Received 20 May 2024; revised 30 August 2024; accepted 21 October 2024. Date of publication 28 October 2024; date of current version 11 November 2024.

Digital Object Identifier 10.1109/OJUFFC.2024.3487147

Direct Digital Simultaneous Phase-Amplitude Noise and Allan Deviation Measurement System

MARCO POMPONIO^{®1,2}, ARCHITA HATI^{®1} (Member, IEEE), AND CRAIG NELSON^{®1} (Member, IEEE)

¹National Institute of Standards and Technology, Boulder, CO 80305 USA ²Department of ECEE, University of Colorado Boulder, Boulder, CO 80305 USA CORRESPONDING AUTHOR: M. POMPONIO (marco.pomponio@nist.gov; marco.pomponio@colorado.edu)

ABSTRACT In this paper, we present a direct digital measurement system capable of simultaneously measuring phase noise, amplitude noise, and Allan deviation with and without cross-correlation. The residual phase noise of the single-channel system achieves \mathscr{L} (1 Hz) = -143 dBc/Hz for a 10 MHz input signal and an Allan deviation noise floor of 3.2×10^{-15} at 1 second averaging time (τ). The system's performance improves as expected with cross-correlation, resulting in an average-limited residual white noise floor of -185 dBc/Hz after only a few minutes of averaging, an improvement of 30 dB compared to a single-channel system. It also reaches an average limited flicker phase noise floor of \mathscr{L} (1 Hz) = -160 dBc/Hz within two days, with an Allan deviation of 5×10^{-16} @ $\tau = 1$ second. To our knowledge, this represents the lowest noise performance ever reported for a digital measurement system. Our solution is based on a pair of high-performance analog-to-digital converters and a single system-on-a-chip (SoC) with multiple processors and a field programmable gate array (FPGA). The architecture allows for processing all data samples in real-time without dead-time between calculation frames, enabling the fastest averaging possible during cross-correlation.

INDEX TERMS Allan deviation, amplitude noise, analog-to-digital converter (ADC), cross correlation, digital down conversion, direct digital measurement, field programmable gate array (FPGA), phase noise.

I. INTRODUCTION

HASE and amplitude noise can have a significant impact on performance in many applications, such as telecommunication [1], [2] and radar [3]. Measuring the phase and amplitude noise of local oscillators and two-port devices becomes critical to predict the behavior and overall capabilities of a particular system. Techniques to measure phase and amplitude noise already exist, especially in the analog domain [4], [5], [6], [7]. However, they require the use of phase-lockable references at the same frequency as the device under test (DUT) for absolute measurements, and the use of additional analog components such as mixers and amplifiers which can limit the measurement noise floor. Additionally, the references' phase-locked loop (PLL) bandwidth limits the minimum frequency offset for the DUT measurement, and other techniques must be used for long-term stability evaluation [8].

Moreover, if ultra-stable signals generated using optical clocks, optical frequency combs [9], or ultra-low noise oven-controlled crystal oscillators are measured, the instrument noise-floor is often a limitation near or far from the carrier. In this case, more advanced techniques such as cross-correlation or carrier-suppression must be used to overcome the measurement system single-channel performance.

In contrast to their analog counterparts, direct digital phase-amplitude noise measurement techniques present many advantages, such as asynchronous reference and DUT frequencies, which removes the need for a phase-locked loop scheme, simplified near-to-the-carrier measurements, simultaneous phase-amplitude noise detection and Allan deviation measurement, and no need for sensitivity calibration [10], [11], [12], [13], [14].

Several products based on this technique are now available since the first commercial introduction [10]. However,

the current commercially available state-of-the-art instrument presents a noise-floor of \mathcal{L} (1 Hz) = -137 dBc/Hz and \mathcal{L} (1 MHz) = -190 dBc/Hz for a 10 MHz signal after 15 minutes of cross correlation averaging (DNA series from NoiseXT).

In this paper, we show a new state-of-the-art measurement system that presents $\mathcal{L}(1 \text{ Hz}) = -143 \text{ dBc/Hz}$ residual flicker and -155 dBc/Hz white noise without the use of cross-correlation, and that can match and surpass the current commercially available instruments. When cross-correlation is used, for a 10 MHz input we achieve an average-limited flicker phase noise floor of $\mathcal{L}(1 \text{ Hz}) = -160 \text{ dBc/Hz}$ within two days of averaging and an average-limited white phase noise floor of $\mathcal{L}(1 \text{ MHz}) = -185 \text{ dBc/Hz}$ within 5 minutes. Lower noise floors can be reached by averaging longer, or by decreasing the size of the fast Fourier transform (FFT) for the same amount of time. Moreover, this instrument is capable of performing phase, amplitude, and Allan deviation measurements simultaneously and without dead-time between calculations.

Section II presents a detailed architecture overview of the main critical components and solutions implemented. Section III shows the built prototype, results and noise-floor achieved, while in Section IV we identify solutions for the current limitations of the first prototype and our plans for the next revision.

II. ARCHITECTURE & HARDWARE

Figures 1, 2 and 3 show a simplified block diagram of the entire system. It consists of analog-to-digital converters (ADCs), a local clock (CLK) and PLLs, the digital signal processing system implemented in field programmable gate array (FPGA) and the software. The DUT signal is power-split and sent to ADC0 and ADC3, while a reference signal (REF) or two independent references are sampled by ADC1 and ADC2. An adjustable clock frequency is used to clock the four ADCs, but instead of being clocked directly, two independent, but loosely locked clocks are used to trigger ADC pairs ADC0 & ADC1 and ADC2 & ADC3 respectively. Each ADC pair is implemented on the same chip, which will later help suppress common-mode clock noise. The dual ADCs used are Analog Devices AD9652 [15] which exhibit exceptional flicker noise performance.

