

Oscillator metrology with software defined radio

Jeff A. Sherman^{a)} and Robert Jördens^{b)}

National Institute of Standards and Technology, Division of Time and Frequency, Boulder, Colorado 80305, USA

(Received 20 November 2015; accepted 6 May 2016; published online 24 May 2016)

Analog electrical elements such as mixers, filters, transfer oscillators, isolating buffers, dividers, and even transmission lines contribute technical noise and unwanted environmental coupling in time and frequency measurements. Software defined radio (SDR) techniques replace many of these analog components with digital signal processing (DSP) on rapidly sampled signals. We demonstrate that, generically, commercially available multi-channel SDRs are capable of time and frequency metrology, outperforming purpose-built devices by as much as an order-of-magnitude. For example, for signals at 10 MHz and 6 GHz, we observe SDR time deviation noise floors of about 20 fs and 1 fs, respectively, in under 10 ms of averaging. Examining the other complex signal component, we find a relative amplitude measurement instability of 3×10^{-7} at 5 MHz. We discuss the scalability of a SDR-based system for simultaneous measurement of many clocks. SDR's frequency agility allows for comparison of oscillators at widely different frequencies. We demonstrate a novel and extreme example with optical clock frequencies differing by many terahertz: using a femtosecond-laser frequency comb and SDR, we show femtosecond-level time comparisons of ultra-stable lasers with zero measurement dead-time. [<http://dx.doi.org/10.1063/1.4950898>]

I. OVERVIEW

Time is best measured by counting periods of natural or manmade oscillators.²² To maximize temporal resolution we must interpolate between integer periods, a task equivalent to determining an oscillator's phase. Consider two oscillators with frequency f , the periods of which can be counted as clocks. Their phase offset $\Delta\phi(t_k)$ (in radians) at a measurement epoch t_k can be interpreted as a time offset,³

$$\Delta T(t_k) = \frac{\Delta\phi(t_k)}{2\pi f}. \quad (1)$$

Resolving whether ΔT is stationary is the most sensitive method for detecting small frequency offsets or fluctuations between the oscillators and thus calibrating or characterizing them as clocks.^{26,49} For continuously running clocks, a linear drift in ΔT defines a (fractional) frequency offset between the oscillators $y = [\Delta T(t_2) - \Delta T(t_1)] / (t_2 - t_1)$ consistent with the notion that frequency is the rate of change of phase.

In this work, we briefly review existing high-resolution techniques for measuring ΔT of radio frequency oscillators. We introduce the software defined radio (SDR) concept in the context of time and frequency metrology, and describe basic demonstration experiments valid for many SDR implementations. Finally, we explore SDR's ability to compare oscillators at dissimilar frequencies and to scale to many-oscillator comparisons. One new SDR application is discussed in some detail: phase-coherent measurement of optical clocks via a femtosecond laser frequency comb.

A. Radio techniques in oscillator metrology

Though clock frequencies may be high, measurement bandwidth need not be for comparing oscillators i and r that are similar in frequency, $f_i \approx f_r$. Since clock oscillators are typically very stable, a signal at $f_i - f_r$ is both low in bandwidth and low in absolute frequency and therefore amenable to high-precision measurement. Such frequency translation is rooted in radio techniques—transmitters shift a signal of low- to moderate-bandwidth upwards to many megahertz or gigahertz for ease of wide-area propagation while receivers spectrally convert the signal back to its original band with no practical loss in information.

The widely applied dual-mixer time-difference (DMTD) technique,³ illustrated in Figure 1(a), is an example of radio frequency translation applied to oscillator metrology. A transfer oscillator is synthesized at $f_r - f_b$, slightly offset from input and reference oscillators f_i and f_r : ($|f_i - f_r| \ll f_b \ll f_{i,r}$). The transfer oscillator is mixed (multiplied) with both f_r and f_i tones, creating two signals with frequencies near f_b after low-pass filtering. A time-interval counter (TIC) counts periods of a fast *timebase* oscillator ν_m , also disciplined by f_r , gated by high slew-rate zero-crossing detectors observing the two heterodyne products near f_b .

As a consequence of the spectral conversion, DMTD methods resolve $\Delta T \leq 1$ ps accurately, despite no component possessing a bandwidth approaching $(1 \text{ ps})^{-1} = 1 \text{ THz}$. DMTD realizations often employ an offset frequency $1 \text{ Hz} \leq f_b \leq 10 \text{ Hz}$ and heterodyne factors $10^5 \leq f_r/f_b \leq 10^7$, so a TIC must only accurately resolve $\Delta T \sim 1 \mu\text{s}$ between the relatively slow oscillations near f_b to discern $(f_r/f_b)^{-1}(1 \mu\text{s}) \sim 1$ ps of oscillator time difference.

While DMTD techniques are already highly refined,⁴⁸ ΔT of 1 ps can be mimicked or masked by fluctuations of

