

A high-speed wide dynamic range time-to-digital converter

M. Lampton and R. Raffanti

Space Sciences Laboratory, University of California, Berkeley, California 94720

(Received 6 June 1994; accepted for publication 14 July 1994)

We describe a time-to-digital converter able to measure intervals as great as 100 ns with a resolution of 4 ps rms. It achieves this large dynamic range by simultaneously sampling four sinusoidal wave forms (sine and cosine waves at 200 and 6.25 MHz) derived from a single quartz oscillator. Twelve-bit analog-to-digital conversion of the 200 MHz waves yields the high time resolution. Eight-bit conversion of the 6.25 MHz samples removes the cycle ambiguity of the 200 MHz data. The digital words are pipelined in a fully parallel data flow architecture. A first-in first-out stage in the pipeline derandomizes the random event arrival times. A subsequent stage in the pipeline uses an arctangent function to convert the sine and cosine pairs into linearized measures of event time. These are subtracted to yield start-stop time interval sizes for individual photoevents. The minimum start-stop interval is 50 ns, set primarily by the cycle time. Because the same processing is employed for the start and stop events, a large class of potential error and drift phenomena are eliminated. The digitizer provides an accurate way to decode the outputs of delay line detectors, offering high event throughput and extremely good long-term timing accuracy. As a side benefit, the pipeline data flow architecture permits simple breadboarding and low-throughput testing of the system stages with the arctangent work implemented in a personal computer. This arrangement is also very convenient for logging diagnostic evaluation data. The same front-end and data flow architecture is directly applicable to very high-speed applications where the event processing is implemented in a digital signal processor.

I. INTRODUCTION

In recent years, position sensitive detectors based on delay lines have supplanted or replaced alternative event read-out methods.¹ In a delay line detector, each detected event produces a start pulse and a stop pulse. The time interval separating these pulses encodes the event position. Present trends in delay line applications involve images in both one and two dimensions that have a large ratio of size to resolution. This dynamic range must be accommodated while retaining full accuracy and good long-term image mapping stability.

To decode time intervals generated by such a detector, a timing circuit is required whose time resolution must not degrade the spatial resolution of the detector and whose maximum time interval capability must be sufficient to accommodate the full image size of the detector. This digitizer must be extremely linear so that the mapping of event position into position code can be reliably determined for each event. Moreover, the conversion slope and offset must remain highly stable over time and over operating conditions so that the detector calibration remains effective over the long term. Finally, in some applications it is important to accommodate high count rates.

Early time-to-digital converters (TDCs) used an analog ramp time-to-amplitude converter²⁻⁷ followed by an analog-to-digital converter (ADC) or pulse-height analyzer. Such a combination is inherently limited in accuracy by its two analog processing stages and by the recycling time of the ramp circuitry; it also needs frequent recalibration against a primary quartz oscillator time base to determine its conversion slope. Alternative approaches are based on fast digital logic and include the direct digital counting of clock pulses in

various vernier and coarse/fine interpolation arrangements.^{8,9} A third technique, by Berry,¹⁰ uses sine and cosine waves for the fine time interpolation in an otherwise basically digital system.

Our approach, the dual arctangent TDC, differs in that it eliminates the high-speed cycle counting and its correction logic. Instead, we generate four spectrally pure, continuous encoding wave forms from a single quartz oscillator at 200 MHz. The oscillator output is filtered and delayed to produce sine and cosine waves at 200 MHz. It is also divided by 32, and filtered and delayed to give continuous sine and cosine waves at 6.25 MHz. Each timing event triggers the immediate digitization of these four wave forms. The digitized data are placed into a first-in first-out (FIFO) buffer pipeline chip that is accessible by a computer. From the four samples we reconstruct the event time downstream using only software.

This pipeline approach offers a number of practical benefits: complete immunity to phase shifts and timing skew in the front end, an absolute minimum of high-speed electronics, and easy flexible software implementation of the pipeline stages for low-speed applications and development diagnostics.

