Jitter Measurements in Serial Data Signals

Michael Schnecker, Product Manager
LeCroy Corporation

Introduction

The increasing speed of serial data transmission systems places greater importance on measuring jitter with higher accuracy. Serial data standards normally require operation at an expected error rate of 10^{-12}. While this represents only one bit error every 2.3 hours at 100 Mb/s, it translates to one error every 4 minutes at 3Gb/s. It is, therefore, important to understand the characteristics of the jitter in order to maintain system performance.

Jitter is characterized by the relative variation in the location in time of the transitions of the signal level across a specific level (Figure 1). For clock signals, the relative time between threshold crossings (rising-to-rising or falling-to-falling) is measured. Data signals, on the other hand, generally require the measurement of the relative positioning of the data signal to that of the sampling clock, which is related to setup-and-hold time. Because of its random nature, jitter is normally described in terms of its probability density function or PDF.

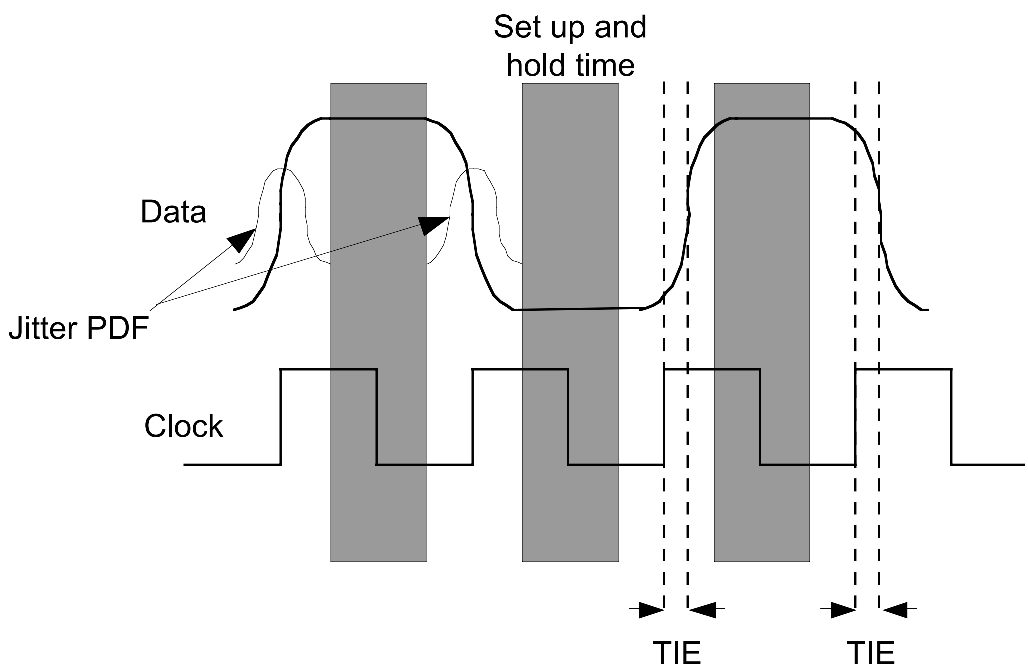

Figure 1 Setup and hold time requirement for error-free operation. Data transitions within the setup-and-hold time (gray area) will result in bit errors. Time interval error (TIE) is the time difference between clock and data edges, and the PDF of TIE is a measure of the probability of an edge occurring during the setup-and-hold time.

The processes that make up jitter are complex and come from many different random and non-random (deterministic) sources. The PDF of the jitter is the convolution of all individual component PDFs. Measurements are able to estimate the jitter PDF, but are not able to determine the distributions of the random and deterministic parts of the overall distribution. The lack of exact measurements for the jitter distributions of Rj and Dj has lead to the use of a simplified model for the total jitter. Equation 1 describes this model, which was first presented in the Fibrechannel MJSQ document.

$$Tj = \mathrm{N}(BER)*Rj + Dj \qquad \text{(Eq. 1)}$$

Equation 1 is a heuristic that describes total jitter as a function of bit error rate (BER) and is related to a distribution consisting of a Gaussian convolved with a pair of impulses, as shown in Figure 2. The constants Rj and Dj represent all of the components of random and deterministic jitter. The function N(BER) is the total peak-to-peak jitter of a unit normal distribution (i.e., a Gaussian with zero mean and a standard deviation of 1) at the specified bit error rate. The process of determining Rj and Dj involves finding the "best fit" values that solve equation 1. There are many possible ways to fit Rj and Dj to Equation 1, and since it is a simplification, no single set of solutions can completely describe the behavior of actual jitter completely. It is for this reason that the SDA uses two separate methods to measure Rj and Dj — effective and direct — and presents these to the user.

