TIME INTERVAL MEASUREMENT

The function of a time interval measurement device is to produce a quantitative measure to the length of a time interval at the input of the device. Typically this interval is presented by short, logic level start-and-stop pulses that have a relation to some physical phenomena whose time behavior is to be studied by the device. Its output is given either by an analog quantity, such as a voltage change in a capacitor or as a digital number, or an output of a counter, for example (Fig. 1). Typically the final output of the device, independent of the method that is being used, is characterized by a number, which gives the length of the time interval as a multiple of the least resolvable time interval of the device. This can be achieved by an analog-to-digital converter (ADC) in case of analog realizations, for example. The least resolvable time interval characterizing the measurement system being used can be considered the value of its least significant bit (LSB) analogically to ADCs. It defines the resolution of the device. Accordingly, a time interval measurement device is often called a time-to-digital converter (TDC).

TDCs have manifold applications. It is one of the critical components of a pulsed time-of-flight laser radar, for example, where the distance measurement accuracy is directly dependent on the time measurement accuracy. Pulsed time-of-flight laser radars can be used in geodesy, space measurements, and various industrial inspection applications (1, 2).

Another application area for TDCs is found in nuclear science, e.g., dE/E and various time-of-flight studies. TDCs are also used in measurement instruments and in the calibration of test equipment for electronic circuits. As in high-speed LSI circuits propagation delays are reduced to 20–200 ps or even below, high speed and accuracy are required in automatic test equipment systems (ATEs). Test timing control with an accuracy to the order of tens of picoseconds is needed, and to maintain such timing performance, the calibration hardware requires even higher performance TDCs with picosecond resolution and accuracy (3–9).

The most straightforward method to realize a TDC is to have an electronic counter counting the pulses of an accurate oscillator during the start–stop time interval. The resolution of this method, as shown later, is defined by the period of the oscillator, which means that a resolution better than in the range of nanoseconds is difficult to achieve because of difficulties in the realization of high-frequency, high-performance oscillators. The purpose of this article is to describe this time interval measurement method with its variations where its fundamental resolution limitation is overcome with some kind of interpolation and where accordingly picosecond-range time interval measurement resolution can be achieved. Some of these interpolator structures, such as time-to-amplitude converters of digital delay lines, can also be used as stand-alone time interval measurement units.

In the following chapters, first some important performance parameters of a time interval measurement device are given. Then the basic counting method is analyzed in more detail. After this, interpolation methods to increase the precision and accuracy of the counting method are described. Next the performance of one interpolation technique, the Nutt method, is analyzed in more detail. This method is selected as it is perhaps the most popular technique in high-performance time-to-digital conversion with many practical realizations. As part of the presentation, some practical realization aspects are also discussed.

PERFORMANCE PARAMETERS IN TIME INTERVAL MEASUREMENT

The main performance parameters used in connection with the time interval measurement method are resolution, precision, and accuracy.

The term "resolution" is used for the smallest time interval that can theoretically be resolved by the TDC in a single measurement, i.e., the quantization step (LSB). The term (single-shot) "precision" is used for the standard deviation (σ) of the measurement result distribution around the mean value (m) when a single time interval is measured repeatedly (Fig. 2). In a practical measurement, the single-shot precision is influenced, besides the quantization error, by nonidealities like jitter in timing signals and power supply noise and especially by the nonlinearities of the possible clock period interpolators. The single-shot precision can be used to estimate the smallest real time interval that can be resolved in a single isolated measurement. As precision is limited by quantization error and statistical error sources, it can usually be improved by averaging (10).

Although precision gives the statistical variation of the measurement result around a mean value, single-shot accuracy is affected by both this statistical variation and any systematic errors in the mean value. The statistical variation (sigma) can be reduced by averaging, but systematic errors cannot; i.e., the precision of the average may be good, but the accuracy may still be poor. Systematic errors include, for example, linearity and stability errors. Integral linearity error (INL) is the deviation of the input–output characteristics from the ideal, straight line input–output characteristics (Fig. 2). Differential linearity error (DNL) is the deviation of each quantization step from the ideal value of the LSB. Note that DNL and INL are related so that the INL of a particular TDC channel is just the sum of the DNLs of all previous channels. The stability of a TDC is defined as the sensitivity of its characteristics with temperature, supply voltage, time, and so on. Clearly, to measure the systematic errors reliably, the statistical variation should be reduced to a negligible level by averaging.

Other parameters that should be considered while choosing time measurement techniques are, for example, the measurement range of a TDC and the conversion time. The range of the device defines the maximum time interval that can be measured or digitized. Conversion time is the time between the end mark of the input time interval and the moment when the measurement result is ready. Sometimes, especially with analog TDCs, a parameter called dead time that typically defines the period during which the system is incapable of accepting a new start

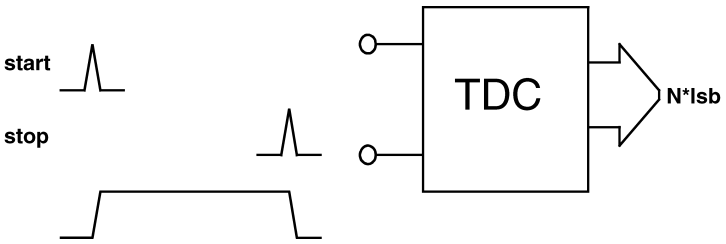

Figure 1. Input–output interfaces of a TDC. LSB is the least significant bit, for example, 10 ps to 10 ns. $N \times$ LSB is the length of the start–stop interval.

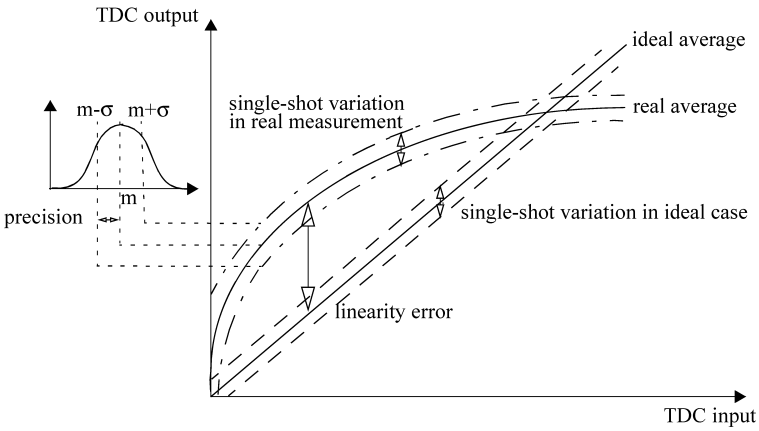

Figure 2. Performance parameters of a time-to-digital converter.

after a registered start signal (conversion time + possible recovery time) is also used. Dead time might be an important parameter, especially in applications where the time intervals to be measured are randomly distributed, which often happens in nuclear measurements. For example, if the start detector efficiency is high and the stop efficiency is low, respectively, a long dead time may result in reduced measurement efficiency because many start signals that trigger the time interval measurement may not be followed by a valid stop signal. In this case, it is important for the TDC to recover from the useless start pulse as quickly as possible.