The leading contemporary architecture for high-speed signal work centers on the use of a dedicated digital signal processor (DSP): a microprocessor whose instruction codes and memory layout facilitate rapid repetitive math operations on pipelined data. Our TDC pipeline interfaces conveniently to such an architecture, and indeed the processing steps that convert the ADC samples into time intervals are likely to be implemented using DSPs in forthcoming flight applications.

The dual arctangent time-to-digital converter is capable of achieving a very large dynamic range, resolving a few picoseconds while accommodating time intervals to beyond




FIG. 1. Block diagram of the TDC. The master 200 MHz clock drives all the sinusoids. Each input event triggers the sample/holds and the 8-bit ADCs. The data flow pipeline begins with the FIFO buffer and includes the two arctangent functions and an event pair time code difference block.

100 ns. (By adding a slow cycle clock, the maximum timing interval could be extended arbitrarily.) Unlike ramp converters, the arctangent converter does not have distinct states in which it is accumulating time or is paused. Its sine waves provide a continuous measure of time and need never be reset. The same circuitry is triggered once for the start event and again for the stop event. The time interval is then evaluated by subtracting the two event times of these identical logic transitions, thereby eliminating a major source of measurement error. The arctangent function is inherently ratio-metric, so that shifts in oscillator amplitude, electronic gain, and ADC scale factors do not appear as output terms. Second-order timing errors are produced by dc shifts, but are removable by a linear correction based on diagnostics derived from the native random-event train. The TDC stability is largely determined by the phase stability of a single quartz-crystal oscillator whose drift can be tightly controlled. Finally, the electrical power consumption is low enough (10 W in its present implementation) to permit the use of this TDC on-board spacecraft for photon-counting space astronomy applications with microchannel detectors. The objective of the work described here is to quantify the sources of actual timing errors for the dual arctangent TDC.

II. OVERVIEW

In Fig. 1, we present a block diagram of the TDC. Here, a single quartz-crystal oscillator runs continually. Its output is bandpass filtered and divided into two branches, one of which is delayed by 1.25 ns to give a 90° phase lag. These two signals are the cosine and sine waves of our 200 MHz timing channel. The clock output is also presented to a digital counter that reduces the frequency by exactly a factor of 32. (An alternative implementation would use a 6.25 MHz

oscillator, whose 32nd harmonic would be generated and filtered for use in the 200 MHz channel.) The 6.25 MHz wave form is bandpass filtered to remove harmonics and subharmonics and split into two signals having a 90° phase difference, establishing the low-frequency cosine and sine pair.

We use a pair of fast analog sample/hold (s/h) circuits to acquire the 200 MHz cosine and sine samples, stabilizing them for 30 ns while two 12-bit ADCs perform the conversion. At the same time, a dual 8-bit ADC samples and digitizes the 6.25 MHz wave pair.

Our TDC makes use of a digital data pipeline to convert the raw ADC samples into time intervals. This pipeline contains stages that detect and eliminate the principal errors accompanying the sampled sine waves: dc offset and drift, gain, oscillator amplitude, reference voltage error, and phase drift. This is accomplished without need for artificially injected test pulses, using instead only the native events normally processed by the system. Such an arrangement makes the TDC highly robust against front-end analog drift, providing a performance that closely approaches the quartz oscillator accuracy and fundamental timing jitter of the trigger wave forms and the sampling process.

Our algorithm for conversion of wave phase into time interval is robust against variations in the relationship between the high-frequency wave pair and the low-frequency pair. This feature makes our TDC tolerant of differing or varying delays in the slow and fast filter, sampler, and ADC. Precision delay matching or compensation of the slow and fast channels is not required.

