IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 14 November 2023, accepted 19 December 2023, date of publication 21 December 2023,
date of current version 29 December 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3345833

APPLIED RESEARCH

Beyond DDMTD—Sub-Picosecond Timestamps for Asynchronous Clocks

ALEXANDER C. UTTER
The Aerospace Corporation, El Segundo, CA 90245, USA
e-mail: alexander.c.utter@aero.org

ABSTRACT Many navigation, sensing, and time-transfer applications call for the high-precision timestamping of events that occur in different digital clock domains. However, state-of-the-art techniques such as digital dual mixer time difference (DDMTD) typically require that the measured clocks be syntonized, i.e., derived from a common reference oscillator and identical in frequency. This article introduces a novel, all-digital technique that can provide accurate collinear timestamps for asynchronous clock domains of arbitrary frequencies. The vernier-referenced digital asynchronous collinear timestamp (VERDACT) system allows timestamps to be generated and compared across any number of arbitrary clock domains with sub-picosecond precision. The novel circuit can be implemented on most field programmable gate arrays (FPGAs) without the need for external components. Under laboratory conditions, an open-source prototype has demonstrated end-to-end precision better than 0.8 ps-rms on a Zynq UltraScale+ FPGA. Potential applications include sub-nanosecond two-way time-transfer using precision time protocol (PTP / IEEE-1588) over an asynchronous multiport Ethernet switch.

INDEX TERMS Clocks, digital integrated circuits, digital dual mixer time difference (DDMTD), field programmable gate array (FPGA), phase estimation, time measurement, time transfer, timing circuits.

I. INTRODUCTION

Many systems require the precise measurement of time, including the timing of events that occur in different digital clock domains. Cross-clock measurements have many applications in time-transfer, frequency-transfer, range-measurement, and remote sensing. Fig. 1 shows a typical example measuring elapsed time between a specific rising edge of CLK_1 and a specific rising edge of CLK_2.

Under the right circumstances, sequential digital circuits can measure time with a resolution finer than the underlying clock period(s). State of the art for high-precision digital time measurements is the digital dual-mixer time-difference (DDMTD) technique [1], [2]. DDMTD measures the phase– or time-offset between two clocks that are syntonized, i.e., derived from a common source and identical in frequency.

DDMTD operates by synthesizing a third clock at a very small frequency offset (i.e., the "beat frequency"), then sampling both input clock signals. As the synthetic clock shifts relative to each input clock, the sampled result is a time-stretched square wave at the beat frequency, whose phase indicates the respective input clock phase. The smaller the beat frequency (i.e., the longer the beat period), the finer the time-resolution.

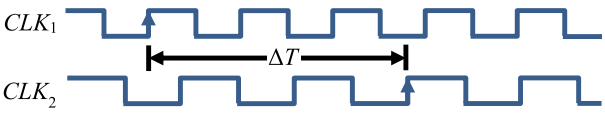

FIGURE 1. Elapsed time between events in different clock domains.

DDMTD is a key enabler for time-distribution applications such as the White Rabbit Synchronous Ethernet switch [3], [4] and can achieve measurement resolution of ±10 femtoseconds under ideal conditions [5].

However, DDMTD has a major limitation: the two input clocks must be syntonized. Because the phase offset is assumed to be constant over each beat period, there is an inherent design tradeoff between maximum resolution and maximum tolerable frequency offset. Syntonization circumvents this tradeoff by forcing the latter term to zero, allowing very long beat periods. However, syntonization is difficult or

The associate editor coordinating the review of this manuscript and approving it for publication was Jiafeng Xie.

impossible in some applications, such as mobile transceivers with unpredictable Doppler effects.

Precise network time monitoring (PNTM) is a related technique that reduces the need for syntonization. PNTM uses relatively short beat periods and interpolates DDMTD measurements over successive beat periods. PNTM achieves end-to-end measurement precision of ±5 picoseconds [6], and tolerates small frequency differences up to ±100 ppm.

To the author's best knowledge, no prior digital techniques offer comparable precision for asynchronous clocks; all require either syntonization or quasi-syntonization.

II. OUTLINE

This article introduces a novel circuit that can be used to timestamp events with sub-picosecond precision. The circuit operates in any combination of asynchronous clock domains over a wide range of target clock frequencies.

Section III presents the concept of mutually collinear timestamps that track a shared value-vs.-time reference. Collinear timestamps allow direct comparison of the absolute or relative timing of events occurring in arbitrary clock domains.

Section IV gives an overview of the novel timestamping circuit and its theory of operation.

Section V shows how a novel bundle of reference signals can be observed digitally with high precision, and Section VI presents a novel, all-digital phase-error detector circuit to make such observations.

Section VII explains the operation of a novel, all-digital closed-loop filter that operates in conjunction with the phase error detector. Section VIII explains how the state of this filter is used to generate best-fit collinear timestamps.

Section IX documents digital synthesis methods required to measure the end-to-end performance of the novel system. Section X discusses empirical observations under laboratory conditions.

Appendix A derives a figure of merit for predicting system performance. Appendices B through E provide solutions for problems that arise in practical implementation.

III. COLLINEAR TIMESTAMP CONCEPT

This section presents the abstract concept of collinear timestamps that track a shared value-vs.-time reference.

Fig. 2 illustrates the three-step process used to generate collinear timestamps. First, a reference counter defines a sequence of time/value pairs. Second, observations of that sequence allow estimation of a local best-fit line. Third, interpolation along that line allows generation of collinear timestamps in any target clock domain.

In its simplest form, the reference counter starts at zero and increments by a fixed amount on each reference clock cycle. The rising edge of the reference clock associates that moment in time with the corresponding numeric value. The numeric units are arbitrary, but for analysis it is simplest to use "seconds" as the unit for both time T and numeric

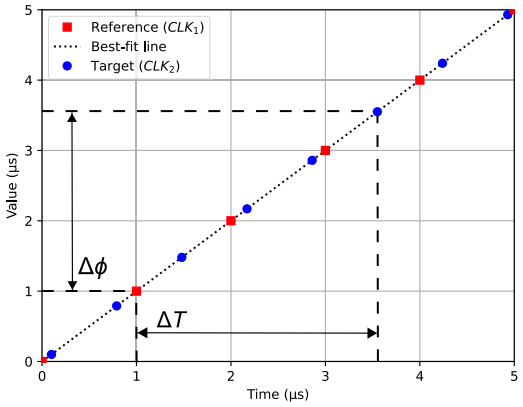

FIGURE 2. Reference sequence of time-value pairs in CLK1 domain (red squares) defines a local best-fit line. Interpolation along this line defines timestamps in target clock domain CLK2 (blue circles).

value ϕ. In practice, the increment is chosen to be a fixed-point representation of the time elapsed per clock cycle.

That series of time/value pairs defines a best-fit line. The reference clock may be imperfect and have its own jitter, phase noise, frequency deviations, etc.; but short-term averaging (e.g., 1 to 100 ms) still allows very precise measurements of the underlying best-fit line over any given interval.

Once a localized best-fit line is defined, the value at any moment can be found by interpolation. For events occurring in sync with a given target clock domain, the timestamp is the projected counter value at the corresponding clock edge.

For an ensemble of estimators observing the same reference, each estimated best-fit line is approximately collinear with the original reference. Therefore, any given estimate is approximately collinear with every other estimate, and elapsed time is the arithmetic difference $\Delta\phi$ between two such timestamps, regardless of the source and destination clock domains. The achievable timestamp or time-difference accuracy is set by the degree of collinearity and by the jitter of each target clock.

For an applied example, consider an Ethernet switch supporting the IEEE-1588 Precision Time Protocol (PTP). For each PTP packet, transparent clock mode requires the device to measure time spent traversing the switch and increment a specific header field accordingly [7]. For most switches, this operation is difficult because ports operate at different line rates, and each port has asynchronous transmit and receive clock domains. With collinear timestamps, this operation is trivially calculated by subtracting the ingress timestamp from the egress timestamp, regardless of the source and destination clocks.

IV. SYSTEM OVERVIEW

This section gives a brief overview of the vernier-referenced digital asynchronous collinear timestamp (VERDACT) circuit and its subsystems.

In a mechanical vernier caliper, subsidiary-scale markings are inscribed at 9/10th the scale of the primary markings.

