Jitter Separation - 50 Mb/s to Over 40 Gb/s Using the Agilent 86100C Infiniium DCA-J

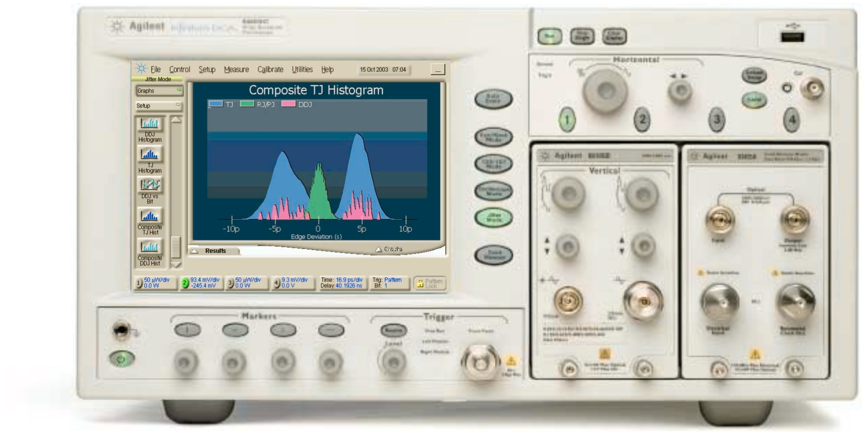

Abstract

Jitter is one of the most difficult measurement problems facing digital design and signal integrity engineers. As data rates increase, bit periods become very small, forcing allowable system level jitter to be as small as a few picoseconds and lower. Separation of jitter into its constituent components (random, determinisitic, data dependent, etc.) is a powerful tool for quickly diagnosing root causes of jitter problems. It allows system and component designers to precisely know the sources and causes of jitter leading to efficient troubleshooting when problems occur. Several key standards define jitter performance requirements based on specific values of constituent components. A number of tools exist for separation of jitter at bit rates up to approximately 3 Gb/s. However, many new designs are pushing well beyond 3 Gb/s.

This paper will provide a brief guide to making jitter measurements with the Agilent 86100C DCA-J. The hardware and firmware algorithms of the DCA-J and how they work together to perform jitter separation at virtually any data rate to 10 and even 40 Gb/s are described in detail. The procedure is straightforward and simple to perform. Results are easy to obtain and interpret. There is an extensive description of the significant changes that have been made to the classic equivalent time (wide bandwidth) oscilloscope that allow it to be an effective jitter measurement solution. Measurement limitations are reviewed.

Table of Contents

Table of Figures

Introduction

Virtually every high-speed communications design must deal with the issue of jitter. Definitions of jitter as well as its impact on system performance have been well documented but can be generally summarized as follows: When data are misplaced from their expected positions in time, receiver circuits will make mistakes in trying to interpret logic levels. Bit error ratio (BER) is degraded.

As speeds increase, jitter problems tend to be magnified. What might have been considered a small and tolerable time deviation at a lower data rate appears to be large and intolerable at high data rates. Consider that the bit period of a 10 Gb/s signal is only 100 picoseconds. When combined with signal impairments such as attenuation, dispersion, and noise, just a few picoseconds of timing instability can mean the difference between achieving or failing to reach BER objectives. The problem is further aggravated by the difficulty presented in making accurate measurements of jitter. A variety of approaches exist for making jitter measurements. While the various methods to measure jitter are well thought out and are based upon sound architectures, there has been frustration within the industry around the complexity of setting up a measurement, getting repeatable results, and perhaps more important getting the various techniques to agree with each other. The measurement problem is severely compounded when the 3 Gb/s data rate threshold is crossed. Above this rate, the list of solutions for jitter measurements becomes very short. The "equivalent time" sampling oscilloscope, with over 80 GHz of bandwidth and extremely low levels of intrinsic jitter, is a candidate for jitter measurements at very high data rates. However, some fundamental limitations of the wide-bandwidth oscilloscope have historically prevented it from being more than just a coarse jitter measurement tool. These limitations have now been overcome with the new Agilent 86100C DCA-J. Patented architectural changes have resulted in a flexible and easy to use jitter solution based upon a familiar laboratory instrument. Perhaps most important is that this tool provides thorough and accurate jitter analysis at rates from 50 Mb/s to over 40 Gb/s.