^{a)}Electronic mail: jeff.sherman@nist.gov

^{b)}Electronic mail: rj@m-labs.hk

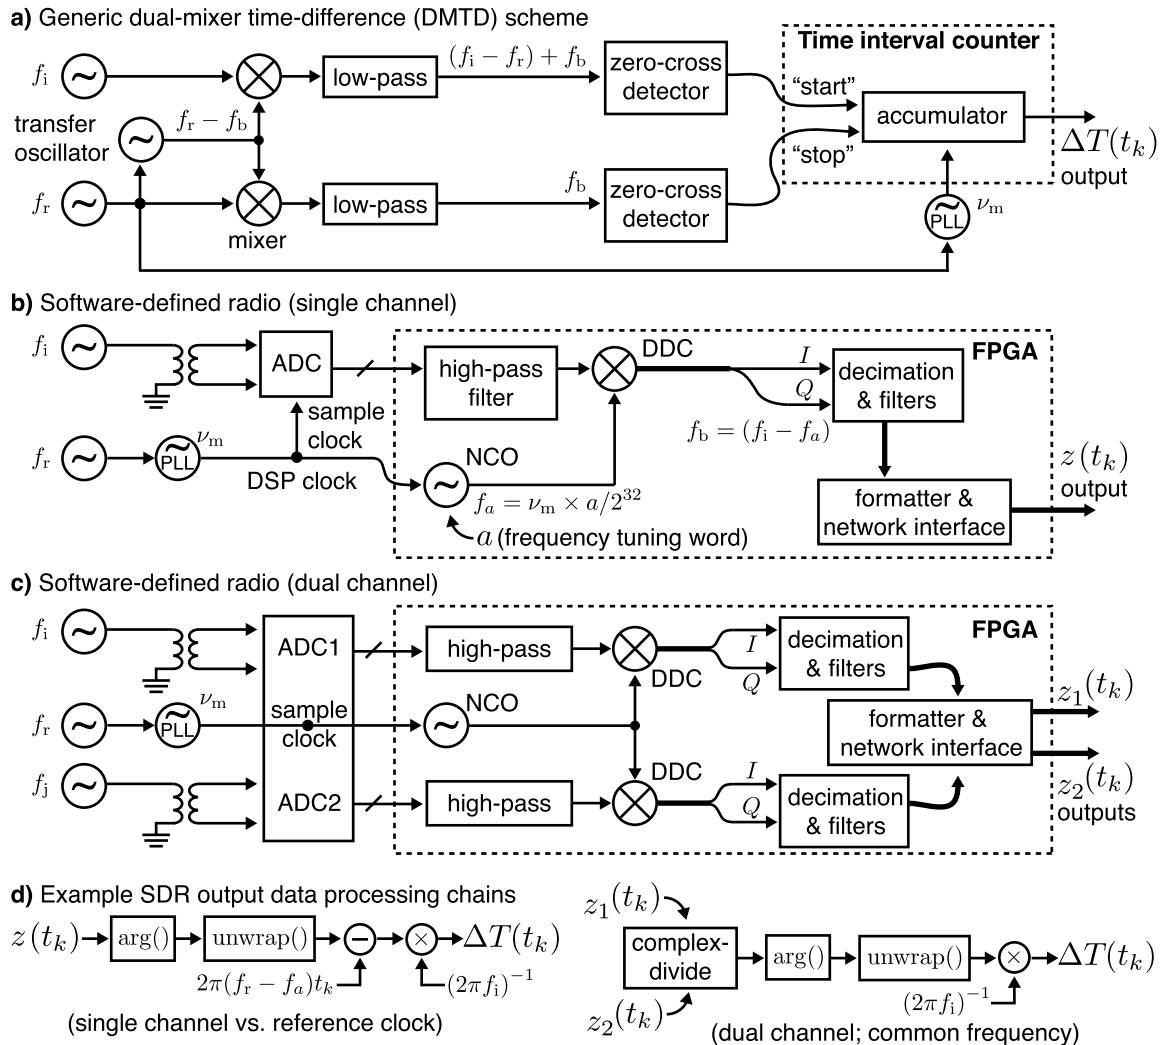


FIG. 1. Schematic comparison of (a) generic dual-mixer time-difference (DMTD), and (b) a software defined radio (SDR) described here. Dashed-lines surround digital processing sections. In both concepts, f_i is the frequency of an oscillator under test. A reference oscillator f_r disciplines a digital clock at fast frequency ν_m through a phase-locked-loop (PLL). Both methods gain resolution by spectrally shifting f_i to a low frequency f_b ; in SDR, the mixer analogue is digital downconversion (DDC) with a synthesized numerically controlled oscillator (NCO). (c) Two oscillators f_i and f_j are compared in two channels of a single ADC, suppressing noise due to the ν_m PLL, its reference tone f_r , and the ADC's aperture jitter. (d) While DMTD directly outputs time-offset data; further processing is performed on the SDR sampled waveform $z(t_k)$ with a computer to determine time offset. We illustrate two simplified processing chains for single- and dual-channel measurements; see text for details.

$\approx 200 \mu\text{m}$ in electrical length or group delay, so even cables contribute to instability (the temperature dependence is of order $0.5 \text{ ps m}^{-1} \text{ K}^{-1}$). Analog components (including “digital” mixers) can contribute flicker-phase noise,³¹ amplitude-to-phase-modulation conversion,⁹ sensitivity to interference (e.g., channel crosstalk, ground loops, wideband ambient rf), and coupling to the environment (e.g., temperature, humidity). The TIC start- and stop-inputs require high-bandwidth, high slew-rate triggers, but the signals following the mixers are slow sinusoids. Zero-crossing detectors must therefore boost signal slew rates by $\sim 10^6$ while accurately preserving phase.⁷ Without additional synthesis steps, DMTD requires f_r and f_i to be very similar, and among a small set of frequencies compatible with the analog processing components. Mixer and filter non-linearity and frequency-dependent group delay complicate maintaining a whole-system ΔT calibration over arbitrary signal frequencies. Finally, DMTD schemes cannot resolve fluctuations over time scales shorter than $1/f_b$.

B. Related work

Some limitations in DMTD can be addressed by replacing certain analog processing steps with digital implementations. The TIC can be dramatically redesigned³⁴ with much higher effective ν_m . One group⁴⁴ replaced the TIC by digitizing the mixed and filtered signals at f_b and later eliminated the mixers with high-speed direct sampling of the input signals.²⁹ Others have replaced mixer-based spectral down-conversion with aliasing through under-sampling.³⁶ Early consideration of a direct-sampling system²⁵ very similar to the present work showed plausible limits due to quantization effects alone can be $\Delta T < 1 \text{ fs } (\tau/1 \text{ s})^{-1/2}$ for averaging intervals τ . While high-speed samples can be processed entirely in software³⁹ or with custom hardware,¹⁷ this work explores oscillator metrology using an inexpensive, commercially available, unmodified software defined radio (SDR). We note a similar approach for characterizing ADCs.⁸ We employ an ADC noise cancellation

technique in the time domain, which perhaps is analogous to cross-spectral analysis³⁸ in the frequency domain.

II. SOFTWARE DEFINED RADIO

A. “Sample first, ask questions later”

In SDR,²⁸ signals of interest are sampled by a fast, high-resolution analog-to-digital converter (ADC) with little or no analog processing, amplification, or filtering. A numerically controlled oscillator (NCO), computed synchronously with ADC sampling, takes the place of the local oscillator tone in analog radio reception. A digital multiplication of the sampled signal and NCO phasor performs the role of signal mixer. Filtering and sample rate decimation are also performed digitally, reducing noise bandwidth while conserving signal information. Here we focus on SDR receiver functions, but many SDRs are capable of transmission as well. Since the signal processing chain in SDR is highly configurable, it has applications in radar, spread-spectrum, and multiple-input multiple-output (MIMO) communication, and advanced protocol demodulation and simulation.

SDR seems to suffer a significant disadvantage: noise figures of high-speed ADCs are much worse than a collection of radio frequency filters, amplifiers, and mixers. On the other hand—especially in the context of precision metrology—analogue components are subject to strict impedance matching requirements and exhibit long-term sensitivity to shock, vibration, supply voltage, temperature, humidity, aging, interference, and signal crosstalk. A low ADC signal-to-noise ratio (SNR) is at least amenable to averaging and process gain, while environmental sensitivities are more pernicious sources of stochastic noise and drift over long durations. In contrast, digital processing steps are stable, deterministic, and environmentally insensitive.

B. Technical details

At the time of writing, the techniques presented here ought to apply to SDRs from at least ten manufacturers. While we attempt to consider SDR generically, Figure 1(b) illustrates relevant components in the SDR receiver studied here (an Ettus USRP N210 except where noted¹). Field programmable gate array (FPGA) hardware description code and circuit schematics are available for inspection and customization.^{11,12}

A receiver daughterboard couples a ground-referenced input signal (1–250 MHz) into a differential ADC via a transformer. The ADC (Texas Instruments ADS62P44) has an analog input bandwidth of 450 MHz (−3 dB), a full-scale range of ±1 V, and 14-bit resolution. The ADC specifications include an aperture jitter $t_{ap} = 150$ fs, a significant technical timing uncertainty between an idealized sample trigger and actuation of the converter’s sample-and-hold circuitry. The sample timebase, a voltage-controlled crystal oscillator (VCXO) at $\nu_m = 100$ MHz, drives the ADC sampling trigger and the FPGA’s digital signal processing (DSP) pipeline. The SDR includes phase-locked loop (PLL, bandwidth ≈ 3 kHz) components, which we often use to discipline ν_m to a +14 dBm

signal at $f_r = 10$ MHz derived from an active hydrogen maser. In our configuration, this SDR consumes about 10 W of dc power.

SDR’s three important DSP tasks are frequency translation, filtering, and data decimation. After a high-pass filter suppresses the ADC’s zero-offset, the input signal undergoes digital down-conversion (DDC), or frequency translation by an NCO tuned to

$$f_a = \nu_m \times \frac{a}{2^{32}}, \quad (2)$$

where a is an integer $0 \leq a < 2^{32}$. As in direct digital synthesis (DDS), a phase register accumulates the frequency-tuning word a upon every ν_m clock cycle, the most significant bits of which are used to derive complex NCO phasor components. However, unlike many DDS implementations, SDR often does not use the phase register as an index in a large lookup table of precomputed trigonometric values. Instead, SDR often implements *coordinate rotation digital computer* (CORDIC)⁴⁷ to compute NCO phasors in fixed-point arithmetic. Exploiting the equivalence between angle rotation and phase accumulation, CORDIC is a successive approximation algorithm built from logical operations well suited to a DSP pipeline: comparisons, bit shifts, and binary addition. After inspection of the two most significant accumulator bits fixes the phase quadrant, this SDR implements $K = 20$ CORDIC iterations on 24-bit phase words for an approximate angle resolution of $\tan^{-1} [2^{-(K-1)}] = 1.9 \mu\text{rad}$. CORDIC approximates resampling the real input signal into a complex frame rotating at f_a , adding negligible quantization noise (approximately equivalent to $\sigma_x(\tau) = 0.3 \text{ fs} (\tau/1 \text{ s})^{-1/2}$ at 10 MHz). The SDR ultimately truncates the resulting signal to 16 bits of resolution for each of the in-phase (I) and quadrature-phase (Q) components.