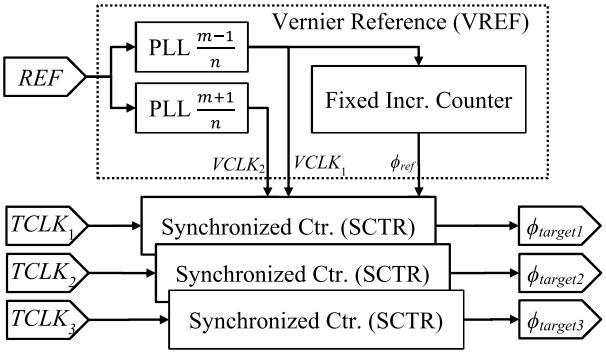

FIGURE 3. The complete VERDACT system consists of one VREF driving any number of SCTRs, one for each target clock of interest.

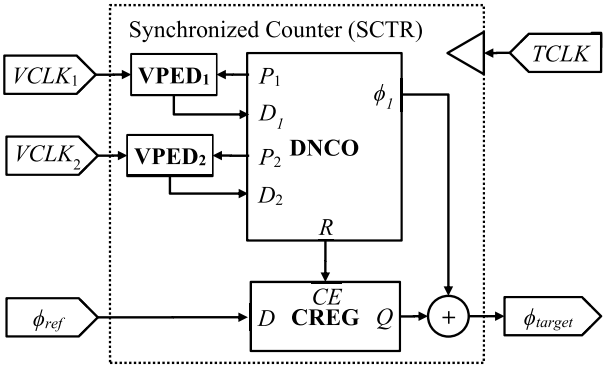

FIGURE 4. The SCTR circuit generates collinear timestamps ϕ_{target}. The VPED pair compares the sampled VCLK signals against predictions from the DNCO (P1, P2). The DNCO accumulates early/late indicators (D1, D2) to form a closed-loop estimation process. CREG allows calculation of ϕ_{target} from the DNCO state.

The subtle alignment offsets between the two scales enables precise length measurements. By an analogous process, a pair of vernier reference clocks enables extremely precise digital measurements of time.

Fig. 3 shows a block diagram for VERDACT, which generates collinear timestamps in each target clock domain. A single vernier reference (VREF, Section V) provides signals to several synchronized counter units (SCTR, Fig. 4), one for each target clock ($TCLK_n$). The VREF circuit defines the reference line. Each SCTR independently generates a timestamp sequence (ϕ_{target}) in its target clock domain, collinear with the original reference defined by VREF.

The vernier reference (VREF) generates three signals that define the best-fit line. The first and second are synthetic reference clocks with a carefully selected frequency ratio ($VCLK_1$, $VCLK_2$). These clocks are synthesized from a shared external reference using one or more phase-locked-loops (PLLs). For reasons discussed in Section V, discrete-time observations of the two clock signals allow digital circuits to measure time-differences with picosecond-scale resolution. The third signal is a fixed-point reference counter (ϕ_{ref}) operating in the $VCLK_1$ domain. The value of this counter at each rising edge of $VCLK_1$ defines the reference line.

Each synchronized counter (SCTR) is a closed-loop system that operates entirely in the associated *TCLK* domain.

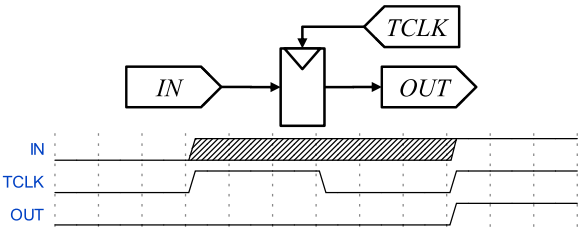

FIGURE 5. For a simple D flip-flop, the "dead zone" is equal to the clock period, because a rising-edge transition anywhere in the shaded region produces the same discrete-time output.

It estimates the instantaneous phase of $VCLK_1$ and $VCLK_2$ relative to *TCLK*. These parameters allow the calculation of ϕ_{target}, which is the estimated timestamp at each rising edge of *TCLK*, collinear with the time-value pairs defined by VREF.

Fig. 4 shows a block diagram of one SCTR, comprising:

- Two vernier phase error detectors (VPED, Section VI). The VPEDs sample $VCLK_1$ and $VCLK_2$ in the *TCLK* domain, then compare against predicted signals (P_1, P_2) to generate early/late delay indicators (D_1, D_2).
- A dual-output numerically controlled oscillator (DNCO, Section VII). The DNCO accumulates early/late indicators from the VPED to predict the phase of each reference clock. The output ϕ_1 indicates the time offset between each *TCLK* rising edge and the preceding $VCLK_1$ rising edge.
- A counter register (CREG, Section VIII), which brings ϕ_{ref} into the *TCLK* domain to allow calculation of ϕ_{target}. Its clock-enable signal *CE* is driven by the modulo-arithmetic rollover events *R* of DNCO output ϕ_1.

V. VERNIER REFERENCE

This section explains how the paired clocks generated by the vernier reference (VREF) circuit allow for picosecond-scale time measurements.

Consider a simple circuit that samples a slow clock signal (*IN*) using a D flip-flop operated by a faster target clock (*TCLK*). As shown in Fig. 5, the sampled output (*OUT*) cannot resolve time-differences smaller than the target clock period; i.e., for any zero-to-one transition in the discrete-time output sequence, it is impossible to determine if the continuous-time rising edge occurred near the beginning or end of the given cycle of *TCLK*.

This ambiguity causes a nonlinear "dead zone" equal to the period of *TCLK,* since a small timing difference causes no observable change in the discrete-time output.

Dead zones are present in all discrete-time measurements of asynchronous signals; the dead zone resolution ΔT is defined as the smallest resolvable timestep, i.e., the smallest change in the continuous-time offset that produces a change in the discrete-time output.

Using two clocks of nearly the same frequency is a common and effective mitigation strategy for decreasing the size of dead zones in various configurations [8], [9], [10], [11].

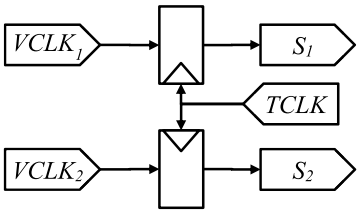

FIGURE 6. Discrete-time sampling of the vernier clock.

The two clocks are called a vernier pair[1] because they operate at a carefully chosen frequency ratio, derived from the same reference using fractional PLL(s). The ratio is typically chosen to be nearly 1:1.

As the two clocks of a vernier pair shift in and out of phase alignment, they form a pattern that results in a very small dead zone. That pattern repeats with an overall period equal to the least common multiple (LCM) of the two clock periods. For small frequency differences, the LCM is inversely proportional to the difference in frequency.

Vernier clock pairs can be used in many ways. In VERDACT, the continuous-time vernier pattern is sampled in the *TCLK* domain, as shown in Fig. 6. The resulting discrete-time vernier sequence (S_1, S_2) is a function of the instantaneous time-offset between (*TCLK*, $VCLK_1$) and (*TCLK*, $VCLK_2$). The LCM sets the number of pseudorandom[2] samples taken over the interval [0, T), where T is the target clock period. Appendix A derives a figure of merit for worst-case dead zone resolution ΔT:

$$\Delta T \leq \frac{M_1 T}{\mathrm{LCM}(M_1, M_2)} \tag{1}$$

where M_1 and M_2 are the vernier clock periods. In practice, the achievable resolution for a given target frequency is typically two to eight times smaller than this bound, but performance is adequate without considering this bonus.

At any given reference and target frequency, the figure of merit (1) is optimized by maximizing the LCM. Since maximizing the LCM results in good performance regardless of T, this allows a single vernier pair to service multiple synchronized counters. In designs with many target clock domains, such as a multiport asynchronous Ethernet switch, this minimizes the use of scarce FPGA resources for clock-generation and clock-distribution.

VI. VERNIER PHASE ERROR DETECTOR

This section presents a novel, all-digital circuit called the vernier phase error detector (VPED), which measures the time offset between a sampled vernier sequence and a predicted vernier sequence.

[1] In a vernier caliper, the human eye's hypersensitivity to subtle alignment offsets between the two scales enables precise length measurements. Every D flip-flop has an analogous hypersensitivity to the timing of inputs changing just before or just after the rising edge of its clock.

[2] Pseudorandom processes are deterministic, but they are complex enough to be treated as random for a given purpose. This assumption is addressed further in "XI."