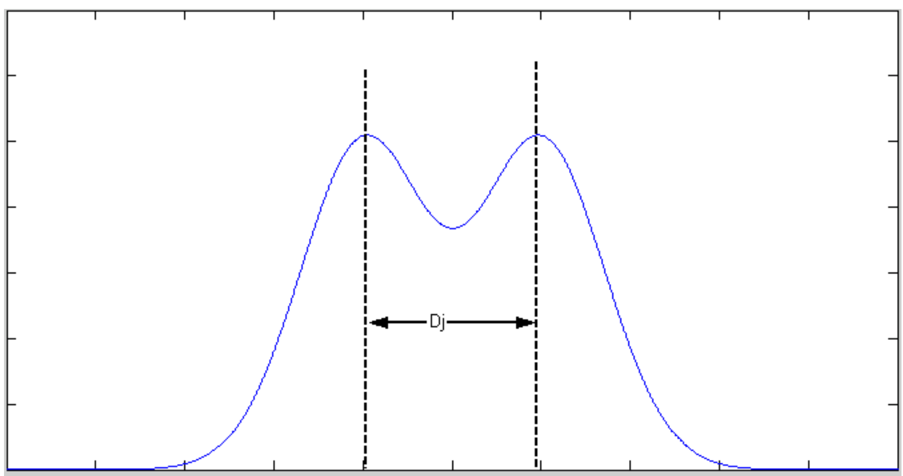

Figure 2 Jitter PDF model corresponding to the heuristic in Equation 1.
The random jitter is modeled by a Gaussian, and the deterministic jitter by a pair of impulses separated by the value of the parameter Dj. The curve shown is the convolution of Rj and Dj.

Bit Error Rate and Jitter

Equation 1 shows that the total jitter is a function of bit error rate. This relationship is based on the effect that jitter has on the bit error rate of a system. The bit error rate is influenced by other parameters of the system, such as noise, so it is not correct to say that BER and jitter are equivalent. It is the contribution to the overall bit error rate caused by jitter that is shown in Equation 1. A bit error will occur when the data signal transitions from one state to another during the setup-and-hold time as shown in Figure 1. Since jitter has a random component, the location in time of the transitions varies over a range of values. The longer the transitions are observed, the greater this range will be. Now, if we think of each transition in the data signal as the change in a bit value, then a transition at the wrong time (i.e., outside the setup-and-hold window) will lead to a bit error. The probability of this event is equivalent to the bit error rate contribution due to jitter. The total jitter gives a confidence interval for the jitter in that it will not exceed a certain value to a confidence of (1-BER). In many specifications, the term "bit error rate" is commonly used in this context to refer to the jitter confidence interval.

Total Jitter

The total jitter is the peak-to-peak jitter in a clock or data signal within a specified confidence equal to 1-BER. An example of a normally distributed jitter PDF is shown in Figure 3. In order to determine the total jitter from the PDF, the probability of the jitter exceeding a certain value t must be evaluated. This is done by integrating the PDF from a time t to $+\infty$ which will give the total probability of an edge occurring at or after this time. The probability can be computed for all values of t by integrating the PDF separately for t > 0 and t < 0.

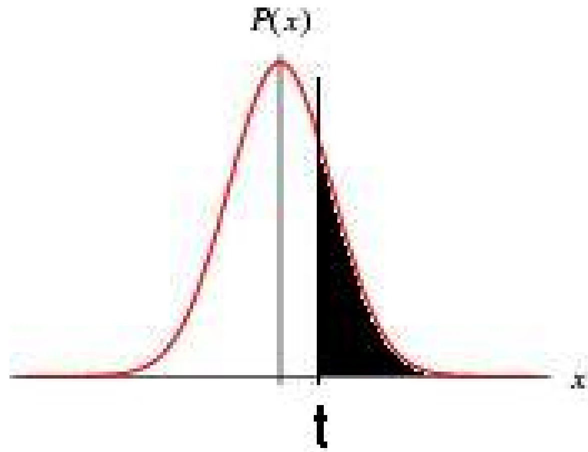

Figure 3 Probability of a data edge displacement greater than time *t* from the sampling clock. The mean value of the distribution is 0, which represents perfect alignment.

Figure 4 Total jitter curve. The vertical values of this curve represent the probability of a data transition occurring at a time represented by the horizontal axis. The horizontal center of the plot is 0 ps. The two markers are placed at the vertical level corresponding to a bit error rate of 10e-12, and the horizontal distance between these two points is the total jitter at this bit error rate.

