by a sophisticated air-conditioning installation. A proportional–integral control guarantees a temperature of 22 °C ± 0.5 °C, with a maximum drift of 0.2 °C/H, and a humidity of 50% ± 10%. The operators are not present in the room during the measurements.

VI. RESULTS

We analyze the results step by step going through consistency checks. This is necessary because the two-sample covariance is still being little used, and thus we cannot take benefit from the general experience.

The first result (see Fig. 9) is the time fluctuation of the differences $(x_{\alpha''} - x_{\beta''})$, $(x_{\beta''} - x_{\gamma''})$, and $(x_{\gamma''} - x_{\alpha''})$ taken from the outputs α'', β'', and γ'' of the TDDS. At the scale shown, the same differences taken from the outputs α', β', and γ' overlap and provide no useful information. A visual comparison between Fig. 9 and a set of simulated random time series indicates that the oscillators are mainly affected by the
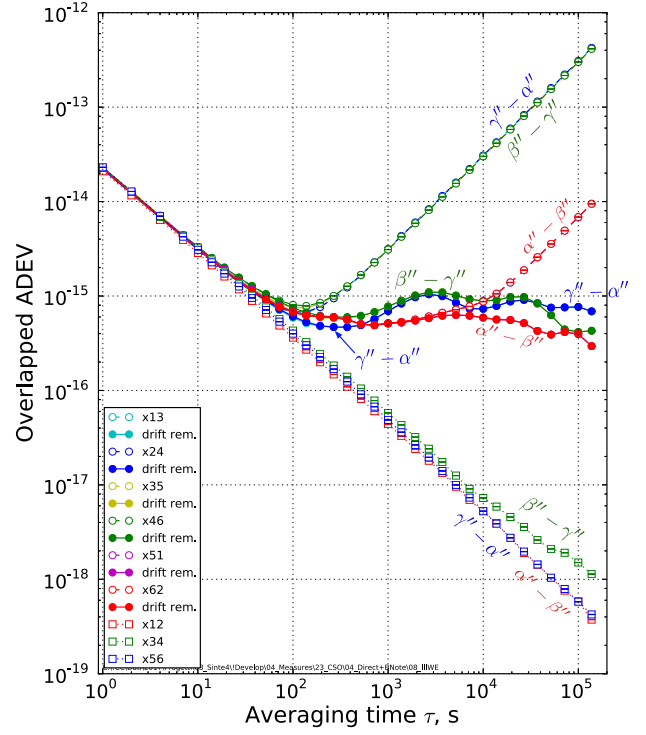
white FM and flicker FM noises, with some drift showing up. The peak-to-peak difference is within 100 ps in the first four days and increases slightly afterward.

The curves are clearly correlated. By visual inspection shown in Fig. 9, the following sums reveal correlation:

$$(x_{\alpha''} - x_{\beta''}) + (x_{\beta''} - x_{\gamma''}) \approx \text{`small'} \quad C, A \text{ anticorrelated}$$
$$(x_{\alpha''} - x_{\beta''}) + (x_{\gamma''} - x_{\alpha''}) \approx \text{`small'} \quad B, C \text{ correlated}$$
$$(x_{\beta''} - x_{\gamma''}) + (x_{\gamma''} - x_{\alpha''}) \approx \text{`large'} \quad A, B \text{ anticorrelated}.$$

We believe that this is still a thermal effect on the oscillators, despite the careful conditioning of the room.

The second result (see Fig. 10) is the overlapped Allan deviation of some relevant differences, taken at the output of the TDDS. In Fig. 10, we identify three groups of plots. The first group (the squares, going down to 10^{-18}) represents $\sigma_{\alpha'-\alpha''}$, $\sigma_{\beta'-\beta''}$, and $\sigma_{\gamma'-\gamma''}$. Each curve is the difference between the two channels measuring the same oscillator. The oscillator fluctuation cancels, and the curve gives the total fluctuation of a pair of channels. The value $\sigma_y = 2.1 \times 10^{-14}$ at $\tau = 1$ s is consistent with the flicker phase noise of the DDS ($b_1 = -110$ dBrad[2]; see Fig. 6), within 1 dB. The curves $\sigma_{\alpha'-\alpha''}$ and $\sigma_{\gamma'-\gamma''}$ match the prediction based on the dominance of flicker phase noise, $\sigma^2 = h_{-1}(3\gamma - \ln(2) + 3\ln(2\pi f_H \tau))/(2\pi\tau)^2$. The third curve, $\sigma_{\beta'-\beta''}$, gets slightly higher beyond $\tau \approx 1000$ s. This is due to a defect in channel β', clearly identified but still not understood.