Transmission and manipulation of output samples $z(t_k) = I(t_k) + iQ(t_k)$ at the physical sample rate ν_m would require ≈ 3 Gbps in network and storage resources. Fortunately, DDC shifts the signal of interest close to baseband, allowing aggressive low-pass filtering and rate-decimation (by up to a factor of 512) in hardware. The SDR filters and decimates in three configurable steps. First, a cascaded integrator-comb filter²¹ divides the sample rate by an integer $1 \leq n_{cic} \leq 128$. This is followed by two optional half-band decimators,⁵ each accomplishing a rate division of 2 and antialias filtering. Within their passbands these filters have a linear phase/frequency dependence and thus are shape-preserving in the time-domain. Figure 2 shows data acquired with typical settings, $n_{cic} = 25$ and both half-band filters enabled, which yields a decimation of $n_{dec} = 4n_{cic} = 100$ and $\nu_m/n_{dec} = 10^6$ samples per second, requiring 32 Mbps of network and buffering resources. ADC quantization noise power, which is nearly uniform in density (“white”) over a Nyquist bandwidth of $\pm \nu_m/2$ is reduced, approximately by n_{dec}^{-1} (see Appendix B). The final DSP section queues and formats $z(t_k)$ along with metadata such as hardware time-stamping and drives their transmission to a general-purpose data acquisition computer. Application programming interfaces are available for several languages, free tools like gnuradio,¹⁵ and commercial data processing packages.




FIG. 2. SDR measurement of a signal at $f_i = f_r = 10$ MHz, spectrally shifted by DDC to $f_b = 10 \text{ MHz} - (429,496,386 \times 2^{-32} \times 100 \text{ MHz}) \approx 8$ Hz. (a) One second (10^6 buffered samples) of $z(t_k) = I(t_k) + iQ(t_k)$ data. (b) The residual amplitude of $I(t_k)$ after removing a best-fit single-tone. The noise is predominately white, but modulation related to f_b and proportional to $|I(t_k)|$ is clearly observed. (c) The instantaneous phase evolves as $2\pi f_b t$; here we plot $\arg z(t_k)$ wrapped into $-\pi < \arg z(t_k) \leq \pi$.

III. DEMONSTRATION EXPERIMENTS

We now outline our study of SDR's suitability for oscillator metrology. We first discuss phase measurements over intervals of a few seconds, the analysis of which includes information about fast fluctuations up to $\nu_m/(2n_{\text{dec}})$ in frequency. Then, we consider measurement noise over several hours to days using methods which average over fast fluctuations. We find that over intervals greater than about 10 ms, ADC aperture jitter is likely a limiting technical noise source. We demonstrate a promising solution available in many SDRs: a second, independent ADC channel is synchronously sampled such that aperture jitter and many other noises subtract in common-mode. We consider application of SDR in a many-clock inter-comparison, and to the problem of optical frequency and phase metrology. Finally, we briefly describe measurement performance of a 6 GHz microwave tone beyond the ADC bandwidth, and the instability of amplitude measurements in two SDR models.

A. Phase of ADC input vs. the sampling timebase

Consider the arrangement in Figure 1(b) where f_i is approximately known and stationary, and f_r is treated as a frequency reference (f_i need not be similar to f_r). We choose the integer a so f_a (see Eq. (2)) is close to f_i . Absent technical noise, $\nu_m = 10f_i$ due to the master timebase's PLL, making f_a exactly computable. The SDR output samples, $z(t_k)$, represent the input signal spectrally shifted to a low frequency $f_b = f_i - f_a$; the sample epochs are $t_k = k \times (n_{\text{dec}}/\nu_m)$. The signal phase, $\arg z(t_k) \equiv \tan^{-1} [\text{Im} z(t_k)/\text{Re} z(t_k)]$ (see Figure 2(c)), is a time-integral of angular frequency $2\pi f_b$ and so evolves in time as $2\pi f_b t_k + \phi_0$, where ϕ_0 includes technical offsets such as cable delays. The \tan^{-1} function is evaluated with independent numerator and denominator arguments, removing

a phase-quadrant ambiguity. Generally, all SDRs are capable of this mode of measurement, though those without the ability to reference ν_m will suffer in accuracy.

To analyze the phase noise floor of this configuration, we split a single 10 MHz oscillator into the f_i and f_r inputs. The amplitude at the ADC is kept near half-scale to avoid distortion (typical input power was ≈ 0 dBm). Ideally, when unwrapped, $\arg z(t_k) = 2\pi(1 - 10a/2^{32})(kn_{\text{dec}}/10)$ (neglecting a fixed ϕ_0). In software, we subtract this deterministic trend, the magnitude of which is made small by an appropriate choice of a , and interpret residual fluctuations as measurement noise. Figure 3(a) shows a typical residual phase signal, a Fourier transform of which yields the phase noise power spectral density (Figure 3(b)).

The black curve in Figure 3(c) depicts a complementary statistical measure: the oscillator time deviation^{2,4} $\sigma_x(\tau) = \frac{\tau}{\sqrt{3}} \text{mod } \sigma_y(\tau)$, where $\text{mod } \sigma_y(\tau)$ is the modified Allan deviation.³⁵ Briefly, $\sigma_x(\tau)$ characterizes the predictability of phase (in time units, see Eq. (1)) as a function of averaging interval τ . Over roughly $20 \mu\text{s} < \tau < 200 \mu\text{s}$ we observe behavior consistent with white-phase noise, $\sigma_x(\tau) \approx 1.2 \text{ fs}(\tau/1\text{s})^{-1/2}$. Regrettably, $\sigma_x(\tau)$ stops decreasing with further averaging, and besides an oscillation peak (related to modulation at f_b) appears limited to a flicker-floor roughly consistent with the ADC's $t_{\text{ap}} = 150 \text{ fs}$. Reducing input power increases the white-phase instability, but otherwise these performance limits persist over many instrument configurations: rf-coupling method (dc-coupled op-amp vs. transformer), choice of heterodyne f_b , decimation factor n_{dec} , stock vs. quiet linear power supply, etc.

B. Phase of one ADC channel vs. another

To do better we must reduce the influence of phase noise in ν_m and the ADC aperture jitter t_{ap} . Fortunately, many