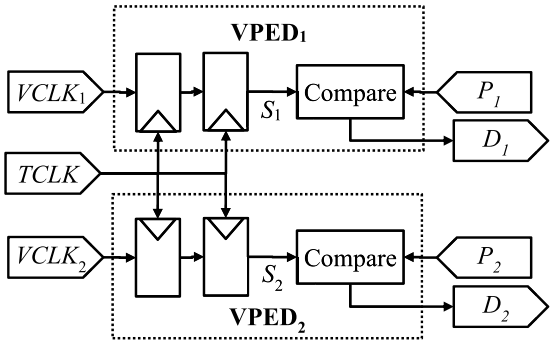

FIGURE 7. A pair of vernier phase error detectors (VPEDs). Each reference signal is sampled by the target clock and compared against a synthetized prediction to generate an early/late indicator.

VPEDs operate in pairs, where each unit samples and compares one of the two vernier reference clocks. Each sampled sequence is compared against the corresponding predicted sequence synthesized by the DNCO described in Section VII.

Each VPED accepts one of the two vernier clock signals $VCLK_1$, $VCLK_2$ and samples it in the target clock domain. The "sampling" circuit is a sequential pair (or triplet) of D flip-flops operated by the target clock, as shown in Fig. 7, forming a standard double-register (or triple-register) synchronizer [12]. The first register acts as a sample-and-hold circuit triggered on each rising edge of the target clock *TCLK* and its state may be metastable. The second register provides settling time to resolve metastability before the signal can be safely used. The third, if present, provides additional settling time. Allowing adequate settling time to resolve metastability is paramount because the vernier clocks exercise all possible phase alignments of the input signal and the target clock; violations of the ordinary setup- and hold-time requirements are inevitable.

Next, each VPED compares the sampled vernier signal (S_1, S_2) against the predicted signal from the DNCO (P_1, P_2) as follows:

- If the two signals are equal, emit $D_n = 0$.
- Otherwise, if S_n has remained constant for at least N clock cycles, emit $D_n = +1$.
- Otherwise, emit $D_n = -1$.

The "recent change" threshold N is set by the ratio of the vernier and target clock periods:

$$N = \lfloor M_1 \div 4T \rfloor \tag{2}$$

For best results, choose the vernier frequency such that N $\geq$ 2 and the ratio is not an integer.

Except for the recent-change state machine, the comparator circuit makes a purely combinational decision during each target clock cycle. Fig. 8 shows comparator operation with $M = 10T$ (i.e., $N = 2$). Each set of three signals shows a scenario with a different time offset, with predicted signal P leading (top) or lagging (bottom) compared to the sampled reference signal S, resulting in changes to the sign and duty cycle of output D.

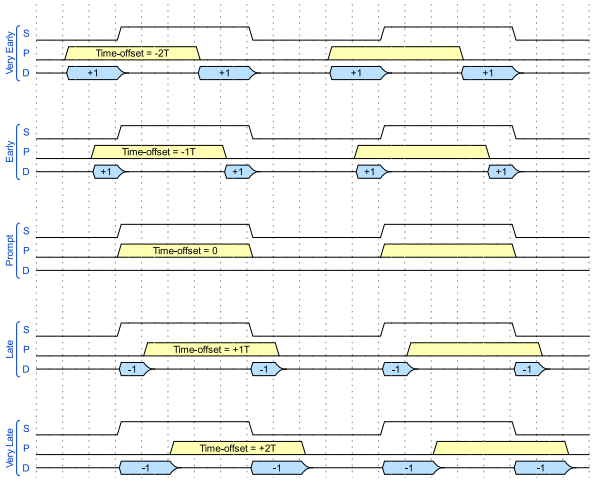

FIGURE 8. Comparator operation in five scenarios, each illustrating a different time offset between signals S and P. As the time offset varies over ±2 target clock cycles, the resulting duty cycle varies accordingly.

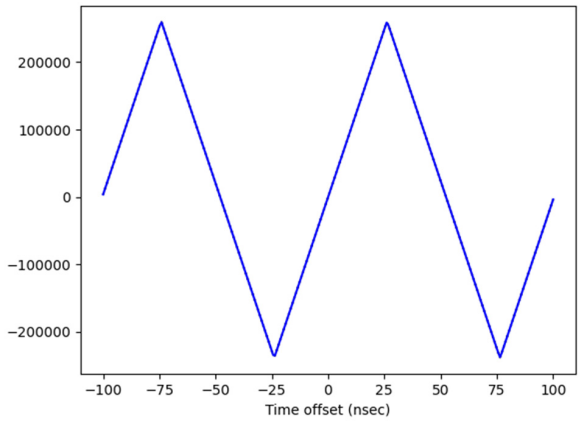

FIGURE 9. Simulated VPED S-curve, showing output vs static time offset. Output is given by sum(D1+D2) over 250,000 TCLK cycles.

Except for subtle secondary corrections described in Appendix B, VPED output signals D_1 and D_2 are added together. Accumulated over time, this gives a simple but effective early/late indicator for aligning the predicted vernier signals to the sampled vernier signals.

In Fig. 8, the duty-cycle of output D_n is proportional to the time offset, i.e., the smaller the magnitude of the time offset, the larger the fraction of clock cycles where S_n and P_n agree and $D_n = 0$. The pseudorandom sampling effect discussed in Section V ensures that, over time, this property remains true even for leading or lagging offsets much smaller than the target clock period T. Though the duty cycle may become very sparse, the remaining bit differences are spread pseudorandomly over the complete vernier sequence, and the number of mismatched bits is approximately proportional to the magnitude of the time offset.

Fig. 9 shows simulation results of the cumulative output from a VPED-pair as a function of the static time offset. The simulation does not include frequency offset or jitter. Each iteration of the simulation selects the static time offset, which is used to synthesize the sampled reference clocks (S_1, S_2) and

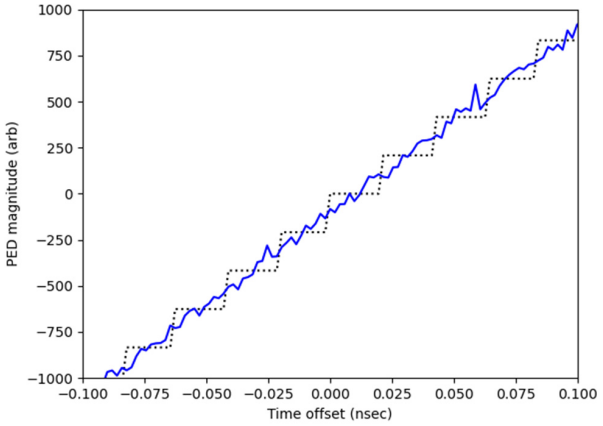

FIGURE 10. The same VPED S-curve, magnified to show small-signal response. Dotted black line shows output with zero jitter, and solid blue line shows output with ±100 ps Gaussian jitter, averaged over 2 ms.

the predicted clocks (P_1, P_2). These discrete-time signals and the corresponding VPED outputs (D_1, D_2) are simulated for each *TCLK* cycle. The output for each iteration is given by $sum(D_1 + D_2)$ over the duration of the simulation.

The simulated vernier clock pair is generated from a 25 MHz external reference scaled by 383/962 and 383/960 (see Appendix E). The simulated target clock is 125 MHz. The simulation duration is 2.0 ms (i.e., 250,000 *TCLK* cycles).

Under these conditions, the VPED output closely approximates a classic triangle-wave S-curve, with a linear range of ±25 nanoseconds (i.e., $\pm M_1/4$) and a safe pull-in range of ±50 nanoseconds (i.e., $\pm M_1/2$). Beyond that limit, each half of the VPED has stable false equilibria that do not correspond to the zero-crossings in the other half. Appendix C discusses methods for obtaining a good initial guess to ensure the system locks to the correct solution.

Fig. 10 shows the same simulation results, magnified to display the small-signal response for time offsets up to ±100 ps. The zero-jitter simulation (dotted black line) shows discrete timesteps at ≈17 ps intervals, consistent with the predicted vernier sequence resolution. A simulation with additive white Gaussian jitter on *TCLK* (blue line, $\sigma = 100$ ps) shows that moderate jitter can be quite effective in mitigating quantization effects, at the cost of added measurement noise.