The second group of plots (the thick dots in Fig. 10) is $\sigma_{\alpha''-\beta''}$, $\sigma_{\beta''-\gamma''}$, and $\sigma_{\gamma''-\alpha''}$ of the 100-MHz signal, drift removed. This two-sample deviation is dominated by the background noise for $\tau < 100$ s. Beyond, the curves converge to σ_{A-B}, σ_{B-C}, and σ_{C-A}, that is, the combined fluctuations of two oscillators each.
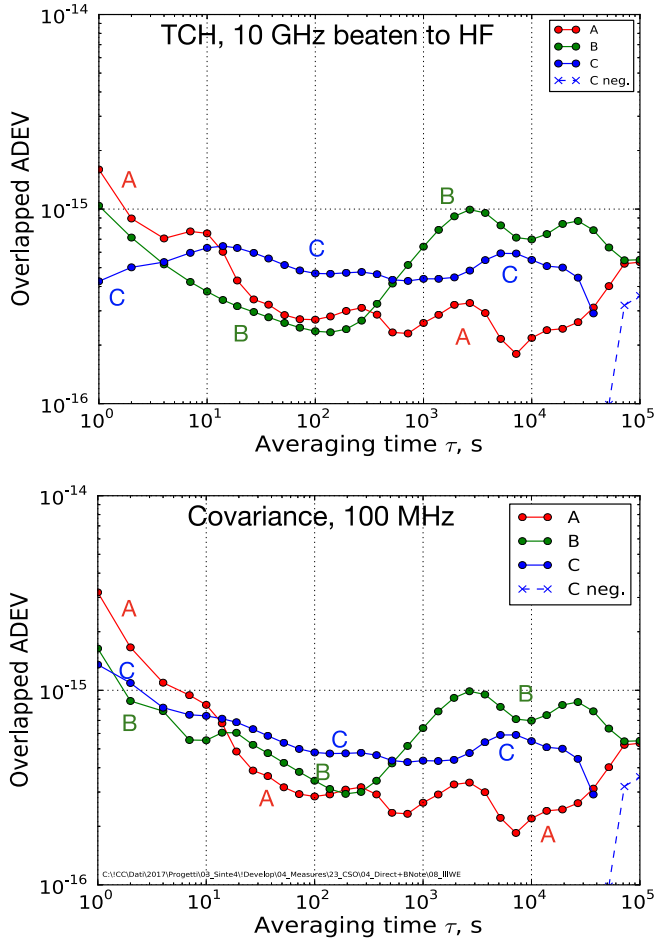
Fig. 11. Allan deviation of the three oscillators measured with the TCH method on the HF beat notes (top) and with the covariance method on the 100-MHz outputs (bottom).
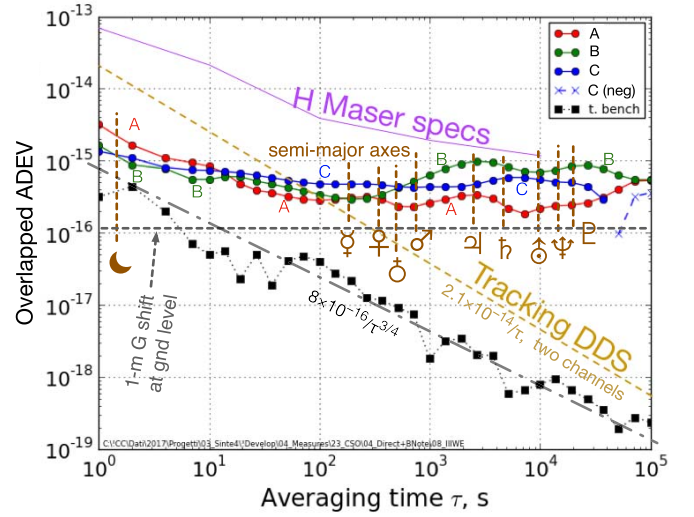

Fig. 12. Allan deviation of our oscillators (100-MHz output) compared with some relevant quantities discussed in the text.

to the insufficient number of data or to the presence of slow correlated terms, probably induced by the temperature.

