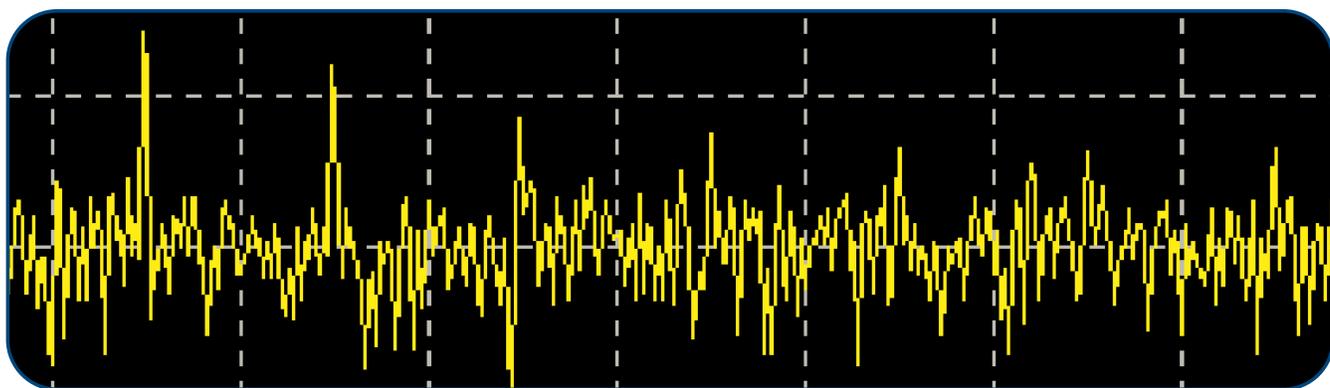


Understanding and Characterizing Timing Jitter



COMPUTING

COMMUNICATIONS

VIDEO

A Guide to Understanding and Characterizing Timing Jitter

► Primer

► Table of Contents

Introduction	3
Section 1: The Consequences of Jitter	3
Computer Bus Design	3
Serial Data Link	4
Section 2: Just What is Jitter?	4
2.1 Defining Short-Term: Jitter vs. Wander	4
2.2 Defining Significant Instants: The Reference Level	5
2.3 Defining Ideal Positions: Clock Recovery	5
2.4 Period Jitter, Cycle-Cycle Jitter, and TIE	6
Section 3: Jitter Measurement and Visualization	8
3.1 Jitter Statistics	8
3.2 Jitter Histogram	8
3.3 Jitter vs. Time (Time Trend)	10
3.4 Jitter vs. Frequency (Jitter Spectrum)	11
3.5 Eye Diagram	12
Section 4: Jitter Separation	13
4.1 Motivations for Decomposing Jitter	13
4.2 The Jitter Model	13
4.2.1 Random Jitter	14
4.2.2 Deterministic Jitter	15
4.2.3 Periodic Jitter	15
4.2.4 Data-Dependent Jitter	16
4.2.5 Duty-Cycle Dependent Jitter	17
4.3 Putting It All Together	18
Section 5: Jitter vs. Bit Error Rate	20
5.1 The Jitter Budget	20
5.2 The Bathtub Curve	21
5.3 The BER Example	22
Section 6: Summary	23
Appendix A: Glossary of Acronyms	23
References	23

Introduction

Timing jitter is the unwelcome companion of all electrical systems that use voltage transitions to represent timing information. Historically, electrical systems have lessened the ill effects of timing jitter (or, simply “jitter”) by employing relatively low signaling rates. As a consequence, jitter-induced errors have been small when compared with the time intervals that they corrupt. The timing margins associated with today’s high-speed serial buses and data links reveal that a tighter control of jitter is needed throughout the system design.

As signaling rates climb above 2 GHz and voltage swings shrink to conserve power, the timing jitter in a system becomes a significant percentage of the signaling interval. Under these circumstances, jitter becomes a fundamental performance limit. Understanding what jitter is, and how to characterize it, is the first step to successfully deploying high-speed systems that dependably meet their performance requirements.

A more thorough definition will be introduced in Section 2, but conceptually, jitter is the deviation of timing edges from their “correct” locations. In a timing-based system, timing jitter is the most obvious and direct form of non-idealness. As a form of noise, jitter must be treated as a random process and characterized in terms of its statistics.

If you have a way to measure jitter statistics, you can compare components and systems to each other and to chosen limits. However, this alone will not allow you to efficiently refine and debug a cutting-edge design. Only by thoroughly analyzing jitter is it possible for the root causes to be isolated, so that they can be reduced systematically rather than by trial and error. This analysis takes the form of jitter visualization and decomposition, discussed in detail in Sections 3 and 4.

Although there are many similarities between the causes, behavior and characterization of electrical and optical jitter, the equipment used to measure jitter in optical systems differs from that used in electrical systems. This paper focuses primarily on jitter in electrical systems.

Section 1: The Consequences of Jitter

“To guess is cheap. To guess wrong is expensive.”

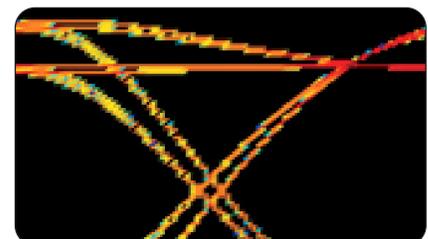
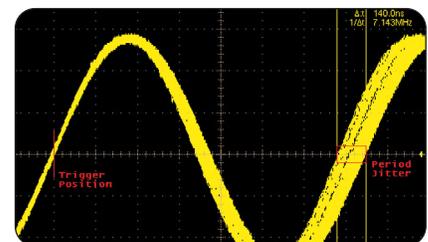
Chinese Proverb

Why should you be concerned about jitter? What impact does jitter have on system performance? In this section, two cases are considered: a high-speed computer bus, and a serial data link. For each case, the specific effects of jitter are considered in some detail.

Computer Bus Design

Let’s say you are bringing up a new embedded processor design, and are seeing occasional data errors when reading flash memory. You suspect that the address decode that generates the flash memory’s chip enable (CE) is not meeting it’s setup time requirement with respect to the rising edge of the write enable (WE). You probe the CE and WE signals with your high-speed scope to observe the timing relationship. After ten single-shot acquisitions, you have measured durations of 87 to 92 nsec, all above the minimum setup time of 75 nsec by what seems like a comfortable margin. But what is an adequate margin? Are these times great enough to remove doubt that you are sometimes violating the setup-time requirement? What percent of the time is a violation likely to occur?

After evaluating several million waveforms in infinite-persistence mode and seeing setup times as short as 82 nsec, you decide that the setup time might be a problem after all. But is the problem due to variations in the system clock period, in the address decoder, or somewhere else?



Serial Data Link

Your Gigabit Ethernet physical-layer transceiver chip is pushing up against deadline and you are a bit nervous about the compliance tests that will soon be run by an external test house. The spec in the standards document calls for a measurement of data jitter relative to the local data clock, and another measurement of the clock's jitter relative to a jitter-free reference, whatever that is. In any event, you'd like to make sure you have enough margin that the parts from the qualification run satisfy the compliance lab.

You start by using an oscilloscope in infinite-persistence mode to check the peak-to-peak jitter on the data clock. Since your scope allows you to define histogram boxes on-screen, you use this feature to create a histogram of the edge locations. You find a peak-to-peak value of 550 psec, and the spec says you must have less than 300 psec. Fortunately, the 300 psec spec is after the jitter has been filtered by a 5 kHz high-pass filter. Unfortunately, you have no way of knowing what part of the jitter in the histogram is due to lower frequencies and can safely be disregarded.

You take a look at the jitter on the data lines relative to the clock, and find that this, too, is dangerously close to the spec limit. However, you suspect that this jitter is not in your chip, but that it has been picked up because of the test board, which doesn't adhere to good design practice on the layout of the differential data lines. You know that the switching supply on the test board may be to blame, but need to determine how much of the measured jitter is coming from this source.

Both this case and the one above are examples of times when even a high-performance oscilloscope may not provide enough muscle to answer all your questions. To really feel confident about your design, you may need two additional tools: some means of advanced jitter analysis, and a good grasp of the basic causes and characteristics of timing jitter.

Section 2: Just What Is Jitter?

This simple and intuitive definition is provided by the SONET specification¹:

“Jitter is defined as the short-term variations of a digital signal's significant instants from their ideal positions in time.”

This captures the essence of jitter, but some of the individual terms (*short-term*, *significant instants*, *ideal positions*) need to be more specific before this definition can be unambiguously used.

In all real applications, jitter has a random component, so it must be specified using statistical terms. Metrics such as mean value and standard deviation, and qualifiers such as confidence interval, must be used to establish meaningful, repeatable measurements. A review of the basic concepts from statistical mathematics that underlie the following material is beyond the scope of this paper, but the bibliography at the end provides some references for the reader who wishes to delve deeper.

2.1 Defining Short-Term: Jitter vs. Wander

By convention, timing variations are split into two categories, called *jitter* and *wander*, based on a Fourier analysis of the variations vs. time. (This kind of analysis is covered in more detail in Sections 3.3 - 3.4.) Timing variations that occur slowly are called *wander*. *Jitter* describes timing variations that occur more rapidly. The threshold between wander and jitter is defined to be 10 Hz according to the ITU² but other definitions may be encountered. In many cases wander is of little or no consequence on serial communications links, where a clock recovery circuit effectively eliminates it.

2.2 Defining Significant Instants: The Reference Level

The significant instants referred to by our definition are the transitions, or edges, between logic states in the digital signal. To be more specific, the significant instants are the exact moments when the transitioning signal crosses a chosen amplitude threshold, which may be referred to as the *reference level* or *decision threshold*. For two-level signal (by far the most common case), the mean signal voltage is frequently used as this reference level. If the signal-of-interest is to be received by a Schmitt-trigger input, it may be desirable to use one reference level when analyzing the rising edges and a different reference level when analyzing the falling edges.

The phrase *digital signal* in our definition is perhaps optimistic, since for high-speed signals, the transitions are analog events that are subject to rise-time and slew-rate limitations. During the small but finite time when the signal is ramping through the reference level, any voltage noise that corrupts the waveform will be converted proportionately into timing jitter.

2.3 Defining Ideal Positions: Clock Recovery

Before a digital signal's deviations from ideal positions can be measured, those ideal positions must be identified. For a clock-like signal (alternating 1's and 0's), the ideal positions conceptually correspond to a jitter-free clock with the same mean frequency and phase as the measured one. More care must be used for a data signal, since no event (transition) occurs when the same bit repeats two or more times in a row. Clock Recovery is the name given to the process of establishing the timing of the reference clock.

