A Multi-Path Gated Ring Oscillator TDC With First-Order Noise Shaping

Matthew Z. Straayer, *Member, IEEE*, and Michael H. Perrott, *Member, IEEE*

***Abstract*—An 11-bit, 50-MS/s time-to-digital converter (TDC) using a multipath gated ring oscillator with 6 ps of effective delay per stage demonstrates 1st-order noise shaping. At frequencies below 1 MHz, the TDC error integrates to 80 fs (rms) for a dynamic range of 95 dB with no calibration required. The 157 $\times$ 258 μm TDC is realized in 0.13 μm CMOS and, depending on the time difference between input edges, consumes 2.2 to 21 mA from a 1.5 V supply.**

***Index Terms*—Noise-shaping, phase detector, quantizer, ring oscillator, time-to-digital converter.**

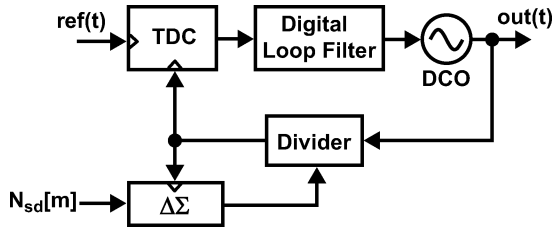

Fig. 1. Basic diagram of a digital PLL.

I. Introduction

As CMOS progresses toward finer scale geometries, there is pressing need for interface circuits to move away from their traditional analog focus and instead leverage digital circuits that take better advantage of Moore's Law. A key focus area for such interface circuits are phase-locked loops (PLLs), which provide provide precise generation and/or alignment of timing signals for a wide variety of applications including frequency synthesizers, clock and data recovery circuits (CDRs), and clock generation. The traditional analog approach to implementing these systems is becoming undesirable due to its need for large loop filter capacitors and the complications of designing analog-intensive components, such as charge pumps, within CMOS processes that are progressively degrading analog process characteristics such as intrinsic gain (i.e., $g_m r_o$), leakage, and matching.

Recently, digital phase locked loops have emerged as a viable alternative to traditional PLLs [1]–[3]. As an example, Fig. 1 displays a digital PLL in the form of a fractional-N frequency synthesizer [4]. As shown in the figure, a digital filter is used in place of a traditional *RC* loop filter, which has the benefit of removing the need for large loop filter capacitance and analog components such as charge pumps. However, there is now a need to use a time-to-digital converter (TDC) to achieve digital encoding of the phase error between reference and divider output, and a digitally-controlled oscillator (DCO) to allow interfacing to the digital loop filter. While there has been much progress on achieving high performance DCO circuits [1], [3], the development of high performance TDC circuits is currently an active research topic in the mixed-signal circuit community [5]–[10].

In this paper, we will focus on the issue of achieving high performance time-to-digital conversion in the form of a structure we call a Gated Ring Oscillator (GRO) TDC [11]. The key performance metrics that we will pursue are high resolution and high linearity in the TDC time-to-digital mapping characteristic, and low power and low area in its implementation. Unlike previous TDC implementations which essentially perform an operation analogous to FLASH [5] or two-step [7], [9] conversion in analog-to-digital converters (ADCs), the GRO structure performs an operation analogous to $\Sigma\Delta$ ADCs by noise shaping the quantization noise of the TDC. Since the loop filter within the PLL acts to filter out the high frequency noise produced by the TDC, such noise shaping yields greatly improved *effective* resolution of the TDC time-to-digital mapping characteristic. To achieve such noise shaping, the GRO must process and store analog information corresponding to the quantization noise produced by the TDC. However, we will show that the GRO structure provides a purely digital circuit implementation to achieve such analog processing and storage.

An overview of the paper is as follows. In Section II, we begin by examining the characteristics of the commonly used delay chain TDC in order to highlight the key challenges to achieving a TDC with high resolution and high linearity. We then quickly review a basic implementation of the GRO TDC in order to explain its ability to perform noise shaping [4] with a purely digital implementation. In Section III, we present the key focus of this paper, which is a proposed multi-path implementation of the GRO [12] that dramatically improves its effective resolution and still preserves its noise shaping properties. Section IV provides details of achieving a robust and efficient implementation of the multi-path GRO, and Section V demonstrates its viability by showing measured results from a custom integrated circuit prototype in 0.13 μm CMOS. In particular, the prototype will demonstrate that less than 100 fs of rms error can be achieved over a bandwidth of 1 MHz given a 50 Msample/s rate on the

Manuscript received September 02, 2008; revised December 02, 2008. Current version published March 25, 2009. This work was supported by MIT-Lincoln Laboratory. The work was performed at the Massachusetts Institute of Technology, Cambridge, MA.

M. Z. Straayer is with Cambridge Analog Technologies, Inc., Bedford, MA 01730 USA (e-mail: mstraayer@cambridgeanalog.com).

M. H. Perrott is with SiTime, Inc., Sunnyvale, CA 94085 USA.

Digital Object Identifier 10.1109/JSSC.2009.2014709

0018-9200/$25.00 © 2009 IEEE

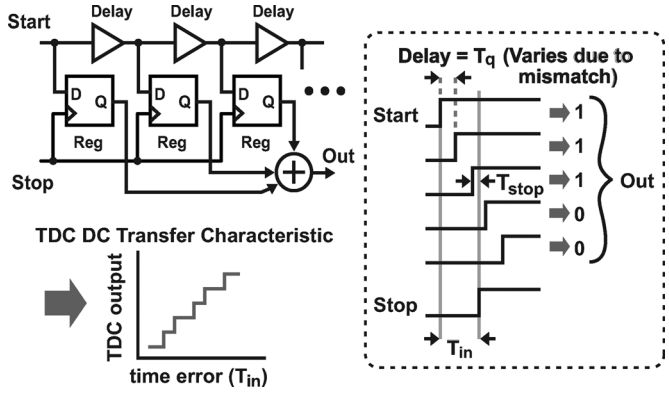

Fig. 2. Classical TDC architecture based on linear delay elements.

TDC. Finally, Section VI provides a summary of the paper and then concludes.

II. Background

Fig. 2 illustrates the commonly used delay chain TDC. To explain the operation of this classical TDC, the rising edge of the start signal, which represents the first event, is successively delayed by a series of inverter gates (polarity is ignored throughout for simplicity), each with average delay T_q. The outputs from each of these inverters are input to a register, which is clocked with the rising edge of the stop signal representing the second event. A thermometer code is then generated at the register output, which corresponds to the number of delay elements that have transitioned within the measurement interval $T_{\rm in}[k]$. The TDC output $Out[k]$ is then simply calculated as the sum of the thermometer code.