In practice, allowing moderate jitter on the synthesized vernier clocks $VCLK_1$ and $VCLK_2$ has the same effect. Any relative jitter between *TCLK* and $VCLK_n$ produces a dithering effect that enables sub-LSB resolution. If that jitter exceeds the quantization step size, then random variation over consecutive iterations smooths the average response, mitigating nonlinear quantization effects.[3] This effect improves linearity at the cost of additional noise. Unlike quantization dead zones, unbiased noise can be mitigated by averaging over time, increasing effective resolution at the cost of a slower

[3]Notably, the prototype discussed in Section X intentionally chooses "suboptimal" clock-synthesis circuit parameters to ensure jitter exceeds this threshold.

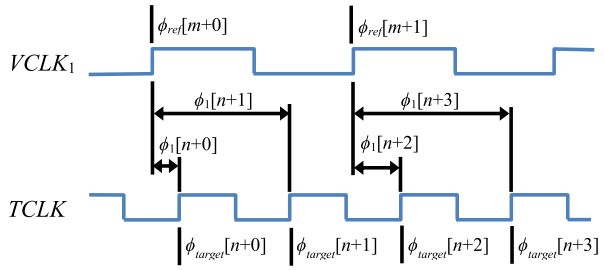

FIGURE 11. Time difference $\phi_1[n]$ from each TCLK rising edge to the preceding VCLK1 rising edge, which corresponds to reference timestamp $\phi_{ref}[m]$. The offset yields the collinear output timestamp ϕ_{target}.

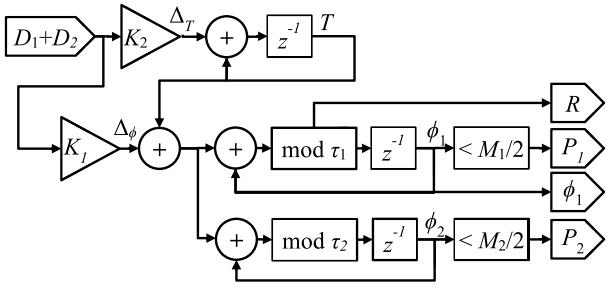

FIGURE 12. DNCO loop filter and numerically controlled oscillator.

response.[4] Jitter scaled to match the quantization step size improves overall system performance.

VII. DUAL-OUTPUT NCO

This section explains the theory of operation for the novel, all-digital dual-output numerically controlled oscillator (DNCO) circuit, which accumulates VPED outputs to perform closed-loop prediction of the sampled *VCLK* signals.

First, define the instantaneous phase ϕ_1 or ϕ_2 as the estimated time difference from the nth rising edge of the target clock (*TCLK*) to the preceding (mth) rising edge of the corresponding vernier clock ($VCLK_1$ or $VCLK_2$). DNCO accumulators ϕ_1 and ϕ_2 are estimates of this time-difference. Fig. 11 shows the relationship between $VCLK_1$, *TCLK*, ϕ_1, ϕ_{ref}, and ϕ_{target}.

Fig. 12 shows the novel DNCO, which incorporates a cross-coupled loop filter. The operation of the interlocked accumulators is roughly analogous to a conventional second-order tracking loop. Inputs D_1 and D_2 are the early/late indicators from each VPED. Outputs R (rollover), P_1, P_2, and ϕ_1 are the signals shown in Fig. 4 and Fig. 7. The DNCO contains three accumulators[5]:

$$\phi_1[n+1] = (\phi_1[n] + \Delta_\theta[n] + T[n]) \bmod \tau_1 \quad (3)$$

$$\phi_2[n+1] = (\phi_2[n] + \Delta_\theta[n] + T[n]) \bmod \tau_2 \quad (4)$$

$$T[n+1] = (T[n] + \Delta_T[n]) \quad (5)$$

where Δ_ϕ, Δ_T, Δ_ε are the outputs from the loop filter; T is the estimated period of the target clock; and τ_1, τ_2 are the nominal period[6] of each vernier reference clock. All ϕ, τ, and T terms have units of time, i.e., picoseconds or some fixed-point representation thereof.

The Δ terms are driven by the VPED outputs D_1 and D_2 as follows:

$$\Delta_\phi[n] = K_1(D_1[n] + D_2[n]) \quad (6)$$

$$\Delta_T[n] = K_2(D_1[n] + D_2[n]) \quad (7)$$

where K_1 and K_2 are scaling factors that configure the loop bandwidth and damping factor of the second-order loop [13].

The final DNCO outputs to the VPEDs are the predicted clock signals P_1 and P_2, which are given by:

$$P_1[n] = 1 \text{ if } \left(\phi_1[n] < \frac{1}{2}M_1\right) \text{ else0} \quad (8)$$

$$P_2[n] = 1 \text{ if } \left(\phi_2[n] < \frac{1}{2}M_2\right) \text{ else0} \quad (9)$$

Quasi-linear operation of the DNCO and VPED require a good initial guess for ϕ_1, ϕ_2, and T. If the initial guess is sufficiently accurate, the SCTR converges and behaves like any other second-order PLL. Thanks to the VPEDs, this process is quite linear and has a worst-case dead zone of less than one picosecond.

Various side-effects of DNCO operation produce high-frequency noise in ϕ_1. The magnitude is typically less than one picosecond, but sensitive applications should apply an IIR filter to ϕ_1 or ϕ_{target} to suppress this noise.

VIII. COUNTER REGISTER AND TIMESTAMP CONSTRUCTION

This section explains how the counter register (CREG) uses DNCO state to calculate best-fit collinear timestamps on each target clock cycle.

Recall that DNCO accumulator $\phi_1[n]$ estimates the time difference indicated in Fig. 11, i.e., the offset from *TCLK* to the preceding $VCLK_1$ rising edge. Reference counter $\phi_{ref}[m]$ is the numeric timestamp associated with that $VCLK_1$ rising edge. Note separate indices m and n.

Therefore, the best-fit estimate of the collinear timestamp is given by the sum of those two quantities:

$$\phi_{target}[n] = \phi_{ref}[m] + \phi_1[n] \quad (10)$$

This process is illustrated in Fig. 13. The process is complicated by the fact that $\phi_{ref}[m]$ and $\phi_1[n]$ exist in different clock domains. The CREG circuit brings ϕ_{ref} into the target clock domain, allowing the calculation to proceed.

To maintain continuity and monotonicity, CREG must be updated concurrently with the rollover of ϕ_1. Consider the *TCLK* rising edge in Fig. 13 that occurs at time 3.0 μs. It is

[4] Because all real clocks wander over time, the speed-vs.-resolution trade-off sets a theoretical bound on VERDACT performance. Systems relying on dither must use filters with a time-constant many times the VREF LCM. Longer time constants decrease residual measurement noise, but degrade the system's ability to accurately track fast changes in VREF or target clock(s), degrading timestamp accuracy.

[5] Since $\Delta_\phi \ll \tau_n$ and $T \ll \tau_n$, the modulo terms will never wrap around more than once in any given clock cycle. Therefore, practical FPGA realizations may replace the modulo operation with a much simpler compare-and-subtract.

[6] Nominally, τ_n is simply the fixed-point representation $\tau_n \tilde{} M_n$. Mitigation for the effect of cumulative rounding error is discussed in Appendix B.

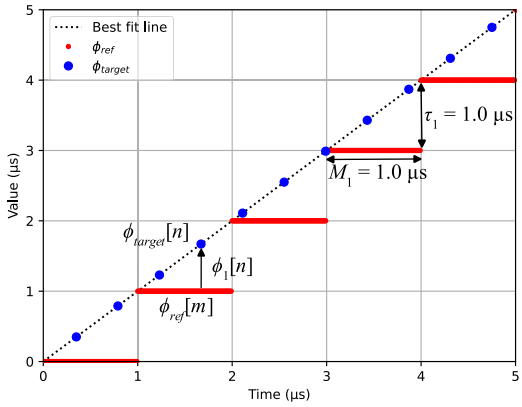

FIGURE 13. Construction of ϕ_{target} (blue dots) from ϕ_{ref} (red lines) and ϕ_1. Reference counter $\phi_{ref}[m]$ increments on each rising edge of VCLK1, indicated by the vertical grid lines. Time offset $\phi_1[n]$ updates on each rising edge of TCLK. CREG coordinates the safe and accurate clock-domain transition of these signals.