Basic procedure for making jitter measurements with the 86100C DCA-J

The procedure for making a jitter measurement using the 86100C is simple and straightforward.

The signal to be measured is connected to the input channel appropriate for the speed of the signal. As a simple rule, the bandwidth should be at least double that of the data rate. For NRZ signals, this will easily pass the third harmonic of the fundamental frequency (which is half of the data rate). For example, the 86112A, the lowest bandwidth electrical channel for the 86100C has a specified bandwidth of 20 GHz with a typical bandwidth in excess of 26 GHz. Thus it can be used for data rates up to and including 10 Gb/s. A variety of other plug-ins is available with optical or electrical channels with up to 65 and 80 GHz bandwidths respectively. Channel bandwidths are adjustable. Verify that the channel bandwidth setting is set correctly. Select **Setup**, **Channels**, **Channel #**, **Advanced >>**, and set the desired bandwidth.

The second requirement is to provide a timing reference as a trigger signal. This must be a clock signal at the data rate, or a divided clock. The allowable clock divisors are 2^N, 5, 10, 20, and 25. Thus for a 10 Gb/s signal, the allowable clock frequencies are 10 GHz, 5 GHz, 2.5 GHz, 1.25 GHz and so on as well as 2 GHz, 1 GHz, 500 MHz and 400 MHz. Whether full rate or divided, this clock must be synchronous with the data signal.

Although not required, it is often a good practice to verify that the signal is visible on the instrument display. This provides a quick check to see that the signal amplitudes and trigger signal are valid for the instrument. Press the **Eye/Mask** mode key and the **Autoscale** key. In one or two seconds, an eye diagram should be visible on the screen. Once the signals have been verified, jitter measurements are activated through pressing the **Jitter Mode** key on the front panel. (Note that the autoscale and Eye/mask steps mentioned above are not required, but are only signal quality verifications. If the signal and trigger are known to be good, the entire measurement setup is achieved by simply pressing the **Jitter Mode** key).

At this point the instrument will automatically perform all the necessary steps to complete the jitter measurement process. The instrument will not display results for two to three seconds as it collects data to detect the required signal parameters. After this, the jitter results page will appear on the display. This page will display graphical as well as tabular jitter results. If the pattern length of the input signal is not in the list of pattern lengths to automatically detect, it must be input manually via the **Trigger, Pattern Lock Setup** dialog (see below). Any pattern length can be added to the auto-detect list by selecting ‘Edit list’ in the pattern length pull-down in the same dialog (see below).