As illustrated by Fig. 2, a practical TDC has a quantized dc transfer characteristic function (i.e., average mapping from input time difference, $T_{\rm in}[k]$, to digital code output, $Out[k]$) which has limited resolution and nonlinearity due to mismatch between delay stages. The limited resolution poses issues for integer-N frequency synthesizers and CDRs since it will cause nonlinear behavior such as limit cycles as encountered with bang-bang feedback systems [13]. This situation can be compared to the classic dead-zone in an analog phase detector, which is well-known to cause erratic limit-cycle behavior in integer-N PLLs. In the case of fractional-N frequency synthesizers, the time error varies in time according to the fractional division operation being performed, which effectively scrambles the quantization noise of the TDC and lowers the possibility of limit cycles occurring. However, the nonlinearity of the TDC dc transfer characteristic due to mismatch will cause noise folding of the quantization noise produced by dithering, which adversely impacts the noise performance of the synthesizer.

Given the above issues, when seeking a high performance TDC, the key goals are to lower the delay per stage so that high resolution is achieved, to achieve high matching between delay stages so that high linearity is achieved, and to prevent nonlinear behavior in the overall PLL such as limit cycles. While the use of future CMOS process will somewhat help the goal of lowering the delay per stage, the issue of mismatch will become increasingly problematic in cases where low power and area are essential. The issue of limit cycles in the PLL can be addressed through scrambling of the TDC quantization noise, as will next be described.

A. The Benefit of Quantization Noise Scrambling

Fig. 3 illustrates the impact of noise and quantization noise scrambling on the dc transfer characteristics of a TDC. As shown in part (a) of the figure, the quantization effects of the classical delay chain TDC yield a dc transfer characteristic that has a staircase function (note that mismatch is being ignored in this case). As shown in part (b), the presence of a small amount of noise in the system, which is manifested as jitter on the start and stop edges as well as the internal signals of the TDC, acts to smooth the dc transfer characteristic but still yields a significant remnant of the staircase TDC nonlinearity. Finally, as shown in part (c), the presence of a high level of noise or purposeful scrambling effectively linearizes the dc transfer characteristic of the TDC. In such case, it is common to assume that TDC quantization process primarily influences the PLL through the introduction of "white" quantization noise. For this to be true, note that the noise or scrambling action must be sufficiently "exciting" and large in comparison to the quantization step size of the TDC.

One might consider some of the traditional approaches to achieving properly scrambled quantization noise, such as intentionally modulating the TDC input with a sufficiently noisy signal in order to improve the randomness of the quantization error. Of course, the presence of such noise adds to the impact of the quantization noise, and can degrade the overall noise performance of the TDC. However, if the "noisy" signal is known and the gain of the TDC is well-characterized, this "noise" can then be subtracted from the TDC output. Unfortunately, this approach is challenged by the presence of mismatch between delay elements in the TDC, which greatly complicates the effort to accurately subtract the added noise from the TDC output. As such, it is worthwhile to consider an alternative approach to quantization noise scrambling in the TDC as offered by the proposed GRO TDC structure.

B. Basic GRO Structure

Fig. 4 illustrates the key concepts of the gated ring oscillator (GRO) TDC [4], [11], which is similar to an oscillator-based TDC [14], [15] in that it measures the number of delay element transitions during a measurement interval as generated by a ring oscillator structure. However, unlike a traditional oscillator-based TDC, the GRO structure only allows the oscillator to have transitions (i.e., to be "gated" on) during a given measurement, and strives to freeze the ring oscillator state between measurements.

The benefit of gating the oscillator is that the residue occurring at the end of a given measurement interval, $T_{stop}[k-1]$, can be transferred to the next measurement interval, $T_{start}[k]$. To see how this results in first-order noise shaping, we first note that

$$T_{start}[k] = T_{stop}[k-1]. \tag{1}$$

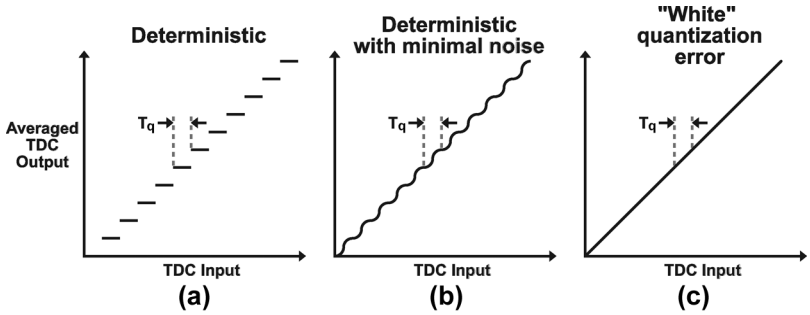

Fig. 3. The dc transfer characteristics for (a) a completely deterministic TDC, (b) a deterministic TDC with small jitter either due to thermal noise or the input, and (c) a TDC with "white" quantization error due to inherent error scrambling or external dithering.

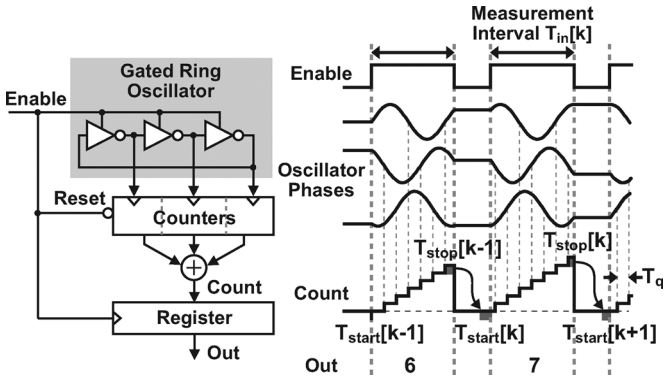

Fig. 4. Concept of the gated ring oscillator TDC.

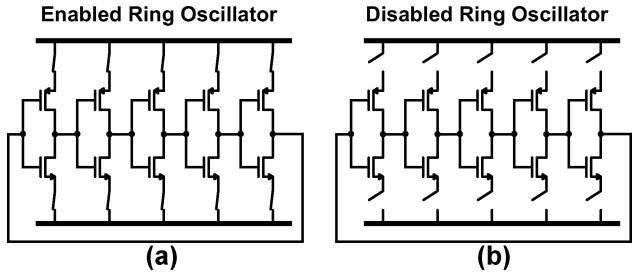

Fig. 5. Conceptual implementation of gating a ring oscillator.

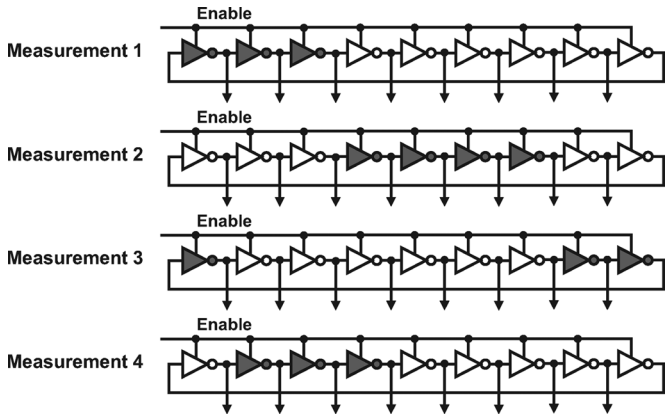

Fig. 6. Barrel-shifting of GRO delay elements to achieve first-order shaping of mismatch error, where the shaded delay elements represent a transition occurring during the measurement.