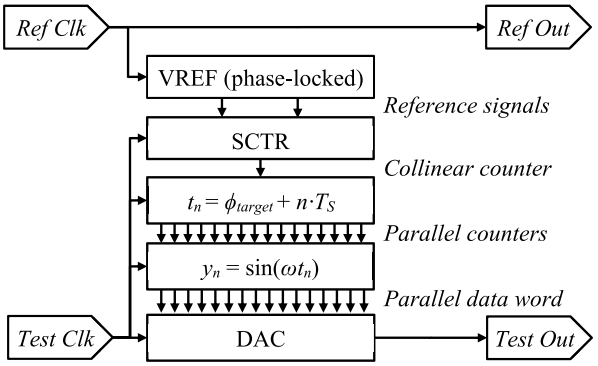

FIGURE 14. Synthesis of a phase-locked sine wave using a DAC.

difficult to determine whether this point corresponds to the earlier ϕ_{ref} (i.e., $\phi_1 \approx \tau_1$) or to the later ϕ_{ref} (i.e., $\phi_1 \approx 0$). Note, however, that the nominal ϕ_{ref} increment is equal to the $VCLK_1$ period, i.e., $\phi_{ref}[m+1] \approx \phi_{ref}[m] + \tau_1$. Therefore, the sum will be nearly identical in either case. As long as ϕ_{ref} updates in sync with the ϕ_1 rollover event, such decisions have negligible effect on the output ϕ_{target}.

To avoid timing violations and metastability problems, CREG must also ensure that it only reads new ϕ_{ref} values when the input signal is stable. Fortunately, this constraint is trivially met once the DNCO has converged.

The operation of the VPED ensures that the rising edge of S_1 occurs shortly after[7] each change to ϕ_{ref}. Once locked, the DNCO predicted output P_1 mirrors the sampled vernier signal S_1 within one $TCLK$ cycle. Therefore, ϕ_1 rollover events (i.e, the rising edge of P_1) occur shortly after updates to ϕ_{ref}. Since $TCLK$ is several times faster than $VCLK_1$, it is safe to infer that the counter value will be stable shortly before and after each rollover event.

As a result, CREG can use the rollover event as a clock-enable for an ordinary flip-flop register $\phi_{ref}[n]$ with no risk of metastability. Once all terms are available in the target clock domain, the output timestamp is calculated by simple addition.

IX. MEASURING PERFORMANCE

This section presents methods for synthesizing sine-waves using an SCTR output sequence, which is required to support empirical measurements in Section X.

Measurement of VERDACT performance is difficult because the output is a sequence of numbers in a specific digital clock domain. To make a verifiable measurement under laboratory conditions, it is necessary to convert the collinear numeric sequence back into a real, measurable signal.

The method in this section accepts an external reference signal, generates a phase-locked vernier reference (see Appendix D), then synthesizes a phase-locked sinusoid in an asynchronous clock domain. There is no analog path from the external reference to the output signal; the collinear numeric output of the SCTR is the only relevant link.

Use of a high-resolution, high-sampling-rate digital-to-analog converter (DAC) allows synthesis of a sine wave whose phase accurately reflects very small changes in the numeric SCTR output. For example, consider a 125 MHz sine wave generated by a 14-bit DAC operating at 6.4 GSPS. A time shift of one picosecond yields a phase shift of about 0.05°, corresponding to a maximum value change of just 6 LSBs. This change is small but measurable with specialized instruments such as a signal-source analyzer or lock-in amplifier.

The logic shown in Fig. 14 synthesizes a phase-locked 125 MHz sine wave for the prototype discussed in Section X. The prototype's DAC operates at $F_S = 6.4$ GSPS, yielding time per output sample $T_S = 1/F_S = 156.25$ ps. Because this DAC operates at a higher sample rate than the FPGA digital fabric can support, the FPGA's internal interface to the DAC requires generation of multiple samples per clock that are serialized during transmission. The SCTR generates a collinear counter sequence ϕ_{target} in the parallel clock domain. For each of the N samples in the set of parallel counters, add a time offset that corresponds to the sequential output order of that bit; i.e., for n in $[0..N)$, let $t_n = \phi_{target} + n \cdot T_S$. For each of these parallel counters t_n, use modulo arithmetic to calculate the phase for lookup-table synthesis of a sine-wave.[8]

Because each set of parallel counters spans only a few nanoseconds, this process of linear extrapolation has negligible impact on the quality of the resulting sine wave.

The sine-wave synthesis method can also be used to measure the test clock and DAC performance in isolation.

[7] Recall that ϕ_{ref} updates on the rising edge of $VCLK_1$. If signal S_1 is sampled from $VCLK_1$ through a double-register synchronizer as shown in Fig. 7, then the delay from ϕ_{ref} to S_1 is in the range $[T, 2T)$.

[8] For platforms without a DAC, square wave outputs can be synthesized using a similar technique to drive a high-speed parallel-to-serial converter. This allows effective sampling rates of 10 GHz or higher, depending on the FPGA platform. The disadvantage of this approach is that nearest-neighbor 1-bit sampling introduces uniform-distributed jitter equal to the sample interval T_S.

To synthesize a free-running sine wave, the same parallel logic is driven by a simple free-running counter instead of the SCTR. Such a system is no longer coupled to VERDACT or to the external reference clock in any way, so it can be used to measure the phase noise and jitter performance of the DAC clock itself. In the following section, this method is used to establish a performance baseline for characterization of noise added by VERDACT.

X. LAB MEASUREMENTS

This section describes the apparatus used to measure VERDACT performance under laboratory conditions, and briefly discusses the implications of those measurements.

The objective of this test is to demonstrate the collinearity of the VERDACT output to an analog reference. The demonstration locks VREF to the analog reference, then uses SCTR timestamps and a DAC to synthesize a phase-locked 125 MHz sine-wave. Because the DAC operates in an asynchronous 6.4 GHz clock domain, the measured phase-noise and phase-offset of the synthesized signal indicate the collinearity of the SCTR timestamps.

All measurements in this section were performed using the Xilinx ZCU208 development board, which contains a Zynq UltraScale+ RFSoC (XCZU48DR). VHDL and C++ source code for VERDACT and for the ''zcu208_clksynth'' design are available on GitHub: https://github.com/the-aerospace-corporation/satcat5.

First, the FPGA accepts a 125 MHz external clock for the vernier reference. The reference counter is numerically phase-locked to the 125 MHz external clock using the technique described in Appendix D.

Separately, the development board's CLK104 clock-synthesis daughtercard is configured to synthesize a 400 MHz clock derived from its onboard TCXO.[9] From this reference, the ZCU208's RF-DAC synthesizes a 6.4 GHz sample clock (for the RF-DAC itself) and a 200 MHz data clock (for the FPGA's parallel interface to the RF-DAC). Use of a free-running TCXO ensures that the RF-DAC clocks are asynchronous to the external reference. Separately, synthesis of a free-running 125 MHz sine-wave allows phase-noise measurements of the RF-DAC clock itself.

The SCTR under test operates in the 200 MHz RF-DAC clock domain. It tracks VREF with a time-constant of 50 ms, followed by an auxiliary IIR filter with a time-constant of 0.5 ms. The numeric SCTR timestamps drive the synthesis of a 125 MHz sine-wave as described in Section IX.

Fig. 15 shows the instrument configuration used for phase noise and phase offset measurements. The VREF's 125 MHz external reference is sourced by an Agilent E4428C signal generator. The ZCU208's RF-DAC clock is used to synthesize free-running and phase-locked DAC outputs. The DAC outputs are measured by a Rohde & Schwarz FSUP signal source analyzer. The phase-locked DAC output is also compared against the external reference using a Stanford

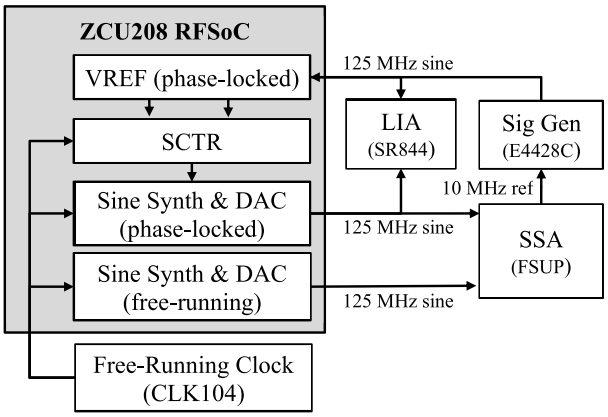

FIGURE 15. Synthesis and measurement of free-running and phase-locked sine-waves using two DACs on the ZCU208.