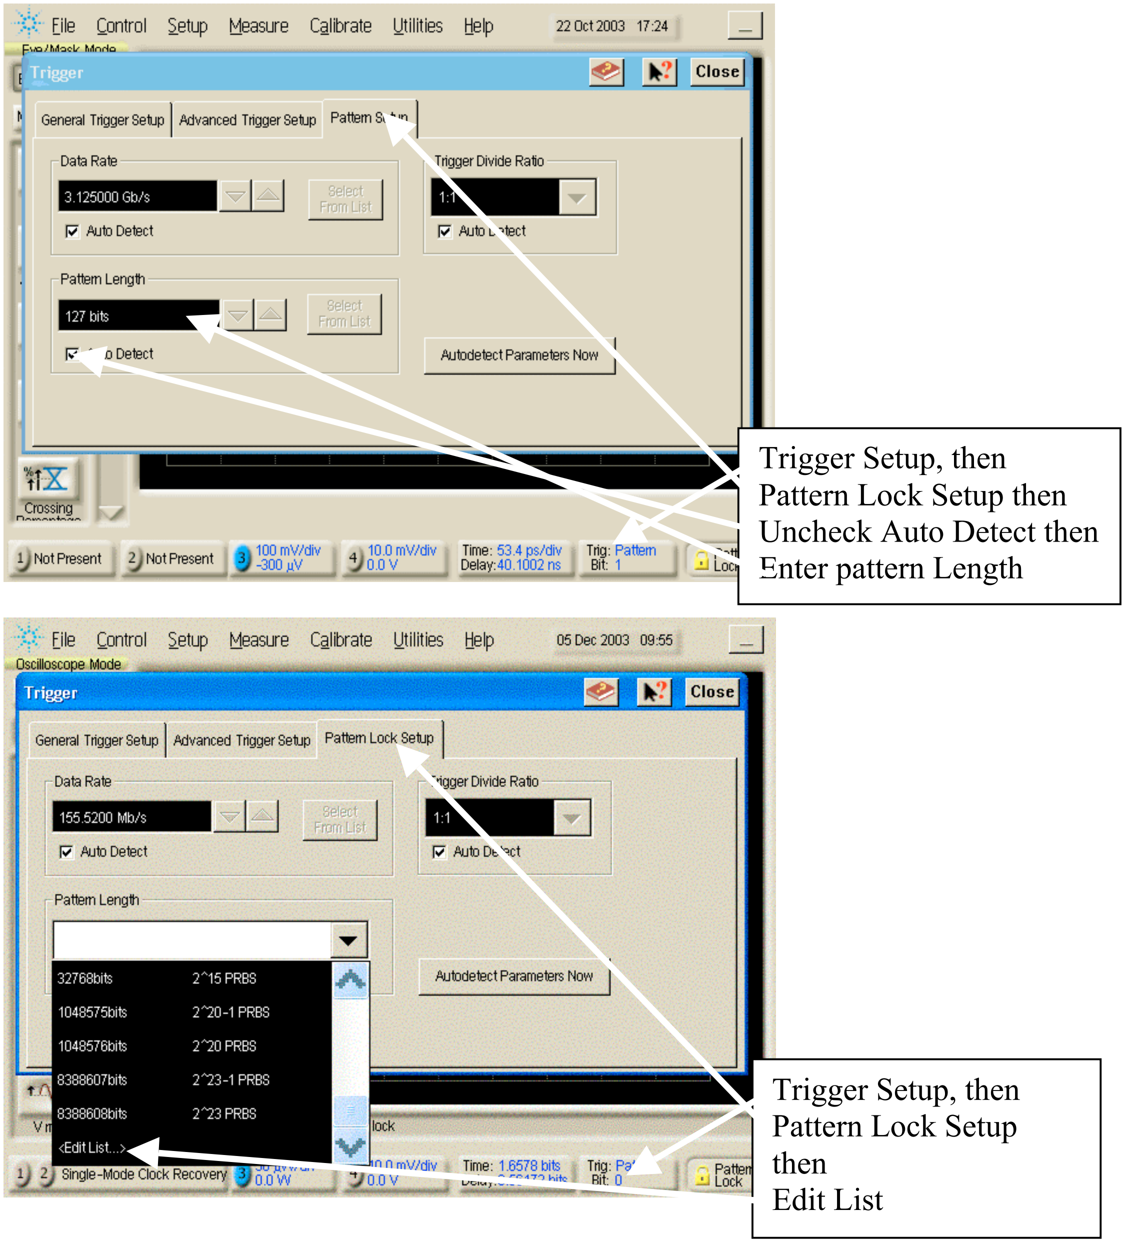

Figure 1. Pattern Lock trigger setup

There is a jitter setup menu available as one of the selections on the left side toolbar in jitter mode. This allows the user to reconfigure a variety of parameters including how data is presented as well as how data is acquired. In most cases, the default state of the instrument is valid, and jitter measurements are performed by selecting the jitter mode key alone.