A. ANALOG TO DIGITAL CONVERTERS

To avoid adding additional noise that could compromise the AD9652 performance, we kept the analog front-end as simple as possible. Wilkinson power splitters have been used for DUT and REF, and the signals are sent directly to two evaluation boards EVAL-AD9652 [16]. After some preliminary measurements, we found out that the residual close-to-the-carrier phase noise performance was highly dependent on temperature fluctuation of the two unbalanced-to-balanced (balun) transformers present in the EVAL-AD9652 ADC front-ends. For this reason, the EVAL-AD9652 front-ends have been modified by removing T2 and T6 (designators

from EVAL-AD9652 schematic [17]). Further modifications were made to disable the switching regulators featured on the evaluation boards. They have been replaced with linear regulators based on the LT3045 [18] chip to avoid leakage of the main power supply switching frequency and additional noise into measurements. Dithering in the AD9652 channels is also enabled to reduce spurious products due to intermodulation and aliasing. Thanks to the internal dithering correction of the AD9652, signal-to-noise ratio is degraded by only 0.5 dB. Anti-aliasing filters are external to this prototype and should be chosen depending on the DUT, REF and ADCs clock frequencies.

As shown in Figure 1, the phase information extracted from the two ADCs on the same chip is subtracted. This effectively suppresses the residual phase noise introduced by the ADC clock since it is common to both ADCs [19]. However, since this suppression is most effective at low frequency offsets, the second pair of ADCs are clocked via a loosely locked PLL so that uncorrelated clock noise far from the carrier can be removed via cross-spectrum analysis. Despite this requirement, the two ADC chips should be clocked with the same frequency to allow for Groslambert codeviation measurements. Fig. 2 represents a block diagram of the clocking solution implemented. The FPGA evaluation board used (ZCU102 [20]) contains a Skyworks Si570 [21] programmable oscillator connected directly to one of the FPGA's transceivers. This oscillator is used as the main local clock, and its output is sent to the SMA connectors present on the ZCU102 board by transmitting alternating 0s and 1s on the transceiver. This solution removes the need for additional hardware, and regenerating the clock in the transceiver section helps clean up the spurious nature of the Si570's power supply.

One of the two transceiver's complementary outputs is power split and sent to two synthesizer evaluation boards (EVAL-ADF4002 [22]). These boards feature an ADF4002 [23] phase detector/frequency synthesizer indicated as PD in Fig. 2 and a Crystek CVCO55CW-0250-0450 [24] as a cleanup oscillator. Finally, these boards are configured to lock the cleanup oscillator to the input frequency with a loop bandwidth of about 200 Hz. The low loop bandwidth and identical input-to-output frequency are achieved by setting the R and N counters in the ADF4002 to the same value, choosing a low charge-pump current, and selecting components for the loop filter accordingly.

Therefore, with the described solution we achieved two in-phase, relatively spur-free, identical frequencies with uncorrelated far-from-the-carrier phase noise for each AD9652 chip.

B. DIGITAL DOWN CONVERTER

Figure 3 shows a simplified block diagram of the core of a direct digital measurement system, the digital down converter (DDC). A numerically controlled oscillator (NCO) driven by the ADC sample rate is then used to down convert the DUT

UEE Open Journal of Ultrasonics, Ferroelectrics, and Frequency Control



FIGURE 1. Simplified block diagram of the Digital Phase-Amplitude Noise Measurement System. In yellow the analog components, in green the digital block implemented in the FPGA, in red the bare-metal software running on the Real-time Processor Unit (RPU), and in blue the software running on the main Application Processor Unit (APU). A single Reference or two independent References can be used. DUT = Device Under Test; REF = Reference; ADC = Analog to Digital Converter; DDC = Digital Down Conversion (Fig. 3); FIR \downarrow = Finite Impulse Response filter and decimation; FFT = Fast Fourier Transform; FIFO = First In First Out Memory; SW MEM = APU Software Memory; MEM = Averaging Memory; PLL = Phase Locked Loop; CLK = Main ADC Clock; s = Differentiation; U \uparrow D \downarrow = Up-sampling and Down-sampling to change sampling rate.



FIGURE 2. Clocking of the analog to digital converters. The Si570 programmable crystal oscillator present on the ZCU102 board clocks an FPGA transceiver, the transceiver signal is then power split and used to loosely phase lock two independent voltage controlled oscillators (VCO). Phase/Frequency Detector (PD) = ADF4002, VCO = CVC055CW-0250-0450.

signal to baseband with a digital in-phase (I) and quadrature (Q) mixer scheme. After decimation and low-pass filtering, the I and Q components can be used to calculate the instantaneous phase and amplitude via the two-argument arctangent and the modulus. The instantaneous phase φ will include the phase noise of the DUT, of the ADC clock and the residual phase noise of the ADC. On the other hand, the modulus of the signal will contain information from the amplitude noise of the DUT, noise of the ADC voltage reference, and the residual noise of the ADC. These last two noise contributions are considered as one. The same is repeated for the REF signal.



FIGURE 3. Digital Down Converter block diagram. The analog components are highlighted in yellow, and the digital ones are in green. Low-pass filtering and decimation is done in the CIC and FIR filters. DUT = Device Under Test; CLK = Clock; ADC = Analog to Digital Converter; NCO = Numerically Controlled Oscillator; CIC = Cascaded Integrated-Comb filter; FIR = correction Finite Impulse Response filter.

The DDC outputs can be represented as

$$\varphi_k = 2\pi \left(\frac{\nu_D - \nu_{NCO}}{\nu_C}\right) k + \varphi_{D_k} - \varphi_{C_k} + \varphi_{A_k} - n_k \pi \quad (1)$$

with $n_k \in \mathbb{Z}$ constraining φ_k in the interval $[-\pi, \pi)$, effectively taking into account the atan2 phase wrapping.

$$A_k = A_D \left(1 + \alpha_{D_k} + \alpha_{A_k} \right) \tag{2}$$

where φ_k is the DDC phase output samples; $k \in \mathbb{Z}$ is the sample number; ν_D and ν_C are the frequencies of the DUT and CLK respectively; $\nu_{NCO} = w\nu_C$ is the NCO frequency which depends on ν_C through a scaling factor (tuning word) w; φ_D , φ_C and φ_A are the phase noise samples of the DUT,

ADC clock and residual phase noise introduced by the ADC, respectively; A_k represents the modulus output of the DDC; A_D is the amplitude of the DUT signal; and α_D and α_A are the amplitude noise samples of the DUT and the residual amplitude noise introduced by the ADC respectively. Quantities α_D and α_A are assumed $\ll 1$.