Counting Method

Analysis of the Method. The counting method where clock pulses are counted during the input time interval is perhaps the most simple time measurement method (Fig. 3). Provided that the reference clock is accurate, a crystal oscillator, for example, the counter method has a wide linear range and good stability. In asynchronous measurement, i.e., when the measurement begins in a random phase with respect to the clock, the maximum error in one measurement is $\pm T_{clk}$, where T_{clk} is the clock period. For each input interval, the counter will measure either N or $N + 1$ counts, and for this binomial distribution, it can be shown that for an input interval of $(Q + F) \cdot T_{clk}$, where Q is an integer and $0 \leq F < 1$, the expected value of the measurement result is $Q + F$ (counter reading divided by the total amount of measurement results) and the standard deviation is (11)

$$\sigma = \sqrt{F \cdot (1 - F)} \tag{1}$$

Thus, the measurement precision varies with the input interval, and the worst case value for $F = 0.5$ is $0.5\, T_{clk}$. For a 1-GHz clock, for example, the maximum single-shot error is $\pm$ 1 ns and the worst case value for σ is 500 ps.

To improve precision, averaging can be used in applications where several samples per measurement can be taken. In asynchronous measurement and for N_{av} samples, the precision is (11)

$$\sigma = \frac{\sqrt{F \cdot (1 - F)}}{\sqrt{N_{av}}} \tag{2}$$

which is shown in Fig. 4.

In asynchronous measurement, the phase of the input interval with respect to the clock is continuously distributed over the clock period or even several periods. It means that the signal repetition rate is not coherent with the counter clock, which is an essential requirement as the time relationship between the signal and the counter clock must be such as to sweep through the full range of the $N/N + 1$ count ambiguity in a random manner to satisfy the statistical requirement of averaging.

Synchronous measurement is also possible, which can be realized if the start timing signal is repeated at a constant rate and if the oscillator of the system is locked to this rate. However, if the measurement is completely synchronous with respect to the clock, averaging does not improve precision. On the other hand, if the phase of the synchronous measurement can be controlled, it is possible to achieve a faster precision improvement rate than in asynchronous measurement. For example, if the phase of the start signal and oscillator has M discrete values evenly distributed within the clock period, the measurement

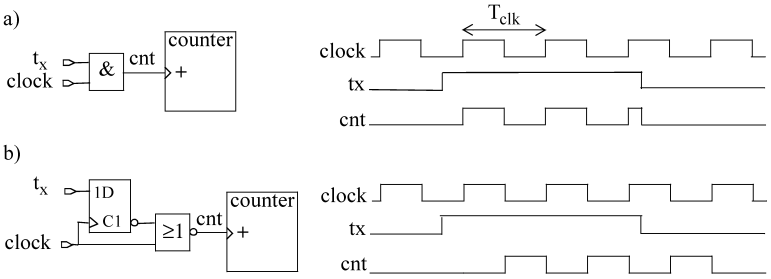

Figure 3. Basic counting method, where clock pulses are enabled to increment the counter during the input time interval. The gating logic of (b) is preferable to the one in (a) because it allows only integral clock pulses to reach the counter. The method in (a) can produce short pulses that may or may not be detected by the counter.

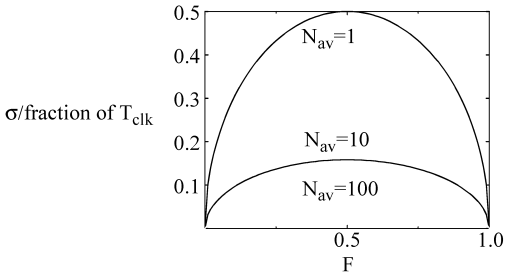

Figure 4. Precision of the counting method as a function of the fractional part F (by which the length of a time interval exceeds an integral number of clock periods) and number of averaged samples N_{av}.

precision can be improved in averaging with a rate proportional to $1/M$ (rather than to $1/\sqrt{M}$ as in synchronous measurement), but the improvement is limited by factor $1/M$. The synchronous method using discrete phase steps is, however, more difficult to realize than the asynchronous method (11).

The basic counting method has good accuracy, because with a stable crystal oscillator reference, the linearity and both the short-term and the long-term stabilities are good. The measurement range can be increased simply by increasing the number of bits in the counter. With standard commercially available high-performance crystal oscillators, the maximum frequency is limited to a few hundred megahertz, which corresponds to a single-shot resolution of several nanoseconds. With advanced clock generation techniques, single-shot resolutions of about 500 ps to 1 ns have been achieved. Here the idea typically is to multiply the clock frequency of a reference source by a phase-locked loop, which includes a divider and a VCO in its loop. In integrated time-to-digital converter realizations, a VCO based on a ring oscillator is especially attractive, because the measurement resolution can be further improved by using all the clock phases of the ring oscillator (12).

Input Synchronization. In the counting method, the gating of the digitization clock may have a great impact on the measurement accuracy. If the input interval (asynchronous with respect to clock) is directly used for gating the counter, short clock pulses appear frequently at the input of the counter (Fig 3a). These may or may not be long enough to increment the counter and cause an unpredictable error in the averaged result. Therefore, a synchronizer, where the input interval is synchronized to the clock is needed. An example is shown in Fig. 3b, where only full-size clock pulses are allowed to the counter. Note that because of the use of edge-triggered flip-flop in the synchronization circuit, the clock becomes effectively a train of "zero-width" pulses (impulses) (13). The advantages following are thus that the measurement result will be unbiased, and that time intervals much less than the clock period can also be reliably measured.

Increasing Precision and Accuracy

The precision of the counting method can be improved by averaging but at the cost of measurement time. In some applications, however, averaging is not possible because of the single-shot nature of the measured phenomena. Typically this is the case in the experiments made in nuclear physics. The aim of this article is to discuss the variations of the basic counting method, which enable one to achieve an improved single-shot precision with realistic clock rates. These methods typically use accurate clock generators in connection with other time interval measurement techniques, which enable a resolution markedly better than the clock period to be achieved. In a way these methods digitize the basic clock period so that the resolution of this digitization (interpolation) defines the system resolution.