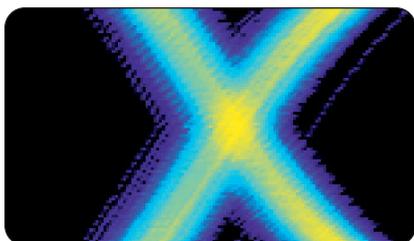
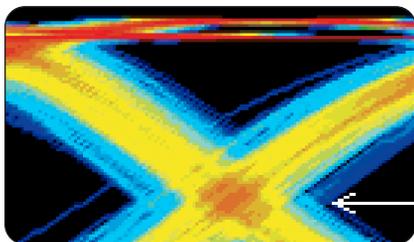
One method of clock recovery is to use the constant-frequency clock that best fits the measured events, in the least-squares sense. This means that a reference clock of the form:

$$A \cdot \sin(\omega_c \tau + \phi_c)$$

is assumed, where ω_c and ϕ_c are constants. The constants are chosen such that the sum of the squares of the time errors between the reference clock and the measured clock is minimized. This is an excellent approach when a finite-length block of contiguous data is to be analyzed. If the duration of the data is sufficiently long, the resulting jitter measurement may include wander as well as jitter. In this case, a high-pass filter may be used subsequently to remove the wander component.

Another effective method of clock recovery is through the use of a phase locked loop (PLL). A PLL constantly tracks slowly-varying changes in the symbol rate of the measured data. Consequently, it acts as a high-pass filter with respect to the jitter that remains on the signal. Since most datacom links use PLLs in their receivers, this measurement approach has the advantage of modeling the behavior of the system in which the measured device will be used.

For consistency and repeatability of measurement, many datacom specifications define a *golden PLL*. In this context, *golden* simply means that the PLL characteristics are precisely defined and tightly controlled. If a PLL conforming to these specs is used to measure jitter on multiple devices, the jitter may be objectively compared and logically related to the system in which the devices are to be used.



2.4 Period Jitter, Cycle-Cycle Jitter, and TIE

There are several ways in which jitter may be measured on a single waveform. These are *period jitter*, *cycle-cycle jitter*, and *time interval error* (TIE). It is important to understand how these measurements relate to each other and what they reveal.

Figure 2.4a shows a clock-like signal with timing jitter. The dotted lines show the ideal edge locations, corresponding to a jitter-free version of the clock.

The *period jitter*, indicated by the measurements P1, P2 and P3, simply measures the period of each clock cycle in the waveform. This is the easiest and most direct measurement to make.

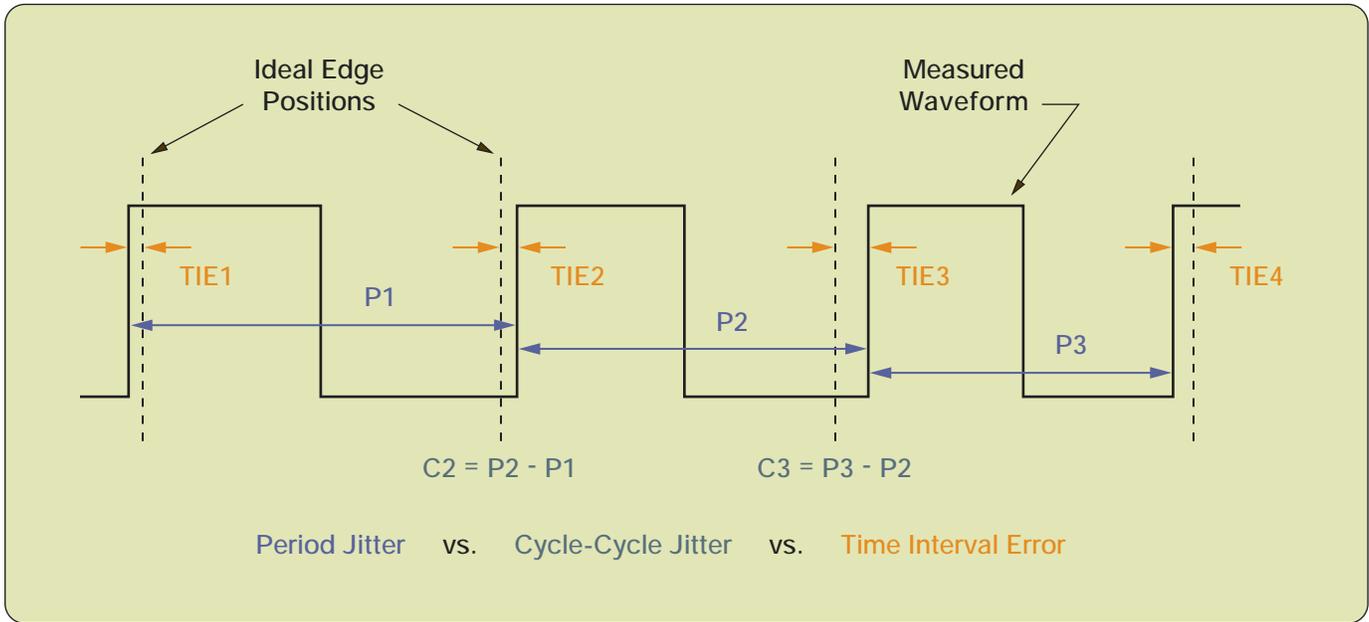
Its peak-to-peak value may be estimated by adjusting an oscilloscope to display a little more than one complete clock cycle with the display set for infinite persistence. If the scope triggers on the first edge, the period jitter can be seen on the second edge. This is shown in Figure 2.4b.

The *cycle-cycle jitter*, indicated by C2 and C3 in Figure 2.4a, measures how much the clock period changes between any two adjacent cycles. As shown, the cycle-cycle jitter can be found by applying a first-order difference operation to the period jitter. This measurement can be of interest because it shows the instantaneous dynamics a clock-recovery PLL might be subjected to. Notice that no knowledge of the ideal edge locations of the reference clock was required in order to calculate either the period jitter or the cycle-cycle jitter.

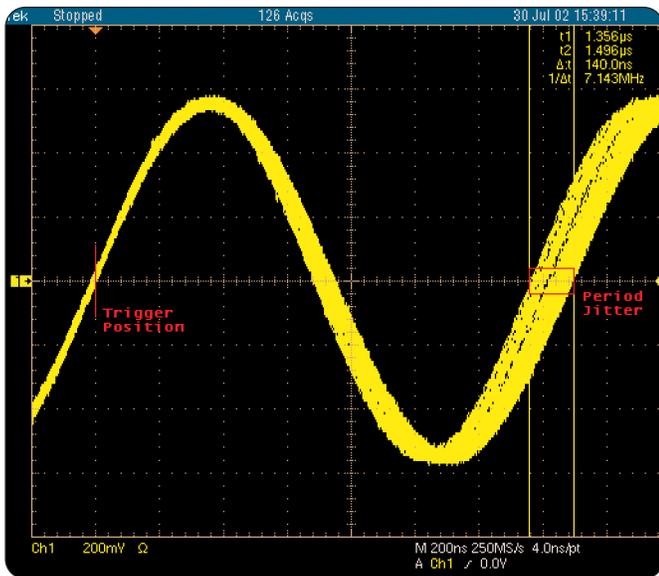
The time interval error is shown in Figure 2.4a by the measurements TIE1 through TIE4. The TIE measures how far each active edge of the clock varies from its ideal position. For this measurement to be performed, the ideal edges must be known or estimated. For this reason, it is difficult to observe TIE directly with an oscilloscope, unless some means of clock recovery or post-processing is available.

The TIE may also be obtained by integrating the period jitter, after first subtracting the nominal (ideal) clock period from each measured period. TIE is important because it shows the cumulative effect that even a small amount of period jitter can have over time. Once the TIE reaches ± 0.5 unit intervals, the eye is closed and a receiver circuit will experience bit errors.

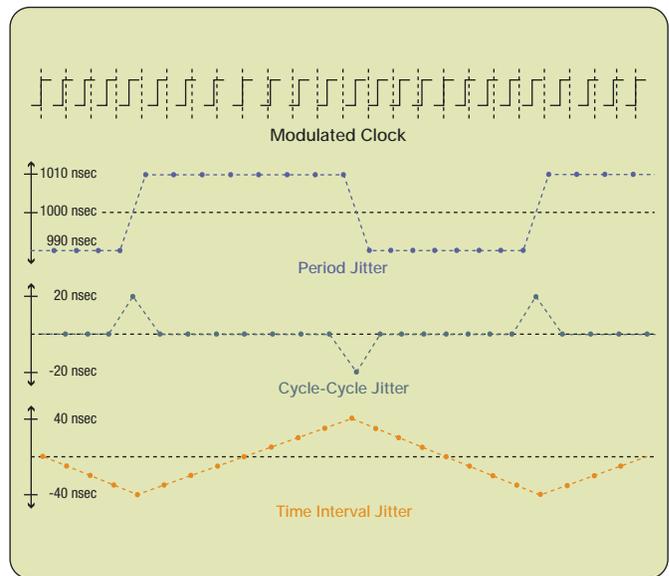
Figure 2.4c gives an example of how these three jitter measurements compare to each other for a specific waveform. For this example, the waveform has a nominal period of 1 μsec , but the actual period follows a pattern of eight 990 nsec cycles followed by eight cycles of 1010 nsec.



► Figure 2.4a



► Figure 2.4b



► Figure 2.4c

Section 3: Jitter Measurement and Visualization

The paragraphs in this section discuss some of the tools and techniques that are useful for quantifying and/or analyzing jitter.

3.1 Jitter Statistics

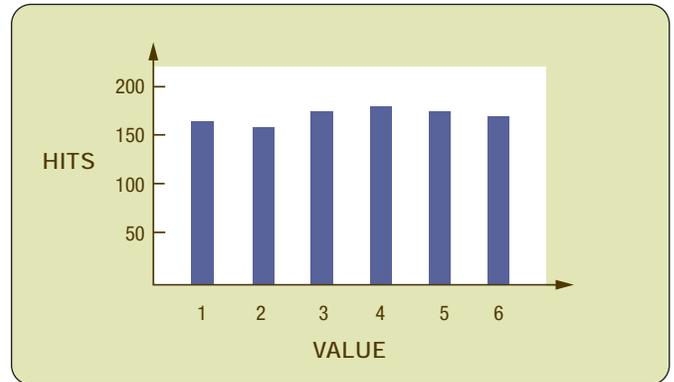
Since all known signals contain jitter that has a random component, statistical measures are required to properly characterize the jitter. Some of the commonly used measures are:

Mean Value: The arithmetic mean, or average, value of a clock period is the nominal period. This is the reciprocal of the frequency that a frequency counter would measure. The mean value of the TIE is theoretically zero, although a small residual value may result depending on the measurement technique.