As it can be seen from (1), φ_k has a saw-tooth behavior because the atan2 function can only output numbers from $-\pi$ to $+\pi$ and it represents the continuously evolving and wrapping phase difference between the DUT and CLK. The discontinuity of this phase wrap can be problematic to further signal processing, so it has been addressed by representing $[-\pi]$ with the minimum usable value -1 in fixed point binary math, $+\pi$) with the maximum value, and by taking the first difference of consecutive samples. This effectively computes a scaled time-derivative of φ_k , converting the output from a phase demodulator to a frequency demodulator, thus avoiding the problematic wrapped phase evolution completely (Fig. 1, s blocks). Setting v_{NCO} really close to v_D , and a phase-unwrap function can solve the problem for short term measurements. However, for long term ones, because of the drifting nature of oscillators, there is a non-zero chance that the unwrapped phase will hit the maximum or minimum binary representable values, and therefore an increased number of bits should be used. Steering v_{NCO} to follow v_D is not an option because it will compromise close-to-the-carrier measurements. To guarantee this type of derivative to function, the difference of consecutive atan2 outputs needs to be constrained to $|(\varphi_{k+1} + n_{k+1}\pi) - (\varphi_k + n_k\pi)| < \pi$. Note that because of this constraint, to allow the most bandwidth for $\varphi_D - \varphi_C + \varphi_A$ then $\frac{\nu_D - \nu_{NCO}}{\nu_C} \approx 0$. Four DDCs are implemented in the FPGA, and they are

Four DDCs are implemented in the FPGA, and they are designed to process up to 1 GSps at 16 bits. Since the maximum fabric clock frequency of the Zynq Ultrascale+ [25] is 775 MHz, a parallelization by two and a clock frequency of 500 MHz is needed to process up to 1 GSps.

The NCOs are designed with the Xilinx DDS Compiler tool [26] and are configured with 48 bits phase accumulator resolution and 24 bits output in the <24, 23> fixed point format (24 total bits, 23 of which are fractional), giving a spurious free dynamic range of 138 dB. This signal is then multiplied with the ADC output with an assigned fixed-point layout of <16, 15>, and the output is truncated to a fixed point of format <26, 25>. Choosing this fixed-point format, the math guarantees numbers in the [-1, +1) range. To avoid the special case $(-1) \cdot (-1) = +1$, the ADC symbol 8000_{16} is removed and considered the special input over-range case, since +1 cannot be represented with the chosen format.

A combination of custom designed cascaded integratedcomb (CIC) and finite impulse response (FIR) filters are used to low pass and decimate both I and Q paths. The CIC filter is 131 bits wide, 15 stages long, and can decimate by any number from 2 through 128. A 127 tap FIR filter is used to flatten the CIC filter frequency response and to further decimate by 2. The CIC filter alone provides an aliasing rejection > 114 dB when the entire output bandwidth is used and > 210 dB in the first half of the available spectrum. The FIR filter gives an additional suppression > 100 dB. Since the second half of the output spectrum will never be used, the FIR filter can be optimized for a lower number of taps. This high level of suppression is necessary to guarantee an alias-free measurement, especially when using FM demodulation. The output sample rate is defined as $2f_c$, f_c is then the analysis bandwidth and the usable spectrum is $\frac{f_c}{2}$. More aggressive filters can bring the usable spectrum close to f_c .

The two-argument arctangent and square root functions are both implemented in the FPGA using the COordinate Rotation DIgital Computer (CORDIC) algorithm. The Xilinx provided intellectual property is used, and they have been both configured with the maximum available resolution. The atan2 (\cdot , \cdot) output is also designed to represent $-\pi$ as -1, and π as +1.

After the DDC, the phase derivative and the modulus of the I and Q are converted to double floating-point representation, scaled and subtracted from the adjacent DDC channel as seen in figure 1. This operation allows measuring the DUT relative to an external reference signal (REF) and removes any common-mode ADC related clock noise φ_C within the pairs ADC0 & ADC1 and ADC2 & ADC3. For common-mode noise cancellation, it is important to guarantee that ADC0 & ADC1 pair have correlated parameters such as the same clock, voltage references and temperatures. The same applies for the ADC2 & ADC3 pair, and conveniently the AD9652 provides pairs of ADCs co-located on the same silicon chip. The architecture shown in Fig. 1 allows either absolute AM noise measurements of the DUT when $k_1 = k_2 = 0$, or differential/residual AM noise measurements relative to the REF input. In case of the latter, the ADCs voltage reference noise contribution is suppressed during subtraction since it's common to each ADC pair.

C. DECIMATION CHAIN

After scaling and subtraction, a decimation chain is implemented to convert the signals to various sampling rates for analyzing different frequency ranges. This is done because phase and amplitude noise graphs are often represented on a logarithmic frequency offset scale, and a normal FFT would lose graphical resolution at low frequencies. Each filter decimates by 8 and a total of 9 filters per channel are present, giving a total of 10 different sampling rates. The sampling rates for each output are

$$\frac{2f_c}{8^i}$$
 $i = 0, 1 \dots 9$ (3)

where *i* is the decimation chain output.