The resulting curve shown in Figure 4, gives the total probability of an edge being greater than *t* (or less than –*t*). The contribution of jitter to the system BER is given by the probability that an edge occurs at a time greater than *t* as we mentioned earlier. In order to guarantee a BER contribution from jitter below a certain value, the positive and negative values of *t* are chosen so that the probability of an edge at a time greater and less than these times is equal to the desired bit error rate. These jitter values can be measured by finding the intersection between a horizontal line at the bit error rate and the total jitter curve. The horizontal spacing between these two points is the total jitter.

A common way to view the total jitter is by plotting the bit error rate as a function of sampling position within a bit interval. This curve, commonly referred to as the "bathtub" curve is derived from the total jitter curve by scaling it to one bit interval (UI). The right half of the bathtub curve is taken from the left half of the total jitter curve, and the left half of the bathtub curve is taken from the right half of the total jitter curve. The bathtub curve corresponding to the total jitter curve in Figure 4 is shown in Figure 5.

Extrapolating the PDF

Measuring the total jitter requires that the probability density function of the jitter is known exactly. The SDA measures the jitter PDF by collecting a histogram of TIE measurements. This histogram approximates the PDF by counting the number of edges occurring within the time period delimited by each bin in the histogram. In order to accurately measure jitter contributions at very low bit error rates, such as 10^{12}, the histogram must contain measurements with populations that are below 1 in 10^{16} (one TIE measurement out of 10^{16} at a certain value). This number of data transitions would take approximately 38 days at 3 Gb/s. Therefore, measuring this number of edges is clearly impractical.

Figure 5 The bathtub curve is constructed by rescaling the total jitter curve in Figure 4 to one unit interval, and centering the right side of the total jitter curve at 0 UI and the left side at 1UI (the left and right sides of the bathtub curve).

A smaller data set is extrapolated in order to estimate the data for the larger sample size. Extrapolation of the measured histogram of TIE values uses the random nature of jitter at the extremes of the histogram to extrapolate the bins below the 10th percentile and above the 90th percentile. The central part of the distribution is dominated by deterministic jitter while the extremes are entirely random. The bins in these ranges behave as a normal distribution as the jitter range is increased, that is, their populations fall as $\exp(-t^2)$.

Taking the logarithm of the histogram makes this relationship quadratic. The extrapolation, therefore, is simply a quadratic curve fit to the extremes of the log of the jitter histogram (Figures 6 and 7). The extrapolated histogram is used to compute the total jitter curve described above, and it is normalized so that the sum of the populations of all the bins is one. The integrals described above are implemented by summing the extrapolated histogram bins.

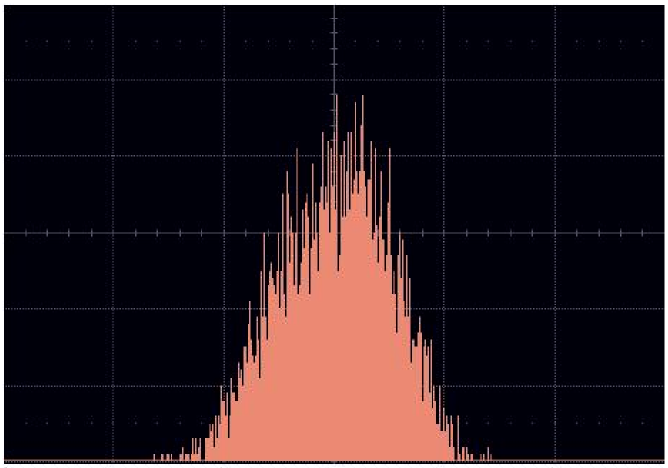

Figure 6 The measured histogram of TIE values is extrapolated by fitting curves to the bins below the 10th and above the 90th percentile. The log of the histogram is used to simplify this process to a quadratic fit.

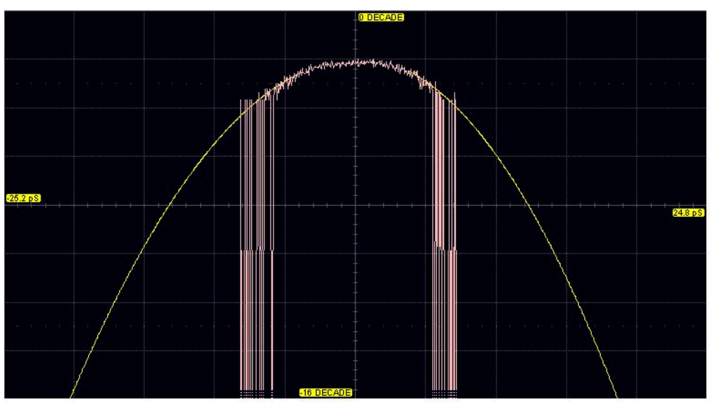

Figure 7 Logarithm of the measured TIE histogram superimposed on the extrapolated curve (in yellow). The extrapolation uses a quadratic curve fit to the histogram bins at the extremes.