The timing error distribution of our TDC has several contributors. First, each cascaded ECL gate in the START and STOP trigger logic contributes about 1 ps rms timing uncertainty; our circuit has three consecutive gates in each path, giving 1.7 ps rms. Second, the Analog Devices AD9100 s/h that we adopted has a quoted timing jitter of 1 ps rms. Third, quartz oscillators have short-term timing errors normally described by the "floor" or level part of the phase-noise sideband spectrum. Our oscillator has not been spectrum analyzed, but a typical quartz noise floor is -140 dBc/Hz or 0.7 mrad in a 50 MHz clock bandwidth. Atop a 200 MHz carrier this phase noise is 0.56 ps rms. Finally, the 12-bit ADC quantization is 0.4 ps per least significant bit or about 0.2 ps rms. Combining these error terms for START and STOP, we expect (and achieve) a 4 ps rms interval timing uncertainty for the system. Our TDC is well oversampled: Its digitization errors are far smaller than the other error magnitudes. This oversampling makes our TDC performance robust against common ADC defects such as missing codes, odd-even effects, and bin-to-bin variations in conversion slope.

Our TDC implementation does, however, pose significant accuracy requirements on the high-frequency front-end circuit components. The sampled fast cosine and sine wave forms must be free of harmonic distortion for the recovered arctangents to be linear in time. We achieve harmonic distortion figures of a few percent by filtering the oscillator wave form and operating the s/h circuit at a low signal level, far below its slew-rate limit. The remaining distortion is periodic and correctable, as described below.

TABLE I. Major circuit elements.

Item	Description
200 MHz quartz oscillator	Standard crystal E1100-AAC
200 MHz bandpass filter	2×PLP-200 and 2×PHP200
Input timing logic	MECL 10 KH gates
1.25 ns delay	25 cm semirigid coax
200 MHz sample/hold	Analog Devices AD9100
Buffer amplifier	Comlinear CLC401 at ×20
12-bit ADC	SPT 7922 TTL, 30 ns
Divide by 32	Plessey SP8743B and MECL
6.25 MHz filter	PLP-10.7
40 ns delay line	Allen Avionics LC040Z50A
Dual 8-bit ADC	Analog Devices AD9058
FIFO	six 8-bit IDT 7203S20

III. IMPLEMENTATION DETAILS

Owing to the rapid improvement rate of electronics accuracy, power consumption, and speed, details of a given implementation are likely to be obsolete before a technical paper such as this can reach its audience. Nonetheless, we offer a snapshot of our current state of TDC development by presenting descriptions of the elements of the system whose performance we report here. Table I lists these elements and how they were implemented.

One design consideration is the fact that the high-speed s/h chip has appreciable nonlinearity if employed with a full scale (2 vpp) sine wave at 200 MHz. On the other hand, it has a very low noise level, amounting to less than 100 μ V rms. It is thus best used at a low level of sine wave drive; we employ <0.2 vpp. We employ a buffer amplifier with a gain of 20 to boost the wave sample to span the 4 vpp working range of the 12-bit ADC chip. The s/h chip also exhibits about -40 dB of unwanted sine wave feedthrough at 200 MHz, which if applied to the 12-bit ADC might cause problems, even though the ADC has its own internal s/h circuit. We limit the bandwidth of the buffer amplifier to 50 MHz with a fairly sharp cutoff to attenuate rather than boost the 200 MHz hold-mode feedthrough. The buffer amplifier also restricts the noise bandwidth of the sampler to about 50 MHz. The combined noise level of the s/h and the amplifier is less than 1 mV rms (1 least significant bit) at the input to the ADC.

IV. ARCTANGENT TIME COMPUTATION

The TDC technique presented here requires the conversion of the cosine and sine samples to phase arguments at two places: The 200 MHz pair needs to be converted with high resolution, while the 6.25 MHz pair needs only enough resolution to resolve the 200 MHz cycle count. We refer to this conversion step as the arctangent stage. Then, we combine the coarse and fine phases to yield a complete clock time reading for START, and another for STOP. Subsequently, we subtract these two readings to estimate the time interval.