Second, by examining Fig. 4 and then combining with (1), the overall quantization error for a given measurement interval, T_{error}, is given as

$$T_{error}[k] = T_{stop}[k] - T_{start}[k] \quad (2)$$

$$= T_{stop}[k] - T_{stop}[k-1]. \quad (3)$$

This discrete-time first-order difference operation on the baseline noise process, $T_{stop}[k]$, corresponds with a first-order noise shaping in the frequency domain.

To gain a practical understanding of the GRO TDC, Fig. 5 displays a conceptual implementation of the oscillator core [11]. In essence, the GRO adds gating functionality to a classical inverter-based ring oscillator by placing transistor switches in series with the positive and negative power supply connections for each inverter. When the switches are closed, oscillation is enabled and the ring of inverters behaves identically to a classical ring oscillator (Fig. 5(a)). Conversely, when the switches are open, the inverter delay element is unable to charge or discharge the parasitic output capacitance, and as a result oscillation is suspended (Fig. 5(b)). The oscillator phase at the end of the enabled state is then held during the disabled state with the charge stored on the parasitic capacitance of the delay elements.

A subtle aspect to the GRO-TDC is that, along with the quantization noise, the delay element mismatch is also first-order shaped. To see this more clearly, let us examine the sequencing of delay elements for successive TDC conversions, as shown in Fig. 6. What is evident in this figure is that the selection of delay elements for a given input is equivalent to the well-known barrel-shift algorithm [16] for dynamic element matching. Similar to the transfer of quantization error, the mismatch errors for one sample are also passed along to and subtracted from the following sample.

Finally, along with the first-order shaping of the quantization and mismatch error, we can also say that the GRO-TDC quantization error is intrinsically scrambled. To explain, if we assume that the ring oscillator is free-running with respect to the time intervals that it is measuring, then $T_{stop}[k]$ from Fig. 4 will take on a uniform distribution over the interval $[0, T_q)$, where T_q is the delay per stage as in Fig. 2. Note that this logic can also be extended to include quantization error from delay element mismatch.

In summary, the GRO TDC can achieve intrinsic scrambling of its quantization and mismatch error, as well as first-order noise shaping. As such, there is no need to perform calibration

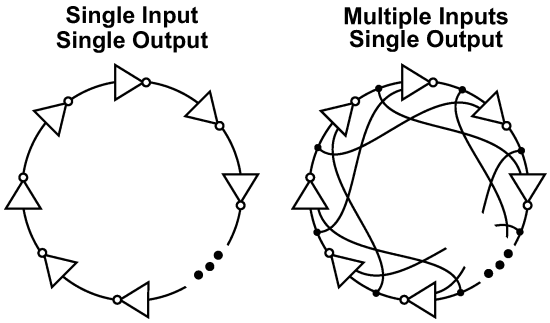

Fig. 7. Traditional versus multi-path ring oscillator implementation: (a) traditional structure, (b) multi-path structure.

of the TDC to achieve high resolution and linearity, even in the presence of large mismatch. As with other delay-based TDCs, the GRO may require calibration of its overall gain due to the influences of process and temperature variations on the delay per stage. However, such TDC gain calibration can be achieved in a continuous manner without significant difficulty, as discussed in [1] and [4].

III. Multipath Gated Ring Oscillator

While the basic GRO structure discussed in the previous section (Fig. 5) offers scrambling and noise shaping benefits over classical TDC structures, its effective resolution is still related to the delay per stage that is achieved in the structure. As such, reduction in the delay per stage is the chief means of improving its resolution. In this section, we will present a multi-path implementation technique for the GRO which achieves over a factor of 5 improvement in resolution over the basic GRO structure discussed in the previous section. The resulting quantization noise of the GRO-TDC is effectively reduced to the point where thermal and $1/\mathrm{f}$ noise in the GRO becomes the dominant noise source at low frequencies, and the impact of the first order shaped quantization noise can be completely eliminated by the lowpass filtering action of the PLL.

The key idea of the proposed GRO structure is to leverage multi-path techniques that have previously been demonstrated for ring oscillators to lower the effective delay per stage of the oscillator [17], [18]. As shown in Fig. 7, a classical ring oscillator, which is shown in part (a) of the figure, taps only the previous delay stage output for its input. In contrast, a multi-path ring oscillator, which is shown in part (b) of the figure, taps several previous delay stages in order to speed up the transition time at the output of the delay stage, and thereby lowers the effective delay per stage.

Note that there are other possibilities for achieving an effective reduction in the delay per stage of the GRO. One such possibility is to realize the GRO as M coupled gated oscillators, each with N stages, which theoretically reduces the effective delay per stage by a factor of M [19], [20]. One issue that must be carefully considered for a system of coupled oscillators, and especially for the GRO, is how to initialize and maintain oscillation within the primary mode. Because the very premise of the GRO is that it will be stopped and started at the same phase state with no intervention or reset operation, we need to achieve a well-defined oscillation through the gating operation. As such, the multi-path structure offers a potentially simpler implementation than this coupled oscillator structure, the details of which are described in the following two subsections.

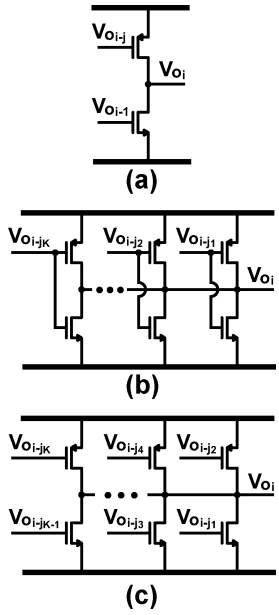

Fig. 8. Techniques to reduce effective delay by modifying the standard inverter. (a) Asymmetrically skewed inverter; (b) multiple skewed inverters; (c) unrestricted connections.

A. Multi-Path Delay Cell

To achieve the multi-path GRO implementation, we must modify the basic inverter delay structure so that it allows multiple inputs. There are several means of accomplishing this, as shown in Fig. 8. The first is to asymmetrically connect the pMOS and nMOS inputs to different delay stage outputs (Fig. 8(a)) [17], which allows an earlier arrival of the input transition to the slower pMOS transistor, thereby speading up the average delay per stage. Alternately, multiple connections to entire inverter cells can be made (Fig. 8(b)) [18] to improve the delay per stage while maintaining stable primary mode oscillation. In both cases, the effective delay through the stage has been reported to be reduced by a factor of 2 in the context of realizing a ring oscillator [17], [18].