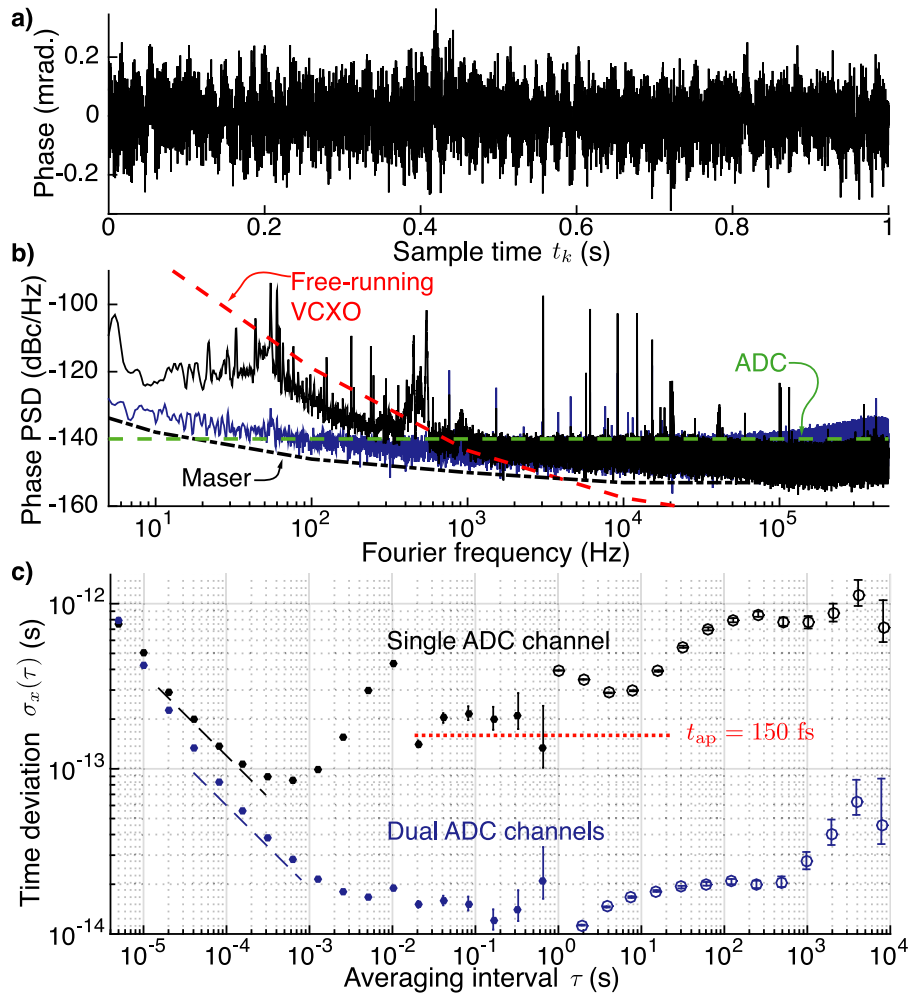


FIG. 3. (a) Unwrapped phase signal $\arg z(t_k)$ when f_i and f_r derive from the same 10 MHz oscillator (0.1 mrad corresponds to $\Delta T \approx 1.6$ ps); a deterministic ramp caused by the choice of f_a is removed. Here, $n_{\text{dec}} = 100$. (b) A Fourier transform of $\arg z(t_k)$ (black data) yields a single-sided phase noise power spectral density (PSD). At high Fourier frequencies we observe a white noise of ≈ -140 dBc/Hz (green dashed line), consistent with the ADC's SNR, signal power, and decimation filtering. At low Fourier frequencies we observe technical noise roughly tracking the rising noise density of the ν_m VCXO (red dashed line), which the PLL cannot fully suppress. A hydrogen maser noise specification (black dotted-dashed) provides context. The relative PSD between two ADC channels (blue data) has much improved flicker noise. (c) Time deviation $\sigma_x(\tau)$ in one-channel (black) and two-channel (blue) modes; $n_{\text{dec}} = 500$. Solid points derive from short streams of $\arg z(t_k)$ samples without averaging. Open circles result from pre-averaging streams in 1 s chunks. White phase noise of $1.2 \text{ fs}(\tau/1\text{s})^{-1/2}$ (black dashed line) is equivalent to ≈ 86 dB SNR (see Appendix A). The blue dashed line represents a further 6 dB improvement. A red dashed line marks the ADC's aperture jitter t_{ap} . For $\tau \gg 10$ s, we expect environmental coupling to dominate both measurement modes. See text for further detail.

SDRs can process two independent ADC channels which are sampled synchronously (specifically, the two ADC channels exist on the same chip). To examine residual noise in this differential configuration, we split the same oscillator into the three inputs (f_i , f_j , and f_r) as shown in Figure 1(c), though it is not crucial that the f_r input be identical to either of the others. Since $f_i = f_j$, the same NCO frequency f_a is used to DDC both channels, giving the same deterministic trend to both output phase signals. The phase signals should be unwrapped before subtraction because, due to noise and small phase offsets, 2π -discontinuities can appear at different sample epochs. Example single- and dual-channel signal processing chains are illustrated for comparison in Figure 1(d).

Blue curves in Figures 3(b) and 3(c) show the significant improvement in phase noise and time deviation from dual-channel operation. A flicker-floor of $\sigma_x \approx 20$ fs now appears roughly an order-of-magnitude below t_{ap} and persists over

$1 \text{ ms} < \tau < 500 \text{ s}$. It also improves by an order-of-magnitude upon the typical noise floor of the DMTD instrument ($\sigma_x \approx 300$ fs). While we lack detailed knowledge of the ADC, we posit that each channel's sample-and-hold circuitry shares a trigger-input threshold-detector. After this element, circuit paths, component/process variation, and environmental non-uniformity on the ADC chip are likely minute. ADC voltage-reference fluctuations and phase noise in the ν_m PLL (and its reference, f_r) are also highly common to both sampled channels. Remaining non-common elements include off-chip transmission lines, coupling transformers, and on-chip ADC sampling circuitry. We show later that similar common-mode suppression is present in a different ADC with much larger $t_{\text{ap}} = 1$ ps.

In this mode of operation, it is less important that ν_m be locked to a high-quality oscillator because phase noise in ν_m will be highly common between the two sampled inputs. Noise is not completely suppressed, however. We found slightly

better performance, at the level of 20% in σ_x , when ν_m was referenced to a hydrogen maser versus the SDR's quartz oscillator. We hypothesize that parasitic coupling of the digital sample clock at ν_m is slightly imbalanced between the two ADC inputs. This feature is likely specific to the SDR model and circuit layout.

C. Instability over long averaging intervals

Maximum decimation in an SDR still results in several megabits per second of data per channel. As a practical matter for long-duration measurements, we reduce this data stream as it is acquired to one recorded ΔT value per second. This step reduces measurement bandwidth to ≈ 1 Hz. We tested two simple averaging methods with similar performance: uniform weight (“rectangular window”) averaging of $\arg z(t_k)$ over groups of $N = \nu_m/n_{\text{dec}}$ samples per second, and the phase estimation routine discussed in [Appendix A](#). Issues related to windowed averaging here are analogous to those in frequency meters.³⁷

The SDR measurement stability does not degrade much over intervals of several hours, an important requirement for an atomic-clock measurement system.²⁶ We undertook no special environmental stabilization beyond standard laboratory conditions (ambient temperature control of ≈ 0.5 K). The SDR operated in its original enclosure with a continuously active cooling fan. For these tests, matched cables were flexible, double-shielded (RD-316), and SMA terminated. Open circles in [Figure 3\(c\)](#) show typical long-term performance of the one- and two-channel SDR techniques. In terms of frequency instability, the two-channel ADC residuals at 10 MHz typically average as flicker-phase noise with $\sigma_y(\tau) = 7 \times 10^{-14}(\tau/1 \text{ s})^{-1}$ through $\tau = 10^3$ s.

D. Clock comparison with software radio

The tests described above demonstrate the low instability of the SDR technique; here we discuss time accuracy. Two 5 MHz signals, sourced by hydrogen masers (NIST masers ST0010 and ST0014), are input into the two SDR ADC channels. A non-linear frequency doubler converted one of these to create the $f_r = 10$ MHz PLL reference. The maser signals were measured simultaneously by a commercial system based on DMTD. [Figure 4\(a\)](#) shows excellent agreement between the methods. The time-series of the difference between the data sets ([Figure 4\(b\)](#)) reveals technical noise in one or both measurement systems, some details of which are not yet understood. An initial transient of about 15 ps in magnitude is a repeatable “warm-up” SDR characteristic lasting several minutes. Key component temperatures, measured with platinum resistive thermometers attached with thermally conductive epoxy, increase by 5 K to 10 K in these first several minutes of operation. We also observe a periodic variation (of roughly 31 h) with an amplitude of order 5 ps. Such a variation would contribute $< 10^{-16}$ to fractional frequency instability, which is of marginal significance in the inter-comparison of maser clocks. [Figure 5](#) shows the frequency instability of the maser comparisons and typical SDR and DMTD residual instabilities. At averaging intervals of $\tau \approx 10^3$ s, the single-channel SDR technique is comparable with the commercial DMTD instrument; the two-channel SDR technique outperforms both by almost an order-of-magnitude.

E. Multi-channel operation

Some commercial DMTD instruments accept 16 or more input oscillators, where one channel is permanently assigned