Vernier Method. In the vernier method (14), two startable oscillators with slightly different frequencies are used to achieve an LSB equal to the difference of the oscillator periods (dt in Fig. 5). The start mark of the input interval enables the oscillator with a lower frequency $f_1 = 1/T_{clk}$, and the stop mark enables the oscillator with a higher frequency $f_2 = 1/(T_{clk} - dt)$. A counter records the pulses from the oscillators until, since $f_1 < f_2$, at some point the two oscillators will be in phase. Then, from Fig. 5,

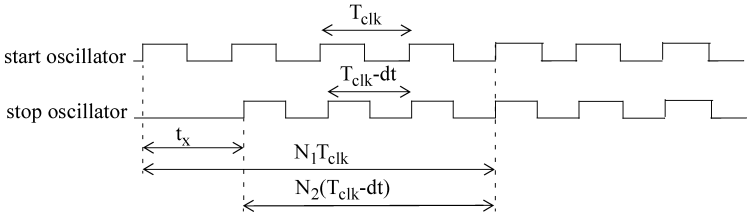

Figure 5. Principle of the vernier method in time measurement. Two oscillators with different periods are triggered by the timing pulses. The counting is terminated at coincidence. The time interval t_x is effectively digitized with a resolution defined by the period difference.

$t_x = (N_1 - N_2) \cdot T_{clk} + N_2 \cdot dt$. If $t_x < T_{clk}$, $N_1 = N_2 = N$ and $t_x = N \cdot dt$. The maximum value of N_2 is T_{clk}/dt, so the maximum conversion time (time from the stop mark to the coincidence) is $(T_{clk}/dt)(T_{clk} - dt)$. To measure input intervals longer than T_{clk}, two counters recording both N_1 and N_2 are needed; see Ref. 14.

Functionally the vernier method can also be considered a time interval multiplier where the multiplication factor is T_{clk}/dt. When the multiplied time interval is digitized with a clock having a period T_{clk}, the effective resolution is equal to dt. The accuracy and the resolution of this method can be high if f_1 and f_2 are stable and dt is made small (typically 1% of the period).

In the dual vernier method (Fig. 6), the startable start-and-stop oscillators have an equal frequency $f_0 = 1/(T_{clk} + dt)$ and their phase crossover points with a third, continuously running oscillator with a frequency $f_{clk} = 1/T_{clk}$ are detected (15). Similarly to the basic vernier method of Fig. 5, the LSB of the measurement is dt and the maximum conversion time is $(T_{clk}/dt)(T_{clk} + dt)$. Using the three counter values N_c, N_1, and N_2, the input interval can be calculated from

$$t_x = N_c \cdot T_{clk} + N_1 \cdot (T_{clk} + dt) - N_2 \cdot (T_{clk} + dt) \qquad (3)$$

Good single-shot resolution can be achieved with the vernier techniques. In Ref. 15, the LSB is 20 ps with a clock frequency of 200 MHz, and in Ref. 16, the measured standard deviation is 2.3 ps in a 2-ns range with a clock frequency of 500 MHz. The time measurement circuitry in Ref. 15 is a discrete implementation for a commercial counter, whereas the TDC in Ref. 16 is an integrated implementation, but the start-and-stop clock sources are external. This circuit is used for tester timing calibration, where the start input is the reference clock and the stop input is a tester output.

Nutt Method

Analysis of the method. A powerful method for measuring time intervals is to combine an accurate digital clock and an analog or digital interpolation circuit as shown in Fig. 7 (17). The input time interval is roughly digitized by counting the reference clock periods during this interval. The counter is enabled at the first clock pulse following the start mark, and it is disabled at the first clock pulse following the stop mark. The resulting time interval T_{12} is synchronized to the clock and is, therefore, accurately measured. The fractions T_1 and T_2 are digitized separately with interpolators to improve single-shot resolution. For an n-bit interpolator, the LSB of the measurement is equal to $T_{clk}/2^n$ and the input interval t_x can be calculated from the following equation:

$$t_x = T12 + T1 - T2 = N_c \cdot T_{clk} + N_1 \cdot T_{clk}/2^n - N_2 \cdot T_{clk}/2^n \qquad (4)$$

Note that if the system clock is asynchronous with respect to the time intervals to be measured, the length of time fractions T_1 and T_2 change randomly in a repeated measurement although their difference has only two discrete values (dt if the main counter result is $N_c - 1$ or $dt - T_{clk}$ if the main counter result is N_c). This means that in the averaged results the nonlinearities of the interpolators are also in a way averaged so that the accuracy of the system is not limited by them. Note also that the drifts of the interpolators tend to cancel as it is the difference $T_1 - T_2$ that counts in the final result (18).

The interpolators are generally based on analog time-to-voltage conversion or on digital delay lines. These structures can achieve good single-shot resolution in a limited dynamic range. Practically, they can be used for measuring time intervals from tens to some hundreds of nanoseconds with a resolution of 10 ps–1 ns depending on the measurement range. The Nutt method can thus be considered a technique that combines the inherently good single-shot resolution of an analog time interval measurement method such as time-to-amplitude conversion or of a digital delay line, for example, with the accuracy and wide linear range of the counting method.

Synchronization. For the Nutt method, it is typical that the timing signals and the clock are deliberately asynchronous, which produces a synchronization problem in the generation of time intervals T_1, T_{12}, and T_2. In Figure 8a, a simplified scheme of one possible control block of a TDC based on the Nutt method and analog interpolation is shown. The end mark of T_1 (T_2) is taken from the output of the flip-flop D2a (D2b). However, in asynchronous measurement, the start (stop) pulse arrives at a random phase with respect to the reference clock. When the setup time requirement of the flip-flop D2a or D2b is not fulfilled, the propagation delay of that flip-flop will increase (Fig. 8b), and in the extreme case, the flip-flop will enter a metastable state. Thus, if the start (stop) pulse occurs near the rising clock edge, an erroneous measurement is possible. Unlike in the basic counting method, even a small excess delay affects the measurement result directly through T_1 and T_2.

The probability of a synchronization error can be reduced by waiting before sampling the output of flip-flop D2a (D2b). In the scheme of Fig. 9, the measurement will be accurate if the flip-flop D2a (D2b) settles in less time than T_{clk}. This comes from the fact that even though the

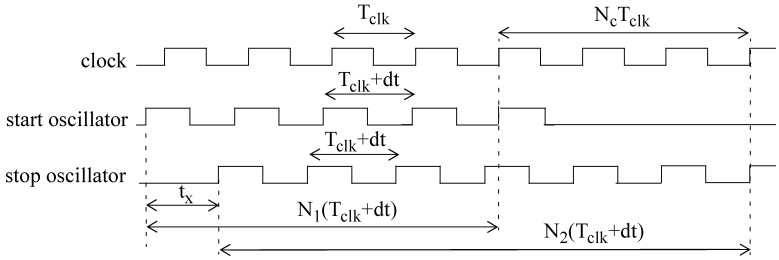

Figure 6. Principle of the dual vernier method in time measurement.