Standard Deviation: The standard deviation, represented by the greek character sigma (σ), is the average amount by which a measurement varies from its mean value. It is particularly useful in describing Gaussian processes, for which the distribution is completely specified by the mean and standard deviation. This is discussed at greater length in Section 4.3.

Maximum, Minimum and Peak-Peak Values: The Max and Min values generally refer to values actually observed during a measurement interval, and the Peak-Peak value is simply the Max minus the Min. These measurements should be used judiciously. For a deterministic signal, these values may equal the true values even after a relatively short measurement interval. But for a random signal with a Gaussian distribution, there is theoretically no limit on the max and min values, so the observed peak-peak value will generally grow over time. For this reason, the peak-peak value should be used in conjunction with the population size and some knowledge of the type of distribution.

Population: The population is the number of individual observations included in a statistical data set. For a random process, a high population intuitively gives greater confidence that the measurement results are repeatable. If the characteristics of the distribution are known or can be estimated, it is possible to calculate the population needed to reduce the measurement uncertainty below a desired point.



► **Figure 3.2a**

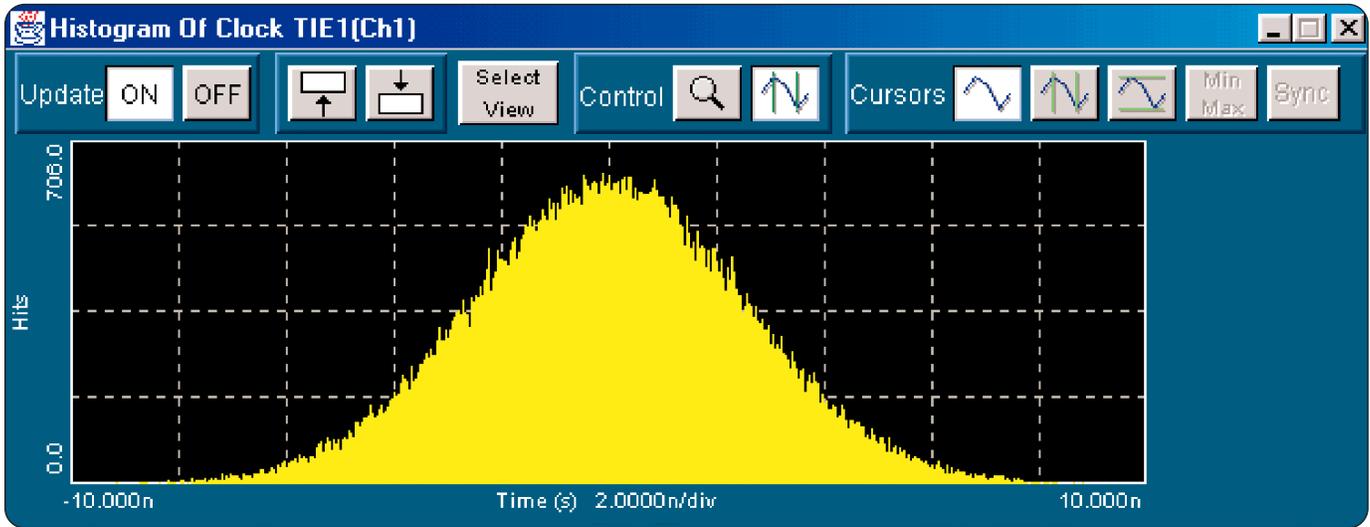
3.2 Jitter Histogram

A histogram is a diagram that plots the measurement values in a data set against the frequency of occurrence of the measurements. If the number of measurements in the data set is large, the histogram provides a good estimate of the probability density function (pdf) of the set. For example, if you rolled a fair die 1000 times and recorded the results, they might be as depicted in Figure 3.2a, where the HITS axis shows the number of times each value occurred.

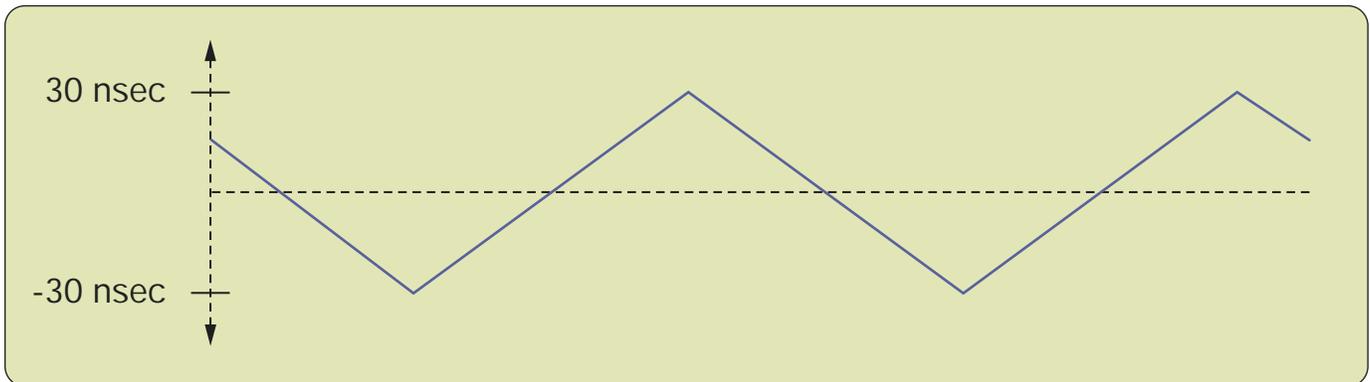
Note that a histogram provides no information about the order in which the observations occurred.

An example of a jitter histogram for a TIE measurement is shown in Figure 3.2b. In this case, the continuous variable is mapped into 500 bins and the total population of the data set is 100,000. Since this is a TIE measurement, the mean value is 0 nsec. For this plot, the distribution is approximately Gaussian with a standard deviation of 2.63 nsec.

Consider a second example, in which the TIE of the signal in Figure 3.2b is modulated by the triangle wave depicted in Figure 3.2c.