The differentiated phase signal, used to overcome phase wrapping, requires additional aggressive filtering when decimated. The signal being decimated is $2\pi v_D y(t)$ where y(t) are the fractional frequency fluctuations, and the signals' power spectral density $4\pi^2 v_D^2 S_y(f)$ is computed instead of the typical $S_{\varphi}(f)$. The attenuation requirements for the antialiasing decimation filters can be determined as follows. For residual

measurements, one can assume phase noise to be comprised by only white and flicker components $S_{\varphi}(f) = W \left(1 + f_{i/f}/f\right)$ where *W* is the white noise level and $f_{i/f}$ is the corner frequency where flicker and white noise intercepts. Thus, the system will measure

$$4\pi^2 \nu_D^2 S_y(f) = 4\pi^2 f^2 W\left(1 + \frac{f_{1/f}}{f}\right) \quad 0 < f < f_c.$$
(4)

When decimating by an even number, the frequency content near f_c will alias close to 0. To prevent an aliasing bias at offset frequencies above f_m the following inequality must be maintained.

$$4\pi^{2}v_{D}^{2}S_{y}(f_{m}) \gg K4\pi^{2}v_{D}^{2}S_{y}(f_{c}-f_{m})$$
(5)
$$K \ll \frac{f_{m}^{2}+f_{m}f_{1/r}}{(f_{c}-f_{m})^{2}+(f_{c}-f_{m})f_{1/r}} \approx \frac{f_{m}f_{1/r}}{f_{c}^{2}}$$
$$f_{c} \gg f_{1/r} \gg f_{m} \gtrsim 0$$
(6)

where $10 \log_{10}(K)$ is the attenuation for the decimation filter used in dB. Inequality (6) shows that with just white and flicker phase features, it is impossible to prevent aliasing as f_m approaches 0. However, K can be selected to make measurements valid from a practical point of view. Fortunately, when measuring oscillators, higher order noise is present (white and flicker frequency modulation and random walk) and this will relax the requirement on K. When white frequency noise is considered, $S_{\varphi}(f) = W (1 + f_{1/f}/f + f_{1/f^2}/f^2)$ with f_{1/f^2} defined as the corner frequency where white frequency and white phase noise intercepts. In this case, (6) becomes

$$K \ll \frac{f_m^2 + f_m f_{1/r} + f_{1/r^2}}{(f_c - f_m)^2 + (f_c - f_m) f_{1/r} + f_{1/r^2}} \approx \frac{f_{1/r^2}}{f_c^2}$$

$$f_c \gg f_{1/r}, f_{1/r^2} \gg f_m \gtrsim 0.$$
(7)

Inequality (7) allows finding the minimum attenuation needed to measure any close-to-the-carrier frequency and Allan Deviation averaging time without aliasing data corruption for spur-less spectra when oscillators are measured. If spurs are present, a smaller K might be needed.

Each decimation filter is implemented with a 255 tap FIR clocked with a 250 MHz signal, and the transfer function has > 200 dB of stop band suppression to avoid aliasing. In fact, by using (7) and $f_c = 62.5$ MHz our measurements will be accurate as long as $f_{1/2} \gg 39.1 \mu$ Hz. Because these filters need to run on double floating-point math, their FPGA implementation is heavily optimized, and utilizes resource sharing techniques as much as possible.

D. FFT, CROSS-CORRELATION & AVERAGING

For each decimation chain output a fast Fourier transform must be computed, but to minimize the required FPGA resources, we implemented only one FFT processing chain per channel per phase/amplitude analysis. This is possible because the total maximum sample rate out of all decimation chain outputs is,

$$\sum_{i=0}^{9} \frac{2f_c}{8^i} \approx 142.9 \text{ MSps}$$
(8)

meaning that we can use one shared Hanning windowing, FFT and averaging module per channel per phase/amplitude analysis as long as this section can process more than 142.9 MSps. Eq. 8 can be extended to infinite frequency ranges with the use of one processing chain as long as $\frac{16f_c}{7} < f_{clk}$, with f_{clk} being the processing chain clock speed. FPGA resources will limit the maximum number of frequency ranges allowed.

We implemented an FFT with size 1024, and this section was designed to process 1024 sample bursts at a 250 MHz clock speed. Because of the filters' shape half of the spectrum is usable and because the decimation chain reduces the sampling rate by 8 at each stage, each frequency span will show 224 bins. The lowest frequency span can show more bins. To create bursts of data, the samples from each decimation chain output are collected into first-in first-out (FIFO) memories. When at least 1024 samples are present in the FIFO for both channels for a given decimation chain output, a signal is generated and an arbiter (represented as the multi-selector switch in Fig. 1) selects its output for processing. In the case of multiple signals from different decimation chain outputs, the arbiter prioritizes the ones with higher sample rate.

In the worst case, the fastest sample rate output FIFO reaches 1024 samples right after the arbiter selected another output for processing. By the time all 1024 samples are extracted, the FIFO will continue to fill. The extra filled amount will depend on the burst length, how fast data is extracted from the FIFO, and how fast it fills. The FIFOs will then need a depth of at least

$$1024\left(1 + \frac{2f_c}{250 \text{ MHz}}\right) = 1536.$$
 (9)

This minimum depth applies for the decimator output i = 0, other outputs can be designed with smaller FIFOs since the sample rate at which they fill is lower. Equation (9) can be modified to

$$1024\left(1+\frac{1}{8^{i}}\frac{2f_{c}}{250 \text{ MHz}}\right) \qquad i=0,1\dots9.$$
 (10)

In our case, since FPGA memory was sufficient, a depth of 2048 was instead implemented for all outputs to eliminate the possibility of data loss.

Cross-correlation is the most critical part for this section. In fact, if data taken at different moments is cross-correlated, the result will be meaningless since no correlation will be present. If there is a time delay offset t_d between the two channels, the cross correlated result will be modified according to

$$C_{S}(f, t_{d}) = C_{S}(f, 0) e^{j\pi f t_{d}}$$

$$C_{S}(f, 0) = \mathcal{F}_{1}\{\cdot\} \mathcal{F}_{2}^{*}\{\cdot\}$$
(11)

where $C_S(f, t_d)$ is the cross-spectrum in function of frequency and time offset, and $C_S(f, 0)$ is computed by multiplying the FFT of one channel with the complex conjugate of the other. With enough averaging $C_S(f, 0)$ will tend to a real function, and the real and imaginary parts of $C_S(f, t_d)$



FIGURE 4. Simplified block diagram of the timestamp balancing logic.

will result as

$$\Re \left(C_S \left(f, t_d \right) \right) = C_S \left(f, 0 \right) \cos \left(\pi f t_d \right) \tag{12}$$

$$\Im \left(C_S \left(f, t_d \right) \right) = C_S \left(f, 0 \right) \sin \left(\pi f t_d \right) \tag{13}$$

and we see that nulls are created in $\Re(C_S(f, t_d))$ and $\Im(C_S(f, t_d))$ functions, and their meaning will be periodically inverted, compromising the measurement.