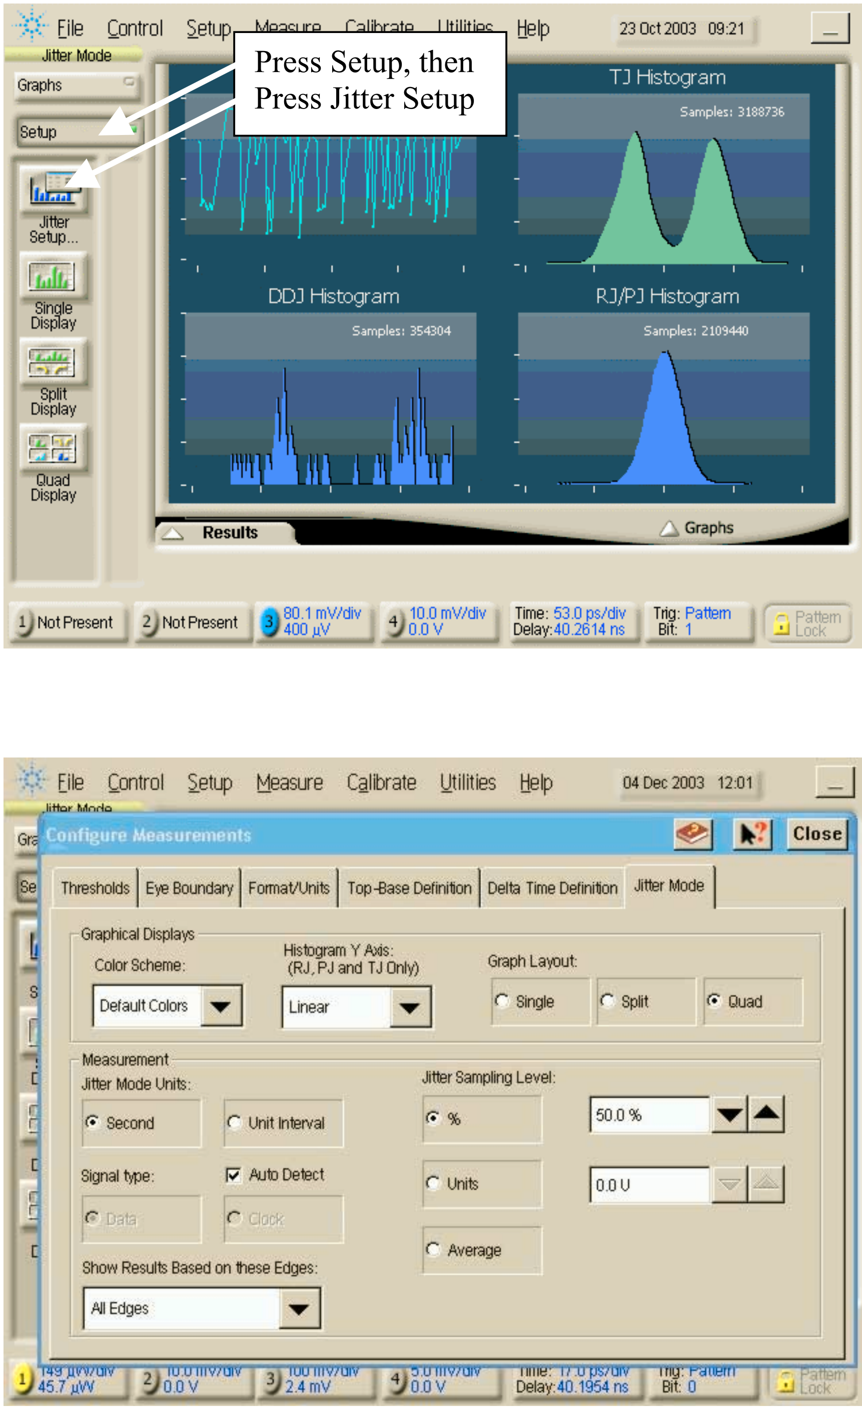

Figure 2. Jitter setup menu for customizing the measurement process

The tabular results include DCD, ISI, DDJ, DDJ versus bit position, PJ, RJ, DJ, and TJ.

One, two or four graphical displays of jitter are provided depending upon the user-defined configuration (available on the toolbar on the left side of the screen or the jitter setup page of the configure measurements menu). The following graphical displays are available:

RJ, PJ histogram: This is the probability distribution function (PDF) of random and uncorrelated periodic jitter. The histogram represents all jitter that is uncorrelated from the data pattern.