Separating Rj and Dj – Two methods

The total jitter curve is the basis for estimating the magnitude of Rj and Dj. Since the total jitter curve is derived directly from the signal under test, its value is the most accurate representation of the jitter for a given bit error rate. There are basically two ways of separating the random and deterministic jitter. The first method, which models the growth of total jitter as BER is decreased, leads to the effective jitter parameters Rje and Dje. These values are effective in the sense that they provide an equivalent total jitter model for low bit error rates. Starting with the total jitter curve, the growth in the total jitter as a function of decreasing BER is plotted. The curve described by Equation 1 is fitted to the measured curve by selecting the Rj (slope) and Dj (intercept) to minimize the error in the fit.

The second method of fitting Rj and Dj to the measured data is based on direct measurements of the deterministic jitter. Random jitter is the difference between this value and the total jitter at the selected bit error rate measured from the total jitter curve. This, of course, exactly matches the measured total jitter at the selected bit error rate, but is a poor predictor of the jitter for bit error rates below this level. The motivation behind employing this method is to better represent the contribution of deterministic jitter in the overall jitter at the specified bit error rate.

Each method of measuring Rj and Dj results in different values for the standard deviation and spacing between the Gaussian curves in the distribution in Figure 2. The total jitter at the specified bit error rate, however, is the same for either distribution.

Effective Random and Deterministic Jitter

The effective jitter components Rje and Dje represent the best fit values for Equation 1 to the behavior of the measured total jitter as the observation time is increased or, equivalently, the bit error rate is decreased. For a given bit error rate, the total jitter is measured from the width of the total jitter curve. The value of the total jitter as the bit error rate is decreased can be plotted as shown in Figure 8. The vertical axis of the left-hand plot is the log of bit error rate. The Gaussian nature of the jitter at the extremes of the distribution results in a total jitter that grows approximately linearly with the log of BER, as shown in the upper curve of the right-hand plot. The function N(BER) in Equation 1 represents the width of a normal distribution with a variance of one at a given confidence level equal to 1-BER. The lower curve of the right-hand plot shows the variation of N(BER) with the log of BER, which is approximately linear. The values of Rje and Dje are chosen so that the lower curve lies on top of the upper one. From Equation 1, it can be seen that Rje is a slope parameter, while Dje adjusts the intercept point.

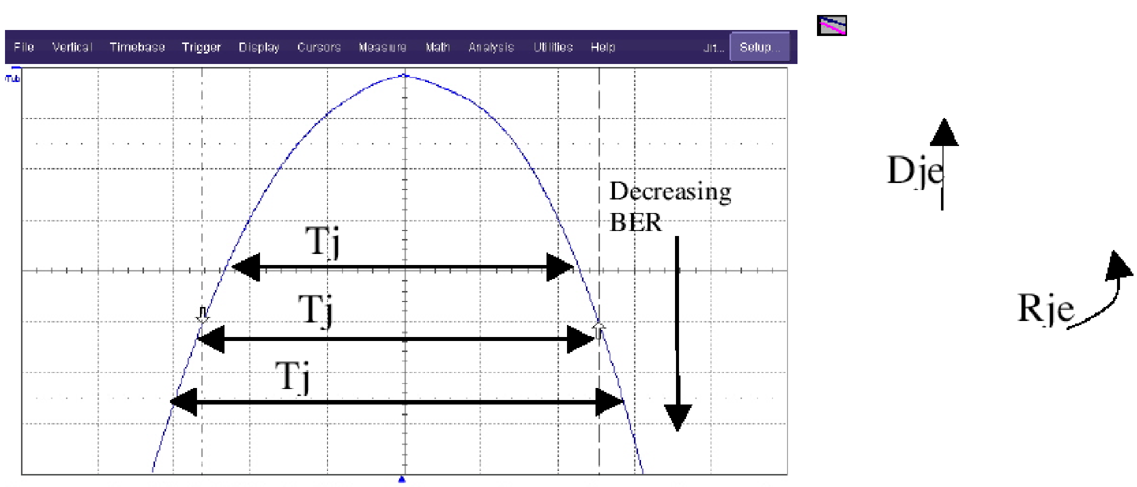

Figure 8 The change in total jitter with BER is represented by the values Tj_n in the total jitter curve on the left. These values are plotted vs. bit error rate in the upper line of the chart on the right, while the lower line shows the variation vs. BER for a pure Gaussian. Rje and Dje are chosen such that the curves line up.