In contrast, consider that each of the above techniques can be combined together to result in an unrestricted set of transistor connections as shown in Fig. 8(c). In the proposed topology, K *transistors* connect to a set of output stages $\{V_{o_{(i-j_1)}}, V_{o_{(i-j_2)}}, \ldots, V_{o_{(i-j_K)}}\}$, which gives the designer a larger optimization space compared with Fig. 8(a) and (b). Specifically, the connection and size of *each* transistor in the delay cell can be *independently* adjusted, and the overall design can be *fully* optimized. We will show later in the paper that the prototype GRO decreases the effective delay per stage by a factor of 5–6 while maintaining a stable and robust oscillation.

B. Design of the Proposed Multi-Path GRO

Fig. 9 displays the overall multi-path GRO that is achieved by cascading the unrestricted delay structures presented in Fig. 8(c)

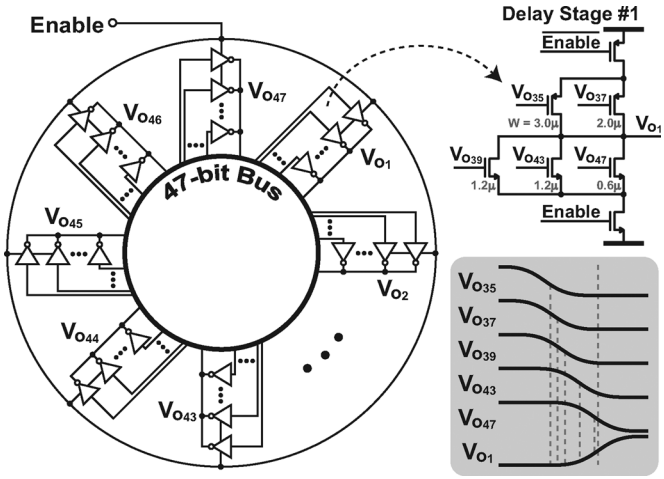

Fig. 9. Schematic of the proposed multi-path GRO.

within a ring topology. To achieve the desired gating functionality in the same manner as previously described for the basic GRO structure, each delay cell is modified by appropriately placing switches above and below the inverter core. As revealed by the figure, there are numerous design choices including the number of stages employed, and the choice of each delay stage connection as well as its weight (i.e., the size of the connection device). We will briefly discuss such design issues in the next few paragraphs.

Let us first consider the number of stages N that is appropriate for use in the GRO-TDC application. Counting and measuring the GRO outputs with standard digital logic places an upper bound on the oscillation frequency of 2 GHz in 0.13 μm CMOS, which along with the achievable delay per stage places a bound on the minimum value for N. Another issue for choosing N is that a prime value inherently has better rejection of undesirable modes than does a value of N with large odd factors, such as $45 = 3 \cdot 15 = 5 \cdot 9$. As a result of these considerations, we chose $N = 47$ for our prototype.

To set the delay cell transistor connections and sizes, we used a soft design procedure based on empirical simulation results in combination with the desire to minimize layout complexity and area. A useful metric for evaluating designs is the power-delay product, which can achieve a local minimum for a well-designed multi-path oscillator. Another useful indicator of stability, albeit somewhat qualitative, is the steady-state start-up time of the oscillator when given a minor charge injection onto one of the oscillator nodes.

Although simulation is used for final assignment of connections and weights, there are a number of guidelines that are also useful to generate a first-pass design that is relatively close to optimal. As discussed earlier in regard to Fig. 8(a), connecting the slower pMOS transistors to stages farther away than the nMOS connections results in better efficiency. Depicted in Fig. 9, the pMOS connections are then made the 13th and 11th preceding stages $\{V_{o(i-13)}, V_{o(i-11)}\}$. Another consideration is to properly distribute the weight of connections across the entire set, which aids in establishing the primary oscillation mode. Therefore, nMOS connections are made to $\{V_{o(i-9)}, V_{o(i-5)}, V_{o(i-1)}\}$.

In terms of sizing, a larger weight is assigned to the connections with longest distance, with only a small drive from the immediately preceding stage. To facilitate a compact and simple layout, the same number of equal-width fingers (5) is used for each of the four levels in the transistor stack (PMOS switch, pMOS inverter, nMOS inverter, nMOS switch). For optimal power and speed, the switch transistors are sized wider than the total width of the inverter core transistors.

IV. Measurement Circuits

While we have thus far focused on achieving high performance for the core gated oscillator portion of the GRO by lowering the effective delay per stage, the function of counting its edges during a given measurement interval becomes challenging when trying to achieve a reliable implementation with low power and area. Aside from the issue of power and area, one must also be careful to avoid corruption in the counting process as this will easily destroy the noise shaping properties of the GRO and therefore degrade its performance. With such issues in mind, we now present an efficient approach to counting the GRO edges during each measurement inverval. To simplify our discussion, we consider a basic GRO structure first, and then show additional techniques that address the added complexity of the multi-path GRO structure.

A. Illustration With a Simple GRO

The task of the measurement circuitry is simply to calculate the number of delay element transitions that occur during each sample, which could simply be implemented with $2N$ counters (N for each the positive and negative transitions). However, for the GRO, *counting* all of the delay element output transitions is much more expensive in terms of area and power than is *sampling* the outputs with a digital register. Fig. 10 illustrates the basic concept of calculating the GRO-TDC output by quantizing the oscillator state with registers, mapping the state into a discrete phase count, and then performing a first-order difference on each sample.

In sampling the phase information of the oscillator state, we need to account for wrapping that occurs when the transitions travel farther than the number of delay cells within a given measurement interval. Assuming N stages, the quantized phase state of the GRO is limited to $2N$ levels (which corresponds to 2π in phase of the oscillator). To account for wrapping, we separate the phase into two components: a fine phase residual that is calculated from the registers, and a coarse phase that accumulates $2N$ (in the figure $N = 15$) each time the oscillator phase wraps around. The coarse phase accumulation can simply be implemented by counting the positive transitions of a chosen delay element, with an overflow bit to correct for exceeding the counter range.