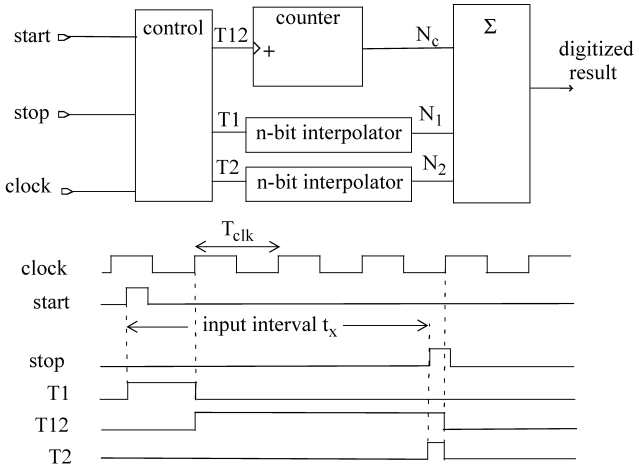

Figure 7. Block diagram and operating principle of a TDC based on the Nutt method. The start–stop time interval is digitized coarsely by the counter. In addition, time fractions from start and stop to the next following clock pulses (or next but one), respectively, are digitized by interpolators.

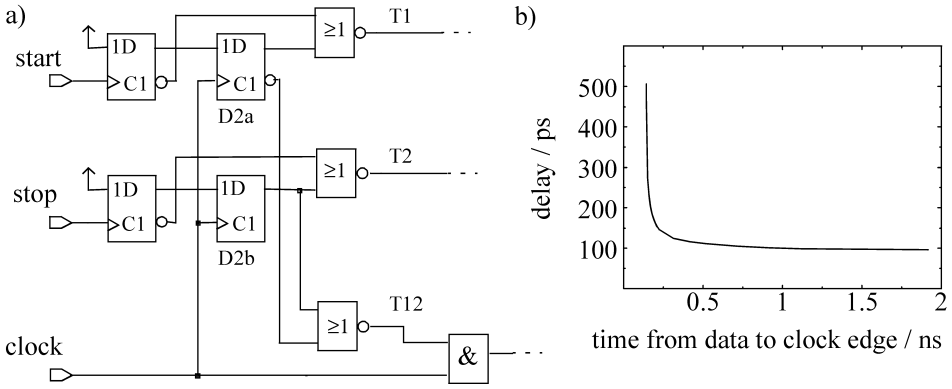

Figure 8. (a) Synchronization of start and stop inputs in a TDC based on the Nutt method. Problems may occur when the clock edge and output of first flip-flops occur within the set-up time of D2. (b) Propagation delay of a CML-flip-flop (current mode logic) as a function of the data-clock edge time interval.

delay of the flip-flop D2a (D2b) might increase because of coincident clock and data edges, this has no effect to the length of T_1 (T_2) because the back edge of $T_1(T_2)$ is defined by the next following clock edge. Note also that even in the case if the timing pulse would be completely missed with a particular clock edge, the measurement is correct because this simply means that $T_1(T_2)$ would be longer by one T_{clk}, but on the other hand, T_{12} would now be correspondingly shorter by one T_{clk} (18).

The scheme of Fig. 9 increases the measurement range required of the interpolators, which is a disadvantage in digital interpolators, where the integral nonlinearity caused by delay element mismatch increases. In time-to-voltage conversion, a short time offset is an advantage, because the nonlinear part of the characteristics from switching effects is then not used for measurement. It should also be understood that the detailed construction of the synchronization mechanism and circuitry used depends

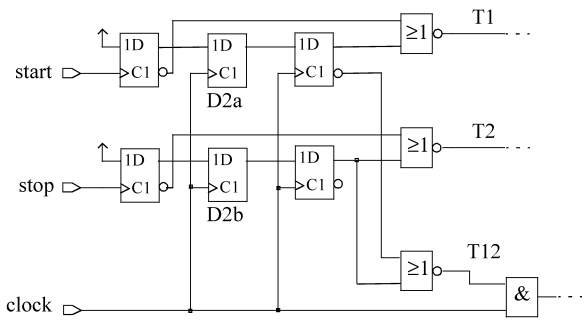

Figure 9. Improved synchronization of start and stop inputs in a TDC based on the Nutt method. Here the setup time violation in D2 does not produce error in T_1 or T_2.

largely on the particular interpolation method and architecture. The above example is shown just to demonstrate the typical design challenges met is designing interpolating, especially analog or double-level digital, time-to-digital converters.

Analog Interpolators. A straightforward method to digitize the start and stop time fractions T_1 and T_1 is to apply time-to-amplitude conversion (TAC) followed by an A/D converter. In time-to-amplitude conversion, a capacitor is discharged with a constant current during the input time interval. A schematic diagram of a time-to-amplitude converter is shown in Fig. 10. It consists of an accurate current generator, a current switch based typically on a BJT or MOS differential pair, a reference voltage source, and the conversion capacitor. In the operation cycle, the capacitor floating in a reference voltage is discharged by a rate defined by the current of the current source during the time interval to be measured. The change in the capacitor voltage is thus proportional to the input time and can be converted into digital form with an A/D converter. After the conversion, the capacitance is again charged to the reference voltage and the cycle can be repeated. Time-to-digital converters based on this technique are presented in Refs. 18 and 19.

A fairly common choice for the analog interpolator is the dual-slope converter (4,20–22). In the basic dual-slope converter, a capacitor is first discharged with a constant current (I) during the input time interval (t_x) and then charged back with a smaller current (I/N). Clock pulses are counted during the charging time ($N \cdot t_x$), which gives a measurement resolution equal to the case in which t_x is directly digitized with a clock frequency N times higher. (Fig. 11).

Increasing the stretch factor N not only improves precision but also increases conversion time. To shorten the conversion time of the TDC, dual interpolation or multiple interpolation methods can be used (23). In multiple interpolation, a stretching operation with a stretch factor of K_s is repeated N times to achieve an effective stretch factor of ${K_s}^N$.

Interpolators based on an analog interpolation technique easily achieve sub-nanosecond single-shot resolution. With a clock frequency of 100 MHz and a 1-V dynamic range in the time-to-voltage conversion, a time resolution of 10 ps corresponds to a voltage resolution of 1 mV. However, lowering of the supply voltage has a direct impact on the operation of these interpolators because it limits the linear dynamic range and thus the single-shot resolution. The power consumption and conversion time of the analog interpolator depend on the chosen A/D conversion architecture.

The above analog interpolator devices can, of course, be used as stand-alone time interval measurement units. The TAC method is especially useful if good single-shot resolution is needed and the measurement range is modest (less than about 100 ns). Typically they give an excellent single-shot resolution in the picosecond range but suffer from limited linearity (INL typically 0.1% of the range). In connection with the Nutt method, the linearity problem can, however, be avoided in averaging measurement as explained in detail later.