The jitter computed using this method allows Equation 1 to accurately model the jitter behavior of systems as a function of bit error rate. This model is especially useful when computing jitter margins in system applications.

Direct Measurement of Deterministic Jitter

Deterministic jitter can be completely characterized by measuring the threshold crossing times of the data signal over a finite time period. The two classes of deterministic jitter are periodic and data dependent.

Data-dependent jitter is caused by system effects that are dependent on the data pattern. A common source of data-dependent jitter is the frequency response of the channel through which the serial data signal is transmitted. In this case, data patterns with many transitions, such as a 101010… pattern, contain more high frequencies in their spectrum than patterns containing fewer transitions (11001100… for example). The patterns with higher frequency content will be attenuated and phase shifted relative to the lower frequency patterns. In addition to data-dependent jitter, the rise and fall times of the data bits can be different. The detection threshold in the receiver is normally set to the 50% amplitude (midway between the '1' and '0' levels); therefore, unequal rise and fall times will generate jitter. This type of jitter is known as Duty Cycle Distortion (DCD).

The SDA uses a patent-pending method to measure both forms of data-dependent jitter. The method uses the history of a number of bits in the waveform to determine their effect on the transition of a given bit. The measurement uses a user-selectable number of bits (from 3 to 10). The acquired waveform is processed in segments the size of the selected number of bits. For example, if 5 bits are selected, segments 5 UI long are examined. For each segment, the value of the 5 bits is determined and each group of 5 bits is averaged with segments of similar value. When the entire waveform is scanned, a set of up to 32 (for 5 bits) waveforms are created. The averaging process removes all random jitter, noise, and periodic jitter from the segments. The waveform segments are overlaid by lining up the first transition of each of the segments and measuring the relative timing of the transitions to the last (5^{th} in this example) bit.