VII. CONCLUSION

Fig. 12 summarizes the relevant results obtained with the TDDS and the covariance at the 100-MHz output of the CSOs.

The gold dashed line shows the background noise of the TDDS, accounting for two channels. This is fair because the simplest measurement, with no correlation, takes two channels. The background noise is low enough for the measurement of H masers and the other classical atomic standards. In addition, the CSOs can be measured in the correlation mode, using two channels per oscillator and the covariance method.

The statistical limit is set by the duration $\mathcal{T}$ of the time series according to (15). Taking $\sigma_0 = 2.1 \times 10^{-14}$, which is the combined background of two channels, we find $\sigma = 8 \times 10^{-16}/\tau^{3/4}$. The theoretical prediction (gray dashed-dotted line) fits well the observed residuals (dotted line with black squares).

The RF/microwave hardware can be a serious limitation at this level. The phase fluctuation of the oscillator C, $\sigma = 1.5 \times 10^{-15}$ at $\tau = 1$ s, corresponds to a time stability of 1.5 fs or a length stability of 400 nm on a coaxial cable (velocity factor of 0.88 for a 1/2" Heliax cable).

The CSO instability is in the low 10^{-15} at 1 s and below 10^{-15} for $\tau \geq 10$ s up to one day. For reference, the gravitational shift g/c^2 is of 1.09×10^{-16} when a clock is raised by 1 m at the ground level.

The highest stability is seen for τ between seconds and hours, which is the travel time of the light at interplanetary distances (semimajor axes of the planet orbits in Fig. 12). This feature makes the CSO an ideal clock for applications related to the exploration of the solar system.

The third group of plots (the circles in Fig. 10) is the same as the second, but the drift is not removed. The drift is clearly visible on the right-hand side of the plot, where the curves are proportional to τ. Two of these curves, σ_{A-B} and σ_{B-C}, reveal a significantly higher drift, due to the oscillator B. This is a consequence of the power control not working. Using $\sigma^2 = (1/2)D_y^2\tau^2$, the drift of B is $D_y = 4.5 \times 10^{-18}$/s or 3.9×10^{-13}/day.

Fig. 11 shows the ADEV of the three oscillators measured with the TCH method on the beat notes (top) and with the covariance method on the 100-MHz outputs (bottom). The drift is removed in both the cases. The results are substantially equivalent for $\tau > 100$ s, while the noise of the synthesizers shows up for shorter τ. The Allan deviation plot of the oscillator B has a bump at $\tau = 20$ on the covariance plot [see Fig. 11 (bottom)] and is not present on the TCH plot [see Fig. 11 (top)]. This is probably due to the frequency divider, which has no thermal shield. The oscillator B also has irregular behavior and significantly higher instability than the other oscillators for $\tau > 1000$ s. This is ascribed to the failure in the power control. The blue line (Oscillator C) is dashed for $\tau > 5 \times 10^4$ s because a negative covariance appears in the evaluation, and the result makes no sense. This can be ascribed

REFERENCES

[1] J. G. Hartnett, N. R. Nand, C. Wang, and J.-M. Le Floch, "Cryogenic sapphire oscillator using a low-vibration design pulse-tube cryocooler: First results," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 57, no. 5, pp. 1034–1038, May 2010.

[2] C. Fluhr, S. Grop, B. Dubois, Y. Kersalé, E. Rubiola, and V. Giordano, "Characterization of the individual short-term frequency stability of cryogenic sapphire oscillators at the 10^{-16} level," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 6, pp. 915–921, Jun. 2016.

[3] N. R. Nand, J. G. Hartnett, E. N. Ivanov, and G. Santarelli, "Ultra-stable very-low phase-noise signal source for very long baseline interferometry using a cryocooled sapphire oscillator," *IEEE Trans. Microw. Theory Techn.*, vol. 59, no. 11, pp. 2978–2986, Nov. 2011.