Depending on implementation details these computational steps may be carried out in a general-purpose microprocessor, in a dedicated digital signal processor, or in a dedicated wired-logic system. In this section we describe a

software approach in which a 64-wire parallel direct memory access communications channel inserts the raw ADC information directly from the FIFO into a block of computer memory and sets a flag. The computations are then done using software.

The arctangent function of two arguments is a common library function in high-level computer languages; for example, it is an essential part of the algorithm that converts a Cartesian coordinate pair into a polar coordinate pair. Let the function $\arctan(y,x)$ represent such a function expressed in cycles (the trigonometric arctangent divided by 2π). Let X and Y be the fine high-frequency cosine and sine samples, and let U and V represent the coarse wave samples. Define

$$F = \arctan(Y,X) \quad (1)$$

and

$$C = \arctan(V,U), \quad (2)$$

where F and C are expressed in cycles so that $0 < F < 1$ and $0 < C < 1$. These fine and coarse phases must be combined to create an event-time estimate that is continuous and monotonic. Because the coarse phase is not precisely coordinated with the fine phase, the fine cycle count must be driven by the passage of the fine phase through zero, not simply by the advance of the coarse phase. In a clock analogy, hours start when the minute hand passes 0, not when the hour hand crosses an hour mark.

We adopt the time formula

$$T = F + \text{int}(32 * C - F + K), \quad (3)$$

where T is expressed in cycles of the fast clock. In this formula, the $\text{int}()$ function returns the integer part of its argument; here, it provides the added number of cycles of coarse time. The argument of the $\text{int}()$ function varies only slightly while F increases, since the $32 * C$ balances the $-F$ term on the average. However, the abrupt reduction of F , unaccompanied by any significant change in C , is what increments the $\text{int}()$ function. The constant K centers the transition in the fine range to provide a generous margin against phase drifts between the slow and fast timing channels. K can be determined in software. K is approximately equal to 0.5 in our circuit.

To obtain the time interval separating the START and STOP events, the start time must be subtracted from the stop time using modulo arithmetic. If the difference is negative, a full coarse period ($T=32$, which is 160 ns in our system) has to be added to the result.

V. PIPELINE DATA DIAGNOSTICS AND MANAGEMENT

Because the event data are processed sequentially by software, it is convenient to implement data diagnostics at various points in the processing flow, both before and after the arctangent conversion of raw ADC data into finished event interval timings. It has also proven instructive to accumulate diagnostic histograms of raw and processed data values generated by random and periodic event trains. Some of these results are presented in the Sec. VI below.

A key requirement of the TDC described here is that the arctangent function used to decode the fine cosine and sine

samples must remain centered on the zero voltage points of the sinusoids. Drift in the amplitudes of the sinusoids, or in the ADC reference voltages, is cancelled by the arctangent function, but drift in the zero points of the s/h or ADC units will directly impact the precision with which the derived time codes are constructed.

To eliminate zero-point drift, our pipeline includes a preliminary stage that monitors and corrects for any shift in the range of ADC codes. Because the codes are triggered by events that are unsynchronized with the TDC clock, the event codes are essentially random. Averaging the X codes to ascertain a zero point accurate to 1 part in 4000 would require millions of events. Instead, we use software peak detectors that sense the positive and negative extreme values of the X and Y that are found in the pipeline. The method makes use of the fact that random samples of a noiseless clock wave form give a distribution with a low central density but a very high density at the extremes: A useful fraction of the events lies in the extreme bins. For example, if a noiseless encoding circle has a diameter of 4000 bins, about 1% of all events lie near bin 0 and another 1% lie near bin 3999. Consequently, in a sequence of 500 random event timings (250 start-stop pairs), five extreme-bin events of each kind are expected. The probability of having no samples representing the upper extreme bin in such a sequence is $\exp(-5)$: less than 1%. The effect of a few bins of radial noise in the encoding oval is a slight increase in the jitter of the amplitude determination.