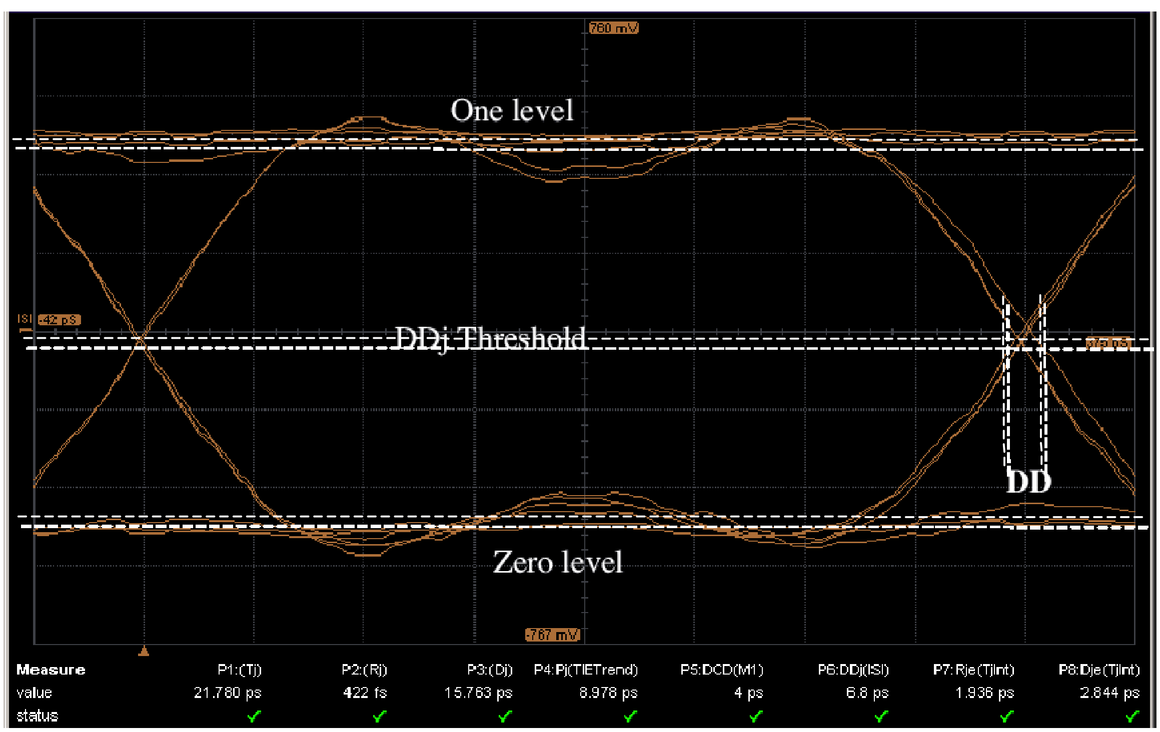

Figure 9 DDj measurement procedure. Averaged waveform segments for each pattern in the data stream are overlaid by lining up the first data transition. The curve above shows all of the transitions between

the next-to-last bit and the last bit on the right-hand side. The DDj is measured by examining the width of this crossing point at the selected threshold level.

Periodic jitter is measured by examining the spectrum of the trend of TIE values. The time interval error is measured for each edge in the data stream. Where no edges are present, as is the case for consecutive '1' or '0' values, edges are inserted at the expected data transition times. These inserted edges do not add any additional jitter since they are placed at the ideal edge locations for the data rate. The trend of TIE is continuous and the spectrum can be computed. Periodic jitter is the complex sum of the spectral components, excluding those associated with the repetition frequency of the data pattern and its harmonics.

The deterministic jitter is computed by adding up the periodic (Pj) and data-dependent (DDj) components. The random jitter is computed using Equation 1 and subtracting the measured deterministic jitter from the total jitter at the selected bit error rate:

Rj = (Tj(BER) – Dj)/_ (BER) (Eq. 2)

Comparing Models

Equation 1 is plotted for Rj and Dj computed from both methods along with the measured total jitter in Figure 9. The plot shows the bathtub curves for the measured values, as well as for both estimates. Viewed in this way, it is clear why both the effective and direct measurements for Rj and Dj are used. Both estimates arrive at the same total jitter at the specified BER (10^{-12} in this case) but they give different values of Tj at other BER values. The effective jitter values give a very accurate prediction of total jitter for bit error rates below about 10^{-10}, which is where they are fitted. The direct measurement underestimates the total jitter at error rates below the specified one, and overestimates the jitter above this. Note that the effective parameters underestimate the jitter for high bit error rates. The three curves in figure 9 below show the resulting bathtub curves from the measured signal (blue line) and the two models; Hj(BER) for the direct Dj method (red line) and Hje(BER) for the effective jitter method (green line).

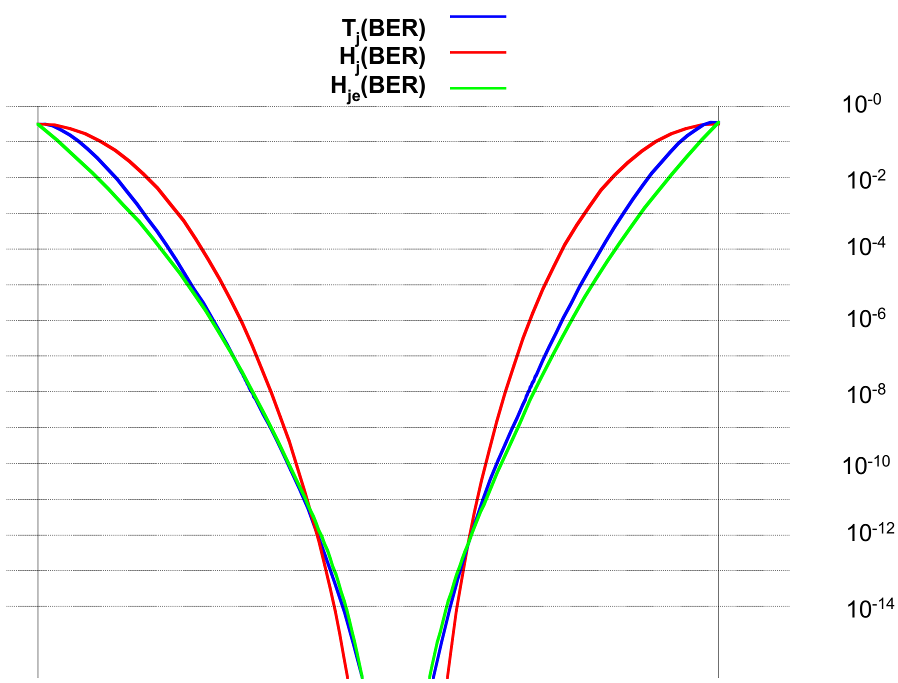

Figure 10 Measured jitter bathtub curve (blue curve) based on the extrapolation of the measured TIE histogram. The red and green curves are the estimated bathtub curves based on the direct Dj and effective measurements respectively. Note that the direct Dj method underestimates the total jitter below the BER at which it is computed.