Digital Interpolators. An example of a digital delay line is shown in Fig. 12. In a delay line, the time measurement unit is the propagation delay of a logic gate, usually of an inverter. In Fig. 12, the start mark of the input interval travels along the delay line. When the stop mark (clock edge) arrives, it stores the status of the delay line into the flip-flops. From this data, the time interval between the start and stop can be coded as a multiple of one gate delay. The delay of the element must be controllable to compensate for the effects of process variations and temperature and supply changes. Alternatively, digital calibration methods can also be used. As shown in Fig. 13, the control parameter can be the current or the number of the load elements, for example. The control voltages (pbias, nibias, bias) are commonly created in analog PLL or DLL control loops, but also digital control has been implemented.

However, if the propagation delay of a logic gate is directly used as the measurement unit, the achievable resolution is limited by the minimum gate delay that, in turn, depends on the technology and operating temperature range. This limitation can be avoided by using the difference of two gate delays as the quantization step. The delay difference can be implemented with two delay lines having slightly different gate delays (24–26). A more compact solution is the pulse-shrinking delay line presented in Ref. 27. In a pulse-shrinking delay line (Fig. 14), the time resolution depends on the difference between two de-

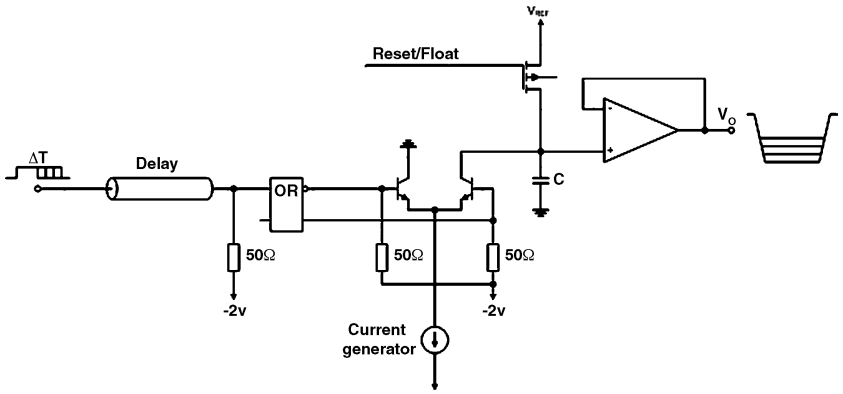

Figure 10. A schematic diagram of a time-to-amplitude converter, which gives out a pulse whose amplitude is linearly proportional to the length of the input pulse.

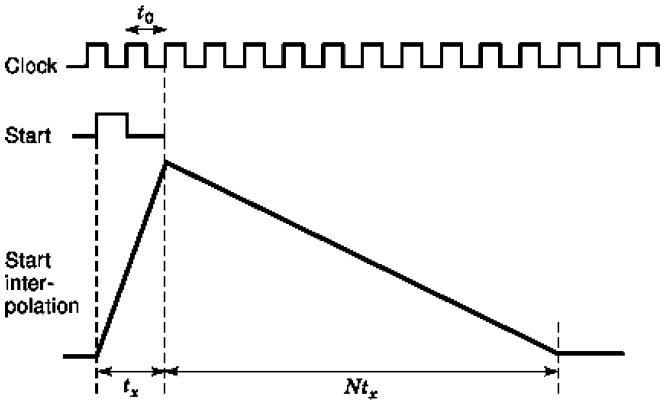

Figure 11. Operation principle of dual-slope converter, which discharges a capacitor with a constant current I and then charges back with a smaller current I/N.

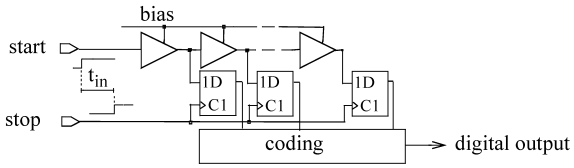

Figure 12. Delay line as a time measurement unit. The output of the flip-flops is latched by the stop signal. First flip-flop not set indicates the length of the start–stop interval.

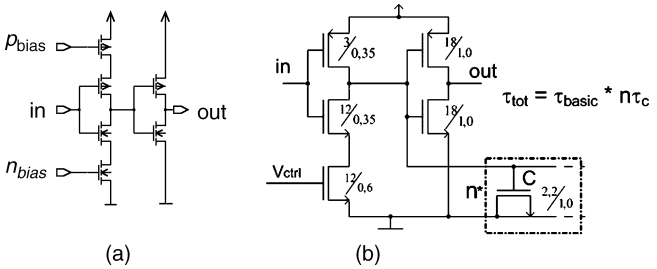

Figure 13. Controllable delay elements used in delay lines. In (a), the delay is controlled by adjusting the current of the inverter with the bias voltages nbias and pbias. In (b), the delay is controlled by varying the number of load capacitances seen by the inverter.

lays of a single element. The input pulse t_{in} propagates in a chain of delay elements. The pulse shrinks by a constant amount in each element until it disappears entirely. The rs flip-flops are set by the propagating pulse until the pulse vanishes, after which the following flip-flops are left reset. The address of the first flip-flop not set is coded to the output. The propagation of the rising edge of the input pulse is slowed down by the current starving transis-

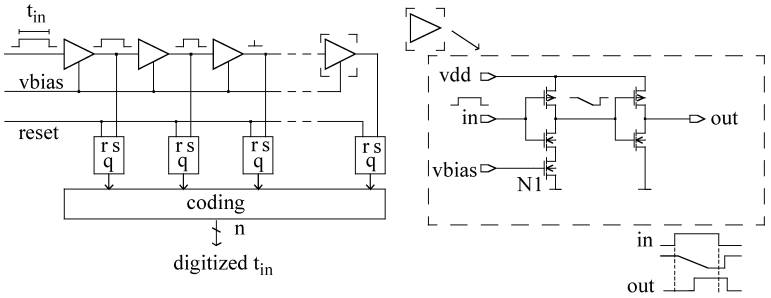

Figure 14. Pulse-shrinking delay-line and schematic of the delay element.

tor N_1, whereas the falling edge travels fast. Thus, in each element, the pulse shrinks with an amount equal to the delay difference. The amount of shrinking, i.e., the LSB of the delay line is controlled by the bias voltage v_{bias} at the gate of the current starving transistor N_1. V_{bias} is generated in the delay-locked loop, which stabilizes the length of the delay line to be equal to the reference clock period. Basically, the bias voltage is adjusted so that a pulse with a width of one clock period, which is equal to the maximum input time interval of the interpolator, just disappears in the last delay element. Note also that the delay line method where the difference of the two absolute delays defines the resolution is equivalent to the vernier method where resolution was defined by the difference of the oscillator periods.