[4] S. Doeleman, T. Mai, A. E. E. Rogers, J. G. Hartnett, M. E. Tobar, and N. Nand, "Adapting a cryogenic sapphire oscillator for very long baseline interferometry," *Astronom. Soc. Pacific*, vol. 123, no. 903, pp. 582–595, May 2011.

[5] R. Bara *et al.*, "Generation of 103.75 GHz CW Source With 5.10^{-16} frequency instability using cryogenic sapphire oscillators," *IEEE Microw. Wireless Compon. Lett.*, vol. 22, no. 2, pp. 85–87, Feb. 2012.

[6] G. J. Dick and R. T. Wang, "Stability and phase noise tests of two cryocooled sapphire oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 47, no. 5, pp. 1098–1101, Sep. 2000.

[7] V. Giordano *et al.*, "New-generation of cryogenic sapphire microwave oscillators for space, metrology, and scientific applications," *Rev. Sci. Instrum.*, vol. 83, no. 8, p. 085113, Aug. 2012.

[8] M.-S. Heo *et al.*, "Drift-compensated low-noise frequency synthesis based on a cryocso for the KRISS-F1," *IEEE Trans. Instrum. Meas.*, vol. 66, no. 6, pp. 1343–1348, Jun. 2016.

[9] M. Abgrall *et al.*, "High-stability comparison of atomic fountains using two different cryogenic oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 8, pp. 1198–1203, Aug. 2016.

[10] T. Ikegami, K. Watabe, S. Yanagimachi, A. Takamizawa, and J. G. Hartnett, "Autonomous cryogenic sapphire oscillators employing low vibration pulse-tube cryocoolers at NMIJ," *J. Phys. Conf. Ser.*, vol. 723, no. 1, p. 012032, 2016.

[11] C. E. Calosso, "Tracking DDS in time and frequency metrology," in *Proc. Eur. Freq. Time Forum*, Prague, Czech Republic, Jul. 2013, pp. 747–749.

[12] D. W. Allan and H. Daams, "Picosecond time difference measurement system," in *Proc. Int. Freq. Control Symp.*, Atlantic City, NJ, USA, May 1975, pp. 404–411.

[13] D. W. Allan, "Report on the dual mixer time difference system (DMTD) built for the time-domain measurements associated with phase 1 of GPS," NBS (now NIST), Boulder, CO, USA, Tech. Note 75-827, Jan. 1976.

[14] G. Brida, "High resolution frequency stability measurement system," *Rev. Sci. Instrum.*, vol. 73, no. 5, pp. 2171–2174, May 2002.

[15] C. Massimo, G. A. Costanzo, and C. E. Calosso, "6/12-channel synchronous digital phasemeter for ultrastable signal characterization and use," in *Proc. Eur. Freq. Time Forum*, Denver, CO, USA, Apr. 2015, pp. 681–683.

[16] J. Grove, J. Hein, J. Retta, P. Schweiger, W. Solbrig, and S. R. Stein, "Direct-digital phase-noise measurement," in *Proc. Int. Freq. Control Symp.*, Montreal, QC, Canada, Aug. 2004, pp. 287–291.

[17] A. C. Ecker, "A digital method for phase noise measurement," Ph.D. dissertation, Univ. Washington, Seattle, WA, USA, Apr. 2014.

[18] J. A. Sherman and R. Jördens, "Oscillator metrology with software defined radio," *Rev. Sci. Instrum.*, vol. 87, no. 5, p. 054711, May 2016.

[19] R. G. DeVoe, "Measuring the allan variance by sinusoidal fitting," *Rev. Sci. Instrum.*, vol. 89, no. 2, p. 024702, Jan. 2018.

[20] S. J. Yu, E. Fajeau, L. Q. Liu, D. J. Jones, and K. W. Madison, "The performance and limitations of FPGA-based digital servos for atomic, molecular, and optical physics experiments," *Rev. Sci. Instrum.*, vol. 89, no. 2, p. 025107, Feb. 2018.

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

[22] A. C. Cárdenas-Olaya *et al.*, "Noise characterization of analog to digital converters for amplitude and phase noise measurements," *Rev. Sci. Instrum.*, vol. 88, no. 6, p. 065108, Jun. 2017.

[23] A. C. Cárdenas-Olaya, "Digital instrumentation for the measurement of high spectral purity signals," Ph.D. dissertation, Univ. Bourgogne Franche Comté, Politecnico di Turin, Besançon, France, Jul. 2018.