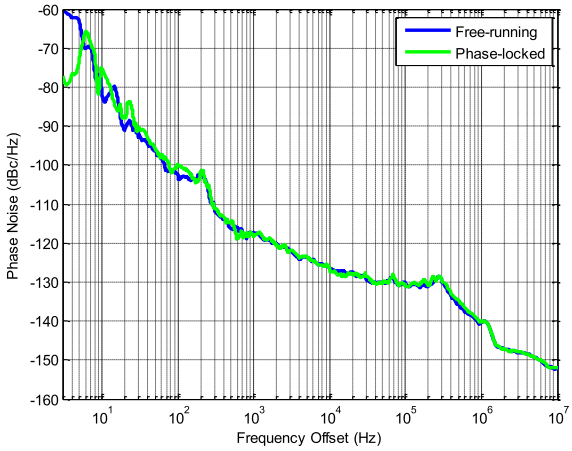

FIGURE 16. Measured phase noise at RF-DAC output (125 MHz sine). Above 10 Hz, phase noise contributions from VERDACT are below the measurable floor set by the RF-DAC clock.

Research Systems SR844 lock-in amplifier (LIA). Passive components such as baluns and power-splitters are omitted for clarity.

Fig. 16 shows the measured phase-noise spectrum from each DAC output, averaged over two trials, measured from 3 Hz to 10 MHz. For the phase-locked output, the FSUP reports that jitter over this span is 1.53 ps-rms. The blue curve shows the baseline phase noise[10] from the free-running output; the green curve shows the phase noise from the phase-locked output. The two curves are largely identical above 10 Hz, indicating that high-frequency phase noise and jitter are dominated by the performance of the RF-DAC clock itself and not by VERDACT. Therefore, further discussion will be focused on low-frequency effects (i.e., 0 to 10 Hz) measured in the time domain, to provide a more accurate bound on VERDACT performance.

Fig. 17 shows the phase offset vs. time reported by the SR844. The LIA tracks the phase difference between the synthesized, phase-locked DAC output and the original 125 MHz reference. The LIA is configured to filter the baseband signal

[9]Part number: Connor Winfield DOT050F-010.0M

[10]Phase noise contributions of the E4428C and FSUP are well below this baseline over the full frequency range.

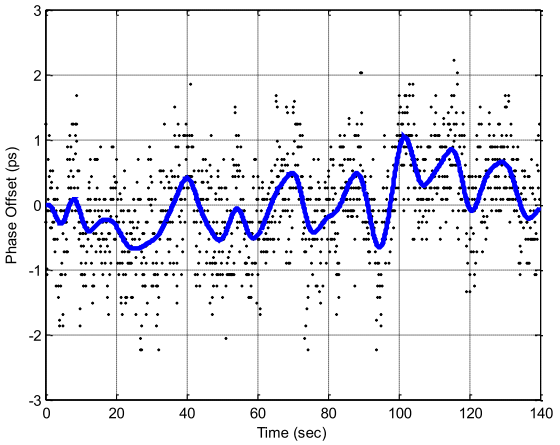

FIGURE 17. Measured phase offset vs. time for the phase-locked output. Black dots show raw 10 Hz measurement (0.78 ps-rms); blue trendline shows the result after application of a low-pass filter (0.42 ps-rms).

(time-constant 100 ms), measure the angle difference, and present this value as an analog output (scaling factor 1V = 1° = 22.2 ps). This signal is captured by an oscilloscope sampling every 100 ms for several minutes, then analyzed digitally. Black dots show the raw captured data, which has a standard deviation of 0.78 ps-rms. The blue trendline shows the result after application of a low-pass filter,[11] which yields a standard deviation of 0.42 ps-rms.

Many factors influence the end-to-end jitter of the synthesized phase output, including jitter in the DAC clock, VERDACT resolution limits, VERDACT synchronization errors, and instrumentation noise. To the extent practical, the measurement of interest is the portion caused by VERDACT itself.

As previously discussed, jitter above 10 Hz is attributed entirely to the DAC clock. Phase-noise measurements conclusively indicate that the DAC clock is the dominant jitter source at carrier offsets above 10 Hz, and these contributions will be ignored.

Jitter below 10 Hz is caused by a mixture of contributions from VERDACT, the DAC clock, and other sources. If we conservatively assign all such jitter to VERDACT, then this yields an upper bound jitter equal to that of the raw LIA output (i.e., the black point-cloud of Fig. 17), yielding a maximum estimated VERDACT jitter of 0.78 ps-rms.

Conversely, setting a very low cutoff frequency allows a lower bound on the total jitter attributed to VERDACT. Below 0.125 Hz, the residual jitter contribution from the DAC clock is negligible. This corresponds to the digitally filtered LIA output (i.e., the blue trendline of Fig. 17), yielding a minimum estimated VERDACT jitter of 0.42 ps-rms.

XI. CONCLUSION

VERDACT produces timestamps with sub-picosecond collinearity, allowing the comparison of events occurring in arbitrary combinations of asynchronous digital clock domains.

[11] Forward-backward filtering is used to eliminate group delay. Each pass is a 4th-order Butterworth low-pass filter with -3 dB cutoff of 0.125 Hz.

The novel, all-digital system circumvents syntonization or quasi-syntonization requirements common to prior state-of-the-art approaches, such as DDMTD and PNTM. The entire VERDACT system can be constructed with building blocks that are commonly available in off-the-shelf FPGAs, and it requires no external components except stable reference and target oscillators.

A proof-of-concept VERDACT prototype running on a Xilinx Zynq Ultrascale+ RFSoC has demonstrated end-to-end jitter between 0.4 and 0.8 ps-rms.

The Aerospace Corporation has applied for a patent on VERDACT technology: USPTO Application #18/066,542 filed December 15, 2022, titled "Vernier Phase Locked Loop."

APPENDIX A
PSEUDORANDOM VERNIER SAMPLING

This appendix describes the pseudorandom behavior of the sampled vernier sequence and derives a figure of merit for predicting its dead zone resolution.

Recall that a pair of vernier clocks forms a repeating continuous-time pattern. In the VPED, both vernier signals are then sampled in the target discrete-time domain.

The resulting discrete-time vernier sequence is a function of the reference parameters, the time-offset, and the period of the target clock. Over a moderate time-window, slow-varying frequency and time-offset parameters can be assumed to be constant. Under this condition, the sampling operation is directly analogous to the operation of a linear congruential generator (LCG), a simple but effective pseudorandom number generator [14]:

$$\mathrm{LCG}_k[n] = (n{\cdot}T + T_0) \bmod M_k \tag{11}$$

$$S_k[n] = 1 if \left(LCG_k[n] < \frac{1}{2}M_k\right) \text{ else } 0 \tag{12}$$

where T is the target clock period, T_0 is the target clock's initial time offset, M_k is the period of the kth vernier clock, LCG_k is the sampling phase for that clock, and S_k is the sampled vernier signal.

A conventional LCG's scale, offset, and modulo parameters are directly analogous to T, T_0, and M_k, respectively. As with the conventional LCG, choosing $T \ll M_k$ breaks the pseudorandom illusion, resulting in long stretches of obvious linear outputs. To avoid this, the period of each reference clock should be no more than 30 times the period of the target clock. The pseudorandom sequence repeats with an interval N_k that depends on the precise interplay of these three parameters:

$$N_k = \frac{\mathrm{LCM}(T, M_k)}{T} \tag{13}$$

Because the sequence eventually repeats, the sampling phase LCG_k takes on discrete values. In effect, each pseudorandom sequence shuffles the order of the points on this lattice, which is evenly distributed over the range [0.. M_k).

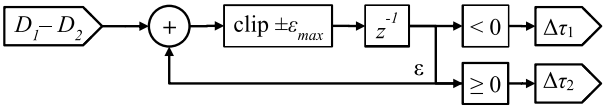

FIGURE 18. Accumulator for ratiometric balance correction.

Since the lattice spacing sets the smallest observable time offset, longer sequences are always preferred.