Delay line interpolators based on inverter gates are easily implemented in a CMOS process, have low-power consumption, and can relatively easily be transferred to a low supply voltage environment. Also, the conversion time of a TDC based on delay line interpolators is short, because the conversion time depends on the propagation delay of the delay element chain and on the coding logic. The main factor limiting the performance is the nonlinearity caused by random mismatch of the individual delay elements, which determines the achievable precision. Several realization variants of these kind of time-to-digital converters are presented in Refs. 28–38, for example.

At best published integrated CMOS time-to-digital converters based on multilevel interpolation and reference recycling give a single-shot precision of about 10 ps with intervals of less than 200 μs and an external reference of 5 MHz (38). This performance is achieved in a 0.35-μm CMOS technology by using a double-level interpolation structure as shown in Fig. 15. The first interpolation stage consists of a 16-element delay line that effectively multiplies the reference clock by 16 so that with a reference clock of 65 MHz, for example, the internal clock frequency is 1 GHz. Actually the frequency of the external reference clock can be N times lower because the clock edge is circulated N times within the delay line before accepting a new "jitter-free" clock edge into the line from the external clock. In this particular realization, the external reference frequency is 5 MHz and N is 32. The time interval between the edges of the virtual clock is interpolated with four parallel delay elements so that the delay difference between the interpolated edges is about 100 ps. Fine interpolation is achieved by hooking to the timing pulse (start, stop) additional timing edges with a delay difference of about 10 ps. These edges are produced by parallel delay lines. The position of the timing pulse with respect to the interpolated edges of the virtual clock is determined by the coincidence of these edges resulting in a single-shot resolution of 10 ps. The precision of the device is shown in Fig. 16 as a function of F ($\Delta T = NT_{clk} + F$). As shown by the figure, a precision of about 8 ps can be achieved by using a look-up table to correct the nonlinearities of the interpolators.

Characteristics of the Interpolation Method

Single-Shot Precision. Similarly to a counter, in asynchronous measurement (i.e., the start pulse arrives at a random phase with respect to the reference clock), the worst-case single-shot precision of an ideal TDC based on the Nutt method is, according to Eq. (1), 0.5 LSB and the precision improves with averaging according to Eq. (2). The maximum single-shot quantization error is ± 1 LSB and the measurement result distribution is binomial (N or $N+1$ counts). In practice, however, several error sources from the timing logic and interpolators deteriorate precision. In asynchronous measurement, the precision is often limited by gain error or, more generally, by the nonlinearity of the interpolators.

In Ref. 22, the effect of interpolator nonlinearity on the accuracy and precision of the TDC is analyzed. Assuming a general form $e(x)$ for the interpolator error (nonlinearity and/or gain error), using subscripts 1 and 2 for the start and stop interpolators, respectively, and by normalizing T_{clk} to 1, the TDC measurement error for input time interval $Q + F$ (Q is an integer and $0 \le F < 1$) in a single measurement can be calculated as

$$\begin{aligned} e(x) &= e_1(x) - e_2(x + F), \quad x + F < 1 \\ e(x) &= e_1(x) - e_2(x + F - 1), \quad x + F \ge 1 \end{aligned} \tag{5}$$

The maximum single-shot measurement error is now equal to the maximum difference $e_1(x) - e_2(x)$, and the measurement result distribution is no longer binomial, but more than two results are possible. Since in asynchronous measurement x varies randomly and with equal probability between $0 \le x < 1$, the mean value of the TDC measurement

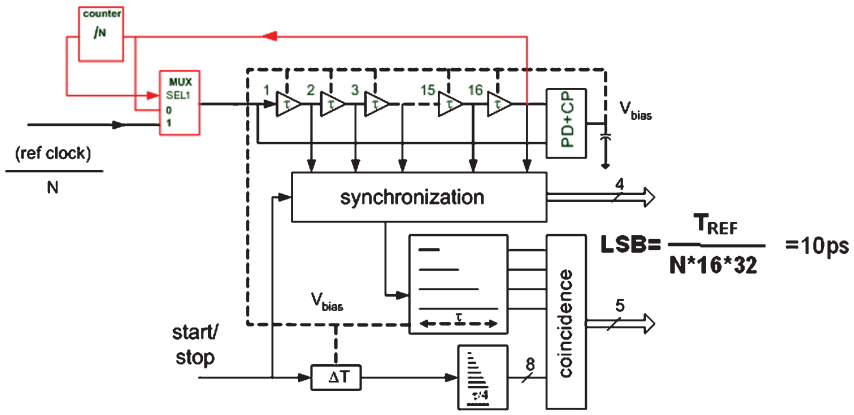

Figure 15. Architecture of a time-to-digital converter using reference recycling and dual-level interpolation with delay lines.

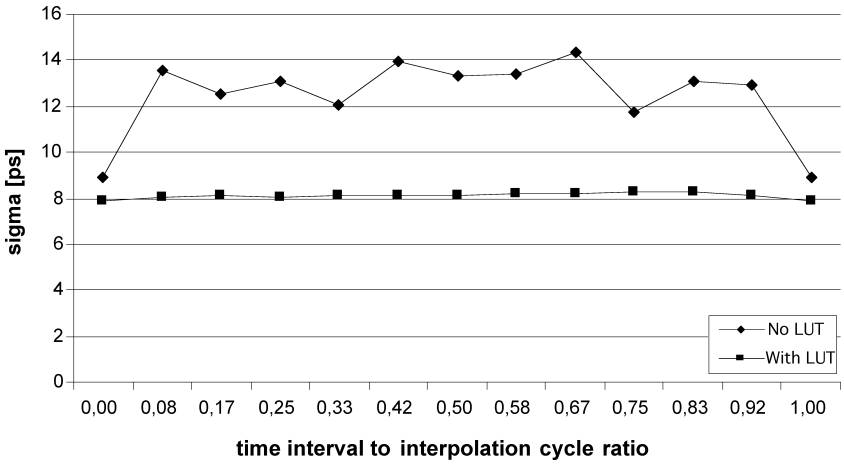

Figure 16. Precision of the TDC of Ref. 38 as a function of the fractional part of the time interval to be measured.