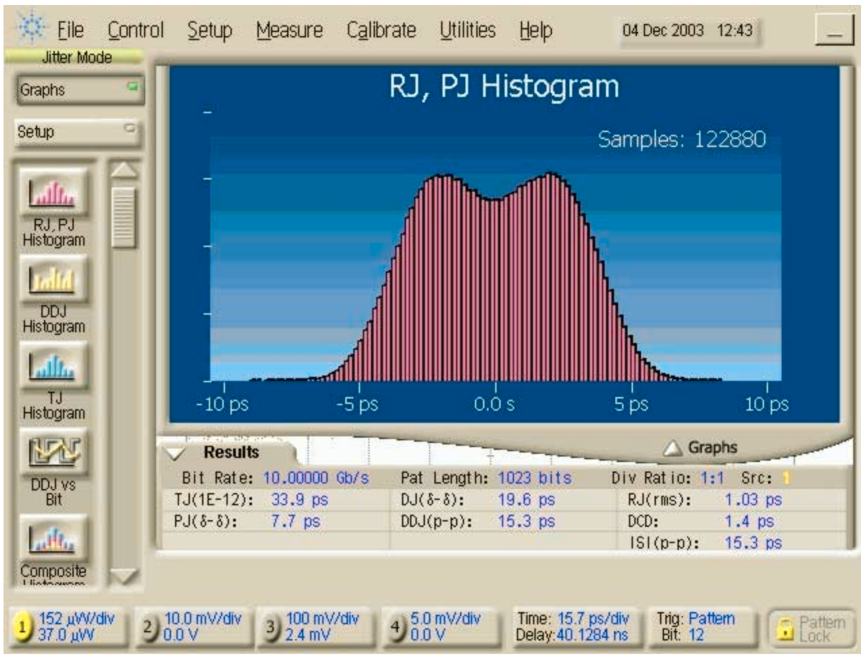

Figure 3. RJ, PJ histogram display

DDJ histogram and Composite DDJ histogram: The DDJ histogram is the PDF of the correlated jitter including ISI and DCD. The DDJ histogram display shows the DDJ data associated with all of the edges. In the Composite DDJ histogram display the different colors represent the rising edge data, falling edge data, and the composite data for all edges.