Measurement Example

The following example is for a 3.125 Gb/s serial data signal. The signal was connected in line with 16 inches of FR4 backplane, which added approximately 60 ps of data-dependent jitter. The eye pattern shown in Figure 10 below shows the added jitter in the zero crossings. The histogram superimposed on the eye pattern shows the jitter, and the markers indicate the approximate peak-to-peak jitter caused by the backplane.

Figure 11 shows a jitter measurement of this signal. The Rj and Dj numbers listed below the grid are from the direct jitter measurement of deterministic jitter. The periodic jitter DCD and DDj are listed to the right of the Dj parameter. The DDj parameter indicates the 60 ps of jitter added by the backplane while, in this case, there is no significant periodic jitter. The total deterministic jitter is the sum of DDj and Pj. The effective deterministic jitter parameter, on the other hand, does not directly indicate the data-dependent jitter in the signal. The effective random jitter parameter is also larger than the value of Rj. This is so because the effective jitter parameters indicate the jitter at and below the specified bit error rate. The total jitter for either set of parameters is the same since they are both computed using Equation 1.

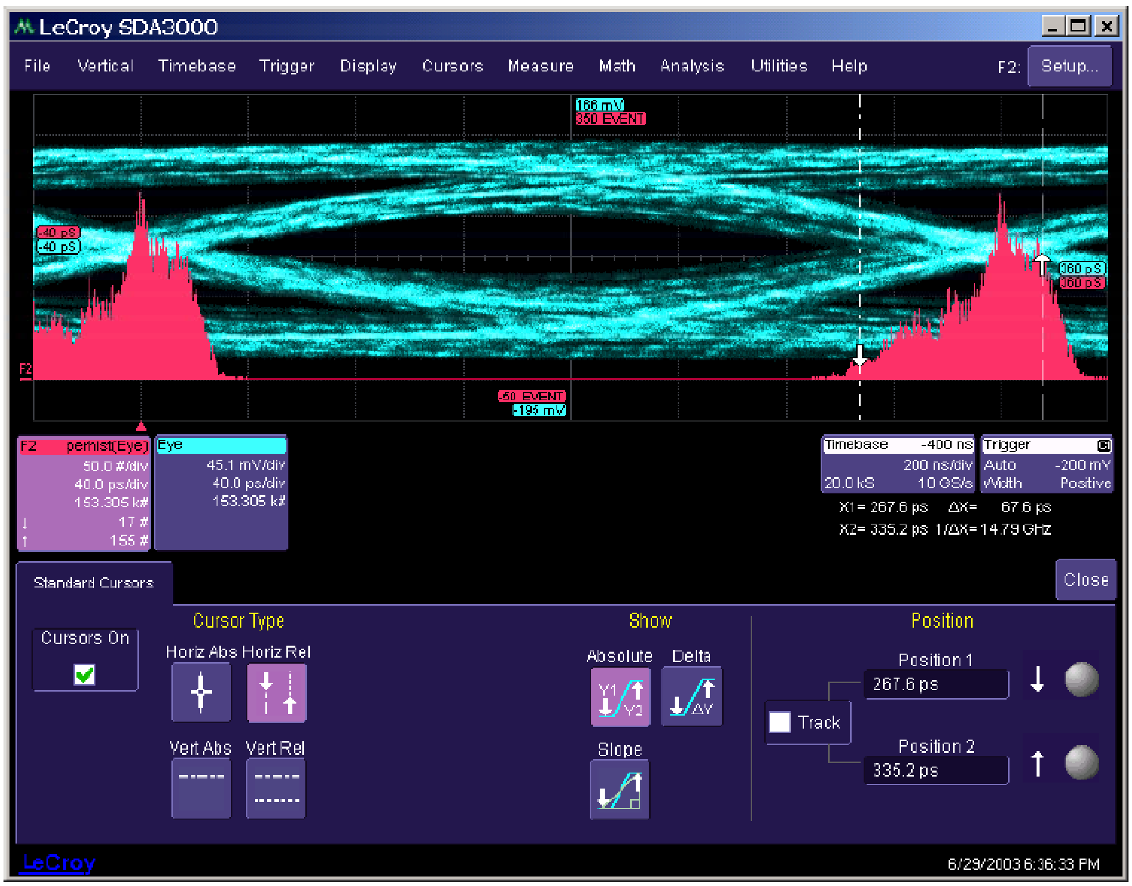

Figure 11 Eye pattern of 3.125 Gb/s data signal through 16 inches of backplane.
The superimposed histogram shows the jitter at the zero crossings.
The markers show the approximate size of the data-dependent jitter added by the backplane.