Table 1. Manufacturers of commercially available frequency/time-interval counters

Manufacturer	Homepage
Agilent Technologies	http://www.home.agilent.com
BK precision	http://www.bkprecision.com
Blue Sky Electronics	http://www.blueskyelectronics.com
CAEN S.p.A. (Costruzioni Apparecchiature Elettroniche Nucleari)	http://www.caentechnologies.com
EADS North America Defense Test and Services, Inc. (Earlier Racal Instruments)	http://www.racalinstruments.com
FAST ComTec GmbH	http://www.fastcomtec.com
Fluke Corporation	http://us.fluke.com
Hameg Intstruments	http://www.hameg.com
Highland Technology Inc	http://www.highlandtechnology.com
Ortec	http://www.ortec-online.com
Pendulum Instruments	http://www.pendulum-instruments.com
Standford Research Systems	http://www.thinksrs.com

error is

$$m = \int_0^1 e(x)\cdot dx = \int_0^1 e_1(x)\cdot dx - \int_0^{1-F} e_2(x+F)\cdot dx - \int_{1-F}^1 e_2(x+F-1)\cdot dx \qquad (6)$$
$$= \int_0^1 e_1(x)\cdot dx - \int_0^1 e_2(x)\cdot dx = m_1 - m_2$$

where m_1 and m_2 are the mean values of the start-and-stop interpolator errors. Thus, the gain error or the INL of interpolators reduces to a constant bias error independent of the time to be measured. This result is extremely important as it indicates that the error caused by the interpolator nonlinearities is effectively averaged out in a repeated measurement from the randomization process present in the asynchronous measurement. This source provides the excellent linearity of the method.

The nonlinearities of the interpolators do have an effect on the precision of the system, however. The variance (σ^2) of the TDC measurement result is

$$\begin{aligned}\sigma^2(x) &= \int_0^1 (e(x) - m)^2 \cdot dx \\ &= e_{1rms}^2 + e_{2rms}^2 - (m_1 - m_1)^2 \\ &\quad -2\int_0^{1-F} e_1(x) \cdot e_2(x + F) \cdot dx \\ &\quad -2\int_{1-F}^1 e_1(x) \cdot e_2(x + F - 1) \cdot dx\end{aligned} \tag{7}$$

The value of the two last terms in the above equation depends of F, the fractional part of the time interval to be measured, and thus, the standard deviation of the measurement result is a function of the fractional part (F) of the input time interval ($Q + F$). Since $0 \leq F < T_1$, the measurement precision caused by interpolator nonlinearity or gain error has now a period of T_{clk} instead of LSB.

Nonlinearity of the interpolators can only be compensated for by measuring the linearity error and using these data for correction of the measurement result. However, since this is technically more complicated than in the case of gain error, minimization of interpolator nonlinearity is important.

Also jitter in the input timing signals, clock, and timing logic deteriorates precision with respect to the theoretical value. Their effect is usually small compared with the nonlinearity of the interpolators. However, power supply noise or other types of noise coupling can significantly increase jitter.

Linearity. In asynchronous averaging, the linearity of a TDC based on the Nutt method is basically as good as the linearity of the counting method. Gain error or nonlinearity of the interpolators increases the single-shot error from the theoretical quantization error, but for averaged results, their effect is a constant bias error offset independent of the time to be measured (as shown above). Thus, interpolator gain error and nonlinearity have no effect on the linearity of the TDC, provided that enough samples are being averaged. If the interpolators were exactly identical, their errors would cancel completely (zero offset). In practice, the magnitude of the offset depends on the matching of the interpolators. If the nonlinearity is caused by systematic errors, the interpolators usually have very similar nonlinearities and only a small offset error can be expected to remain. However, if the nonlinearity is caused by random mismatch (as in well-designed delay lines, for example), similar cancellation cannot be assumed.

A prerequisite for the randomization of the systematic errors in averaging measurement is that the system clock and the time intervals are asynchronous; i.e., the length of the time fractions T_1 and T_2 vary randomly in the operating range of the interpolator in a repeated measurement (although their difference has only two values).

Stability. The stability error of a TDC based on the interpolation method can be divided into two components. The first component is an offset-type error that is independent of the input time interval and arises in the control logic (differences in delays) and interpolators (mismatch). If the start-and-stop signal paths in the control block and in the interpolators were identical, including the loading and layout, then according to Eqs. (4) and (6), the measurement results would have zero offset. However, random mismatch cannot be avoided, and when this offset changes with temperature, for example, stability error is created. As a result to achieving highest possible stability, it is important to symmetrize the timing paths. The effect of the mismatch in the interpolators can be minimized by using interpolators alternately in the start and stop channels; see Ref. 29.

The second error component arises from the temperature or time dependency of the reference oscillator. This gain error is dependent on the input interval. For example, a stability specification of ± 25 ppm inclusive of temperature dependency ($0 - +50^\circ$C) and aging corresponds to an error of $< \pm$ 25 ps in a measurement range of 1 μs.

As a summary, it can be the concluded that the nonidealities of the interpolators to large extent cancel out in averaged results. The difference between the nonidealities of the interpolators remains, but in many cases (systematic error source), it can be expected to be smaller than the nonideality itself. Furthermore, this error is constant for all input intervals so it can be corrected simply by subtracting it from the measurement results provided that it does not change with temperature, supply voltage, and so on. Single-shot accuracy, on the other hand, is affected by all nonidealities.

Commercial Time-Interval Counters. Table 1 lists certain manufacturers of universal frequency/time-interval counters and time-to-amplitude converters. The list is not exhaustive, but it contains devices having best promised performance.