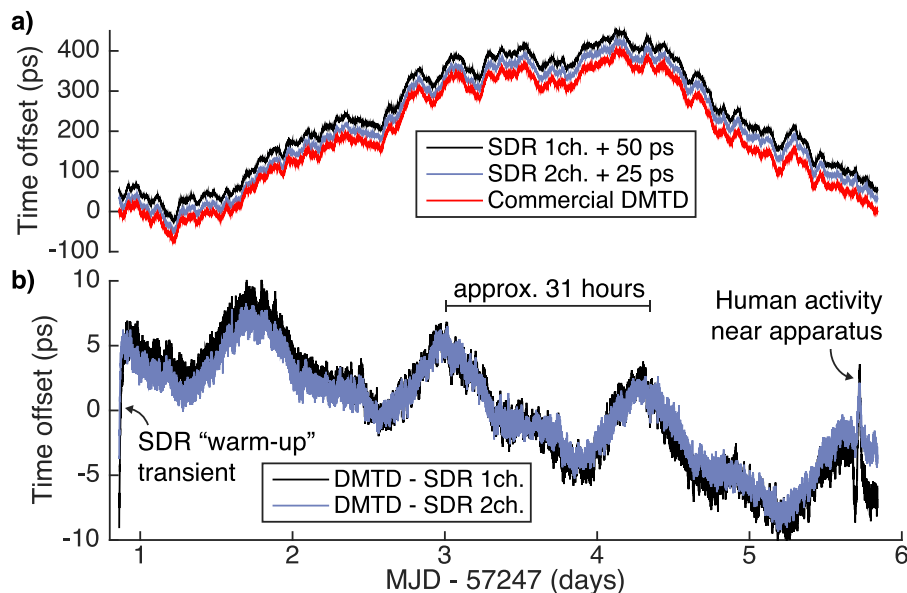


FIG. 4. (a) A comparison of two hydrogen masers over five days (MJD is the modified Julian date), using a commercial instrument based on DMTD (red) and the SDR described here (black/blue for single-/dual-channel mode, respectively). From each time series we subtract a linear phase trend corresponding to the masers' frequency difference of $y = 8.85 \times 10^{-14}$. We introduce 25 ps and 50 ps offsets for visual clarity. (b) We show the differences of each SDR measurement with that of the DMTD. Some technical noise features are understood and annotated; it is not yet known whether the DMTD, SDR, or both systems contribute to the ~ 31 h periodic modulation.

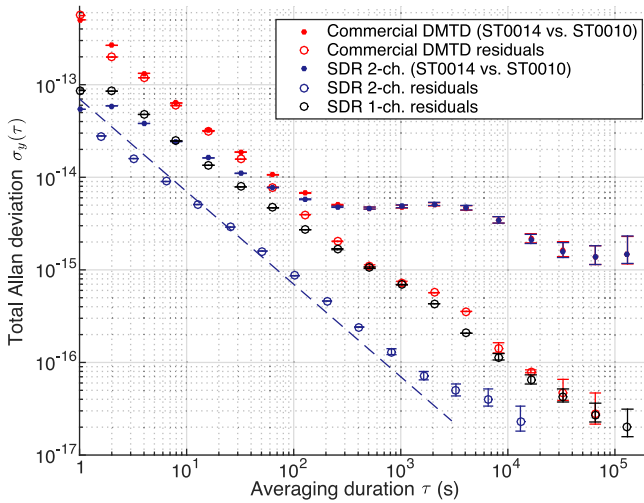


FIG. 5. Fractional frequency instability $^{16} \sigma_y(\tau)$ of hydrogen masers (NIST masers ST0014 vs. ST0010) as measured by a DMTD commercial instrument (red, solid) and the SDR two-channel technique (blue, solid) described here. From $\tau \geq 200$ s both techniques become identically limited by maser frequency fluctuations. Open points show typical residual instabilities of the DMTD instrument (red), the single-channel SDR method (black), and the two-channel SDR method (blue). The blue dashed line is an eye guide placed at $\sigma_y(\tau) = 7 \times 10^{-14} (\tau/1 \text{ s})^{-1}$. Both DMTD and SDR methods yield one datum per second, but the effective measurement bandwidth of the DMTD instrument is known to be $\gg 1$ Hz.

a special role as reference for the TIC timebase ν_m . The two-channel SDR scheme presented here is scalable to an unlimited number of channels, and it is possible but not necessary that one oscillator be assigned a special role. Figure 6(a) sketches a scheme whereby multiple SDR instruments are arranged in a “ring,” immune to any single oscillator or SDR fault. A “hub” model (Figure 6(b)), where one oscillator is distributed to all measurement nodes is also possible. Simultaneous implementation of the one-channel SDR technique using a distinct f_r oscillator provides a “backup hub” mode of operation with degraded performance. In a scaled deployment, it may be desirable to increase the decimation performed in hardware, perform the phase computation and averaging itself in the FPGA, and/or distribute the software data processing among multiple connected computers. We estimate that, per measurement channel, the material cost of a SDR solution is a factor of two or more below competitive multi-channel DMTD instruments.

F. Optical oscillator measurement

Optical atomic frequency references now exceed the performance of official primary standards based on microwave frequencies by factors of 1000 in stability²⁰ and potentially 100 in accuracy.^{6,19,46} Generally, optical frequency references operate by disciplining a pre-stabilized laser oscillator to an atomic resonance in neutral atoms or single trapped ions.³² Direct phase and frequency comparisons between two lasers at f_α and f_β are only possible if they are sufficiently close to create a heterodyne beatnote on a photodiode or other transducer. Otherwise, a now standard technique employs a broadband femtosecond laser frequency comb (FLFC, or comb) as a common heterodyne oscillator spanning hundreds

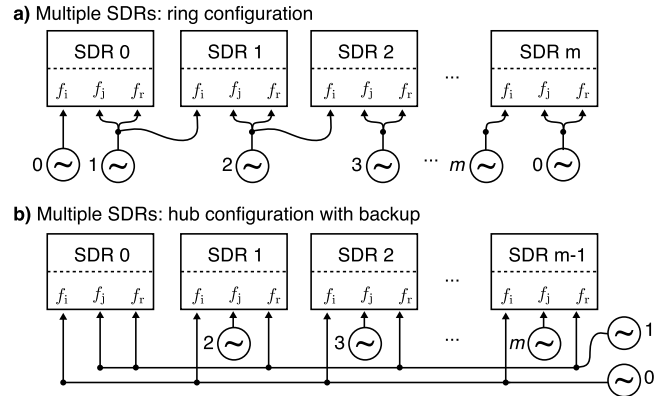


FIG. 6. Multiple SDRs can scale for coherent many-oscillator comparisons in flexible arrangements, the choice of which will depend on which failure modes are judged most likely. (a) For example, in a “ring” configuration, each SDR node produces the two-channel differential signal $f_j - f_i$, a one-channel signal $f_i - f_r$, and unique one-channel residual $f_j - f_r$. Phase data collection for all oscillators is uninterrupted with any single node failure. (b) In a “hub” configuration, the oscillator indexed “0” is distributed to an ADC channel in each node as part of a two-channel differential measurement. To protect the network against failure of oscillator “0,” oscillator “1” provides a shared PLL reference to all nodes, enabling one-channel measurements of all oscillators as a “degraded backup.” Here, junctions imply distinct distribution amplifier channels; differential amplifier and cable delays must be accounted for when comparing oscillator phase differences.

of terahertz.^{30,45} A FLFC spectrum consists of many optical modes, whose absolute frequencies can be expressed as $f_n = n f_{\text{rep}} + f_{\text{ceo}}$ for many thousands of consecutive integers n . The comb’s pulse repetition rate, f_{rep} , scales inversely with the laser resonator length, and $|f_{\text{ceo}}| < f_{\text{rep}}$ depends on the details of the intra-cavity dispersion. For our purposes, it is sufficient to note that both degrees of freedom correspond to radio frequencies controllable by phase-lock techniques.

We measured and tracked phase fluctuations between two laser oscillators using a FLFC and the SDR. A Ti:sapphire FLFC¹³ with $f_{\text{rep}} \approx 1$ GHz was stabilized by locking a comb mode 640 MHz offset from an ultra-stable optical frequency $f_\alpha \approx 259$ THz. We used self-referencing interferometry¹⁰ to stabilize f_{ceo} . A second ultra-stable laser,¹⁴ $f_\beta \approx 282$ THz, interfered with another comb mode to make a heterodyne tone f_o near 160 MHz on an amplified photodiode. Independent characterizations have determined frequency instability floors of $\leq 2 \times 10^{-16}$ for laser α and 1×10^{-15} for laser β . Due to the comb’s phase locks, fluctuations of f_o are directly related to the fluctuations between the α and β laser oscillators; the required comb mode integers for absolute determinations can be obtained by low-resolution wavemeter measurements of f_α and f_β .

Traditionally, only gated frequency measurements are made of f_o , discarding information about phase fluctuations. A DMTD scheme to track phase is impractical: generally, f_o can appear at any frequency up to $f_{\text{rep}}/2$ depending on FLFC preparation, and f_o fluctuations and drift are typically too large. In contrast, the SDR has a high input bandwidth, a tunable NCO for down-converting arbitrary f_o , and tracks phase information over very short intervals $\nu_m/n_{\text{dec}} \leq 5 \mu\text{s}$ with no dead time.