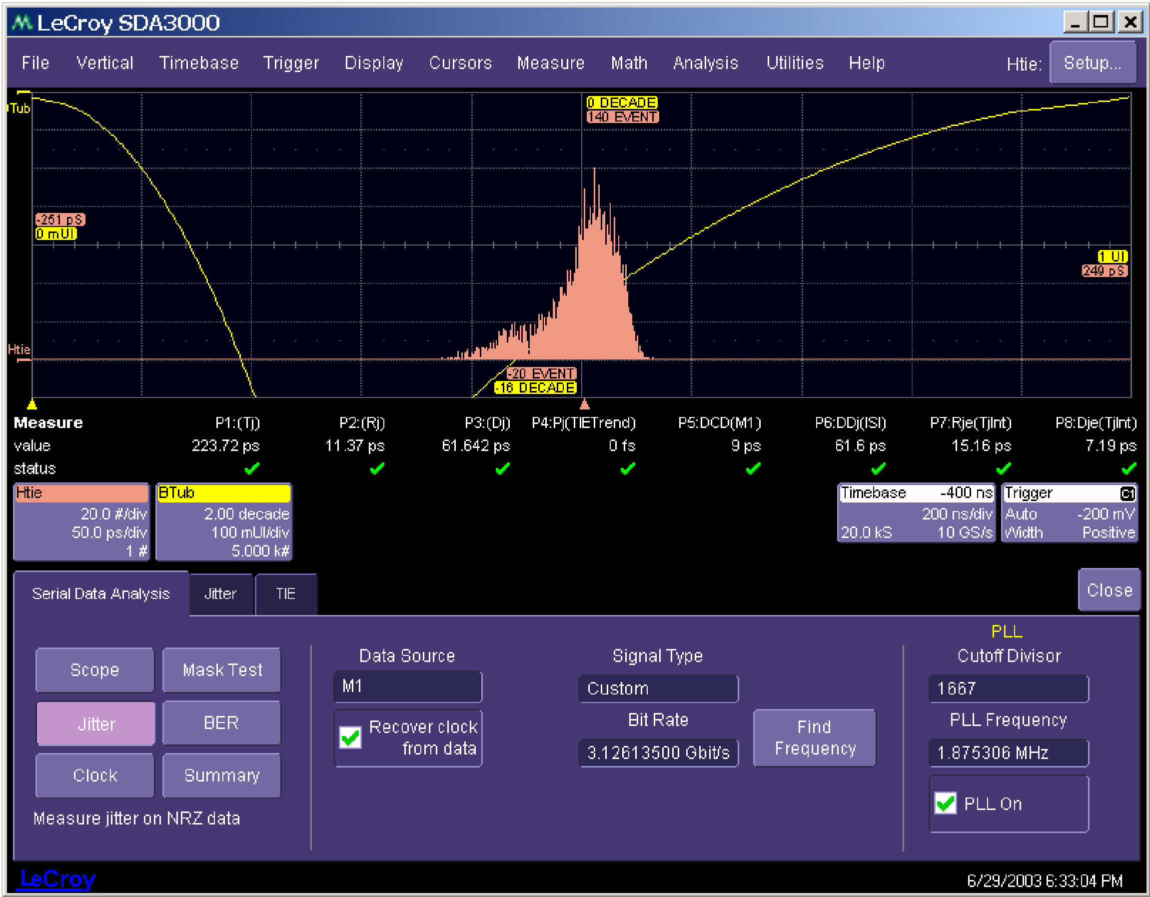

Figure 12 Jitter measurement of the 3.125 Gb/s data signal using both the direct and effective methods. Note that the DDj and Dj parameters agree with the value measured off the eye diagram. The effective Rj and Dj are displayed at the far right of the plot.

Conclusion

Jitter measurement is very important in high-speed serial data systems and detailed analysis of the jitter including the separation of random and deterministic components is essential to characterizing system performance. Because of the complex nature of the underlying processes that lead to jitter, a simplified model is used to describe them. Two methods exist for estimating the parameters Rj and Dj in this equation. Both of these measurements are useful, and each one gives separate information about the system under test. Both estimates are necessary to give a more complete picture of the total jitter for all bit error rates.