[24] F. Vernotte, C. E. Calosso, and E. Rubiola, "Three-cornered hat versus Allan covariance," in *Proc. Int. Freq. Control Symp.*, New Orleans, LA, USA, May 2016, pp. 282–287.

[25] V. Giordano, S. Grop, C. Fluhr, B. Dubois, Y. Kersalé, and E. Rubiola, "The autonomous cryocooled sapphire oscillator: A reference for frequency stability and phase noise measurements," *J. Phys., Conf. Ser.*, vol. 723, no. 1, p. 012030, Jul. 2016.

[26] C. R. Locke, E. N. Ivanov, J. G. Hartnett, P. L. Stanwix, and M. E. Tobar, "Design techniques and noise properties of ultrastable cryogenically cooled sapphire-dielectric resonator oscillators," *Rev. Sci. Instrum.*, vol. 79, no. 5, pp. 051301-1–051301-12, May 2008.

[27] S. K. Jones, D. G. Blair, and M. J. Buckingham, "Effect of paramagnetic impurities on frequency of sapphire-loaded superconducting cavity resonators," *Electron. Lett.*, vol. 24, no. 6, pp. 346–347, Mar. 1988.

[28] V. Giordano, C. Fluhr, S. Grop, and B. Dubois, "Tests of sapphire crystals manufactured with different growth processes for ultra-stable microwave oscillators," *IEEE Trans. Microw. Theory Techn.*, vol. 64, no. 1, pp. 78–85, Jan. 2016.

[29] C. Wang, "Numerical analysis of 4 K pulse tube coolers: Part II. Performances and internal processes," *Cryogenics*, vol. 37, no. 4, pp. 215–220, 1997.

[30] R. V. Pound, "Electronic frequency stabilization of microwave oscillators," *Rev. Sci. Instrum.*, vol. 17, no. 11, pp. 490–505, Nov. 1946.

[31] Z. Galani, M. J. Bianchini, Jr., R. C. Waterman, Jr., R. Dibiase, R. W. Laton, and J. B. Cole, "Analysis and design of a single-resonator GaAs FET oscillator with noise degeneration," *IEEE Trans. Microw. Theory Techn.*, vol. MTT-32, no. 12, pp. 1556–1565, Dec. 1984.

[32] S. Chang, A. G. Mann, A. N. Luiten, and D. G. Blair, "Measurements of radiation pressure effect in cryogenic sapphire dielectric resonators," *Phys. Rev. Lett.*, vol. 79, no. 11, pp. 2141–2144, Sep. 15, 1997.

[33] *VoCore2, the Coin-Sized Linux Computer*. [Online]. Available: http://vocore.io/v2.html

[34] C. A. Greenhall, "A method for using a time interval counter to measure frequency stability," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 36, no. 5, pp. 478–480, Sep. 1989.

[35] J. E. Gray and D. W. Allan, "A method for estimating the frequency stability of an individual oscillator," in *Proc. Int. Freq. Control Symp.*, Atlantic City, NJ, USA, May 1974, pp. 243–246.

[36] D. Fest, J. Groslambert, and J.-J. Gagnepain, "Individual characterization of an oscillator by means of cross-correlation or cross-variance method," *IEEE Trans. Instrum. Meas.*, vol. IM-32, no. 3, pp. 447–450, Sep. 1983.

[37] Lange-Electronic GmbH. Gernlinden, Germany. *Lange-Electronic GmbH Web Site*. Accessed: Jul. 20, 2018. [Online]. Available: http://lange-electronic.de/

[38] G. Kramer and W. Klische, "Extra high precision digital phase recorder," in *Proc. Eur. Freq. Time Forum*, Guildford, U.K., Apr. 2004, pp. 595–602.

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

[40] 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.

Claudio E. Calosso, photograph and biography not available at the time of publication.

François Vernotte, photograph and biography not available at the time of publication.

Vincent Giordano, photograph and biography not available at the time of publication.

Christophe Fluhr, photograph and biography not available at the time of publication.

Benoît Dubois, photograph and biography not available at the time of publication.

Enrico Rubiola, photograph and biography not available at the time of publication.