Since the ADC sample clock $\nu_m = 100$ MHz, $f_o \approx 160$ MHz appears in the third $\pm \nu_m/2$ Nyquist zone, aliased to




FIG. 7. Tracking optical phase with SDR (see text for details). (a) The SDR down-converts a heterodyne between a femtosecond laser frequency comb (FLFC) stabilized to reference laser f_a , and laser f_b to an audio tone of ≈ 140 Hz. We plot the complex components of the SDR output $z(t_k)$. (b) Laser β is transmitted to the FLFC heterodyne via an uncompensated fiber optic link. By shaking the fiber, we observe and can coherently track resulting phase fluctuations. (c) Dividing $\arg z(t_k)$ by $2\pi f_b$, we cast phase fluctuations as time instability of the optical oscillator β . A constant phase and frequency offset are suppressed in the plot.

$-40.005\,860$ MHz. We set the NCO $f_a = -40.006\,000$ MHz in order to obtain an audio beat note $|f_b| \approx 140$ Hz. Figure 7(a) shows the output sample data under normal conditions; Figure 7(b) shows directly observable phase noise created by vigorously shaking the uncompensated²⁷ fiber optics coupling laser β to the comb. It is important to appreciate that a radian of optical phase remains unscaled by mixing with the comb to make f_o , nor is it scaled by the DDC process $f_o \rightarrow f_b$ in the SDR. So, treating laser α as a reference, we can derive the time fluctuations of laser β by unwrapping and dividing the f_b phase $\arg z(t_k)$ by a factor $2\pi \times 282$ THz, following Eq. (1). Figure 7(c) shows the result: well-resolved femtosecond-level temporal instability between two would-be optical clocks, lasers α and β . In measurement of optical heterodyne tones, the SDR noise floor is negligible.

A multi-channel SDR arrangement monitoring several FLFC heterodyne beat notes could form the measurement basis for an *optical time scale*, meaning an ensemble of optical oscillators statistically weighted to produce a robust and reliable “average clock.”^{26,43} Related technology is approaching a high level of readiness, including robust fiber-FLFC designs,⁴⁰ stabilized “flywheel” lasers^{18,24} with frequency instabilities $\sigma_y \leq 1 \times 10^{-16}$, and optical atomic standards characterized at the 10^{-18} uncertainty level.⁶

G. Microwave frequencies

Microwave frequencies far beyond the ADC input bandwidth are measurable by SDR models that incorporate an analog mixer and microwave local-oscillator (LO) synthesizer referenced to the same source as ν_m . In a separate investigation,

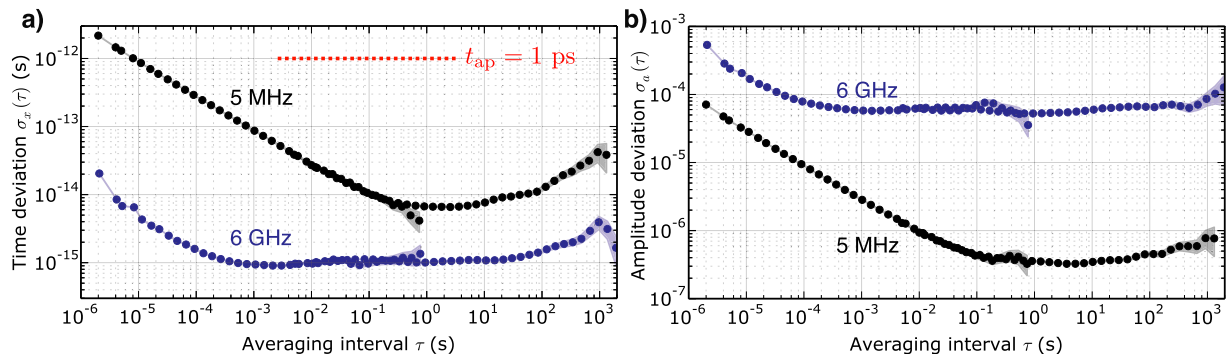


FIG. 8. (a) Time deviation of differential phase measurements of 5 MHz and 6 GHz signals. For 5 MHz, the SDR featured a 12-bit ADC with 1 ps aperture jitter. The instability floor is more than two orders of magnitude lower than t_{ap} , indicating excellent common-mode suppression of technical noise. For 6 GHz, the SDR featured a LO synthesizer and analog mixer front-end to translate the signal into the ADC bandwidth. Though phase-noise performance is made worse by these elements, the high signal frequency leads to a time stability floor of 1 fs, roughly an order-of-magnitude better than the results at 10 MHz (Figure 3(c)). (b) We also investigated amplitude measurement instability (normalized to input amplitude) of two-channel signals in these SDR models. In both plots, data at longer τ are obtained by additional software decimation by a factor of 2500 prior to storage. These data were acquired in an unstabilized office environment and with the sample clock ν_m un-referenced. Shaded bands indicate standard statistical uncertainties.

we tested a SDR (Ettus USRP B210) featuring such a front-end (Analog Devices AD9361) capable of down-converting two ≤ 6 GHz signals before sampling them at 12-bit resolution. To characterize its phase-noise performance, we input a 6 GHz (-22 dBm) signal and set the SDR's programmable amplifiers to 49 dB to use the full ADC range. ν_m was set to 30.72 MHz, and n_{dec} to 32. Due to the mixer front-end, we observed significantly higher phase noise than the results in Section III B: a white noise floor at -123 dBc/Hz and flicker noise of -90 dBc/Hz $(f/1 \text{ Hz})^{-1}$. However, given the much higher carrier frequency, the equivalent time deviation limits were $\sigma_x(\tau) = 20$ as $(\tau/1 \text{ s})^{-1/2}$ over short intervals and a flicker floor of 1 fs, as shown in Figure 8(a).

H. Amplitude metrology

Though we have so far ignored it, the amplitude of a complex sampled SDR signal is also available as $\sqrt{I^2(t_k) + Q^2(t_k)}$. In a separate investigation, we studied the relative amplitude instability limit of signals input into two ADC channels. We tested a SDR (Ettus USRP B100) with a 12-bit ADC (Analog Devices AD9862), $\nu_m = 64$ MHz, $n_{\text{dec}} = 64$, and $f_i = f_j = 5$ MHz. As shown in Figure 8(b), we observed a relative amplitude instability floor of 3×10^{-7} over the averaging interval $0.1 \text{ s} \leq \tau \leq 100 \text{ s}$. The 6 GHz configuration, described in Section III G, achieved an amplitude instability floor of 5×10^{-5} .

IV. CONCLUSIONS

Generally, SDR receivers are little more than high-speed signal samplers followed by a series of digital filters designed to reduce data rate and noise bandwidth. However, these few ingredients are sufficient for several recipes in high-precision time and frequency metrology. Phase/time-offset measurements using unmodified SDR hardware can exceed the stability performance of a commercially available instrument based on the classic DMTD design while offering increased flexibility. SDR measurement of phase using two input channels differentially reduces the influence of technical timing noise and has demonstrated a maser clock frequency resolution $\sigma_y \leq 10^{-16}$ within 10^3 s of averaging. Over several days of continuous hydrogen maser measurement, the SDR technique appears highly accurate, with relatively low environmental noise coupling in a typical laboratory environment. SDR hardware is scalable to coherently measure any number of oscillators at almost any radio or microwave frequency. We have shown the SDR can resolve relative oscillator amplitude fluctuations below the part-per-million level. Finally, we have demonstrated that SDR can be usefully employed in the comparison of ultra-stable optical clocks and oscillators by measuring heterodyne products of clocks with a femtosecond laser frequency comb.

Useful extensions of this work could include a long-term frequency comparison of atomic-clocks' output signals at multiple frequencies (e.g., 5 MHz and 100 MHz), and integration of a many-channel fast ADC into an SDR architecture for better multi-channel scalability. Alternatively, the transmission functions of the SDR could be employed in active

phase-noise compensation in optical or FLFC interferometry applications.

NIST's Time and Frequency Division funded this investigation. The work is a contribution of NIST and not subject to U.S. copyright. The authors thank Judah Levine for helpful discussions, Joshua Savory (maser comparisons), Franklyn Quinlan (FLFC measurements), and Roger Brown for careful reading of the manuscript.