If a double-counting error in the counter is made in this topology, it is likely that the TDC output will be wrong by at least $2N$, since the counter output is amplified by this value. This magnitude of error lacks noise-shaping and would likely be very disruptive at the application level. To prevent such errors, we therefore propose the use of a de-glitch circuit at the counter input that employs two techniques to avoid double-counting. As shown in Fig. 11, the proposed de-glitch logic combines the

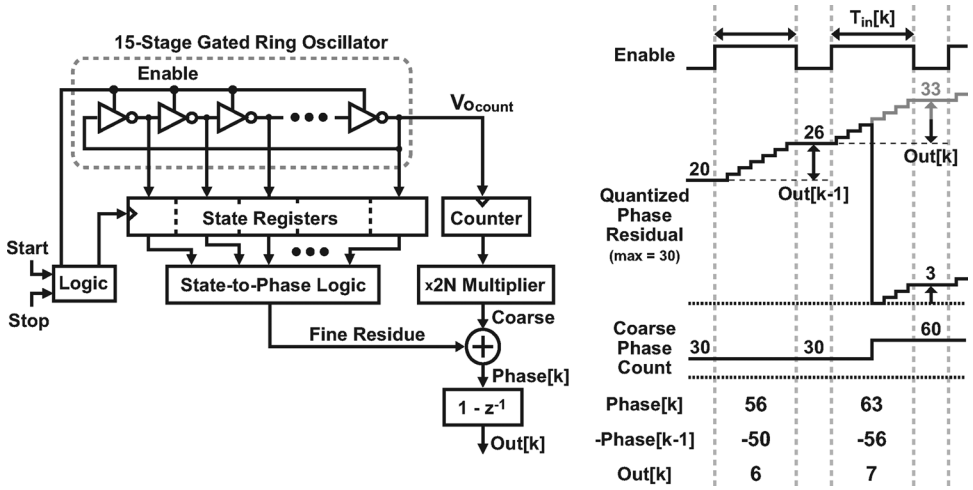

Fig. 10. Basic concept of calculating the GRO-TDC output by differentiating phase.

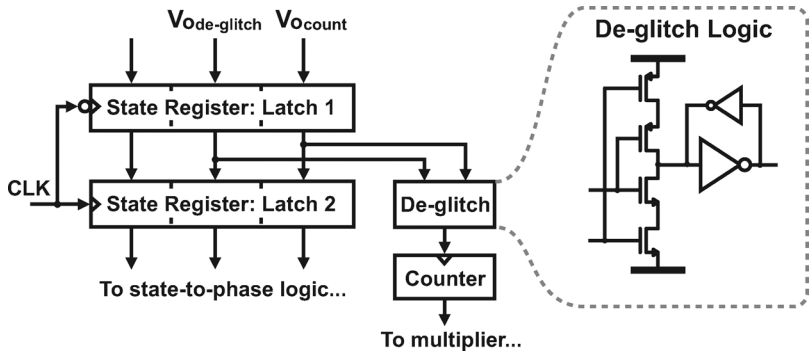

Fig. 11. Implementation of a de-glitch circuit that achieves hysteresis by relying on the sequence of oscillator states.

latching action from the first half of the state register along with logic that only allows the counter to observe phase rotation in a single direction.

To explain further, in addition to the original GRO output to be counted, $V_{o_{count}}$, the de-glitch circuit monitors a second GRO output, $V_{o_{de-glitch}}$, whose transitions occur just before $V_{o_{count}}$ and in the same direction. When the two de-glitch inputs share the same logic level, the counter observes their state, but when the inputs do not have the same value, the counter input is held constant with positive feedback. In this way, a logically ambigous state of $V_{o_{count}}$ during the gating operation is resolved to ensure measurement accuracy and to preserve the GRO noise-shaping.

B. Multi-Path GRO-TDC Implementation Details

Although we can leverage a similar phase-sampling approach for counting edges in the multi-path GRO structure, we must modify it appropriately to account for the fact that its phase state is much more complicated. For example, in contrast to the inverter-based topology, each delay element in the multi-path oscillator may begin transitioning well before the preceding stage is close to completing it own transition. In this case, the sequence of transitions may be deterministic within a specific realization, but predicting this exact sequence prior to measurement is extremely challenging, and in addition there is the possibility that two transitions could cross their respective logical thresholds in a random order. Because a predictable relationship between the transition sequences is required for the state-to-phase logic, the GRO state ambiguity becomes a primary challenge.

One potential way to solve this issue is to create an algorithm that populates a dynamic look-up table based on observing the TDC output, but this approach is cumbersome and inefficient. Alternately, we could simply revert back to counting each of the delay element outputs independently, but we have already discussed the associated drawbacks in this case. Fortunately, there is a compromise between having a single counter and having $2N$ counters.

Fig. 12 illustrates the concept of partitioning the entire GRO state into 7 smaller measurement cells. Here we choose enough cells and distribute the cell inputs so that instead of having multiple ambiguous inputs in the state-to-phase logic, there is at most one delay element in transition per cell at any given time. The tradeoff in this approach is the increased power of having one counter for each cell instead of one counter for the entire GRO. This small penalty is far outweighed by re-establishing the predictable sequence of states, at least with respect to each individual cell. The measurement cells can then independently calculate their outputs, which are then simply summed together in the final step to result in the overall TDC output.

Finally, we show a system block diagram for the proposed 47-stage multi-path GRO-TDC in Fig. 13. Although we have

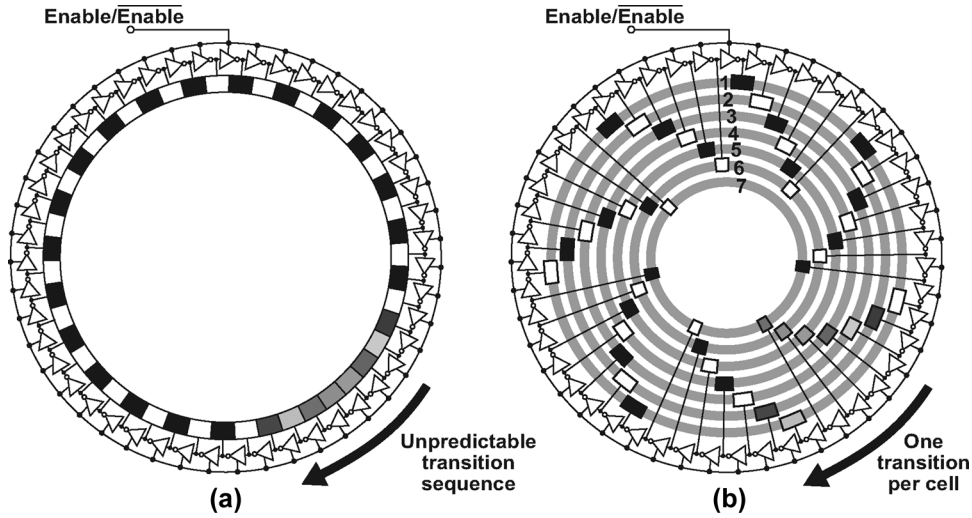

Fig. 12. A geometric view of an example multi-path GRO state illustrating (a) the unpredictable transition sequence considering the entire multi-path oscillator, and (b) a partitioned approach that re-establishes predictable transition sequences within each of the 7 independent measurement cells.

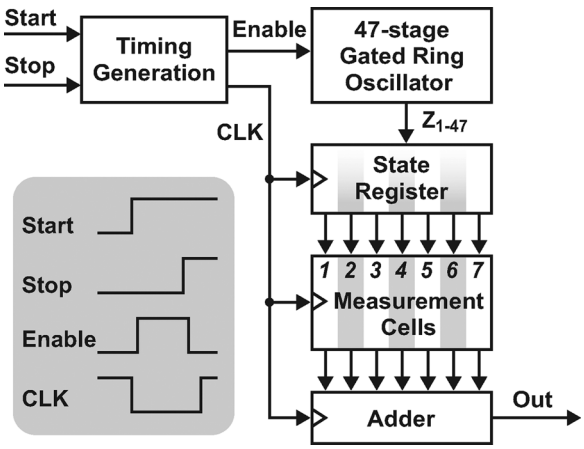

Fig. 13. Overall system block diagram for the proposed 47-stage multi-path GRO-TDC.