Therefore, time offsets between the two channels must be minimized as much as possible, and this is the reason clocks for ADC0, ADC1 and ADC2, ADC3 are designed to be in phase. Before the measurement starts, all decimation counters in the various filters are reset, and a "timestamp" is assigned to each sample as soon as it has been taken. The "timestamp" is essentially a 48-bit counter running at the highest available clock frequency (500 MHz) and incremented at each clock cycle. When a sample is available from an ADC, the current value of the counter is associated with it and treated as a "timestamp". This value flows through the entire data processing (DDC, scaling, etc.) until it gets saved in the decimation chain FIFOs. Here, the arbiter checks the timestamps from the two channels for each FIFO output i = $0, 1 \dots 9$ and makes sure their difference is within a certain maximum tolerance. If the tolerance is exceeded, samples are discarded from one FIFOs until the two channels are again synchronized to within the set threshold. Moreover, when a burst of data is sent for FFT analysis, the timestamp difference for each sample pair is integrated, and this information is used to further discard samples. Figure 4 shows a simplified block diagram of the described delay balancing logic.

A check on the timestamp difference ensures an upper and lower limit for t_d , while keeping $\int t_d dt = 0$ reduces the effect of t_d on the cross-spectrum imaginary part during averaging. The idea is that since (13) is an odd function, the sign of t_d is alternated when averaging, therefore suppressing the effect of t_d .

Note that the described balancing circuit is not actively used if the system is well-built and all delay discrepancies between the two channels are accounted for. However, this logic opens the possibility to use two independent almost-identical clock sources for ADC0, ADC1 and ADC2, ADC3. If the two ADC clocks frequencies are within a few tens of parts-per-million apart, it's possible to have the two channels completely uncorrelated without the use of loosely locked PLLs. However, this type of system won't be able to compute a proper 2-sample codeviation since the samples will be taken at slightly different τ_0 .

Finally, after FFT, cross-correlation and vector modulus computation, an averaging circuit is implemented, with the data saved into FPGA memory. The averaging circuit can compute either normal or exponential averaging. Since most of the processing is done in double-precision floating point, in an ideal case, where the same number (such as 1) is added over and over, the maximum number of averages allowed are 2^{52} where 52 is the number of bits in the fractional part. From a practical point of view, time is the main limitation instead. In fact, with the maximum $f_c = 62.5$ MHz and 1024 sample FFTs, the 2^{52} limit will be reached in about 1169 years of measurement. This allows the averaging circuit to safely work in non-ideal averaging conditions without a practical limit on the maximum number of averages.

E. ALLAN DEVIATION

The data from phase decimation chain, can be easily converted to fractional frequency fluctuations and used to calculate the Allan deviation. Since the overall decimation ratio is a power of 2 and the clock frequency is adjustable, the Allan deviation is generated with an inconvenient, non-round value for a starting τ_0 . To address this problem, a real-time sample rate conversion is used on the decimator output before the Allan deviation to generate a new stream of samples that can produce a τ_0 's with more traditionally used values. Two identical bare-metal programs running on each core of the real-time processing unit (RPU) of the Zynq System On a Chip (SoC) are dedicated to this sample-rate conversion. Each program retrieves samples from the FPGA, then it up-samples them by *N* by adding 0s, and then down-samples the result by *M* with the use of low-pass decimation FIR filters.

Firstly, the software finds the lowest decimation chain output with spectral information that supports the desired τ_0 . Note that because of the filters' frequency response, only the first half of the spectrum content is usable. Therefore, we find

$$i_{\rm A} = \arg\min_i \frac{2f_c}{8^i}$$
 $i = 0, 1...9$ s.t. $\frac{2f_c}{8^i} \ge \frac{2}{\tau_0}$ (14)

Secondly, we find the best ratio $\frac{N}{M}$ that brings the sample rate as close as possible to $\frac{2}{\tau_0}$ with $1 \le N \le N_{\text{MAX}}$, $N \in \mathbb{Z}^+$ and $M \in \mathcal{D}_M$. N_{MAX} is the maximum up-sampling viable, and \mathcal{D}_M contains all the available implementable values for M.

Up-sampling is trivial to implement, and it consists of adding N - 1 zeros between samples and multiplying them by N. The multiplication by N is necessary to maintain the correct DC value. Down-sampling is instead executed with the use of some prime number decimation FIR filters. In our implementation we have low-pass FIRs that can decimate by 2, 3, 5 and 7, and for each type of filter up to 4 are available.

IEEE Open Journal of Ultrasonics, Ferroelectrics, and Frequency Control

The domain \mathcal{D}_M is then defined as

$$\mathcal{D}_M = \left\{ x \in \mathbb{Z}^+ | x = 2^{n_2} 3^{n_3} 5^{n_5} 7^{n_7}, \quad n_{2,3,5,7} = 0, 1 \dots 4 \right\}$$
(15)

After sample rate conversion by $\frac{N}{M}$, the result can be within a few percent from the $\frac{2}{\tau_0}$ sample rate, which is still double the desired $\frac{1}{\tau_0}$. Here, the first half of the spectrum is usable, while the second half is not flat and can contain aliased power. For this reason, a final decimation by two FIR filter, designed to have a steep and flat response, is used to get the desired and final $\frac{1}{\tau_0}$ sample rate. The finally derived samples from both channels are then saved in the main processor's memory, and used for Allan Deviation and Groslambert co-deviation (GCov_A) computation.

F. SOFTWARE

The SoC present on the ZCU102 board features, besides the FPGA and RPU, also a quad-core ARM processor called the application processor unit (APU). The APU is used to run the Ubuntu 18.04 operating system and the main program controlling this project. The software and firmware development tools used were based on the Koheron SDK [27] with significant modifications to support the Zynq Ultrascale+ SoC and the ZCU102 development board.