BIBLIOGRAPHY

1. P. Palojärvi, K. Määttä, and J. Kostamovaara,“ Pulsed time-of-flight laser radar module with mm-level accuracy using full custom receiver and TDC ASICs,” *IEEE Transactions on Instrumentation & Measurement*, **51**(5): 1102–1108, 2002.
2. K. Karadamoglou et al.,“ An 11-bit high-resolution and adjustable-range CMOS time-to-digital converter for space science instruments,” *IEEE Journal of Solid-State Circuits*, **39**(1): 214–222, 2004.
3. Y. Arai and M. Ikeno,“ A time digitizer CMOS gate-array with 250 ps time resolution,” *IEEE Journal of Solid-State Circuits*, **31**(2): 212–220, 1996.
4. E. Gerds, J. Van der Spiegel, R. Van Berg, H. Williams, L. Callewaert, W. Eyckmans, and W. Sansen,“ A CMOS time to digital converter with 2 level analog CAM,” *IEEE Journal of Solid-State Circuits*, **29**(9): 1068–1076, 1994.
5. B. K. Swann et al.,“ A 100-ps time-resolution CMOS time-to-digital converter for positron emission tomography imaging applications,” *IEEE Journal of Solid-State Circuits*, **39**(11): 1839–1852, 2004.
6. P. M. Levine and G. W. Roberts,“ High-resolution flash time-to-digital conversion and calibration for system-on-chip testing,” *Proc. IEE Proc.-Comput. Digit. Tech.*, **152**(3): 415–426, 2005.
7. K. Koch et al.,“ A new TAC-based multichannel front-end electronics for TOF experiments with very high time resolution,” *Proc. IEEE Transaction on Nuclear Science*, **52**(3): 745–747, 2005.
8. C. Herve and K. Torki,“ A 75 ps rms time resolution BiCMOS time to digital converter optimized for high rate imaging detectors,” *Proc. Elsevier Science, Nuclear Instruments and Methods in Physics Research*, **A 481**: 566–574, 2002.
9. T. Xia and J-C. Lo,“ Time-to-voltage converter for on-chip jitter measurement,” *Proc. IEEE Transaction on Instrumentation and Measurement*, **52**(6): 1738–1748, 2003.
10. Institute of Electrical and Electronics Engineers Inc., *IEEE Standard Dictionary of Electrical and Electronics Terms*, 1984.
11. Hewlett-Packard Inc., Time Interval Averaging. Application Note 162-1, USA.
12. M. Loinaz and B. Wooley,“ A BiCMOS time interval digitizer based on fully-differential, current-steering circuits,” *IEEE Journal of Solid-State Circuits*, **29**(6): 707–713, 1994.
13. J. Rabaey, *Digital Integrated Circuits*. Prentice Hall, Englewood Cliffs, NJ, 1996.
14. D. Porat,“ Review of sub-nanosecond time-interval measurements,” *IEEE Transactions on Nuclear Science*, **NS-20**(1): 36–51, 1973.
15. D. Chu, M. Allen, and A. Foster,“ Universal counter resolves picoseconds in time interval measurements,” *Hewlett-Packard Journal*, August: 2–11, 1978.
16. T. Otsuji,“ A picosecond-accuracy, 700-MHz range Si-bipolar time interval counter LSI,” *IEEE Journal of Solid-State Circuits*, **28**(9): 941–947, 1993.
17. R. Nutt,“ Digital time intervalometer,” *The Review of Scientific Instruments*, **39**(9): 1342–1345, 1968.
18. J. Kostamovaara and R. Myllylä,“ Time-to-digital converter with an analog interpolation circuit,” *Rev. Sci. Instrum.*, **57**(11): 2880–2885, 1986.
19. K. Määttä, J. Kostamovaara, and R. Myllylä,“ Time-to-digital converter for fast, accurate laser rangefinding,” *SPIE Proceedings of the International Congress on Optical Science and Engineering*, **1010**: 60–67, 1988.
20. B. Turko,“ A modular 125ps resolution time interval digitizer for 10MHz stop burst rates and 33ms range,” *IEEE Transactions on Nuclear Science*, **26**(1): 737–745, 1979.
21. E. Räisänen-Ruotsalainen, T. Rahkonen, and J. Kostamovaara,“ A low-power time-to-digital converter,” *IEEE Journal of Solid-State Circuits*, **30**: 984–990, 1995.
22. J. Kalisz, M. Pawlowski, and R. Pelka,“ Error analysis and design of the Nutt time interval digitiser with picosecond resolution,” *Journal of Physics E: Scientific Instruments*, **20**: 1330–1341, 1987.
23. J. Kalisz, M. Pawlowski, and R. Pelka,“ A multiple-interpolation method for fast and precise time digitizing,” *IEEE Transactions on Instrumentation and Measurement*, **IM-35**(2): 163–169, 1986.
24. T. Rahkonen, J. Kostamovaara, and S. Säynäjäkangas,“ Time interval measurements using integrated tapped CMOS delay lines,” *Proceedings of the 32nd Midwest Symposium*, **1**: 201–205, 1989.
25. J. Genat,“ High resolution time-to-digital converters,” *Nuclear Instruments and Methods in Physics Research*, **A315**: 411–414, 1992.
26. J. Christiansen,“ An integrated high resolution CMOS timing generator based on an array of delay locked loops,” *IEEE Journal of Solid-State Circuits*, **31**(7): 952–957, 1996.
27. T. Rahkonen and J. Kostamovaara,“ Pulsewidth measurements using an integrated pulse shrinking delay line,” *Proc. 1990 IEEE Int. Symposium on Circuits and Systems*, **1**: 578–581, 1990.
28. T. Rahkonen and J. Kostamovaara,“ The use of stabilized CMOS delay lines for the digitization of short time intervals,” *IEEE Journal of Solid-State Circuits*, **28**: 887–894, 1993.
29. E. Raisanen-Ruotsalainen, T. Rahkonen, and J. Kostamovaara,“ A low-power CMOS time-to-digital converter,” *IEEE Journal of Solid-State Circuits*, **30**(9), 984–990, 1995.
30. A. Mäntyniemi, T. Rahkonen, and J. Kostamovaara,“ An integrated digital CMOS time-to-digital converter with sub-gate-delay resolution,” *Proc. Kluwer Academic Publisher, Analog Integrated Circuits and Signal Processing*, **22**: 61–70, 1999.
31. M. Mota and J. Christiansen,“ A high-resolution time interpolator based on a delay locked loop and an RC delay line,” *IEEE Journal of Solid-State Circuits*, **34**(10): 1360–1366, 1999.
32. P. Dudek, S. Szczepanski, and J. Hatfield,“ A high-resolution CMOS time-to-digital converter utilizing a Vernier delay line,” *IEEE Transactions on Solid-State Circuits*, **35**(2), 2000.
33. R. Szplet, J. Kalisz, and R. Szymanowski,“ Interpolating time counter with 100 ps resolution on a single FPGA device,” *Proc. IEEE Transaction on Instrumentation and Measurement*, **49**(4): 879–883, 2000.
34. P. Chen, S-I. Liu, and J. Wu,“ A CMOS pulse-shrinking delay element for time interval measurement,” *Proc. IEEE Transaction on Circuits and Systems-II: Analog and Digital Signal Processing*, **47**(9): 954–958, 2000.
35. C-C. Chen et al.,“ A precise cyclic CMOS time-to-digital converter with low thermal sensitivity,” *Proc. IEEE Transaction on Nuclear Science*, **52**(4): 834–838, 2005.
36. J. Song, Q. An, and S. Liu,“ A high-resolution time-to-digital converter implemented in field-programmable-gate-arrays,” *Proc. IEEE Transaction on Nuclear Science*, **53**(1): 236–241, 2006.
37. R. B. Staszewski et al.,“ 1.3 V 20 ps time-to-digital converter for frequency synthesis in 90-nm CMOS,” *Proc. IEEE Transac-*

tion on Circuits and Systems-II: Express Briefs, **52**(3): 220–224, 2006.

38. J-P. Jansson, A. Mäntyniemi, and J. Kostamovaara," A CMOS time-to-digital converter with better than 10 ps single-shot precision," *IEEE Journal of Solid-State Circuits*, **41**(6): 1286–1296, 2006.

JUHA KOSTAMOVAARA
KARI MÄÄTTÄ
Department of Electrical and
Information Engineering,
Electronics Laboratory,
University of Oulu,
Linnanmaa, Finland