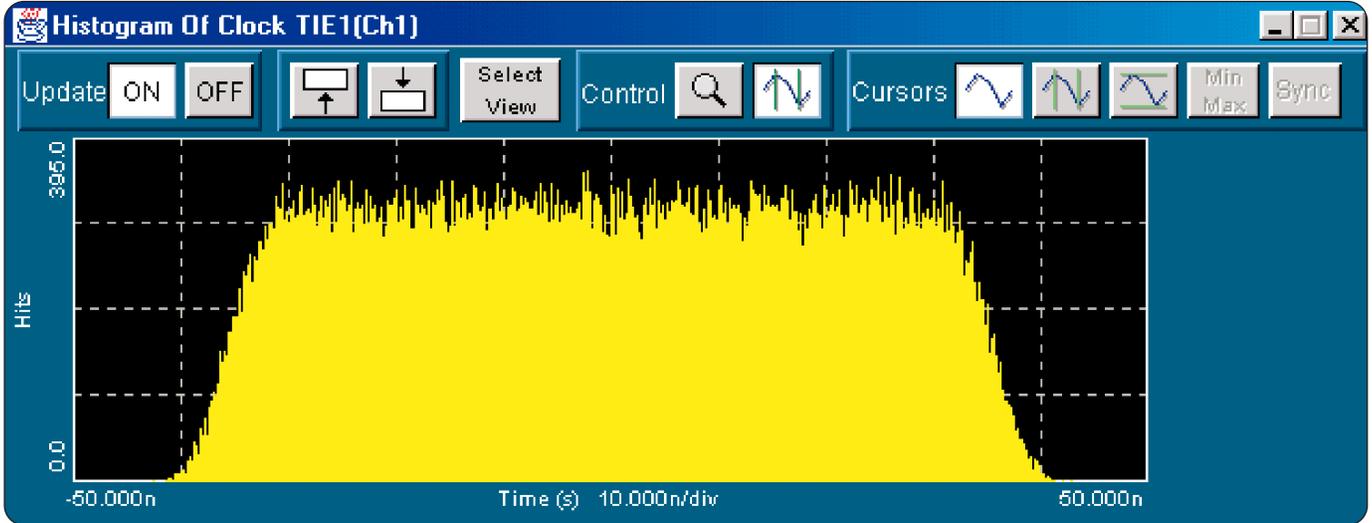
▶ Figure 3.2b



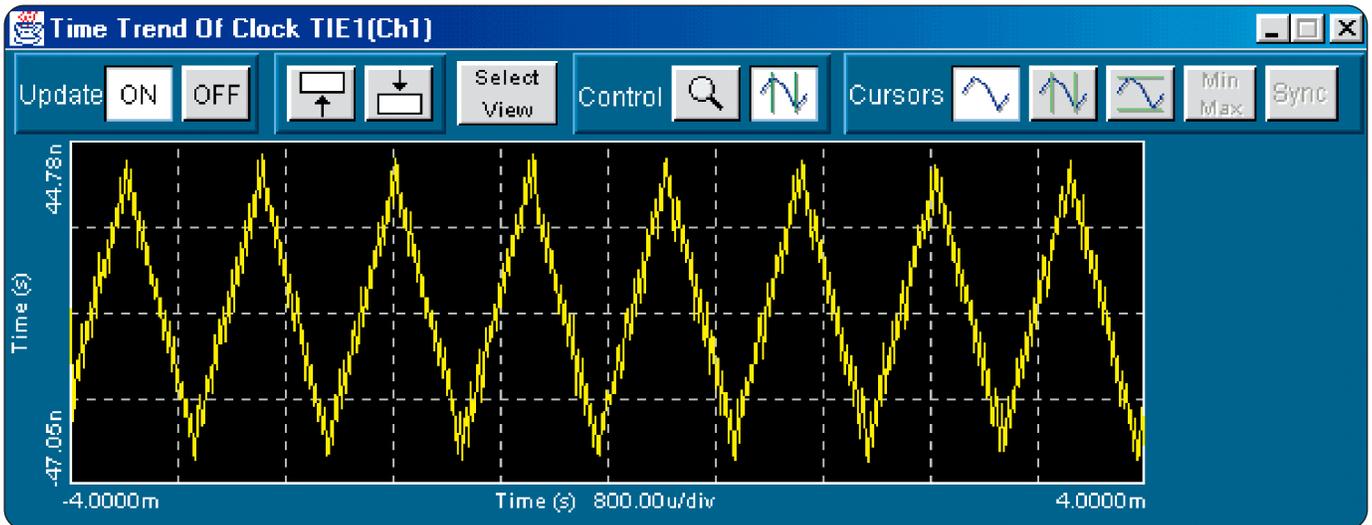
▶ Figure 3.2c

A Guide to Understanding and Characterizing Timing Jitter

► Primer



► **Figure 3.2d**

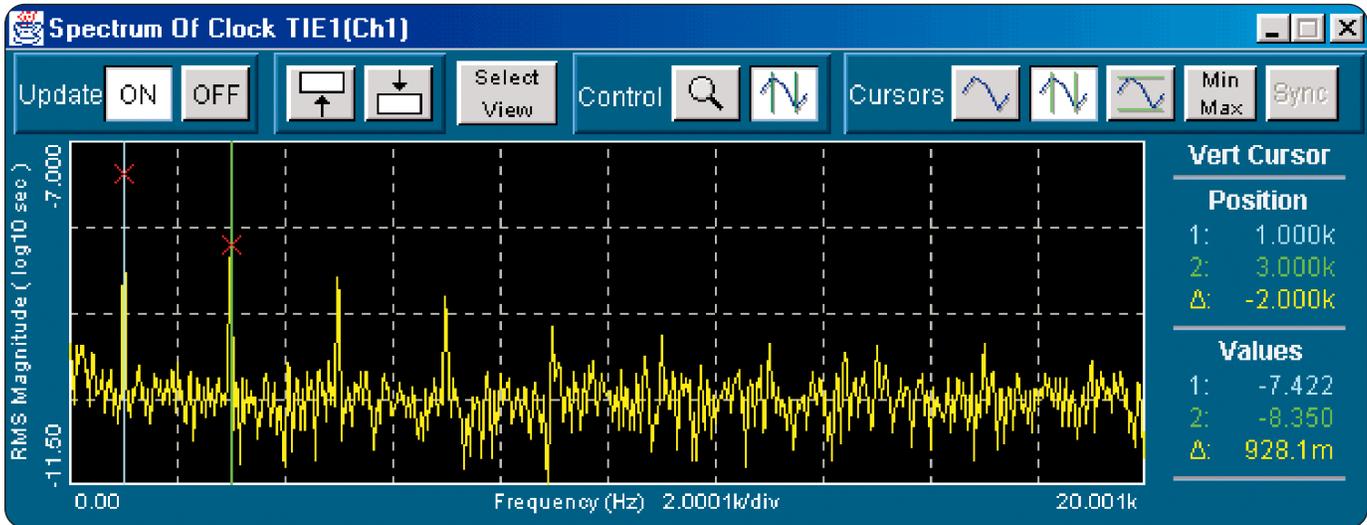


► **Figure 3.3a**

If this triangle wave were sampled at an arbitrary instant, there is an equal likelihood that the sample value would lie anywhere between -30 nsec and 30 nsec. Hence, the TIE histogram of the modulated signal is spread over ± 30 nsec with roughly equal likelihood, as shown in Figure 3.2d. (The sloping tails at the left and right edges of the histogram show that the jitter still has a Gaussian component.)

3.3 Jitter vs. Time (Time Trend)

Since the jitter histogram doesn't show the order in which the measurement observations occur, it cannot reveal repeating patterns that might indicate a modulation or other periodic component. A plot of jitter value versus time can make such a pattern obvious. As an example, the TIE from the phase-modulated signal of Figure 3.2d was plotted against time to produce Figure 3.3a. Now, the pattern of jitter variation becomes apparent, and its correlation with one of several possible sources of coupled noise might become clear. Note that the random component of the jitter can still be seen as noise added to the triangular modulation.



▶ Figure 3.4a

3.4 Jitter vs. Frequency (Jitter Spectrum)

Since the jitter measurements can be plotted versus time, an obvious extension is to apply a Fourier transform to these measurements and display the results in the frequency domain. This results in a jitter spectrum, with the modulation frequency displayed on the horizontal axis and the amplitude of modulation shown on the vertical axis. One of the benefits of spectral analysis is that periodic components that otherwise might be hidden by wideband noise can often be clearly distinguished.

The prior example of a clock with triangular modulation is used again here. Figure 3.4a shows the same TIE measurement as Figure 3.3a, displayed as a TIE spectrum.

Now, the triangular modulation is seen to have a fundamental frequency of 1 kHz, as indicated by the blue cursor. In theory, the Fourier series representation of a triangle wave only has odd harmonics. This is corroborated by the spectrum, where components can clearly be seen at 3 kHz, 5 kHz, 7 kHz, and so on. The random noise still appears, but in this view it is manifested as a broad, flat noise floor.

It was mentioned earlier that timing variations with Fourier components below some limit, usually 10 Hz, are considered by convention to be *wander* rather than *jitter*. Of more practical interest, some other frequency limit (e.g. the loop bandwidth of your system's clock-recovery loop) may determine which noise can be safely tolerated. A spectral view of jitter can reveal whether the noise in a system is of concern or not.

A Guide to Understanding and Characterizing Timing Jitter

► Primer

3.5 Eye Diagram

All of the methods discussed so far rely on edge locations only. These locations are extracted from a waveform by detecting when the waveform crosses one or more amplitude thresholds. The eye diagram is a more general tool, since it gives insight into the amplitude behavior of the waveform as well as the timing behavior.

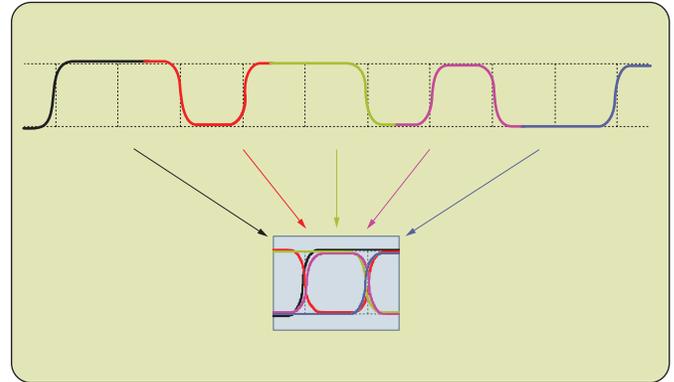
An eye diagram is created when many short segments of a waveform are superimposed such that the nominal edge locations and voltage levels are aligned, as suggested in Figure 3.5a. Usually, a horizontal span of two unit intervals is shown. The waveform segments may be adjacent ones, as shown in the figure, or may be taken from more widely spaced samples of the signal. If the waveform is repeatable, a sampling scope may be used to build an eye diagram from individual samples taken at random delays on many waveforms.

Color has been used in the stylized Figure 3.5a to show how the individual waveform segments are composed into an eye diagram. In practice, eye diagrams are usually either monochrome or use color to indicate the density of waveform samples at any given point of the display. Figure 3.5b shows such a color density display for a waveform that exhibits several types of noise.

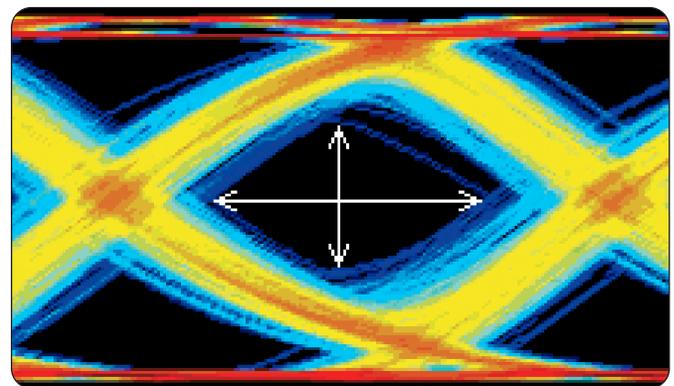
In this diagram, white arrows are used to show the vertical and horizontal extent of the eye opening. As the noise on a signal increases, the eye becomes less open, either horizontally, vertically or both. The eye is said to be closed when no open area remains in the center of the diagram.

The easiest way to create an eye diagram is with an oscilloscope in a long-persistence display mode. You should use caution in how you trigger the scope, if you use this approach. Simply triggering on a waveform edge will result in an eye diagram showing waveforms relative to that edge. This can be very different from an eye diagram relative to the underlying bit clock. To produce a diagram that is relative to the bit clock, you need some form of clock recovery, either in software or hardware. If your oscilloscope does not provide this feature, you may be able to use a trigger derived from an external clock-recovery circuit.

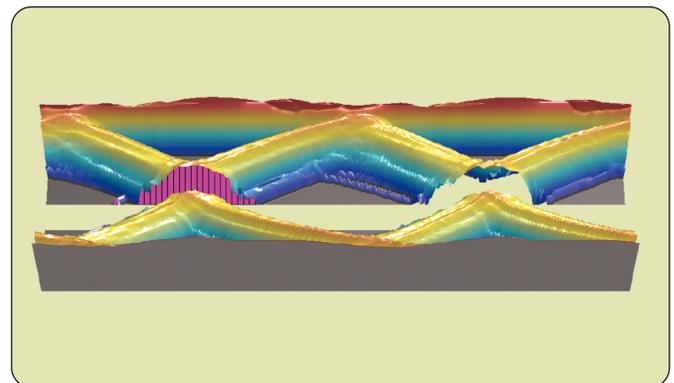
The relationship between the eye diagram and the TIE histogram can be seen by turning the density eye diagram of Figure 3.5b into a three-dimensional figure and slicing it along the decision threshold, as illustrated in Figure 3.5c. The area highlighted in pink is equivalent to a histogram of the first of the two zero-crossings in the eye diagram.



► Figure 3.5a



► Figure 3.5b



► Figure 3.5c

Section 4: Jitter Separation

“All models are wrong; some models are useful”

W. Edwards Deming

Jitter separation, or jitter decomposition, is an analysis technique that uses a parameterized model to describe and predict system behavior. This section explains why this technique is used, and provides details about the jitter model most commonly used today.

4.1 Motivations for Decomposing Jitter

To understand how real systems behave, it is often useful to use a mathematical model of the system. The behavior of such a model can be tuned by adjusting the parameters of its individual components. If the parameters of the model are chosen based on observations of the real system, then the model can be used to predict the behavior of the system in other situations. Thus, one of the motivations for jitter decomposition (also called jitter separation) is to extrapolate system performance to cases that would be difficult or time-consuming to measure directly.

Another motivation for modeling the system this way has to do with analysis. If each of the model components is associated with one or more underlying physical effects, an understanding of the model can provide insight into the precise cause or causes of excessive jitter.