The program sets operational parameters, monitors the FPGA, retrieves the measurement data from the FPGA memory for final processing, and finally computes Allan Deviation and Groslambert co-deviation. The system can be monitored, controlled, and data can be displayed and saved through a web interface or python commands sent via TCP/IP protocol. It is also possible to run a LXI server for standardized Ethernet control; however, this last feature has not been implemented.

Figure 5 shows the designed web interface. On the right side, the user can input DUT and references frequencies, the type of averaging desired, and the Allan Deviation τ_0 . Two buttons start or stop the measurement. Below the measurement settings, there are the main graph controls. Here, the user selects the data to be shown on the interface. Frequency domain options are between Phase/Amplitude Noise for Channels 0/1, real, imaginary, or magnitude of the crossspectrum. Time domain displays containing Allan Deviation computed for the single channels, Groslambert co-deviation and the estimated co-deviation using the phase noise data. The sign information for the cross-spectrum and co-deviation is retained and displayed with black dots. All measurements are calculated and retained in memory, so the user can browse the different graphs without data loss during or after the measurement is completed. Below the graph controls there are units control. The vertical units shown on the graph can be selected between logarithmic scale, root-mean-square (RMS), power spectrum density (PSD), or show directly the phase noise $(\mathcal{L}(f))$ or amplitude noise $(\mathcal{M}(f))$. Finally, below the main interface, all debug/error flags and internal system status are shown.

166

TABLE 1. FPGA utilization.

Resources	Utilization	Available	Utilization %
LUT	171019	274080	62.4
LUTRAM	28034	144000	19.5
FF	248700	548160	45.4
BRAM	717	912	78.6
DSP	1946	2520	77.2
Ю	98	328	29.9
GT	1	24	4.2
BUFG	14	404	3.5
ММСМ	2	4	50.0
LUT-Look Up Table, LUTRAM-Random Access Memory			
implemented in LUT, FF-Flip Flop, BRAM-Random Ac-			
cess Memory Block, DSP-Digital Signal Processing block,			
IO—Input Output pin, GT—GT transceiver, BUFG—Global			
BUFfer, MMCM—Mixed-Mode Clock Manager			

G. HARDWARE

Figure 6 shows pictures of the built prototype system. The ZCU102 development board and its +12 V power supply are located in the bottom section of the case. The same supply is also used to power the analog section after multiple stages of linear regulation based on LM317s [28] and LT3045s to reduce switching noise. The two EVAL-AD9652 and EVAL-ADF4002 are placed on the case's top section and are covered with foam to reduce thermal variation effects on the analog sensitive components. ADCs and FPGA are connected through FPGA mezzanine card (FMC) connectors and custom-made adapters. The front panel has a DUT input port on the left, and a common REF or dual REF inputs on the right. The back panel has the power switch, power connector, Ethernet interface port and a fan exhaust output.

Table 1 shows the FPGA utilization for this first prototype. As it can be seen, the high utilization of LUT, BRAM and DSP combined with high fabric clock up to 500 MHz makes timing closure extremely difficult. The SoC total power consumption is estimated to be around 24 W, making a case exhaust fan mandatory to avoid overheating.

III. EXPERIMENTAL RESULTS

To test the system's performance, a long-term residual phase and amplitude noise measurement has been performed. A 10 MHz 14 dBm signal was power split and connected to both DUT and REF inputs. The ADC clock is set to 297 MHz, which gives the best spurious response for this measurement. The system was brought to its nominal operating temperature by first running a dummy measurement for 24 hours, then, the system was restarted to perform a 4-day measurement. Figure 7 shows the phase and amplitude noise results. The aim was for one million averages, which were reached above 250 Hz frequency offsets. At 1 Hz offset, 5121 averages were completed. From the phase noise graph we can see a residual \mathcal{L} (1 Hz) = -143 dBc/Hz for the single channels which has been improved to -162 dBc/Hz with the use of



FIGURE 5. Web interface used to control/monitor the system and to visualize the live measurements. In the picture, a live 10 MHz instrument residual noise measurement.



FIGURE 6. Pictures of the first prototype.

cross-spectrum averaging. Some spurs have been identified, while others have unknown source and are probably picked up from the ZCU102 digital board. Near the carrier, we see further correlation due to temperature fluctuations, and the wide bump at around 2×10^{-4} Hz in the residual amplitude noise is caused by the air conditioning system. On the other hand, far from the carrier we measure -155 dBc/Hz for the single channels, and we can average below -180 dBc/Hz and reach -185 dBc/Hz with a million averages. For amplitude noise, far-from-the-carrier performance is the same, and some limitations in the cross-correlation are visible due to crosstalk and reduced isolation around the 100 Hz offset area. Our architecture was specifically designed, without dead-time, so that almost no samples are discarded, allowing to complete a million averages in a few minutes far from the

carrier. The time to complete N_A averages for the highest frequency range is for this system, $N_A \frac{1024}{2f_c}$ seconds, where f_c is around a fifth of the lowest input frequency. In Fig. 7 $f_c \approx 2$ MHz (1 MHz usable and shown), allowing to complete far-from-the-carrier one million averages in less than 5 minutes.

Figure 8 shows the Allan deviation of the same 4-day measurement. A bandwidth of $\frac{1}{2\tau_0}$ Hz with $\tau_0 = 1$ s is used. The single channels measure 3.2×10^{-15} @ $\tau = 1$ second, while the Groslambert co-deviation starts at 4×10^{-16} @ $\tau = 1$ second and is average limited (5×10^{-16}). The imaginary part of the cross-correlated phase noise data has been used to estimate the noise floor of the measurement. Note that after long averaging time GCov_A starts to follow the single channels, suggesting that the two channels have





FIGURE 7. Residual phase noise (left) and amplitude noise (right) measured with a 10 MHz signal. The number of averages achieved is indicated on top. CH0 is the channel composed by the ADC0-1 pair, while CH1 is the channel composed by ADC2-3.



FIGURE 8. Residual Allan deviation measurement at 10 MHz signal for averaging time starting at $\tau_0 = 1$ s and BW = 0.5 Hz.

correlated behavior probably due to common-mode temperature fluctuations, which degrade instrument performance.