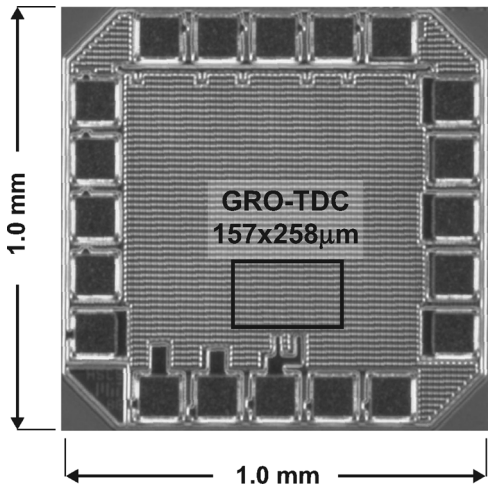

Fig. 14. Microphotograph of the GRO-TDC chip.

discussed at some length the GRO core and the measurement cells, a few other digital circuit blocks are also needed within the TDC. The timing generation block takes a start and stop signal input, generates the differential *Enable* signal, and sufficienctly buffers *Enable* in order to drive the GRO core with modest rise and fall times (and correspondingly modest jitter). In addition, the timing generation block derives the other clocking signals as required by the measurement cells. Last, the output adder receives all of the calculated outputs from each of the measurement cells and sums them to result in the overall GRO-TDC output.

V. Measurement Results

A microphotograph depicting an 11-bit GRO-TDC 1.0 mm $\times$ 1.0 mm die fabricated in 0.13 μm CMOS is shown in Fig. 14. Within the chip, the active silicon area of the larger 11-bit GRO-TDC is 258 μm $\times$ 157 μm.

The 0.13 μm CMOS process used for the prototype supports up to a 1.5 V supply. Since higher supply voltage leads to a shorter delay per stage, we will present measurement results at 1.5 V in order to examine the best practical performance of the GRO. At this voltage setting, the nominal delay per stage is 6 ps. However, proper operation of the prototype GRO was verified across a supply voltage range of 1.0 to 1.6 V, with lower voltages yielding longer delays per stage but also lower power consumption.

The power consumption of the GRO-TDC is a linear function of the width of the enable signal corresponding to its input. At 1.5 V supply and 50 MS/s operation, the power ranges from 2.2 mW to 21 mW as the width of the enable pulse sweeps from a small to large value. The GRO-TDC is operational past 100 MS/s, although here the range becomes constrained by the sampling period. Note that in a typical PLL application, a large input range is only required for obtaining lock, and in normal operation a small enable width can be forced through feedback [4].

The measured multi-path delay of 6 ps represents an improvement factor of over 5 compared to an inverter-based GRO-TDC delay of 30–35 ps under the same voltage supply and operating conditions [11]. This result verifies the significant benefit

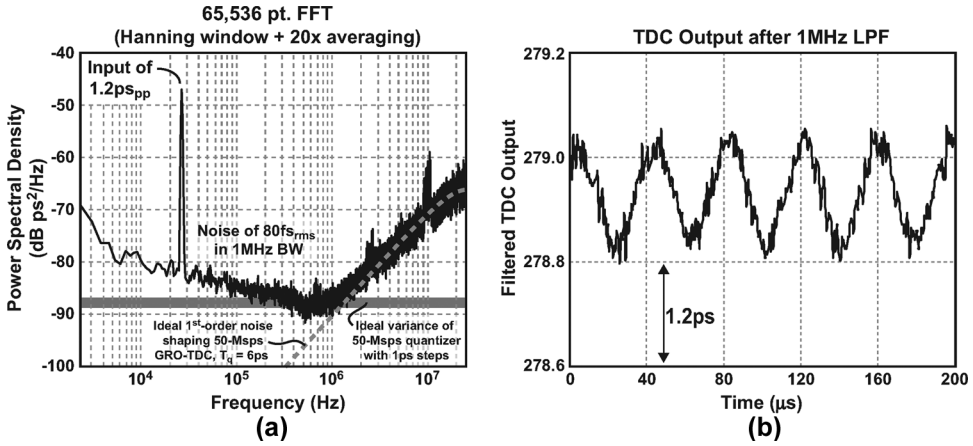

Fig. 15. Measured GRO-TDC output for a 1.2ps_{pp}, 26 kHz input signal. (a) plots the signal and power spectral density in the frequency domain, and (b) is a transient view of the output after digital low-pass filtering with a 1 MHz bandwidth.

in raw resolution that multi-path oscillators can offer for TDC applications.

The GRO-TDC can measure a time difference between edges created by two distinct time sources, however the overall jitter in such a measurement is significantly larger than the internal TDC noise. Therefore, this setup is only used for the large-signal inputs, and for more precise measurements we simply measure the delay of a single off-chip digital buffer (i.e., the relative difference in time between the buffer input and output transitions). In this case, the phase noise of the reference frequency does not contribute to the input signal, and the small noise contribution from the buffer is expected to be comparable to the internal TDC noise. Additionally, the power supply of the buffer can then be modulated to superimpose a very small ac input signal with a dc level.

Fig. 15 shows the both the frequency and time domain GRO-TDC 50-MS/s output with a 26 kHz input of 1.2 ps_{pp} in addition to a dc level of about 1.6 ns. In (a), the 65,536 point FFT is performed with a Hanning window on 20 sequential collects before being averaged to result in the double-sided power spectral density as shown. The measured results have very good agreement with the theoretical first-order noise-shaping, as shown by the thick grey dotted line and given by

$$S_q(f) = \frac{1}{T}\frac{(T_q)^2}{12}|1 - e^{-2\pi fT}|^2 \tag{4}$$

where T is the TDC sample rate. The wide, shaded horizontal line in Fig. 15 shows that the low frequency power spectral density of the GRO-TDC output is dominated by $1/\text{f}$ noise and is comparable to what ideally would be produced by a 50 MS/s classical quantizer (i.e., scrambled quantization errors without noise shaping) with 1 ps steps.

By looking at the time domain output after digitally filtering with a 1 MHz bandwidth in Fig. 15(b), the GRO-TDC is clearly able to resolve a 1.2ps_{pp} signal. In fact, the integrated noise of the GRO-TDC from 2 kHz to 1 MHz is below 80fs_{rms}, which includes the noise of both the GRO and the off-chip buffer delay.