APPENDIX A: SPECTRAL ESTIMATION OF FREQUENCY AND PHASE

In the single-channel setup of Figure 1(b), a f_i known only to within a Nyquist bandwidth $\nu_m/2$ can be quickly acquired by seeking high spectral power while scanning the NCO f_a over its full range. Without loss of generality, we suppose $f_i < \nu_m/2$ and choose the calculable NCO frequency f_a such that $|f_i - f_a| \ll \nu_m/n_{\text{dec}}$; in other words, the DDC frequency must be within the decimated Nyquist zone. The sign of the sampled "beatnote" $f_b = f_i - f_a$ is fixed by the sense of temporal phase rotation in $z(t_k)$, or equivalently, the phase relationship of its real and imaginary components. The problem of high resolution determination of f_i reduces to spectral estimation on groups of N samples of $z(t_k)$ to estimate f_b . Though no closed-form solution exists generally for spectral estimation,²³ our circumstances are unusually favorable: $z(t_k)$ consists of a single, low-frequency tone f_b , with high SNR and little harmonic distortion. Though computationally intensive, an optimal un-biased frequency estimator given these assumptions is the argument \hat{f}_b maximizing the basic periodogram function⁴¹ $|P(f)|^2$, where

$$P(f) = \frac{1}{N} \sum_{k=0}^{N-1} z(t_k) e^{-i2\pi f t_k}. \quad (\text{A1})$$

For signals like ours, $|P(f)|$ is well-approximated by a quadratic polynomial near its maximum. We therefore implemented Brent's method of one-dimension parabolic interpolation³³ to efficiently search for \hat{f}_b . A lower resolution FFT-based spectral estimator seeds this non-linear search with an initial guess. Unlike such FFT-based methods, no windowing function or zero-padding must be applied to the sampled data prior to $|P(f)|^2$ maximization, and there is no need to make N a power-of-2. The search also yields an optimal estimator for the single-tone amplitude, $\hat{A}_b = |P(\hat{f}_b)|$. In the limit of high SNR and spectrally uniform uncorrelated ("white") noise, periodogram maximizing spectral estimates converge with maximum likelihood and non-linear least-squares fit results.

Figure 5 (black open circles) shows measured frequency instability of a $f_i = f_r = 10$ MHz signal, which surpasses that of a commercial frequency meter of comparable cost. Importantly, note that the SDR measurement instability decreases as τ^{-1} , compared to many frequency meters' instability $\propto \tau^{-1/2}$. This difference is attributable to non-zero dead-time and frequency quantization in commercial meter readings. The interval $N(\nu_m/n_{\text{dec}})^{-1}$ is analogous to a "gate interval" of a traditional frequency meter. With SDR, this parameter may be chosen during or after data acquisition since $z(t_k)$ data can be stored. Barring interruption in data transmis-

sion, this method of frequency analysis has zero “dead-time” intervals during which the input oscillations are unmeasured.

We continue the spectral estimation method to determine phase offset measurements from sets of N complex waveform samples $z(t_k)$. If unknown, we first find the \hat{f}_b maximizing the periodogram function $|P(f)|^2$ from Eq. (A1). Then, the optimal estimate of the signal’s phase is

$$\hat{\phi}_b = \tan^{-1} \left[\frac{\text{Im } P(\hat{f}_b)}{\text{Re } P(\hat{f}_b)} \right]. \quad (\text{A2})$$

Successive estimates of phase on continuously sampled data will evolve as

$$\hat{\phi}_b(t_k) = \phi_0 + 2\pi f_b t_k \quad (\text{A3})$$

$$= \phi_0 + 2\pi(f_i - f_r)t_k + 2\pi f_r \left(1 - \frac{v_m}{f_r} \frac{a}{2^{32}}\right) t_k, \quad (\text{A4})$$

where ϕ_0 is the initial phase offset and, for the SDR described in Section II B, $v_m/f_r = 10$. The final term, the result of our choosing a heterodyne offset frequency, is exactly computable in terms of f_r and is removable in post-processing. Subtracting it using complex phase rotation neatly avoids 2π discontinuities, leaving us only with a phase growing linearly with the frequency difference of interest $f_i - f_r$. Phase discontinuities must still be expected and handled over time intervals $\tau \geq 2\pi/(f_i - f_r)$. The variance of a single $\hat{\phi}_b$ estimate using $N \gg 1$ samples is bounded by²³

$$\text{var}(\hat{\phi}_b) \geq \frac{1}{\text{SNR}} \frac{2}{N}. \quad (\text{A5})$$

As $\text{SNR} \propto 1/N$ itself (due to process gain), the bound for variance in the phase estimator is independent of the sample density N under optimal noise conditions, remaining inversely proportional to the SNR and total observation duration. Combining this result with Eq. (1), the resulting theoretical bound on time deviation⁴² is

$$\sigma_x(\tau) = \frac{1}{2\pi f_i} \sqrt{\text{var}(\hat{\phi}_b)} \geq 1.2 \times 10^{-15} \text{ s } (\tau/1 \text{ s})^{-1/2}, \quad (\text{A6})$$

where $f_i = 10$ MHz, $N = 10^6$ samples per second, and the effective $\text{SNR} \approx 86$ dB (see Appendix B). Observations in Figure 3 (black solid points) are consistent with this noise limit over short averaging intervals.

APPENDIX B: DECIMATION FIDELITY IN PRACTICE

Ideally, in the presence of uniform Gaussian noise, the SNR of ADC samples should be improved by a factor of the decimation ratio n_{dec} since the CIC and half-band decimating filters approximate an ideal low-pass filter. Alternatively, with SNR expressed in dB,

$$\text{SNR}_{(\text{ideally observed})} = \text{SNR}_{\text{ADC}} + 10 \log n_{\text{dec}}. \quad (\text{B1})$$

However, finite precision in the numerical filters, and the presence of non-Gaussian noise, such as spurs and input noise near the sample clock v_m , results in slightly worse performance. We observe an approximate $n_{\text{dec}}^{0.8}$ improvement with $20 \leq n_{\text{dec}} \leq 500$ as shown in Table I.

TABLE I. Decimating low-pass filters in the SDR ideally improve SNR proportionally to the decimation factor n_{dec} . We show a slightly worse empirical scaling $\propto n_{\text{dec}}^{0.8}$. Here we compare the measured SNR for a constant, half-scale, $f_i = f_r = 10$ MHz maser-referenced tone under different decimation settings. $f_b \approx 8$ Hz; other choices yielded similar results.

n_{dec}	Expected SNR (dB)	Observed SNR (dB)	Excess noise (dB)
20	81.5	74.5	7.0
40	84.5	76.7	7.8
100	88.5	79.8	8.7
200	91.5	82.1	9.4
500	95.5	85.2	10.3

¹Certain commercial equipment, instruments, or materials are identified in this paper for informational purposes only. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

²D. W. Allan and J. Barnes, “A modified ‘Allan variance’ with increased oscillator characterization ability,” in *Thirty Fifth Annual Frequency Control Symposium. 1981* (IEEE, 1981), pp. 470–475.

³D. W. Allan and H. Daams, “Picosecond time difference measurement system,” in *29th Annual Symposium on Frequency Control* (IEEE, 1975), Vol. 1, pp. 404–411.

⁴D. W. Allan and H. Hellwig, “Time deviation and time prediction error for clock specification, characterization, and application,” in *IEEE 1978 Position Location and Navigation Symposium* (IEEE, 1978), Vol. 1, pp. 29–36.

⁵M. G. Bellanger, J. L. Daguët, and G. P. Lepagnol, “Interpolation, extrapolation, and reduction of computation speed in digital filters,” *IEEE Trans. Acoust., Speech, Signal Process.* **22**(4), 231–235 (1974).

⁶B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye, “An optical lattice clock with accuracy and stability at the 10^{-18} level,” *Nature* **506**(7486), 71–75 (2014).

⁷G. Brida, “High resolution frequency stability measurement system,” *Rev. Sci. Instrum.* **73**(5), 2171–2174 (2002).

⁸A. C. Cardenas-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 *2015 Joint Conference of the IEEE International Frequency Control Symposium and the European Frequency and Time Forum (FCS)* (IEEE, 2015), pp. 676–680.