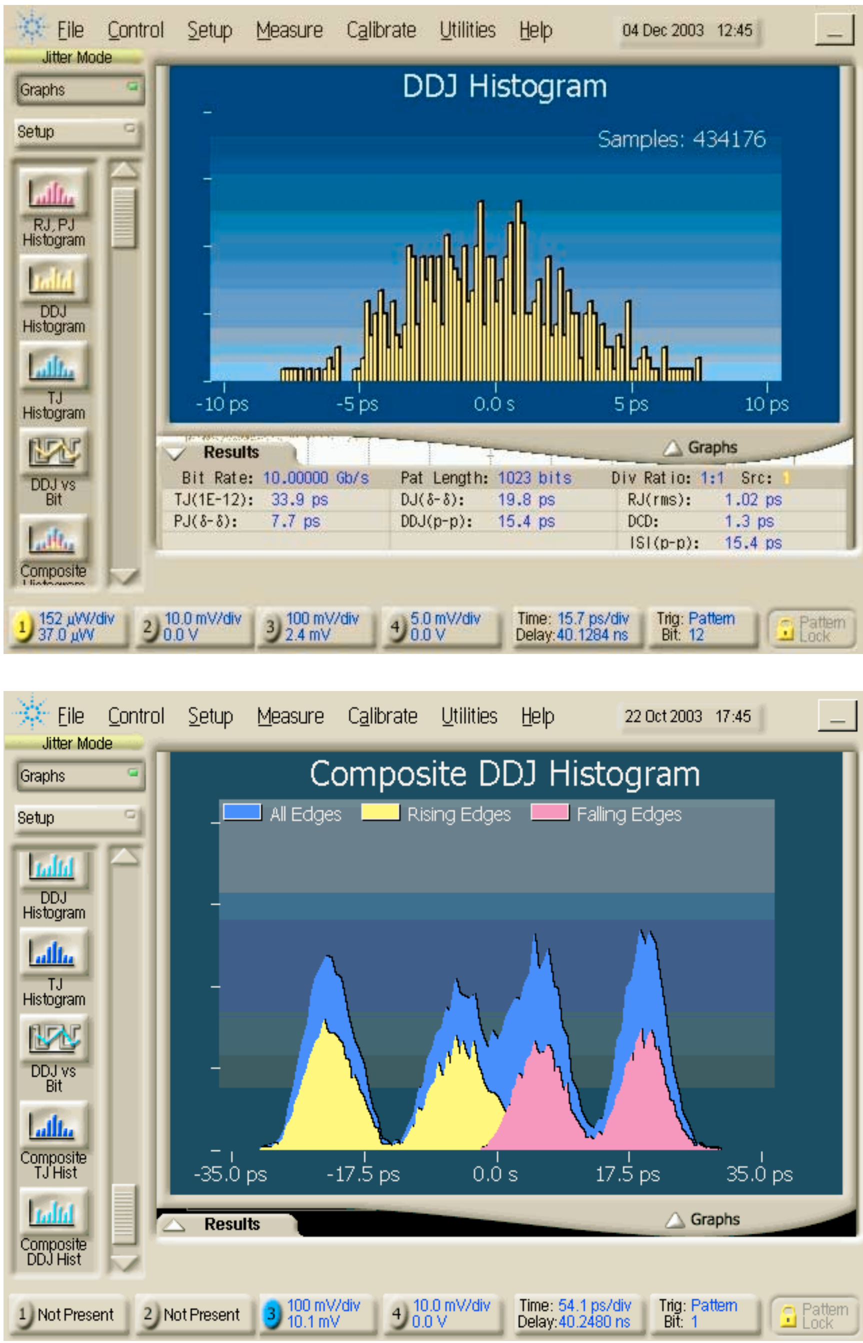

Figure 4. DDJ histograms

DDJ versus bit: This display shows the data dependent jitter (DDJ) for each edge in the entire pattern. The Y-axis is the jitter magnitude while the X-axis is the relative bit number. The data is plotted on top of an ideal representation of the data pattern. This enables quick reference of DDJ data to its position in the pattern. In order to optimize data visualization, the ideal data pattern background is not shown when the screen space to data ratio is less than 2 pixels per bit. This data can easily be assessed with respect to the input signal, as the actual waveform can be precisely viewed by switching to oscilloscope mode. The relative trigger bit can be adjusted in order to view specific sections of the pattern. While the 86100C has locked on to the pattern (using Pattern Lock triggering), the trigger level adjustment knob is reassigned to control the relative bit position of the oscilloscope's trigger in the pattern. The relative trigger bit can also be set directly to a given value in the Trigger Setup dialog. For example, if the DDJ versus bit display indicates that bit 387 has the largest early jitter value, bit 387 can be viewed by adjusting the relative trigger bit-to-bit 387.

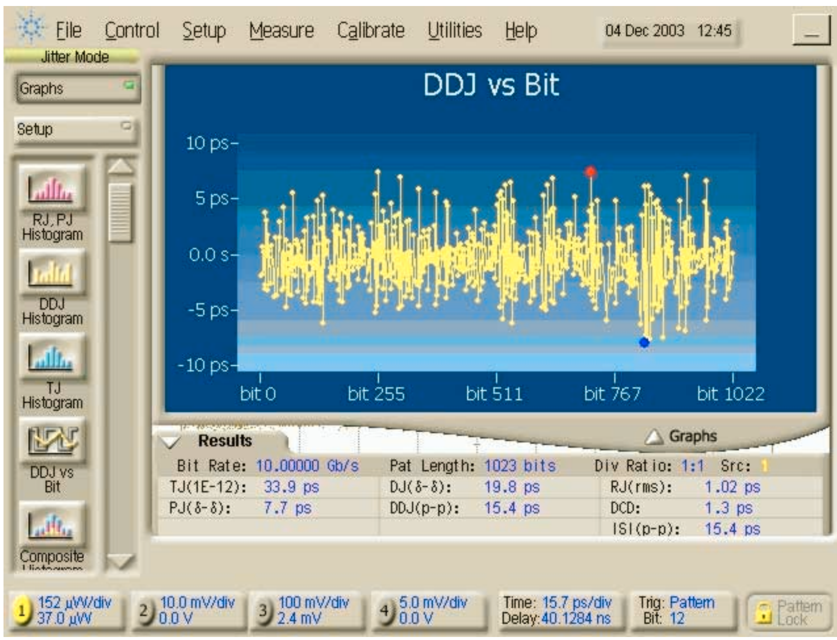

Figure 5. DDJ versus bit display

It is also possible to adjust the span of the DDJ versus bit display. Using a mouse, a box can be created by clicking and dragging. The box width will set the new span of the graph. Also, the oscilloscope's horizontal span and position knobs are reassigned in Jitter Mode to control the span and position of the DDJ versus bit display.