To verify that the residual measurement shown in figures 7 and 8 has not been corrupted by aliasing during decimation, we use inequality (6). From Fig. 7 we see that the single channels' flicker/white corner frequency is $f_{1/r} \approx 30$ Hz, making the measurements valid above ~ 1.3 nHz or below ~ 750 million seconds. Therefore, aliasing is negligible.

The phase noise accuracy of the prototype measurement system was verified by measuring a calibrated noise level from a NIST AM/PM noise standard [29] at 10 MHz. The measured noise is within the noise standard's 2-sigma uncertainty of +/-0.5 dB, as shown in Figure 9(top). The larger statistical variations at lower frequency offsets are due to a smaller number of FFT averages. Furthermore, a comparison



FIGURE 9. Phase noise accuracy of the prototype system at 10 MHz, measured using a calibrated NIST PM/AM noise standard (top). Absolute noise of a commercial signal generator also at 10 MHz measured with the NIST prototype and compared with Microsemi 5125A and Rohde & Schwarz FSWP50 (bottom).

of absolute phase noise of a signal generator at 10 MHz, as measured with our system, and two commercial digital or hybrid measurement systems (Microchip 5125A and Rohde & Schwarz FSWP50) is depicted in Figure 9(bottom). Both measurements show satisfactory agreement.

IV. CONCLUSION

A direct digital simultaneous phase-amplitude noise and Allan deviation measurement system with excellent residual flicker performance has been presented, and its architecture thoroughly described. Using evaluation boards, for a 10 MHz full-scale input signal we achieved $\mathcal{L}(1 \text{ Hz}) = -143 \text{ dBc/Hz}$ and -155 dBc/Hz white noise, which can be improved using cross-correlation to -185 dBc/Hz white noise after a few minutes and $\mathcal{L}(1 \text{ Hz}) = -160 \text{ dBc/Hz}$ in less than 2 days. For residual Allan deviation measurements, the single channels exhibit a stability of 3.2×10^{-15} @ $\tau = 1s$, which can be improved below 5.0×10^{-16} @ $\tau = 1s$ after 4 days of Groslambert co-deviation averaging.

The results obtained are very promising, and there is still room for improvement. In particular, we are going to focus on the following key points:

- careful layout and shielding to address power-line harmonics and other spurious signal pickup.
- more isolation between the digital board/SoC and ADCs by using independent power supplies and signal isolators.
- independent temperature control on the analog sections and ADCs to suppress low offset frequencies correlation between the two channels

Therefore, a custom module will be designed with the above points addressed. This will benefit long term measurements and cross-correlation behavior. Our plan is to also use these custom-made modules for a multichannel digital timescale application in the near future. The multichannel architecture opens the possibility to even faster cross-spectrum measurements.

Moreover, since the entire FPGA digital processing has been designed to support up to 1 GSps, we could investigate the white and flicker noise performance of a newer generation of faster ADCs.

ACKNOWLEDGMENT

The authors would like to thank Jeffrey Sherman and Fabrizio Giorgetta for their useful comments on this manuscript.

DISCLAIMER

Contribution of the U.S. government, not subject to copyright. Commercial products are mentioned in this document for technical and scientific information and does not constitute an endorsement by NIST.

REFERENCES

- G. Colavolpe, "Communications over phase-noise channels: A tutorial review," in Proc. 6th Adv. Satell. Mobile Syst. Conf. 12th Int. Workshop Signal Process. Space Commun., 2012, pp. 316–327.
- [2] A. Spalvieri, "Non-parametric phase tracking in demodulation and decoding of QAM signals affected by phase noise," in *Proc. 11th Int. Conf. Telecommun. Syst. Services Appl. (TSSA)*, Oct. 2017, pp. 1–5.
- [3] K. Siddiq, R. J. Watson, S. R. Pennock, P. Avery, R. Poulton, and B. Dakin-Norris, "Phase noise analysis in FMCW radar systems," in *Proc. Eur. Microw. Conf. (EuMC)*, Sep. 2015, pp. 1523–1526.
- [4] 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*, 1976, pp. 269–274.
- [5] J. Li, E. Ferre-Pikal, C. Nelson, and F. L. Walls, "Review of PM and AM noise measurement systems," in *Proc. Int. Conf. Microw. Millim. Wave Technol.*, 1998, pp. 197–200.