We use a software peak detector to track the actual working values of the extreme bins in the pipeline data flow and to linearly remap the raw data to enforce a uniform data range and zero point before the arctangent function is called. Such a procedure uses a register for each of the positive and negative peak estimates. The procedure compares each sample with the current peak estimate and modifies the current estimate depending on the result of the comparison. Two peak detector routines have been employed:

```

procedure digitalpeak(integer sample,
peak);
  if (sample>=peak)
    then begin
      peak:=sample;
      count:=0;
    end
  else begin
    count:=(count+1) mod 1000;
    if (count=0)
      then peak:=peak-1;
    end;
  procedure analogpeak(integer sample,
real peak);
    if (sample>=peak)
      then peak:=peak+0.1
    else peak:=peak-0.001;
  end;
end;

```

In either of the above two routines, the current peak estimate is adjusted by the stream variable "sample" and is available as a continuously running diagnostic for logging

and data rescaling. The underlying idea is to give the estimate a rapid upward correction rate for samples that exceed the estimate and a slow downward correction rate for the vast majority of samples that are less than the estimate. The "analogpeak" routine above has a 100:1 rate ratio, and therefore adjusts toward the 99th percentile point of the data distribution. It has the feature of continually dithering its estimate by small amounts, which is beneficial in eliminating the small irregular variations in phase probability arising from the "orchard effect" explained below. It is also desirably robust against occasional out-of-range samples.

We use the positive and negative peak values of the sine and cosine wave samples $+X$, $-X$, $+Y$, and $-Y$ to remap the raw ADC data to a unit-radius encoding circle centered on zero:

$$X = \frac{2X - X_{\text{pos}} - X_{\text{neg}}}{X_{\text{pos}} - X_{\text{neg}}} \quad (4)$$

and

$$Y = \frac{2Y - Y_{\text{pos}} - Y_{\text{neg}}}{Y_{\text{pos}} - Y_{\text{neg}}}, \quad (5)$$

where X_{pos} refers to the current positive peak of X , etc.

We use the peak values of $X+Y$ and $-X-Y$ to determine the length of the positive-slope diagonal of the encoding oval, and $X-Y$ and $-X+Y$ to determine the length of its orthogonal diagonal. Comparing these we obtain a measure of the ellipticity of the encoding oval. Let $pp()$ be the peak-to-peak operator. If the encoding phase shift is 90° plus some small additional angle A (radians), the diagonals will be slightly unequal, to first order given by

$$\frac{pp(X-Y)}{pp(X+Y)} = 1 + A. \quad (6)$$

Given a running estimate of the encoding phase shift, subsequent samples of X and Y can be corrected for this phase shift, thereby circularizing the encoding oval. To first order this can be done using the angle A to introduce cross terms:

$$X = X + Y * \frac{A}{2} \quad (7)$$

and

$$Y = Y + X * \frac{A}{2}. \quad (8)$$

These cross terms have the effect of moving each (X, Y) pair from the observed encoding oval to an idealized encoding circle on which the ideal decoding function is the exact arctangent.

In this way pipeline data monitors can make a TDC system robust against electrical drift of its front end: Variations in ADC slope, offset, and phase angle are continually removed by pipeline processing stages. At the same time, diagnostic monitor variables are continually available and can be made part of the data collection stream.

VI. SYSTEM EVALUATION

The TDC described above was built on a four-layer board. The pipeline data buffer was an FIFO memory chip,




FIG. 2. Two-dimensional histogram of fine cosine and sine pairs, binned onto a coarse 256×256 grid. The trajectory is circular to about 1%. The bin count irregularities are not caused by the TDC; they arise from the varying degree of coverage of the trajectory by the checkerboard.

type 7203S20 which accumulates 1024 timings at its half-full point. For evaluation purposes, we implemented the remainder of the data pipeline using a direct transfer to an IBM PC in which a variety of data processing stages were run (written entirely in BASIC). Single events and event pairs were created using a conventional quartz time base and a delay cable.