Considered jointly across all k, the period of the combined discrete sequence will always be at least as long as the continuous-time vernier pattern:

$$M_{joint} = \text{LCM}\,(M_1, M_2, T) \geq LCM\ (M_1, M_2) \tag{14}$$

When that continuous-time pattern is sampled by *TCLK*, this yields the number of discrete points on the joint lattice:

$$N_{joint} = \frac{M_{joint}}{T} \tag{15}$$

Because the lattice points are distributed evenly over $[0..M_k)$, and $M_1 \approx M_2$, the smallest observable time difference (i.e., the dead zone resolution set by the joint lattice) is given by:

$$\frac{M_1}{N_{joint}} \approx \Delta T \approx \frac{M_2}{N_{joint}} \tag{16}$$

Algebraic manipulation yields the performance bound:

$$\Delta T \approx \frac{M_1}{N_{joint}} = \frac{M_1 T}{M_{joint}} \leq \frac{M_1 T}{\text{LCM}(M_1, M_2)} \tag{17}$$

This performance bound is proportional to T. However, all other terms depend only on the vernier reference parameters. Therefore, for any given vernier clock periods M_1 and M_2 where $M_1 \approx M_2$, maximizing the LCM yields excellent time resolution regardless of T.

APPENDIX B
RATIOMETRIC BALANCING OF THE DNCO

This appendix describes a method for mitigating a subtle cumulative error in the DNCO. Without mitigation, the accumulation of error prevents the SCTR from operating correctly.

Recall that DNCO accumulators ϕ_1 and ϕ_2 are calculated modulo τ_1 and τ_2, which are the fixed-point representations of the vernier clock periods M_1 and M_2. Practical implementations of this process introduce some quantization error. Such errors are miniscule, typically less than one femtosecond per modulo wraparound event, but they are a systematic bias that accumulates indefinitely. Without mitigation, the cumulative error gradually drags accumulators ϕ_1 and ϕ_2 in irreconcilable directions until the SCTR loses lock.

To maintain ratiometric balance in the two phase-accumulators, the DNCO requires closed-loop feedback as shown in Fig. 18. The error signal compares the difference in early/late indications from each VPED (D_1, D_2). A new accumulator (ε) tracks the cumulative imbalance:

$$\varepsilon\,[n+1] = \text{clip}\,(\varepsilon\,[n] + D_1\,[n] - D_2[n]) \tag{18}$$

To mitigate windup effects and limit cycles, this accumulator is limited by clipping/saturation to the range $\pm\varepsilon_{\max}$.

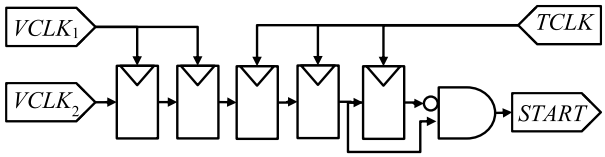

FIGURE 19. DDMTD circuit for detecting vernier phase alignment.

Empirically, a value of $\varepsilon_{\max} = 15$ LSB has been effective under all conditions tested by the author.

Feedback is applied by temporarily adjusting the effective value of τ_1 and τ_2. If cumulative rounding errors are causing ϕ_1 to run slightly faster than ϕ_2, increasing τ_1 or decreasing τ_2 can gradually restore balance. For a fixed-point system, the maximum required adjustment ($\Delta\tau_1$, $\Delta\tau_2$) is one least significant bit (LSB):

$$\Delta\tau_1\,[n] = \text{LSBif}\,(\varepsilon\,[n] < 0)\ \text{else0} \tag{19}$$

$$\Delta\tau_2\,[n] = \text{LSBif}\,(\varepsilon\,[n] \geq 0)\ \text{else0} \tag{20}$$

$$\tau_1'\,[n] = \tau_1 + \Delta\tau_1\,[n] \tag{21}$$

$$\tau_2'\,[n] = \tau_2 + \Delta\tau_2\,[n] \tag{22}$$

APPENDIX C
INITIAL ACQUISITION

This appendix describes methods for reliably initializing the SCTR by iteratively converging on the correct solution.

Due to the false-locking hazard discussed in section VI, the SCTR has a narrow pull-in range for initial acquisition. DNCO variables ϕ_1, ϕ_2, and T must all be initialized carefully to ensure success. DDMTD plays an important role in this process, by giving an indicator of when the two vernier clocks are precisely phase-aligned (i.e., $\phi_1 \approx \phi_2$).

In DDMTD, flip-flops sample the instantaneous value of an input clock signal on every rising edge of a synthesized clock. A small frequency offset in the synthesized clock ensures that, compared to the input clock, its rising edge slips by a small amount on each subsequent clock cycle. The resulting output is a square-wave with a time-stretching factor that is inversely proportional to the beat frequency of the input and synthesized clocks [1].

Since the vernier reference clocks are derived from the same source, all DDMTD prerequisites are met. Therefore, using one of the vernier clocks to sample the other will produce a time-stretched square wave whose rising edge occurs when the two clocks are precisely phase-aligned. The period of the stretched square wave is equal to the period of the vernier pattern, LCM(M_1, M_2).

Since all SCTR logic operates in the TCLK domain, a clock-domain transition is required. Fortunately, the time-stretched square wave changes infrequently, so a simple double-registered buffer is effective. Fig. 19 shows the combined DDMTD and clock-domain-crossing circuit, which generates a "start" strobe shortly after the phase-alignment event.

The phase alignment of the target clock relative to the two vernier clocks remains unknown. However, the delay through

1. Reset T to nominal value, iteration counter = 0.
2. Each iteration:
   a. Set all loop gains based on iteration count.
   b. Set $\phi_1 = \phi_2 = M_1 + 1.5T$ and freeze DNCO.
   c. Wait for start strobe, then unfreeze.
   d. Wait for expected settling time.
3. If locked:
   a. Final iteration: wait a while, then repeat from step 3.
   b. Otherwise, increase iteration counter, then repeat from step 2.
4. If unlocked:
   a. First iteration: repeat from step 1.
   b. Otherwise, decrease iteration counter, then repeat from step 2.

FIGURE 20. Robust iterative acquisition algorithm.

the circuit in Fig. 19 is bounded within the range $[M_1 + T, M_1 + 2T)$, i.e., an average delay of $M_1 + 1.5T$ with worst-case uncertainty of $\pm 0.5T$.

Fortunately, this level of error is accurate enough for a useful initial guess. Recall from Section VI that the VPED's linear pull-in range is $\pm 0.25 M_1$. Since $T \ll M_1$, the SCTR can readily tolerate this level of uncertainty in its initial state.

To compensate for the average delay, adjust the DNCO initial state accordingly. Since the start strobe arrives slightly late, the most accurate initial guess is given by $\phi_1 = \phi_2 = M_1 + 1.5T$. Because both variables are modulo-counters and $M_1 \approx M_2$, this aliases to $\phi_1 = \phi_2 = 1.5T$.

The remaining unknown is the target clock period T. This parameter's uncertainty is usually dominated by the frequency accuracy of the associated external oscillator(s). Depending on size, weight, power, and cost constraints, the actual period may differ from the nominal value by ± 50 ppm or more. Maximum tolerable error for the initial guess depends on many factors, including the SCTR loop bandwidth. Wide loop bandwidths enhance pull-in range but result in higher operating noise and higher probability of cycle slips.

This unfortunate design tradeoff can be circumvented by iteratively refining the estimate of T. This method uses a high loop bandwidth for initial acquisition, then a progressively smaller loop bandwidth producing finer and finer estimates of T, eventually transitioning to steady-state operation. Since it reuses the DNCO circuit, very few additional FPGA resources are required.

The number of iterations and the bandwidth of each iteration are chosen at design time to achieve the desired pull-in range. For example, the prototype design of Section X uses five iterations, each with half the loop bandwidth of the previous iteration. Combined with a lock-detection heuristic, this can produce an extremely robust process for startup and error recovery. An example is shown in Fig. 20.

The lock-detection system is a critical part of robust SCTR initialization. It must accurately warn of false-lock conditions without interrupting normal operation due to excessive false alarms. A simple but effective heuristic is to inspect the inputs

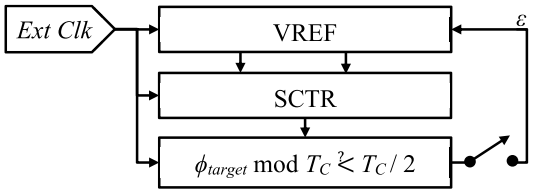

FIGURE 21. Discipline of a phase-locked reference counter.

to each VPED. In a locked state, the sampled and predicted signals (S_n, P_n) should agree almost 100% of the time.

A counter-based state machine allows this accuracy to be measured over very long windows without needing much memory. If the VPED inputs agree, the counter increments by one (up to a designated maximum). If they disagree, the counter decrements by a penalty P (down to a minimum of zero). If the fraction of correct predictions exceeds $P/(P+1)$, then the counter will steadily increase. Comparing the counter to a threshold provides a simple but effective lock indicator.