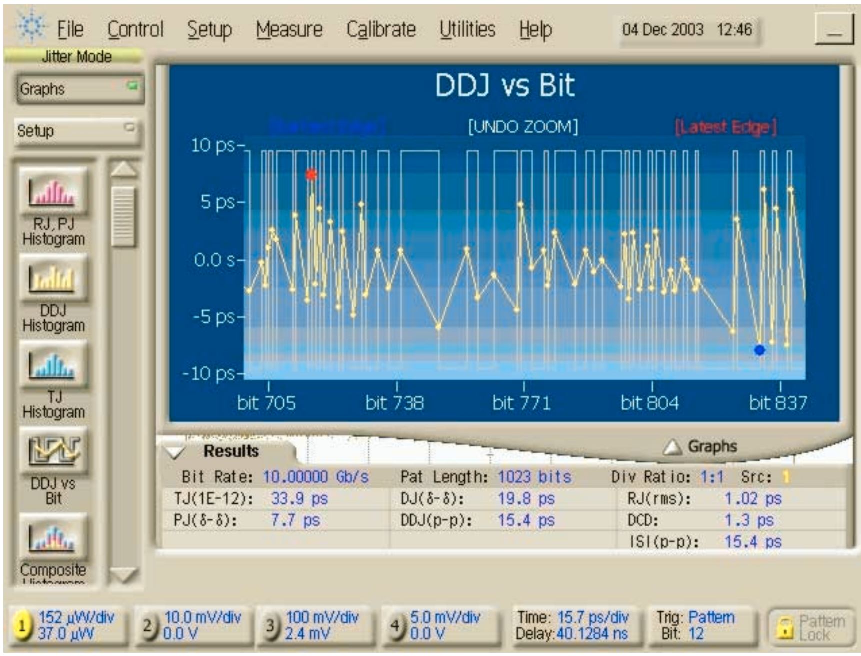

Figure 6. Zooming in on the DDJ versus bit display

Total Jitter histogram: This is the computed histogram derived from all the jitter data. The Total Jitter histogram is constructed by convolving the two directly measured histograms - RJ, PJ histogram with the DDJ histogram.

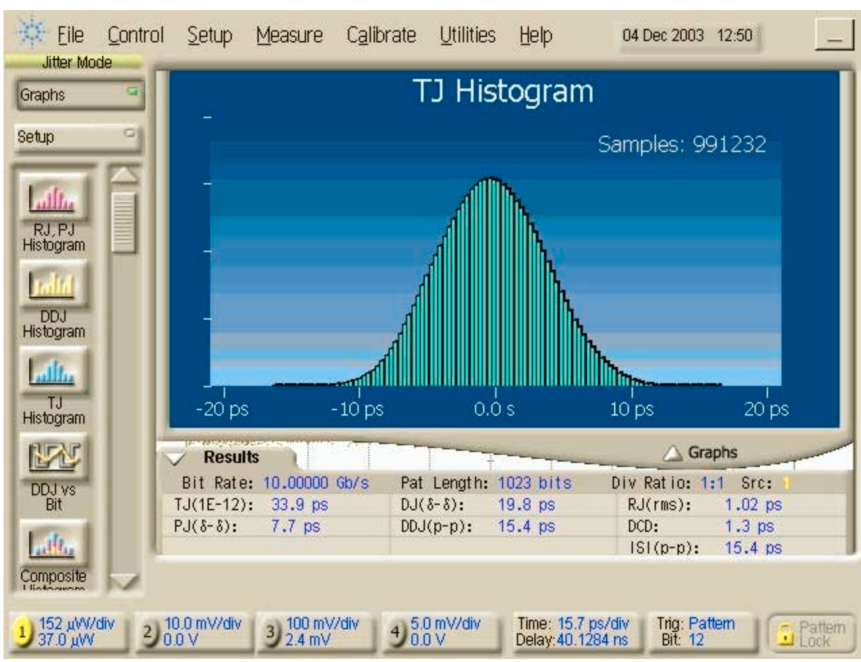

Figure 7. Total jitter histogram

Composite Jitter histogram: This combines the RJ, PJ, DDJ, and TJ histograms on a common X-axis. In order to provide a single histogram view that provides an overall view of the jitter present in the signal.