Our first set of tests explored the distributions of digital codes that the ADCs generated. These tests allowed us to set the sinusoid amplitudes to best span the ADC input voltage ranges. Histograms of each ADC output were prepared; each shows the counting distribution that is expected from a sampled sinusoid. Bin-to-bin count variations on the order of 20% were observed, and one ADC had missing output codes at some major bit transitions. This performance is typical of contemporary high-speed ADCs.

In our second set of tests we examined the two-dimensional distribution of the raw fine cosine-sine pairs. Ideally, these pairs would encode as points on a perfect circle. If the phase relationship between the two channels differed from 90° , the circular Lissajous pattern would become elliptical. If distortion were present, the periodic track would cease to be elliptical. Figure 2 shows an accumulation of 1 000 000 fast-ADC event codes, binned into a coarse grid of 256×256 cells. The pattern is circular to about 1%.

The presentation of Fig. 2 cannot be used to judge small-scale differential nonlinearity because count variations in its large bins result entirely from the geometric variation in coverage of the Lissajous path by the square bin checkerboard rather than by any defect of the TDC. To explore differential timing nonlinearity and time sample density, we conducted a third test exploring the coverage of the (X, Y) plane. Here we accumulated a Lissajous pattern using the full 12-bit resolution of the fine ADCs. The mole run resulting from 567 000 random events is shown in the oblique view of Fig. 3, which spans $1/16$ of the X range and $1/16$ of the Y range. Figure 3 shows that the (X, Y) plane is statistically well covered by samples.

Our fourth set of tests examined the pipeline information for statistical uniformity of the derived fine start-time estimates using a floating point standard BASIC arctangent function. This test is a sensitive indicator of the system's




FIG. 3. Magnified portion of the full resolution 4096×4096 fine cosine and sine trajectory, showing the coverage of the track by the individual data samples. Missing ADC codes occur at some major bit transitions.

freedom from high-frequency wave form distortions. To conduct this test we generated a train of random pulse events, applied Eqs. (4)–(8), performed an arctangent transformation to each fine cosine-sine pair to obtain a measure of fine phase, and accumulated the events in bins of size 1 ps (5 ns/5000). The low-frequency timing channel was ignored as were STOP events.

The resulting data, containing 1 000 000 random starts, are shown in Fig. 4. In this presentation, a fixed phase shift error between the cosine and sine waves would introduce a sinusoidal density variation. Harmonic distortions cause other wavelike density irregularities. What we see is a small-scale bin-to-bin scatter plus an undulation having three cycles per complete cycle of fine phase. The scatter exceeds the 7% rms expected from random-event arrivals. It is largely an "orchard effect" analogous to viewing a periodic grid of trees from a point in an orchard (in contrast to viewing a random population of trees in a forest): The scatter is deterministic, and arises from the limited number of available X and Y codes that contribute counts to a given phase bin. The three-cycle undulation reveals a small (correctable) degree of harmonic distortion at the level of a few percent.

In an ideal TDC, recovered time intervals would be distributed about the true mean start-stop interval with a narrow Gaussian distribution characteristic of the front-end ap-




FIG. 4. Uniformity test for recovered fine phases, using the arctangent function on 1 000 000 random events. Ideal TDC processing would yield a flat probability in phase, subject only to random arrival statistics. The random-looking irregularities arise from the varying numbers of X and Y codes that yield a given phase; we term this the "orchard effect." The waves are caused by harmonic distortion.




FIG. 5. Distribution of TDC outputs for an asynchronous train of 104 ns intervals. Raw ADC data are applied to the arctangent routine. Bin size=1 ps; rms deviation is 39 ps.

erture noise. In practice, the mean of the recovered time intervals accurately corresponds to the true mean to within the precision of the quartz time base. However, the observed probability distribution is broader than predicted from front-end noise considerations alone, and it is non-Gaussian. The reason is the presence of harmonic FM distortion in the sampled 200 MHz clock wave forms. This causes a start sample to be early or late, up to 10 ps, depending on its start clock phase. A stop sample experiences a similar error correlated with the stop phase.