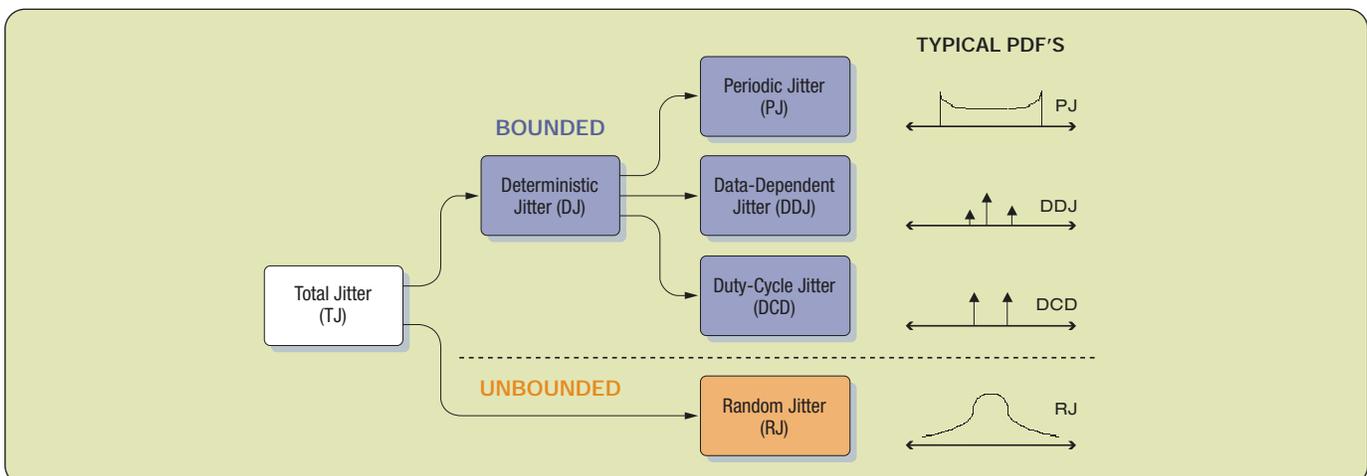
All models of complex systems make assumptions and simplifications, so the fit between the model and the true system behavior will never be exact. And in fact, there is usually some latitude in parameter choice when fitting the model's behavior to the observed measurements. For this reason, jitter separation has some elements of art as well as science, and "four nines" repeatability of measurements is not to be expected.

4.2 The Jitter Model

The jitter model most commonly used is based on the hierarchy shown in Figure 4.2. In this hierarchy, the total jitter (TJ) is first separated into two categories, random jitter (RJ) and deterministic jitter (DJ). It will be seen later that correctly distinguishing between these two jitter types has a dramatic impact on model correctness – more so than any other modeling decision.

The deterministic jitter is further subdivided into several categories: periodic jitter (PJ, also sometimes called sinusoidal jitter or SJ), duty-cycle dependent jitter (DCD), and data-dependent jitter (DDJ, also known as inter-symbol interference, ISI). An additional category (bounded, uncorrelated jitter, or BUJ) is sometimes used.

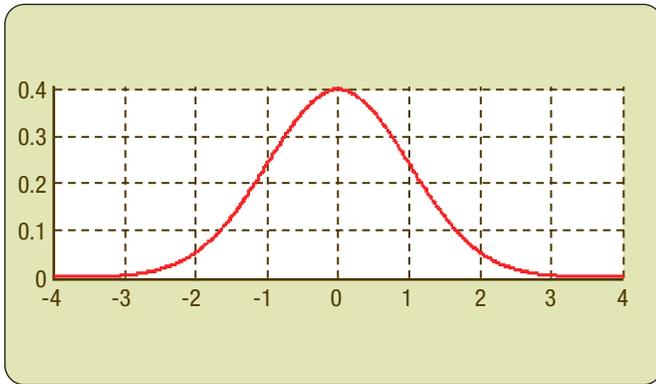
For each of these jitter types, the characteristics and root causes are discussed in the paragraphs that follow.



▶ **Figure 4.2**

A Guide to Understanding and Characterizing Timing Jitter

► Primer

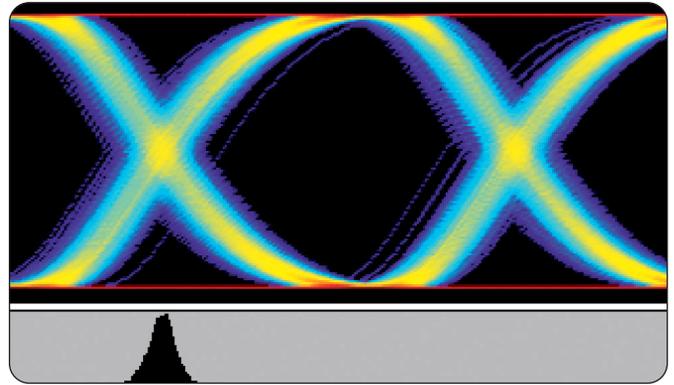


► **Figure 4.2.1a** Gaussian Distribution

4.2.1 Random Jitter

Random jitter is timing noise that cannot be predicted, because it has no discernable pattern. A classic example of random noise is the sound that is heard when a radio receiver is tuned to an inactive carrier frequency. While a random process can, in theory, have any probability distribution, random jitter is assumed to have a Gaussian distribution for the purpose of the jitter model. One reason for this is that the primary source of random noise in many electrical circuits is thermal noise (also called Johnson noise or shot noise), which is known to have a Gaussian distribution. Another, more fundamental reason is that the composite effect of many uncorrelated noise sources, no matter what the distributions of the individual sources, approaches a Gaussian distribution according to the central limit theorem.

The Gaussian distribution, also known as the normal distribution, has a PDF that is described by the familiar bell curve. An example of the distribution, with a mean value of zero and a standard deviation of 1.0, is shown in Figure 4.2.1a. This distribution is covered thoroughly in one of the references in the bibliography⁴, but its single most important characteristic is this: For a Gaussian variable, the peak value that it might attain is infinite. That is, although most samples of this random variable will be clustered around its mean value, any particular sample could, in theory, differ from that mean by an arbitrarily large amount. So, there is no bounded peak-to-peak value for the underlying distribution. The more samples one takes of such a distribution, the larger the measured peak-to-peak value will be.

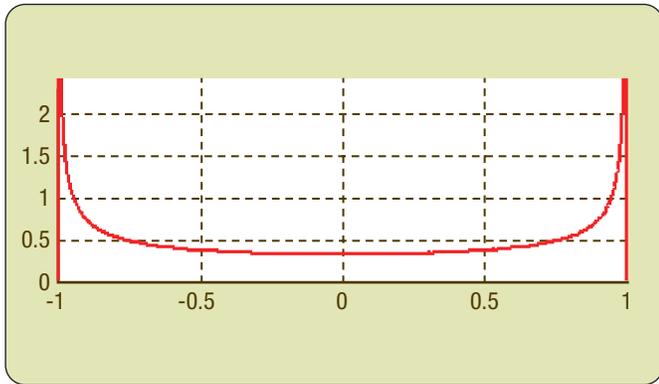


► **Figure 4.2.1b**

Frequently, efforts are made to characterize such a distribution by sampling it some large number of times and recording the peak-to-peak value that results. One should use caution with this approach.

The peak-to-peak value of a set of N observations of a random variable is itself a random variable, albeit one with a lower standard deviation. Using such a random variable as a pass-fail criterion for quality screening, for example, requires that the pass threshold be raised to account for the uncertainty in the measurement, resulting in some acceptable units being failed. A better approach is to fit the N observations to the assumed distribution (in this case Gaussian). Then, the mathematical description of the distribution can be used to predict longer-term behavior to a particular confidence level.

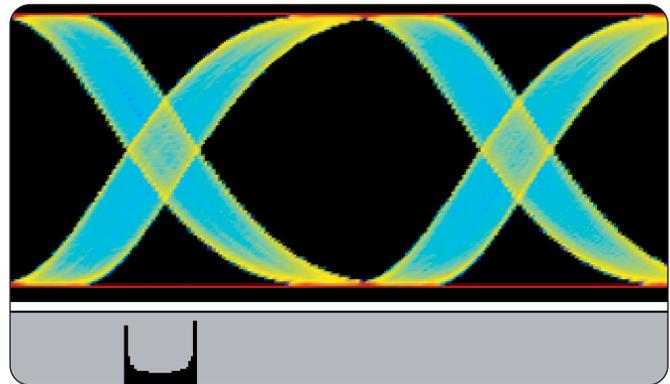
An eye diagram for a signal with only Gaussian jitter, together with the associated TIE histogram, is shown in Figure 4.2.1b.



► **Figure 4.2.3a**

4.2.2 Deterministic Jitter

Deterministic jitter is timing jitter that is repeatable and predictable. Because of this, the peak-to-peak value of this jitter is bounded, and the bounds can usually be observed or predicted with high confidence based on a reasonably low number of observations. This category of jitter is further subdivided in the paragraphs that follow, based both on the characteristics of the jitter and the root causes.



► **Figure 4.2.3b**

4.2.3 Periodic Jitter

Jitter that repeats in a cyclic fashion is called periodic jitter. An example was shown in Figure 2.4c, where the TIE time trend shows a repeating triangle wave. Since any periodic waveform can be decomposed into a Fourier series of harmonically related sinusoids, this kind of jitter is sometimes called sinusoidal jitter. The probability distribution for a sine wave with peak amplitude of 1.0 is shown in Figure 4.2.3a.

By convention, periodic jitter is uncorrelated with any periodically repeating patterns in a data stream. Jitter that is correlated with repeating data patterns is covered in the next paragraph.

Periodic jitter is typically caused by external deterministic noise sources coupling into a system, such as switching power-supply noise or a strong local RF carrier. It may also be caused by an unstable clock-recovery PLL.

An eye diagram for a signal with 0.2 unit intervals of periodic jitter, together with the associated TIE histogram, is shown in Figure 4.2.3b.