When considering the profile of the thermal noise in the output spectrum, recall that the GRO-TDC measures the time *difference* between the start and end of its enable period. This means that thermal noise in the delay stages will only introduce jitter while the GRO is enabled during a given measurement, and it also implies that random jitter will not accumulate across multiple measurements. As such, the effect of such thermal noise is to add *white noise* to the GRO-TDC output (as opposed to the *phase noise* profile of a ring oscillator, which has a $1/\text{f}^2$ frequency dependence). A short enable time will yield a small amount of measurement jitter, or output noise variance due to thermal noise, and a long enable window leads to increased densities of white noise in the output spectrum.

With the noise-shaping operation now verified, let us consider again the ideal dc transfer function from Fig. 3(c). Interestingly, we do observe a small deadzone in the multi-path GRO for the special case when the input time is close to an integer multiple of the GRO period, $T_{\text{in}} = KT_{\text{osc}}$, where K is an integer and $T_{\text{osc}} = 2NT_q$. This result is shown in Fig. 16(a), where the slope of the transfer function becomes zero for this input value.

To briefly explain this deadzone behavior, consider that the process of gating the oscillator will introduce a small gating error to each measurement, with the gating error taking on a value that depends on circuit issues such as charge injection, charge redistribution within the inverter core, and the relative shape of the differential *Enable* signal [21]. If we assume proper scrambling of the GRO, as caused by a sweeping action of the phase error value across progressive measurements, the gating error will be statistically independent of the TDC input, and will therefore simply contribute to the noise floor and overall offset of the TDC. However, if the sweeping action of the phase error is too slow, which occurs under conditions described in [21], the gating error will not be properly scrambled such that its impact becomes dependent on the value of the TDC input. As described in [21], this signal dependence can lead to "injection-locking" of the GRO phase to that of the TDC sample rate, which in turn leads to a deadzone in the overall TDC transfer characteristic. One important note to make is that the multi-path GRO architecture is much less sensitive to this issue than is a classical ring oscillator [11], thanks to the distributed delay elements.

No deadzones are evident for GRO-TDC outputs other than at $2NK$ (e.g., Fig. 16(b)), and the size of the worst-case deadzone is found to be only 1.1 ps. Assuming that the size of this

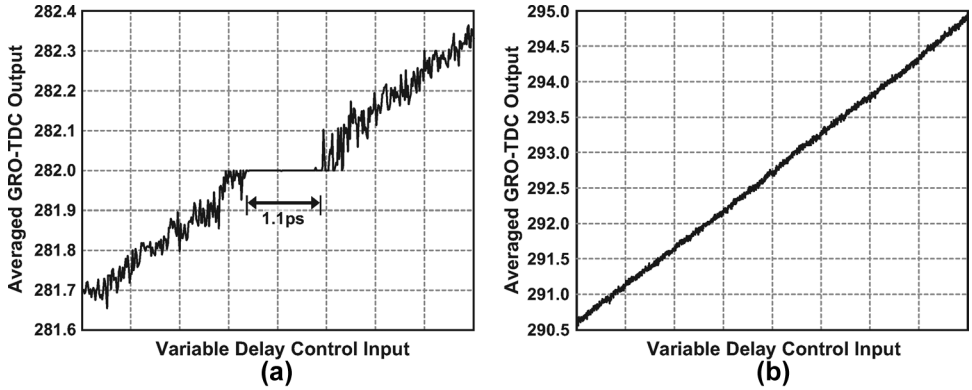

Fig. 16. A measured dc transfer characteristic for the multi-path GRO-TDC that demonstrates (a) the presence of small deadzones for TDC outputs at $2NK$, and (b) linear behavior for integer TDC outputs.

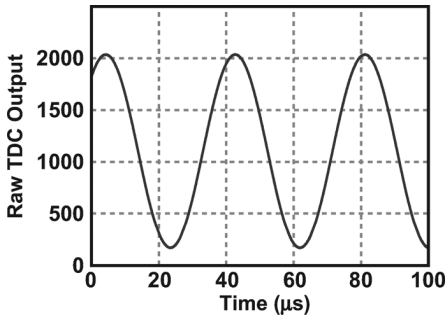

Fig. 17. Raw measured GRO-TDC output for a 26 kHz input signal with an amplitude near full-scale.

TABLE I
SUMMARY OF MULTI-PATH 11-bit GRO-TDC MEASURED PERFORMANCE

Specification	Value
Maximum Sampling Frequency	100 MHz
Range	11-bits
Raw delay resolution	6ps
Effective resolution	1ps @ 50Msps
Integrated noise	80fs, 2kHz-1MHz
Dynamic range	95dB
Power	2.2-21mW (1.5V)
Efficiency	0.2pJ/step
GRO-TDC Area	157μm×258μm
Total Chip Area	1.0mm×1.0mm
Technology	0.13μm IBM CMOS

deadzone corresponds with the peak-to-peak non-ideality for the entire GRO phase state, we can expect that the GRO-TDC output noise will generally be dominated by $1/\mathrm{f}$ and quantization noise as shown in Fig. 15. In a system application with constant input (such as an integer-N PLL), avoiding the small range of GRO-TDC outputs that correspond with $2KN$ is quite straightforward.

To illustrate the full 11-bit operation of the GRO-TDC, Fig. 17 plots raw output data from the chip when a 26 kHz input is applied with amplitude near full-scale (11 bits). The dynamic range in a 1 MHz bandwidth is calculated to be 95 dB, or an equivalent range of 15.5 bits. The TDC efficiency is calculated for a 1 MHz bandwidth to be 0.23 pJ/step, which is almost identical to the efficiency calculated with full bandwidth due to the GRO quantization noise-shaping (0.20 pJ/step). A summary of the performance is shown in Table I.

VI. CONCLUSION

In this work, we have described a prototype GRO-TDC with first-order noise-shaping that achieves state-of-the-art performance in terms of resolution, range, power, and area, with no calibration of differential nonlinearity required. To our knowledge, this is the first demonstration of a noise-shaping time-to-digital converter with the ability to accurately transfer error across a gap of inactivity from one measurement to the next.

Because of the very high resolution that is possible with the GRO-TDC, the applications that will significantly benefit from this technology are likely to be the most demanding in terms of performance. The fundamental architecture of the GRO-TDC is compact, efficient, and simple, and therefore can be easily adapted to many other less demanding applications as well, especially if techniques are used to trade resolution for power. Finally, we anticipate that as TDC become more adapted into integrated systems, the use of digital, high-performance TDC such as the GRO will become more sophisticated, and perhaps lead to the enabling of system architectures that would not be practical in a previous technology.

ACKNOWLEDGMENT

The authors would like to thank the MIT High-Speed Circuits and Systems group for invaluable comments, specifically B. Helal and C.-M. Hsu for providing observations about deadzones and negative skew. The authors are also grateful to MIT-Lincoln Laboratory for supporting this work with packaging, testing, and financial assistance.