We have developed a means of correcting these harmonic errors by accumulating and examining histograms of start phases and stop phases separately. A long accumulation of start and stop phases is made in histogram form. Since the events arrive randomly in time, all start phases and all stop phases have an equal probability of being observed, and, therefore, a long accumulation of events should produce two completely flat histograms. Any deviation from linear time recovery will be revealed as lumpiness in the histograms. The observed histograms are integrated to produce the actual time transfer function, and a list of correction terms is generated to be applied to the START and STOP phases of subsequent events.

In a fifth set of tests we explored the distribution of event interval timings using the combined coarse/fine start-




FIG. 6. Distribution of TDC outputs as in Fig. 5, except raw data are linearly adjusted for offset and span using X and Y peak detectors.




FIG. 7. Distribution of TDC outputs as in Fig. 6, except data are also corrected for phase shift between X and Y , using the ellipticity of the encoding oval.

stop computation described in Sec. IV above. We employed a programmable bench oscillator and a switchable segmented delay line to give reproducible start-stop pairs. The bench oscillator was not synchronized in any way with the 200 MHz TDC clock. Time interval distributions exhibit a statistical uniformity far superior to the start-phase or stop-phase distributions because intervals average the digitizer outputs over all possible TDC time base phases.

Figures 5–8 illustrate the ability of the pipeline processing to accommodate hardware imperfections. Each plot is the result of repeated measurements of a delay line generated time interval of 104 ns. For these accumulations we used a bin size of 1 ps, although a much coarser (or finer) bin size could have been chosen. In Fig. 5, no corrections have been made to the raw ADC values prior to the arctangent stage. The timing error is 39 ps rms. In Fig. 6, peak detectors have been used to linearly center and rescale the ADC code ranges to a uniform diameter. In Fig. 7, the ellipticity correction has been applied as well to remove any systematic interchannel phase error. In Fig. 8, we additionally employed the histogram method for removing harmonic modulation of the start and stop times. The distribution of timing errors is nearly Gaussian. The timing error is 4.1 ps rms. In Fig. 9, we show




FIG. 8. Distribution of TDC outputs as in Fig. 7, except data are also corrected for harmonic distortion using integrated count histograms of start phase and stop phase. Each correction histogram contained 1 000 000 events. Root mean square deviation is 4 ps.




FIG. 9. Distribution of TDC outputs from an accumulation of 10 million pulse pairs from a 101 ns delay cable. Histogram bin size is 125 fs. The bin-to-bin count fluctuations are about 0.3% rms, of the same order as expected from random-event arrival statistics alone. We estimate that our TDC's small-scale differential nonlinearity does not exceed 400 as rms.

a long accumulation with a 101 ns delay cable but with a bin size of 125 fs to show the freedom from fine-structure artifacts. The bin-to-bin count scatter is comparable to the 0.3% rms expected from Poisson event arrivals alone. The presence of small-scale differential nonlinearity would appear here as an additional count variation. We estimate that any such nonlinearity is less than 400 as rms. A test of large-scale differential nonlinearity is provided by binning a large number of random start-stop pairs into time-interval groups. Using 250 intervals of 320 ps each and 75 million start-stop pairs, we obtain the plot shown in Fig. 10, which reveals the presence of some residual differential nonlinearity at the 1% level.

Finally, we have examined the short-term and long-term time measurement stability of our TDC in a laboratory environment with a continually powered double-pulse source. At initial power-on, the TDC timings show an initial upward trend that levels off after 5 min of steady operation. The initial deviation is about 20 ps from the long-term value.




FIG. 10. Differential nonlinearity plot for 75 million random events and 250 bins of 320 ps. Vertical scale: number of events per bin. Horizontal scale: time interval.