A Guide to Understanding and Characterizing Timing Jitter

► Primer

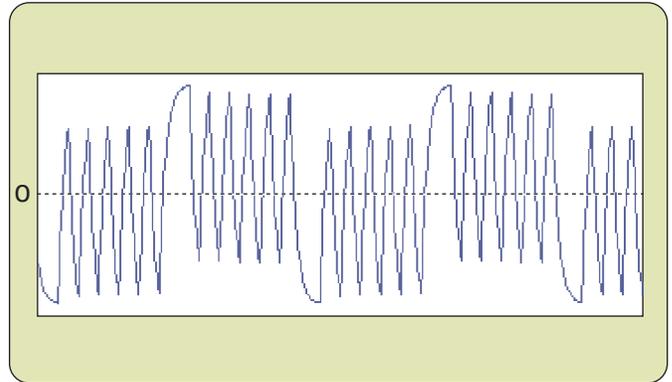
4.2.4 Data-Dependent Jitter

Any jitter that is correlated with the bit sequence in a data stream is termed Data-Dependent Jitter, or DDJ. DDJ is often caused by the frequency response of a cable or device. Consider the waveform in Figure 4.2.4a, in which a data sequence is strongly low-pass filtered. Because of the filtering, the waveform doesn't reach a full HIGH or LOW state unless there are several bits in a row of the same polarity.

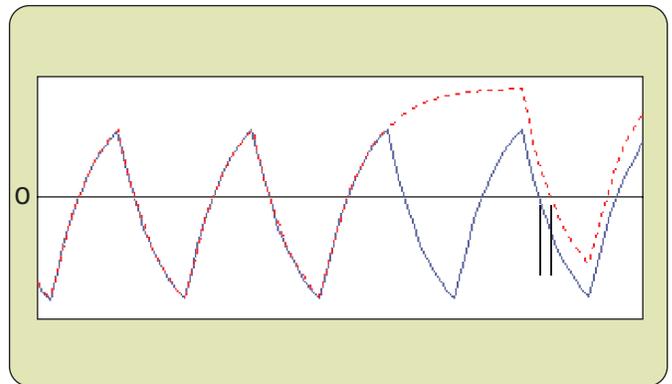
Figure 4.2.4b shows this waveform superimposed on an offset version of itself. You can see that an earlier threshold crossing occurs when a falling transition follows a 1,0,1,0,1,0,1 sequence than when following a 1,0,1,0,1,1,1 sequence.

Since this timing shift is predictable and is related to the particular data preceding the transition, it is an example of DDJ. Another common name is Inter-Symbol Interference, or ISI.

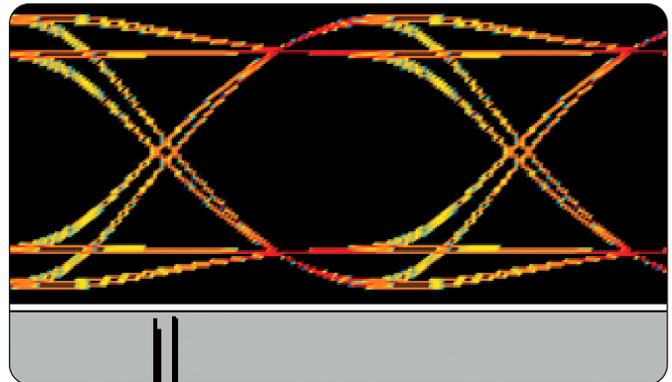
An eye diagram for a signal with 0.2 unit intervals of DDJ, together with the associated TIE histogram, is shown in Figure 4.2.4c. The histogram for ISI jitter always consists solely of impulses.



► Figure 4.2.4a



► Figure 4.2.4b



► Figure 4.2.4c

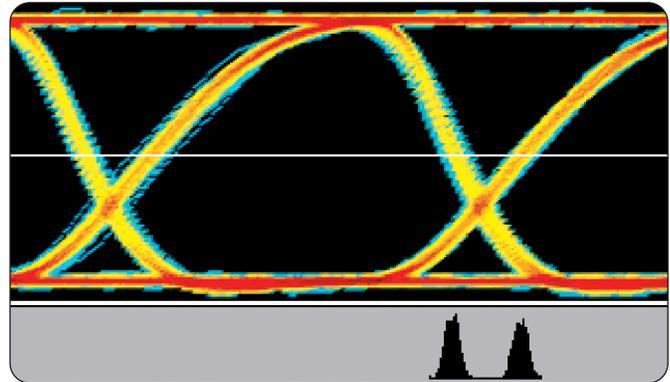
4.2.5 Duty-Cycle Dependent Jitter

Jitter that may be predicted based on whether the associated edge is rising or falling is called Duty-Cycle Jitter (DCD). There are two common causes of DCD:

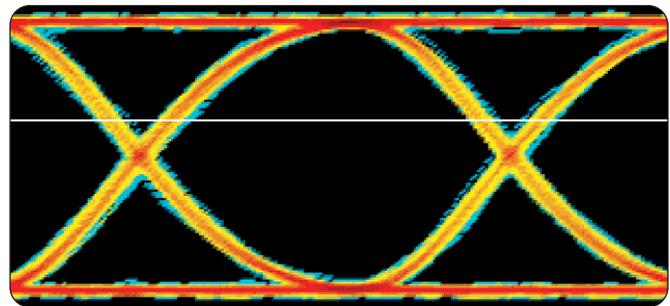
1. The slew rate for the rising edges differs from that of the falling edges.
2. The decision threshold for a waveform is higher or lower than it should be.

Figure 4.2.5a is an eye diagram demonstrating the first case. Here, the decision threshold is at the 50% amplitude point but the slow rise time of the waveform causes the rising edges to cross the threshold later than the falling edges. As a result, the histogram of an edge crossing (gray) shows two distinct groupings. (In contrast to the prior eye diagrams, the waveform in this example has some Gaussian noise in addition to the duty-cycle jitter.)

Figure 4.2.5b is an eye diagram showing the second case, in which the waveform has balanced rise and fall times but the decision threshold is not set at the 50% amplitude point. The edge-crossing histogram would look very much like that from Figure 4.2.5a.



► **Figure 4.2.5a**



► **Figure 4.2.5b**

A Guide to Understanding and Characterizing Timing Jitter

► Primer

4.3 Putting It All Together

The paragraphs above included eye diagrams and TIE histograms for each of the jitter types individually. But what happens to the histogram when more than one jitter type is present at the same time? A useful and applicable result from statistical theory can be paraphrased as follows: If two or more random processes are independent, the distribution that results from the sum of their effects is equal to the convolution of the individual distributions.

A visual example of this concept is much more compelling than the verbal description. Figure 4.3a depicts what happens when duty-cycle jitter (with a histogram consisting of two impulses) is combined with random jitter (with its Gaussian distribution). Another example was shown earlier in Figure 3.2d, where the uniform distribution of a triangle wave (a form of periodic jitter) was combined with random jitter.

A knowledgeable person could look at the final histogram from Figure 4.3a and infer what two types of jitter had contributed to it. However, real examples may involve all the jitter types in various amounts, resulting in a composite histogram that is much less intuitive. One of the objectives of a thorough jitter analysis is to identify all the individual jitter components that contributed to the final result.

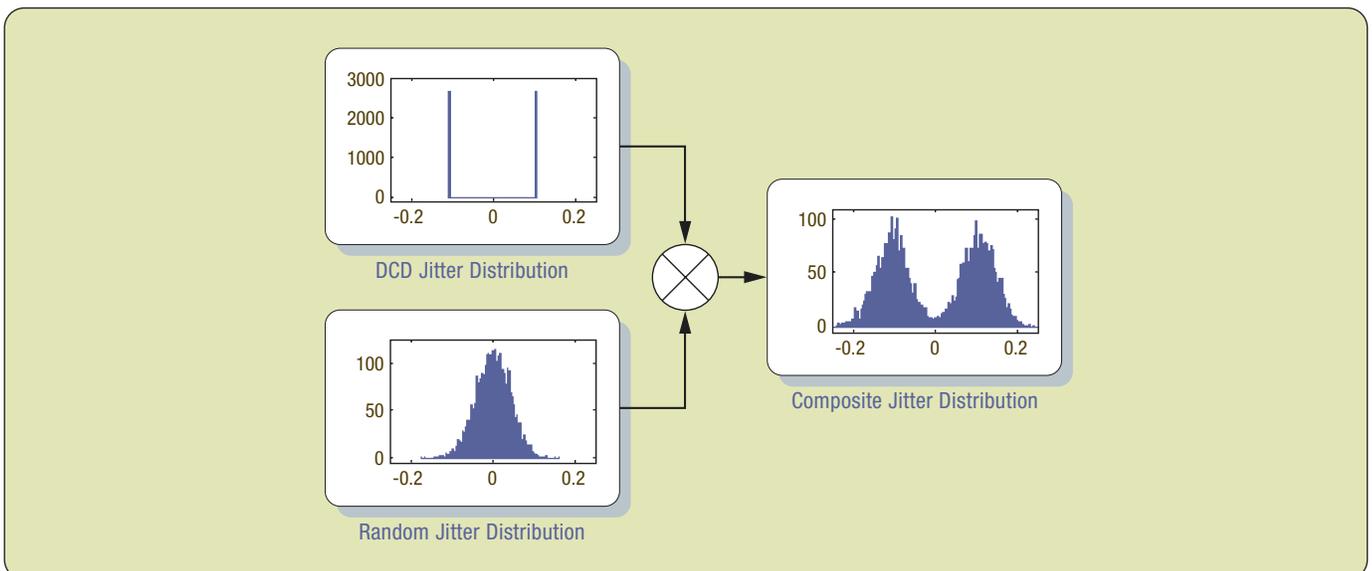
The real power of jitter separation comes when the accurately identified jitter components are inserted back into the jitter model outlined in Figure 4.2 and the model is used to predict how a system will behave in new circumstances. This is the basis for Bit Error Rate estimation, which is covered in more detail in Section 5.

To understand why accurate analysis tools are important, we compare two systems with different jitter characteristics. The first system, identified as System I, has only random jitter with a standard deviation of 0.053 unit intervals (UI). A TIE histogram for this system, which has a data rate of 1062.5 Mbps, is shown in Figure 4.3b.