- [7] 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.
- [8] P. A. Koppang and C. R. Ekstrom, "Degrees of freedom for Allan deviation estimates of multiple clocks," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 63, no. 4, pp. 571–574, Apr. 2016.
- [9] N. V. Nardelli et al., "10 GHz generation with ultra-low phase noise via the transfer oscillator technique," APL Photon., vol. 7, no. 2, Feb. 2022, Art. no. 026105, doi: 10.1063/5.0073843.
- [10] J. Grove, J. Hein, J. Retta, P. Schweiger, W. Solbrig, and S. Stein, "Direct-digital phase-noise measurement," in *Proc. IEEE Int. Freq. Control Symp. Expos.*, Aug. 2004, pp. 287–291.
- [11] T. Imaike, "Full digital phase noise measurement by using two reference oscillators and multichannel ADCs," in *Proc. Joint Conf. Eur. Freq. Time Forum IEEE Int. Freq. Control Symp. (EFTF/IFCS)*, Jul. 2017, pp. 583–586.
- [12] D. A. Howe, A. Hati, C. W. Nelson, and D. Lirette, "PM-AM correlation measurements and analysis," in *Proc. IEEE Int. Freq. Control Symp.*, May 2012, pp. 1–5.
- [13] C. W. Nelson and D. A. Howe, "A sub-sampling digital PM/AM noise measurement system," NCSLI Measure, vol. 7, no. 3, pp. 70–73, Sep. 2012, doi: 10.1080/19315775.2012.11721610.
- [14] P.-Y. Bourgeois, G. Goavec-Merou, J.-M. Friedt, and E. Rubiola, "A fullydigital realtime SoC FPGA based phase noise analyzer with crosscorrelation," in *Proc. Joint Conf. Eur. Freq. Time Forum IEEE Int. Freq. Control Symp. (EFTF/IFCS)*, Jun. 2017, pp. 578–582.
- [15] Analog Devices. 16-Bit, 310 MSPS, 3.3 V/1.8 V Dual Analog-to-Digital Converter (ADC). Accessed: Apr. 2022. [Online]. Available: https://www.analog.com/media/en/technical-documentation/data-sheets/ ad9652.pdf
- [16] Analog Devices. EVAL-AD9652, AD9652 Evaluation Board. Accessed: Jun. 2022. [Online]. Available: https://www.analog.com/en/designcenter/evaluation-hardware-and-software/evaluation-boards-kits/evalad9652.html#eb-overview
- [17] Analog Devices. EVAL-AD9652 Schematic. Accessed: Jun. 2022.
 [Online]. Available: https://wiki.analog.com/_media/resources/eval/ user-guides/14016a_sch.pdf
- [18] Analog Devices. 20V, 500mA, Ultralow Noise, Ultrahigh PSRR Linear Regulator. Accessed: Jun. 2022. [Online]. Available: https://www.analog. com/media/en/technical-documentation/data-sheets/lt3045.pdf
- [19] A. C. Cárdenas-Olaya, E. Rubiola, J.-M. Friedt, M. Ortolano, S. Micalizio, and C. E. Calosso, "Simple method for ADC characterization under the frame of digital PM and AM noise measurement," in *Proc. Joint Conf. IEEE Int. Freq. Control Symp. Eur. Freq. Time Forum*, Apr. 2015, pp. 676–680.
- [20] Xilinx. Zynq UltraScale+ MPSoC ZCU102 Evaluation Kit. Accessed: Apr. 2022. [Online]. Available: https://www.xilinx.com/products/boardsand-kits/ek-u1-zcu102-g.html
- [21] Skyworks. 10 MHZ to 1.4 GHZ 12C PROGRAMMABLE XO/VCXO. Accessed: Feb. 2023. [Online]. Available: https://www.skyworksinc.com//media/skyworks/sl/documents/public/data-sheets/si570-71.pdf
- [22] Analog Devices. EVAL-ADF4002, ADF4002 Evaluation Board. Accessed: Feb. 2023. [Online]. Available: https://www.analog.com/en/designcenter/evaluation-hardware-and-software/evaluation-boards-kits/evaladf4002.html#eb-overview
- [23] Analog Devices. Phase Detector/Frequency Synthesizer. Accessed: Feb. 2023. [Online]. Available: https://www.analog.com/media/ en/technical-documentation/data-sheets/ADF4002.pdf
- [24] Crystek Microwave. Voltage Controlled Oscillator-VCO CVC055CW-0250-0450. Accessed: Feb. 2023. [Online]. Available: https://www.crystek.com/microwave/admin/webapps/welcome/files/vco/ CVC055CW-0250-0450.pdf
- [25] Xilinx. Zynq UltraScale+ MPSoC. Accessed: Apr. 2022. [Online]. Available: https://www.xilinx.com/products/silicon-devices/soc/zynqultrascale-mpsoc.html
- [26] Xilinx. DDS Compiler V6.0. Accessed: Apr. 2022. [Online]. Available: https://docs.xilinx.com/v/u/en-U.S./pg141-dds-compiler
- [27] Koheron. Koheron SDK. Accessed: Apr. 2022. [Online]. Available: https://github.com/Koheron/sdk
- [28] Texas Instruments. LM317 3-Terminal Adjustable Regulator. Accessed: Mar. 2023. [Online]. Available: https://www.ti.com/lit/gpn/lm317
- [29] F. L. Walls, "Secondary standard for PM and AM noise at 5, 10, and 100 MHz," *IEEE Trans. Instrum. Meas.*, vol. 42, no. 2, pp. 136–143, Apr. 1993.

UITrasonics, Ferroelectrics, and Frequency Control



MARCO POMPONIO received the M.Sc. degree in electronic engineering from the Polytechnic of Turin, Italy, in 2017, in collaboration with Italian National Metrology Institute (INRIM). He is currently pursuing the Ph.D. degree with the University of Colorado Boulder in collaboration with the National Institute of Standards and Technology (NIST). In 2018, he presented the M.Sc. thesis work at European Frequency and Time Forum (EFTF). He has been a Research Assistant with

NIST, since 2018. He is also an Electronics Engineer with the University of Colorado Boulder, in collaboration with NIST. His research interests include high-performance digital control loops, field programmable gate arrays (FPGA), signal processing, low-noise electronics, and phase and amplitude noise metrology. He won the student poster competition at EFTF and he won the same competition again in 2021.



CRAIG NELSON (Member, IEEE) is currently an Electrical Engineer and the Leader of the Phase Noise Metrology Group, National Institute of Standards and Technology (NIST). His involvement in this group spans over three decades and research interest are phase and amplitude noise metrology, low-noise electronics, FPGA-based digital control, and instrument control. He has authored over 70 articles and teaches classes, tutorials, and workshops at NIST, the IEEE Frequency

Control Symposium, and several sponsoring agencies on the practical aspects of high-resolution phase noise metrology. He was a recipient of the NIST Bronze Medal in 2012, the Allen V. Astin Measurement Science Award "for developing a world-leading program of phase noise research and measurement services to support industry and national priorities" in 2015, and the IEEE Cady Award "for leadership in the design and development of state-ofthe-art low noise oscillators and phase noise measurement systems" in 2020.



ARCHITA HATI (Member, IEEE) is currently an Electronics Engineer with the Time and Frequency Division, National Institute of Standards and Technology, Boulder, CO, USA, where she is the Calibration Service Leader of the Phase Noise Metrology Group. Her research interests include phase noise metrology, ultra-low noise frequency synthesis, development of low-noise microwave and opto-electronic oscillators, and vibration analysis. She was a recipient of the Allen V. Astin

Measurement Science Award "for developing a world-leading program of phase noise research and measurement services to support industry and national priorities" in 2015. She has been an Associate Editor of IEEE TRANSACTIONS ON ULTRASONICS, FERROELECTRICS, AND FREQUENCY CONTROL, since 2021.