After 15 min of operation, deviations do not exceed 5 ps. These errors are presumably a result of differential gate delays in the start-stop logic and in the s/h circuits.

VII. ALTERNATIVE IMPLEMENTATIONS

The work reported here provides a basis for instrumenting a variety of delay line detectors. In the course of this work, we have considered some alternative implementations which are also attractive. For example, the cosine/sine/arctangent coding is only one of many possible bases for clock generation. Any periodic wave form pair will serve. Triangle waves yield diamond-shaped Lissajous patterns that can be decoded by a suitable table. Trapezoidal waves produce square patterns, which could also work. The cosine/sine pair simply happens to have the least bandwidth, and can be made acceptably pure by applying a reasonable degree of passive narrowband filtering.

The coarse timing channel is far less critical than the fine 200 MHz timing channel, and in considering alternative implementations of the present TDC we note that great liberties can be taken with its design. It is necessary only to resolve individual cycles of the 200 MHz fundamental. For example Berry's¹⁰ design accomplishes this with a digital binary-divider countdown chain driven at a frequency of 2 f. The chain is gated off to determine the coarse count.

The arctangent function on which our present implementation is based could be replaced by an arbitrary lookup table whose elements are chosen in such a way as to minimize the total timing error. In this way, the effect of any sampled clock wave distortion would be eliminated along with phase shifts and nonlinearities in the front end. Since the true distribution of phase estimates is flat for any test signal including (especially) random-event arrivals, it is feasible to create an adaptive algorithm that modifies the phase table continually during production data-acquisition runs in such a way as to maintain a flat statistical distribution of recovered phases.

Finally, we note that the field of microwave analog integrated circuits is advancing rapidly. Specifically, chips for digital sampling oscilloscopes now under development can acquire samples from a full amplitude 5 GHz sine wave.¹¹ Unlike other architectures, our TDCs resolution is limited only by its sampler speed. A 5 GHz arctangent TDC could resolve to better than 1 ps.

ACKNOWLEDGMENTS

We gratefully acknowledge useful conversations with Dr. O. H. W. Siegmund and Professor S. Bowyer. The authors acknowledge the support of NASA Grant Nos. NAGW-3823 and NAGW-1290.

¹M. Lampton, O. Siegmund, and R. Raffanti, *Rev. Sci. Instrum.* **58**, 2298 (1987); M. Lampton, O. H. W. Siegmund, and R. Raffanti, *IEEE Trans. Nucl. Sci.* **37**, 1548 (1990); R. Raffanti and M. Lampton, *Rev. Sci. Instrum.* **64**, 6 (1993).

²R. Nutt, *Rev. Sci. Instrum.* **39**, 3 (1968).

³O. Sasaki, T. Taniguchi, and T. K. Ohoka, *IEEE Trans. Nucl. Sci.* **NS-35**, 1 (1988).

- ⁴B. T. Turko, IEEE Trans. Nucl. Sci. NS-27, 1 (1980).
⁵B. T. Turko, IEEE Trans. Nucl. Sci. NS-31, 1 (1984).
⁶W. F. Kolbe and B. Turko, IEEE Trans. Nucl. Sci. NS-35, 1 (1988).
⁷J. Kalisz, R. Pelka, and A. Poniecki, Rev. Sci. Instrum. 65, 3 (1994).
⁸M. A. Thompson, M. W. Werner, R. R. Egan, P. D. Coan, and P. E. Robl,
IEEE Trans. Nucl. Sci. NS-35, 1 (1988).
⁹S. Kleinfelder, T. J. Majirs, K. A. Blumer, W. Farr, and B. Manor, IEEE
Trans. Nucl. Sci. NS-38 (1991).
¹⁰A. Berry, Rev. Sci. Instrum. 64, 5 (1993).
¹¹S. Naboicheck and S. Ems, Microwave J. 37, 4 (1994).