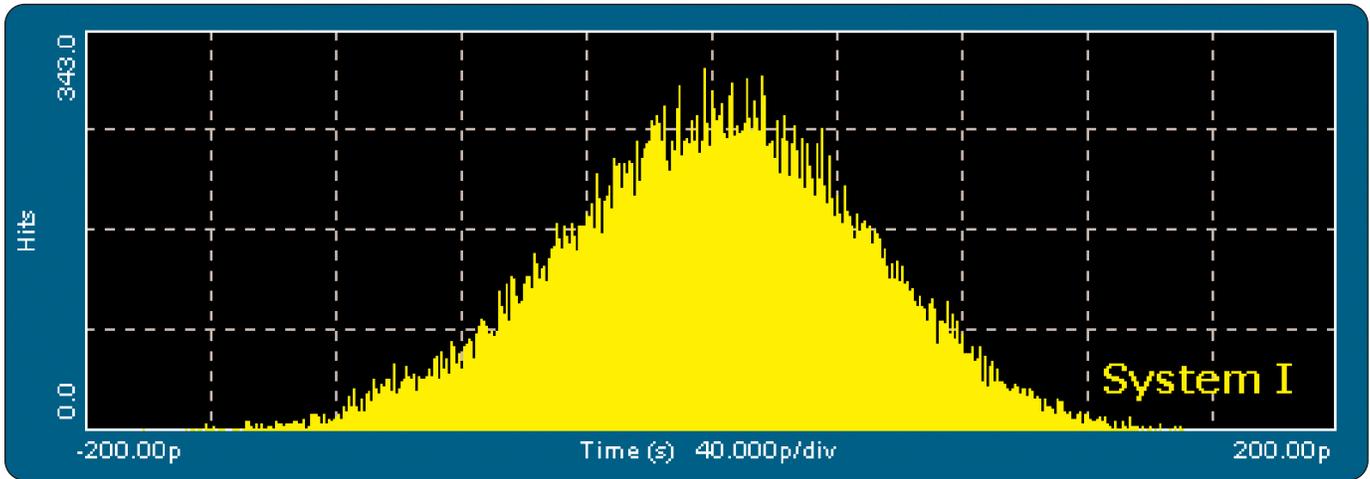
The second system has random jitter with a standard deviation of 0.028 UI, about half that of the first system. However, it also has two uncorrelated components of periodic jitter, each with a peak-to-peak amplitude of 0.14 UI. Figure 4.3c shows the TIE histogram for System II.

The two histograms, which contain identical populations of about 42,000 edges, don't look remarkably different. In fact, for these two sample sets the peak-to-peak value of the time interval error is exactly the same: 430 ps (0.457 UI). If visual comparison were the only tool available, distinguishing between these two cases would be difficult. Accurately predicting how the systems might behave over a longer observation interval would be impossible.

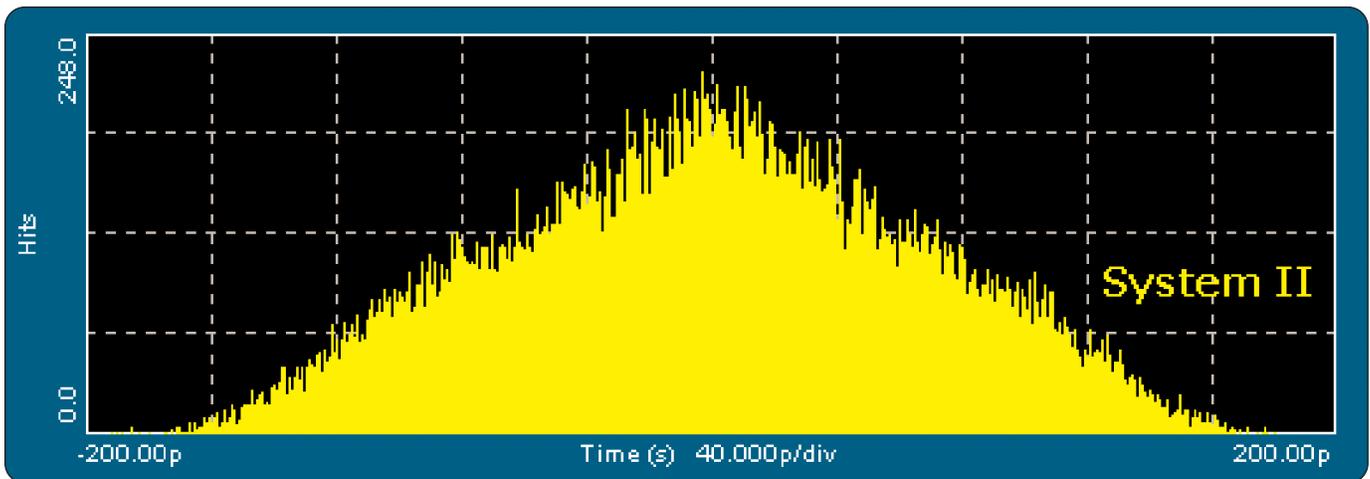
One might ask whether it is important to distinguish between these cases, since they look so similar. We shall return for a second look at this example after a review of Bit Error Rates and Bathtub curves.



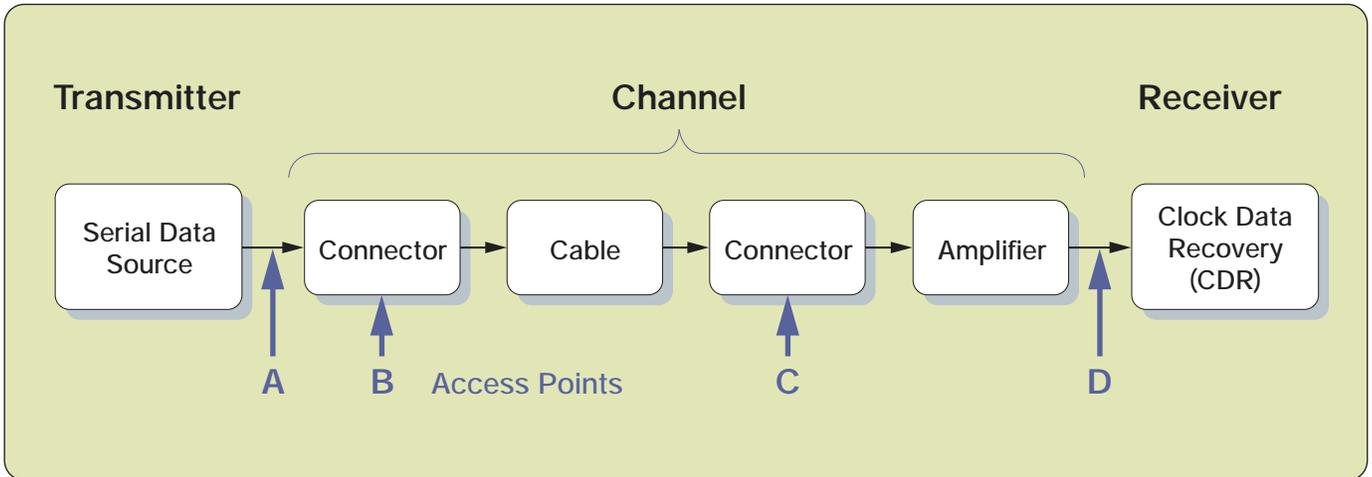
► **Figure 4.3a**



▶ Figure 4.3b



▶ Figure 4.3c



► **Figure 5.1**

Section 5: Jitter vs. Bit Error Rate

“An approximate answer to the right question is worth a good deal more than the exact answer to an approximate problem.”

John Tukey

5.1 The Jitter Budget

A real communications system consists of transmitter, channel and receiver, as modeled in Figure 5.1. The channel may include cabling, interconnects, and active devices that don't perform clock regeneration. The channel may introduce filtering, non-linearity, DC offset, impedance mismatches and additional random jitter. Various circuit test points and system connectors will provide access points where the signal quality may be monitored. As one progresses from the transmitter to the receiver and monitors the signal quality at each access point (perhaps using an eye diagram), the signal jitter generally gets worse.

Simplistically, one might assume that success has been achieved if the eye is still open, even just a little, when the signal reaches the input to the receiver's clock-data recovery circuit (access point D). However, the receiver circuit isn't perfect. The received signal's eye must be sufficiently open horizontally to account for timing uncertainty in the receiver's decision circuit, and open enough vertically to accommodate noise influencing the decision threshold.

Thus, a thorough system design specification will allocate a jitter budget so that each well-defined access point has known jitter limits, and that adequate margin remains when the signal enters the receiver to be resynchronized.

5.2 The Bathtub Curve

The concepts of eye diagrams and eye openings were discussed in Section 3.5. The Gaussian probability distribution, with its theoretically unbounded peak-to-peak value, was covered in paragraph 4.2.1. Considering these two topics together leads to an interesting thought: For any signal that contains some Gaussian jitter, the eye diagram should close completely if you accumulate samples for a long enough time. This would render the concept of eye opening useless as a basis for comparison. Fortunately, the usefulness of the eye diagram is restored if a confidence level is applied to the eye opening.

Consider Figure 5.2a, in which a green ruler 0.5 unit intervals long has been placed horizontally in the center of the eye. Suppose that it is regarded a failure if any waveforms cross this ruler, either rising or falling. In the figure, it appears that this ruler has not been crossed by any waveforms yet, but such a crossing is inevitable if waveform samples continue to accumulate and the signal contains some Gaussian jitter.

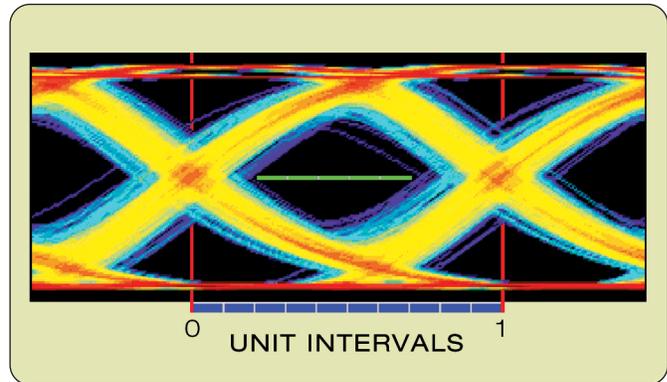
Now suppose instead that the test was considered successful if no more than one waveform out of every 1000 waveforms crossed over the ruler. It would no longer matter how long the test ran. If 50,000 waveforms were allowed to accumulate and 50 or less crossed over the ruler, the test would have passed. One could say that, aside from one waveform in 10^3 , the eye was 50% open. Since each crossing is assumed to represent a bit error, this would be a bit error rate (BER) of 10^{-3} .

If the same signal were tested with a shorter ruler, say, 0.25 unit intervals long, then the ruler would certainly be crossed less frequently. Perhaps only one waveform out of every 100,000 waveforms, on average, would cross this shorter ruler. One could say that, aside from one waveform in 10^5 , the eye was 25% open.