⁹G. Cibiel, M. Régis, E. Tournier, and O. Llopis, “AM noise impact on low level phase noise measurements,” *IEEE Trans. Ultrason. Eng.* **49**(6), 784–788 (2002).

¹⁰S. T. Cundiff and J. Ye, “Colloquium: Femtosecond optical frequency combs,” *Rev. Mod. Phys.* **75**(1), 325 (2003).

¹¹See <http://files.ettus.com/> for Ettus Research circuit schematics and technical information.

¹²See <https://github.com/EttusResearch/uhd> for Ettus Research code repository.

¹³T. M. Fortier, A. Bartels, and S. A. Diddams, “Octave-spanning Ti: Sapphire laser with a repetition rate > 1 GHz for optical frequency measurements and comparisons,” *Opt. Lett.* **31**(7), 1011–1013 (2006).

¹⁴T. M. Fortier, A. Rolland, F. Quinlan, F. N. Baynes, A. J. Metcalf, A. Hati, A. Ludlow, N. Hinkley, M. Shimizu, T. Ishibashi *et al.*, “Digital-photonics synthesis of ultra-low noise tunable signals from RF to 100 GHz,” *Laser Photonics Rev.* (to be published), e-print [arXiv:1506.03095](https://arxiv.org/abs/1506.03095) (2015).

¹⁵GNURadio: The free and open software radio ecosystem, GNURadio Team, April 2015, <http://gnuradio.org/>.

¹⁶C. Greenhall, D. Howe, and D. B. Percival, “Total variance, an estimator of long-term frequency stability,” *IEEE Trans. Ultrason. Eng.* **46**(5), 1183–1191 (1999).

¹⁷J. Grove, J. Hein, J. Retta, P. Schweiger, W. Solbrig, and S. R. Stein, “Direct-digital phase-noise measurement,” in *Proceedings of the 2004 IEEE International Frequency Control Symposium and Exposition* (IEEE, 2004), pp. 287–291.

¹⁸S. Häfner, S. Falke, C. Grebing, S. Vogt, T. Legero, M. Merimaa, C. Lisdat, and U. Sterr, “ 8×10^{-17} fractional laser frequency instability with a long room-temperature cavity,” *Opt. Lett.* **40**(9), 2112–2115 (2015).

- ¹⁹T. P. Heavner, E. A. Donley, F. Levi, G. Costanzo, T. E. Parker, J. H. Shirley, N. Ashby, S. Barlow, and S. R. Jefferts, "First accuracy evaluation of NIST-F2," *Metrologia* **51**(3), 174 (2014).
- ²⁰N. Hinkley, J. A. Sherman, N. B. Phillips, M. Schioppo, N. D. Lemke, K. Beloy, M. Pizzocaro, C. W. Oates, and A. D. Ludlow, "An atomic clock with 10^{-18} instability," *Science* **341**(6151), 1215–1218 (2013).
- ²¹E. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoust., Speech, Signal Process.* **29**(2), 155–162 (1981).
- ²²J. Jespersen and J. Fitz-Randolph, *From Sundials to Atomic Clocks: Understanding Time and Frequency* (Courier Corporation, 1999).
- ²³S. M. Kay, *Modern Spectral Estimation: Theory and Application*. 1988 (Prentice Hall, 1999).
- ²⁴T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye, "A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity," *Nat. Photonics* **6**(10), 687–692 (2012).
- ²⁵G. Paul Landis, I. Galysh, and T. Petsopoulos, "A new digital phase measurement system," in *33rd Annual Precise Time and Time Interval (PTTI) Meeting*, Technical Report, DTIC Document, 2001.
- ²⁶J. Levine, "Invited review article: The statistical modeling of atomic clocks and the design of time scales," *Rev. Sci. Instrum.* **83**(2), 021101 (2012).
- ²⁷L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall, "Delivering the same optical frequency at two places: Accurate cancellation of phase noise introduced by an optical fiber or other time-varying path," *Opt. Lett.* **19**(21), 1777–1779 (1994).
- ²⁸J. Mitola, "The software radio architecture," *Commun. Mag., IEEE* **33**(5), 26–38 (1995).
- ²⁹K. Mochizuki, M. Uchino, and T. Morikawa, "Frequency-stability measurement system using high-speed ADCs and digital signal processing," *IEEE Trans. Instrum. Meas.* **56**(5), 1887–1893 (2007).
- ³⁰N. R. Newbury, "Searching for applications with a fine-tooth comb," *Nat. Photonics* **5**(4), 186–188 (2011).
- ³¹J. Phillips and K. Kundert, "Noise in mixers, oscillators, samplers, and logic an introduction to cyclostationary noise," in *CICC. Proceedings of the IEEE 2000 Custom Integrated Circuits Conference* (IEEE, 2000), pp. 431–438.
- ³²N. Poli, C. W. Oates, P. Gill, and G. M. Tino, "Optical atomic clocks," *Riv. Nuovo Cimento* **36**(12), 555–624 (2013).
- ³³W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C* (Cambridge University Press, 1996), Vol. 2.
- ³⁴I. Prochazka, P. Panek, and J. Kodet, "Note: Precise phase and frequency comparator based on direct phase-time measurements," *Rev. Sci. Instrum.* **85**(12), 126110 (2014).
- ³⁵W. J. Riley, "Handbook of frequency stability analysis," NIST Special Publication 1065 (U.S. Government Printing Office, Washington, DC, 2008).
- ³⁶S. Römisch, S. R. Jefferts, and T. E. Parker, "A digital time scale at the National Institute for Standards and Technology," in *General Assembly and Scientific Symposium, 2011 XXXth URSI* (IEEE, 2011), pp. 1–4.
- ³⁷E. Rubiola, "On the measurement of frequency and of its sample variance with high-resolution counters," *Rev. Sci. Instrum.* **76**(5), 4703 (2005).
- ³⁸E. Rubiola and F. Vernotte, "The cross-spectrum experimental method," e-print [arXiv:1003.0113](https://arxiv.org/abs/1003.0113) (2010).
- ³⁹L. B. Ruppalt, D. R. McKinstry, K. C. Lauritzen, A. K. Wu, S. Phillips, S. H. Talisa *et al.*, "Simultaneous digital measurement of phase and amplitude noise," in *IEEE International Frequency Control Symposium (FCS)* (IEEE, 2010), pp. 97–102.
- ⁴⁰L. C. Sinclair, I. Coddington, W. C. Swann, G. B. Rieker, A. Hati, K. Iwakuni, and N. R. Newbury, "Operation of an optically coherent frequency comb outside the metrology lab," *Opt. Express* **22**(6), 6996–7006 (2014).
- ⁴¹P. Stoica and R. L. Moses, *Spectral Analysis of Signals* (Pearson/Prentice Hall Upper Saddle River, New Jersey, 2005).
- ⁴²D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, "Characterization of clocks and oscillators," in *US Department of Commerce, National Institute of Standards and Technology, NIST Technical Note 1337* (U.S. Government Printing Office, Washington, DC, 1990).
- ⁴³P. Tavella and C. Thomas, "Comparative study of time scale algorithms," *Metrologia* **28**(2), 57 (1991).
- ⁴⁴M. Uchino and K. Mochizuki, "Frequency stability measuring technique using digital signal processing," *Electron. Commun. Jpn. (Part I: Commun.)* **87**(1), 21–33 (2004).
- ⁴⁵Th. Udem, R. Holzwarth, and T. W. Hänsch, "Optical frequency metrology," *Nature* **416**(6877), 233–237 (2002).
- ⁴⁶I. Ushijima, M. Takamoto, M. Das, T. Ohkubo, and H. Katori, "Cryogenic optical lattice clocks," *Nat. Photonics* **9**(3), 185–189 (2015).
- ⁴⁷J. E. Volder, "The CORDIC trigonometric computing technique," *IRE Trans. Electron. Comput.* **EC-8**(3), 330–334 (1959).
- ⁴⁸L. Šojdr, J. Čermák, and R. Barillet, "Optimization of dual-mixer time-difference multiplier," in *EFTF 2004. 18th European Frequency and Time Forum* (IET, 2004), pp. 588–594.
- ⁴⁹F. L. Walls and D. W. Allan, "Measurements of frequency stability," *Proc. IEEE* **74**(1), 162–168 (1986).