REFERENCES

[1] R. B. Staszewski, J. Wallberg, S. Rezeq, C.-M. Hung, O. Eliezer, S. Vemulapalli, C. Fernando, K. Maggio, R. Staszewski, N. Barton, M.-C. Lee, P. Cruise, M. Entezari, K. Muhammad, and D. Leipold, "All-digital PLL and transmitter for mobile phone," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2469–2482, Dec. 2005.

[2] J. Tiero, A. Rylyakov, and D. Friedman, "A wide power supply range, wide tuning range, all static CMOS all digital PLL in 65 nm SOI," *IEEE J. Solid-State Circuits*, vol. 43, no. 1, pp. 42–51, Jan. 2008.

[3] H.-H. Chang, P.-Y. Wang, J.-H. Zhan, and B.-Y. Hsieh, "A fractional spur-free ADPLL with loop-gain calibration and phase-noise cancellation for GSM/GPRS/EDGE," in *IEEE ISSCC Dig. Tech. Papers*, 2008, pp. 200–201.
[4] C.-M. Hsu, M. Straayer, and M. H. Perrott, "A low-noise wide-BW 3.6-GHz digital $\Delta\Sigma$fractional-N frequency synthesizer with a noise-shaping time-to-digital converter and quantization noise cancellation," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2776–2786, Dec. 2008.
[5] R. Staszewski, S. Vemulapalli, P. Vallur, J. Wallberg, and B. P. T. , "1.3 V, 20 ps time-to-digital converter for frequency synthesis in 90-nm CMOS," *IEEE Trans. Circuits Syst. II*, vol. 53, no. 3, pp. 2240–2244, Mar. 2006.
[6] S. Henzler, S. Koeppe, W. Kamp, H. Mulatz, and D. Schmitt-Landsiedel, "90 nm 4.7 ps-resolution 0.7-LSB single-shot precision and 19 pJ-per-shot local passive interpolation time-to-digital converter with on-chip characterization," in *IEEE ISSCC Dig. Tech. Papers*, 2008, pp. 548–549.
[7] M. Lee and A. Abidi, "A 9b, 1.25 ps resolution coarse-fine time-to-digital converter in 90 nm CMOS that amplifies a time residue," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 769–777, Apr. 2008.
[8] C. Weltin-Wu, E. Temporiti, D. Baldi, and F. Svelto, "A 3 GHz fractional-N all-digital PLL with precise time-to-digital converter calibration and mismatch correction," in *IEEE ISSCC Dig. Tech. Papers*, 2008, pp. 344–345.
[9] V. Ramakrishnan and P. T. Balsara, "A wide-range, high-resolution, compact, CMOS time to digital converter," in *VLSI Design (VLSID'06)*, Jan. 2006, 6 pp., online.
[10] J.-P. Jansson, A. Mantyniemi, and J. Kostamovaara, "A CMOS time-to-digital converter with better than 10 ps single-shot precision," *IEEE J. Solid-State Circuits*, vol. 41, no. 6, pp. 1286–1296, Jun. 2006.
[11] B. Helal, M. Straayer, and M. H. Perrott, "A low jitter 1.6 GHz multiplying DLL utilizing a scrambling time-to-digital converter and digital correlation," in *VLSI Symp. Dig. Tech. Papers*, Jun. 2007, pp. 166–167.
[12] M. Straayer and M. Perrott, "An efficient high-resolution 11-bit noise-shaping multipath gated ring oscillator," in *VLSI Symp. Dig. Tech. Papers*, Jun. 2008, pp. 82–83.
[13] R. Walker, C. Stout, J. Wu, B. Lai, C. Yen, T. Hornak, and P. Petruno, "A 2-chip 1.5 gigabaud serial link interface," *IEEE J. Solid-State Circuits*, vol. 27, no. 12, pp. 1805–1811, Dec. 1992.
[14] J. Rivoir, "Statistical linearity calibration of time-to-digital converters using a free-running ring oscillator," in *IEEE Asian Test Symp.*, Nov. 2006, pp. 45–50.
[15] I. Nissinen, A. Mantyniemi, and J. Kostamovaara, "A CMOS time-to-digital converter based on a ring oscillator for a laser radar," in *Proc. IEEE ESSCIRC*, Apr. 2003, pp. 469–472.
[16] R. Baird and T. Fiez, "Linearity enhancement of multibit delta-sigma A/D and D/A converters using data weighted averaging," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 42, no. 12, pp. 753–762, Dec. 1995.
[17] S. J. Lee, B. Kim, and K. Lee, "A novel high-speed ring oscillator for multiphase clock generation using negative skewed delay scheme," *IEEE J. Solid-State Circuits*, vol. 32, no. 2, pp. 289–291, Feb. 1997.
[18] S. S. Mohan, W. S. Chan, D. M. Colleran, S. F. Greenwood, J. E. Gamble, and I. G. Kouznetsov, "Differential ring oscillators with multipath delay stages," in *Proc. IEEE CICC*, Sep. 2005, pp. 503–506.
[19] J. G. Maneatis and M. A. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 28, no. 12, pp. 1273–1282, Dec. 1993.
[20] A. Matsumoto, S. Sakiyama, Y. Tokunaga, T. Morie, and S. Dosho, "A design method and developments of a low-power and high-resolution multiphase generation system," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 831–843, Apr. 2008.
[21] M. Straayer, "Noise shaping techniques for analog and time to digital converters using voltage controlled oscillators," Ph.D. dissertation, MIT, Cambridge, MA, 2008.

Matthew Z. Straayer (S'05–M'09) received the B.S. and M.S. degrees in electrical engineering from the University of Michigan, Ann Arbor, in 2000 and 2001, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 2008.

From 2001 to 2003, he was with Integrated Sensing Systems, Ypsilanti, MI, designing wireless readout ASICs for MEMS sensors. From 2003 to 2008, he was on staff at Lincoln Laboratory, Lexington, MA, working on next-generation RF technologies including low-power transceivers and low-noise frequency synthesizers. Since 2008, he has been with Cambridge Analog Technologies, where he focuses on developing high-performance PLL and ADC IP products.

Michael H. Perrott (M'97) received the B.S. degree in electrical engineering from New Mexico State University, Las Cruces, in 1988, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1992 and 1997, respectively.

From 1997 to 1998, he worked with Hewlett-Packard Laboratories, Palo Alto, CA, on high-speed circuit techniques for Sigma-Delta synthesizers. In 1999, he was a visiting Assistant Professor with the Hong Kong University of Science and Technology, and taught a course on the theory and implementation of frequency synthesizers. From 1999 to 2001, he worked with Silicon Laboratories, Austin, TX, and developed circuit and signal processing techniques to achieve high performance clock and data recovery circuits. He was an Assistant and then Associate Professor in electrical engineering and computer science with the Massachusetts Institute of Technology from 2001 to 2008. He is now with SiTime, a Silicon Valley startup developing silicon timing, clock, and RF chips that incorporate Micro Electro Mechanical Systems (MEMS) timing reference devices.