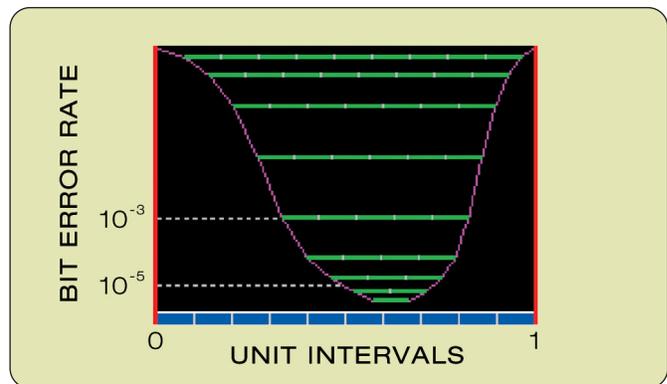
Continuing along these lines and using a series of rulers, one could fully characterize the eye opening versus the bit error rate. (Note that each ruler should be allowed to slide left or right to obtain the best possible fit.) If the rulers are all plotted against their corresponding bit error rates on a single chart, with the ends of the rulers connected, a plot something like Figure 5.2b results.

A plot of this description is called a Bathtub Plot, since the pink lines can be imagined to look like a bathtub. Using such a figure, one can tell what horizontal portion of the eye will remain completely free of signal transitions, for a given confidence level.

Finally, note that it can take a long time to accumulate enough data to directly measure the eye openings near the bottom of the chart. For this reason, the mathematical model covered in Section 4 can be used to predict performance on the basis of a much smaller sample set.



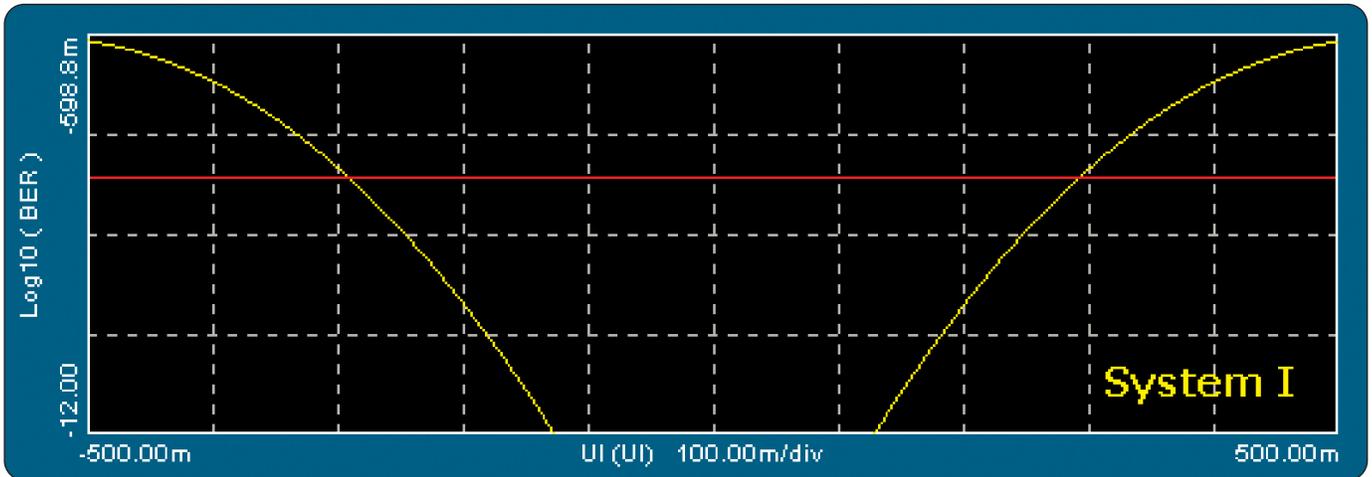
► Figure 5.2a



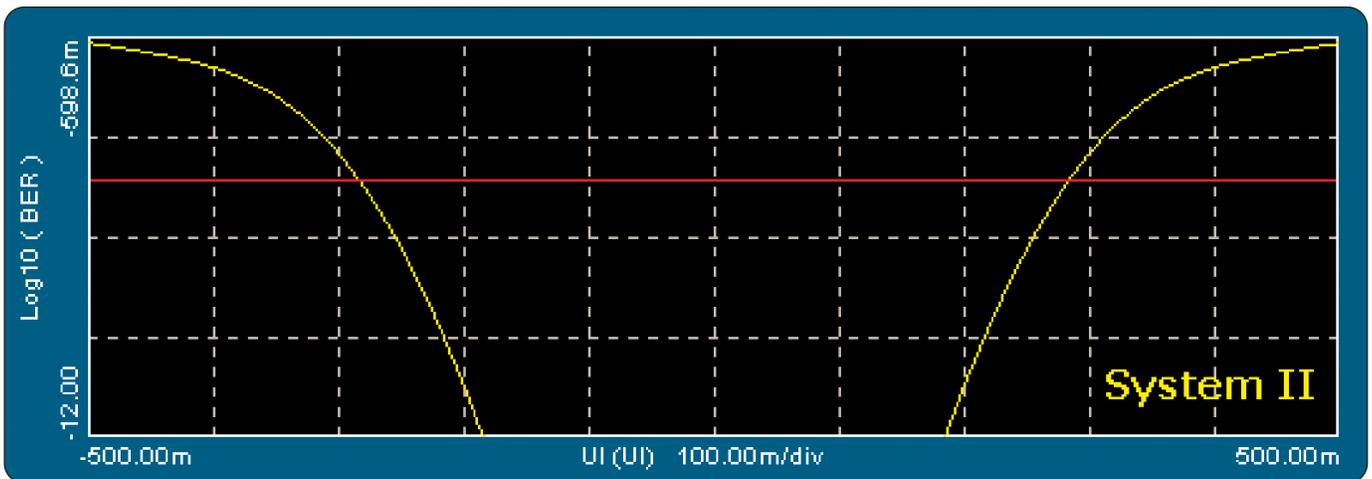
► Figure 5.2b

A Guide to Understanding and Characterizing Timing Jitter

► Primer



► Figure 5.3a



► Figure 5.3b

5.3 A BER Example

Section 4.3 presented an example in which two systems with different jitter characteristics had similar TIE histograms and identical peak-to-peak jitter for a sample size of about 42,000 edges. (See Figures 4.3b and 4.3c.) We now return to that example, and show the associated bathtub curves. Figure 5.3a shows the bathtub curve for System I, and Figure 5.3b shows the plot for System II.

Since both of these plots represent data sets of 42,000 observed edges, a red cursor has been placed on each plot corresponding to a BER of one part in 42,000, or $10^{-4.63}$. At these levels, both of the systems show an eye opening of about 58%, and the eye opening for System I is actually greater than that for System II. However, at a BER of $1e-12$ (which is a frequent specification point for serial communication links) the eye for System I is only open 26%. The eye for System II is open 37%, or about half-again as much as System I. This is a significant difference, and could easily determine whether a system meets its requirements or not.

Section 6: Summary

Timing jitter has always degraded electrical systems, but the drive to higher data rates and lower logic swings has focused increasing interest and concern on its characterization. Characterization is needed to identify sources of jitter for reduction by redesign. It also serves to define, identify or measure jitter for compliance standards and design specifications.

All jitter has random and deterministic components. Because of its random nature, care must be used in specifying acceptable limits for jitter. This is especially true when jitter performance at very high confidence levels is to be estimated based on modest amounts of data. To enable working with random components of jitter, a method that has found favor in jitter analysis and prediction is to use a mathematical model, the parameters of which are adjusted based on observed measurements. The model can then be used to predict performance in other situations. The model can also give insight into the specific causes of the jitter, and therefore aid in understanding how to reduce it.

The oscilloscope has been a traditional tool for observing jitter, using such techniques as histograms and eye diagrams. When augmented by back-end processing that provides additional features like cycle to cycle measurements, trend and spectrum plots, data logging and worst case capture, the scope continues to be the tool of choice for characterizing timing jitter. When further augmented with sophisticated clock recovery, jitter separation algorithms and bit error ratio estimation, the scope becomes the only solution for characterizing and reducing timing jitter.

Many of the figures and examples in this document were taken from the Tektronix TDSJIT3 Jitter and Timing Analysis application, which runs on Tektronix TDS5000-, 6000-, 7000- and CSA7000 Series oscilloscopes to add a wide range of jitter analysis tools.

Appendix A: Glossary of Acronyms

BER	Bit Error Rate
CDF	Cumulative Distribution Function
DCD	Duty-Cycle Jitter
DDJ	Data-Dependent Jitter
DJ	Deterministic Jitter
FFT	Fast Fourier Transform
ISI	Inter-Symbol Interference
PDF	Probability Density Function
PJ	Periodic Jitter
PLL	Phase-Locked Loop
PSD	Power Spectral Density
RJ	Random Jitter
SJ	Sinusoidal Jitter
TIE	Time Interval Error
TJ	Total Jitter
UI	Unit Interval

References

- 1 Bell Communications Research, Inc (Bellcore), *"Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria, TR-253-CORE"*, Issue 2, Rev No. 1, December 1997
- 2 ITU-T Recommendation G.810 (08/96) *"Definitions and Terminology for Synchronization Networks"*
- 3 Papoulis: *"Probability, Random Variables, and Stochastic Processes, Second Edition"*, McGraw-Hill, 1984
- 4 Davenport and Root: *"An Introduction to the Theory of Random Signals and Noise"*, IEEE Press, 1987

Contact Tektronix:

- ASEAN / Australasia / Pakistan** (65) 6356 3900
- Austria** +43 2236 8092 262
- Belgium** +32 (2) 715 89 70
- Brazil & South America** 55 (11) 3741-8360
- Canada** 1 (800) 661-5625
- Central Europe & Greece** +43 2236 8092 301
- Denmark** +45 44 850 700
- Finland** +358 (9) 4783 400
- France & North Africa** +33 (0) 1 69 86 80 34
- Germany** +49 (221) 94 77 400
- Hong Kong** (852) 2585-6688
- India** (91) 80-2275577
- Italy** +39 (02) 25086 1
- Japan** 81 (3) 3448-3010
- Mexico, Central America & Caribbean** 52 (55) 56666-333
- The Netherlands** +31 (0) 23 569 5555
- Norway** +47 22 07 07 00
- People's Republic of China** 86 (10) 6235 1230
- Poland** +48 (0) 22 521 53 40
- Republic of Korea** 82 (2) 528-5299
- Russia, CIS & The Baltics** +358 (9) 4783 400
- South Africa** +27 11 254 8360
- Spain** +34 (91) 372 6055
- Sweden** +46 8 477 6503/4
- Taiwan** 886 (2) 2722-9622
- United Kingdom & Eire** +44 (0) 1344 392400
- USA** 1 (800) 426-2200
- USA (Export Sales)** 1 (503) 627-1916

For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

Updated 20 September 2002

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2003, Tektronix, Inc. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

10/03 ID/BT

55W-16146-1