APPENDIX D
PHASE-LOCKED REFERENCE COUNTERS

This appendix describes a simple technique for phase-locking reference counter ϕ_{ref} to an external reference clock.

The output of the SCTR is a counter that follows the best-fit line defined by the reference counter. For many applications, such as elapsed-time measurement on a network switch, that reference counter need not be tied to any external signal. Other applications, including the laboratory measurements discussed in sections IX and X, require that the counter be phase-locked to an external reference clock.

The first step in producing a phase-locked counter is to derive the vernier reference (i.e., the vernier clock pair and reference counter) from the external reference signal. This ensures that the counter is frequency-locked to the reference. However, the initial phase offset is still effectively random.

The second step is to provide closed-loop feedback to gradually discipline the reference counter to the desired phase. Fig. 21 shows a block diagram for providing such feedback.

When operating in this mode, VREF includes a variable-rate counter that operates in one of the two vernier clock domains:

$$\phi_{ref}[n+1] = \phi_{ref}[n] + M_1 + \varepsilon[n] \tag{23}$$

where ϕ_{ref} is the vernier reference counter, M_1 is the associated vernier clock period, and ε is a time-varying feedback term. To allow phase comparison, an attached SCTR estimates the collinear best-fit line ϕ_{target} in the external clock domain. Modulo-comparison against the external reference period T_C drives a bang-bang feedback term $\pm\varepsilon$ to adjust the counter frequency.

After the initial transient, ε need not be updated frequently and its magnitude should be reduced to minimize impact to the experiment. In the ZCU208 prototype design, ϕ_{ref} has a resolution of 1 LSB = 0.2 attoseconds and operates at 9.961 MHz. Setting the steady-state ε to ± 2 LSB yields a maximum drift rate of 4 picoseconds per second, ensuring

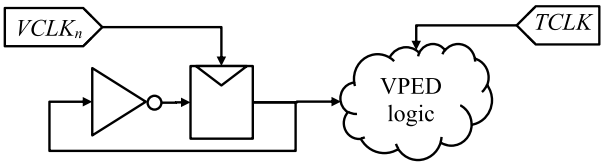

FIGURE 22. Divide-by-two circuit.

adequate margin for loop stability and negligible end-to-end impact even if the polling rate is only 50 Hz.

If preferred, ε can also be driven by a PID loop or any other linear control system.

In other applications, the reference counter could be locked to a broadly distributed time reference, such as a global navigation satellite service (GNSS) receiver. VERDACT timestamps derived from such a reference counter could be compared even if they are generated by devices that are widely separated in time or space. For example, VERDACT timestamps could be used for aligning signals from radioastronomy receivers that are hundreds of kilometers apart, for use in very-long baseline interferometry.

APPENDIX E
OTHER PRACTICAL CONSIDERATIONS

This appendix describes simple mitigations for various problems that arise in the practical implementation of VERDACT, applicable to most FPGA platforms.

One practical consideration is the use of separate distribution circuits for clocks and data. Clock signals on FPGAs typically use a limited number of global buffers that distribute the signal across the entire chip with minimal skew and jitter. Data signals typically use more plentiful but localized resources. In most FPGAs, including the Xilinx UltraScale+ device used in the ZCU208 prototype, the global clock buffers cannot connect to the data input of a flip-flop because such connections never occur in conventional digital logic designs.

However, clock-as-input-signal connectivity is explicitly required to implement DDMTD and VPED circuits. The divide-by-two circuit shown in Fig. 22 is a simple indirect means of making such connections.

The routing of the divide-by-two signal is critical to overall system performance. To ensure routing delays are as short and predictable as possible, the divide-by-two circuit should be replicated at each point of use (i.e., max-fanout $= 2$) to ensure place-and-route keeps the source as close as possible to each destination. The circuit is simple enough that such duplication is inexpensive.

Another consideration is the choice of vernier reference frequencies. In practice, three constraints influence this choice. First, VPED operation requires that the ratio M / T to be at least four. Second, the assumption of pseudorandom behavior (see Appendix A) begins to break down if the ratio exceeds about thirty. Fortunately, the width of this range ($4\times$ to $30\times$) allows a single vernier pair to effectively service many different target frequencies. Choosing $M_n \approx 100$ ns (i.e., $\approx$10 MHz) is a reasonable compromise for target clock frequencies from 40 to 300 MHz and is within the practical output range of most on-chip PLLs. The third constraint is the options available for on-chip fractional PLLs, to maximize LCM subject to the first two constraints.

The ZCU208 prototype uses a 125 MHz external clock. An MMCM on-chip fractional PLL is configured to generate vernier clocks with effective scaling factors of 102/639 and 102/640. After applying the divide-by-two circuit, this yields effective vernier clocks of approximately 9.961 and 9.977 MHz and an LCM of just over 8.0 microseconds. MMCM bandwidth parameters are chosen to ensure jitter is sufficient for full linearization by dithering, as discussed at the end of Section VI.

REFERENCES

[1] P. Moreira, P. Alvarez, J. Serrano, I. Darwezeh, and T. Wlostowski, "Digital dual mixer time difference for sub-nanosecond time synchronization in Ethernet," in *Proc. IEEE Int. Freq. Control Symp.*, Jun. 2010, pp. 449–453.
[2] S. Römisch, T. E. Parker, and S. R. Jefferts, "Novel, all-digital phase measurement system for time scales," in *Proc. 41st Annu. Precise Time Time Interval Syst. Appl. Meeting*, 2009.
[3] M. Rizzi, M. Lipinski, T. Wlostowski, J. Serrano, G. Daniluk, P. Ferrari, and S. Rinaldi, "White rabbit clock characteristics," in *Proc. IEEE Int. Symp. Precis. Clock Synchronization Meas., Control, Commun. (ISPCS)*, Sep. 2016, pp. 1–6.
[4] M. Rizzi, M. Lipinski, P. Ferrari, S. Rinaldi, and A. Flammini, "White rabbit clock synchronization: Ultimate limits on close-in phase noise and short-term stability due to FPGA implementation," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 65, no. 9, pp. 1726–1737, Sep. 2018.
[5] P. Moreira and I. Darwazeh, "Digital femtosecond time difference circuit for CERN's timing system," in *Proc. London Commun. Symp.*, 2011.
[6] F. Girela-López, E. Ros, and J. Díaz, "Precise network time monitoring: Picosecond-level packet timestamping for fintech networks," *IEEE Access*, vol. 9, pp. 40274–40285, 2021.
[7] *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE Standard 1588-2019, Section 7.3.4 and 10.1.1, 2019.
[8] R. Baron, "The Vernier time-measuring technique," *Proc. IRE*, vol. 45, no. 1, pp. 21–30, 1957.
[9] D. Lee, J. Sung, and J. Park, "A 16ps-resolution random equivalent sampling circuit for TDR utilizing a Vernier time delay generation," in *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, Oct. 2003, pp. 1219–1223.
[10] K. Omote, K. Shimizu, and J.-I. Okamura, "A Vernier over sampling and alignment technique for Gb/s serial communication," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2005, pp. 29–32.
[11] Y. Song, D. Zeng, and Y. Tian, "Accurate time interval measurement method based on Vernier caliper principle," in *Proc. Int. Conf. Image Anal. Signal Process.*, Apr. 2010, pp. 553–555.
[12] C. E. Cummings, "Clock domain crossing (CDC) design & verification techniques using SystemVerilog," in *Proc. SNUG*, Boston, MA, USA, 2008, pp. 1–56.
[13] S. A. Stephens and J. B. Thomas, "Controlled-root formulation for digital phase-locked loops," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 31, no. 1, pp. 78–95, Jan. 1995.
[14] G. L. Steele and S. Vigna, "Computationally easy, spectrally good multipliers for congruential pseudorandom number generators," *Softw., Pract. Exper.*, vol. 52, no. 2, pp. 443–458, Feb. 2022.

ALEXANDER C. UTTER received the B.S. degree in engineering from the Harvey Mudd College, Claremont, CA, USA, in 2005, and the M.S. degree in electrical engineering from Stanford University, Stanford, CA, USA, in 2007.

Since 2007, he has been with the Digital Communication Implementation Department, The Aerospace Corporation. He has led the development of FPGA avionics subsystems used on more than a dozen cubesat missions, including the SatCat5 open-source Ethernet switch, the AeroCube camera, and modems for free-space optical communication.