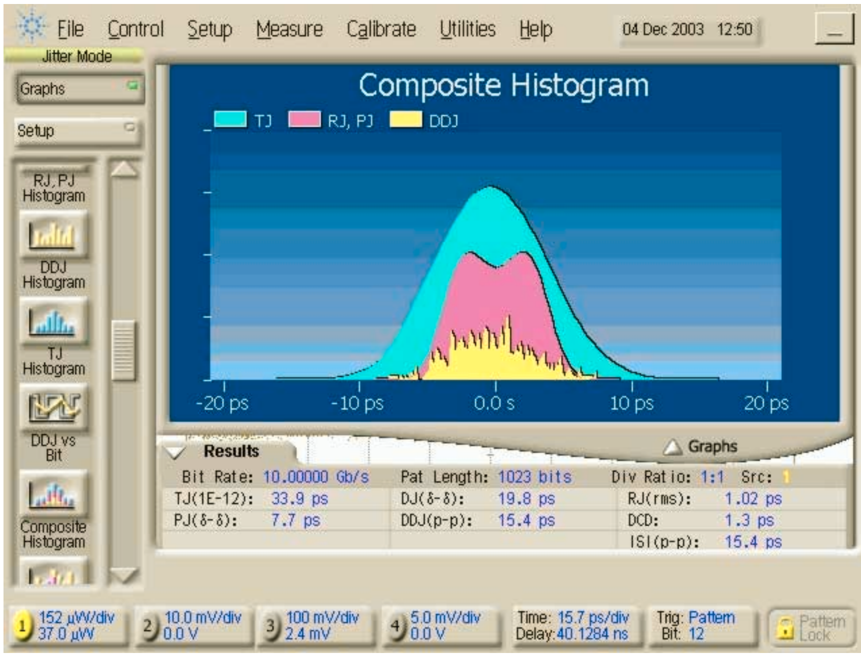

Figure 8. Composite jitter histogram

With the exception of the DDJ versus bit display, all the histogram graphs have no absolute Y-axis scale. The Y-axis magnitude represents the relative population magnitude for X-axis position. For the RJ, PJ, DDJ and TJ histograms, the total number of jitter samples acquired is indicated on the graph.

Interpreting measurement results

The methods the 86100C uses to extract the various elements of jitter are described in detail in subsequent sections of this paper. Certain jitter components are easy to interpret while others require some further explanation. The following is a brief summary of the measurements and what each is intended to describe.

Intersymbol interference (ISI): This value is determined from measuring the average position of each bit in the pattern with all uncorrelated effects removed. It is the difference between the earliest falling edge and latest falling edge, or the difference between the earliest rising edge and the latest rising edge, whichever is larger.

Duty cycle distortion (DCD): Also determined from the same data set as ISI, this is the difference in the mean position of all falling edges and the mean position of all the rising edges, with uncorrelated effects removed.

Data dependent jitter (DDJ): Also determined from the same data set as ISI is the difference in the position of the earliest edge (rising or falling) and the latest edge (rising or falling). Thus the measurement result is dictated by the worst-case bits in the pattern, but does not include the effects of any uncorrelated jitter. Note that as DCD goes to zero, DDJ is equivalent to ISI. As ISI goes to zero, DDJ is equivalent to DCD.

Random jitter (RJ): This value quantifies the jitter that is due to only to random mechanisms. In that it is Gaussian in nature with a magnitude that extends to extreme values (at low probabilities), a root mean square (RMS) value is reported.

Periodic jitter (PJ): This value represents all of the periodic jitter that is uncorrelated from the data pattern. The PJ value is reported in two ways. The PJ (δ–δ) parameter indicates the jitter magnitude required to separate the RJ PDFs to match the dual Dirac-delta model with the acquired RJ, PJ population (see Page 23 for a detailed description of the PJ methodology). The PJ value is also expressed as an RMS value in order to help quickly relate to known amounts of injected jitter. For example, if a sine wave jitter source is injected, its RMS value is known, and one should expect the measured RMS value to be very similar to the known injected amount.
In order to interpret the PJ (δ–δ) value it is important to understand the different impact of different shapes of periodic jitter. Consider the case where the periodic jitter is due to a square wave type signal compared to a triangle wave signal. If the magnitude of the square wave jitter is the same as the magnitude of the triangle jitter, the resulting RJ PJ population density will not be identical. In the square wave case, the combined RJ PJ distribution can be viewed as the random population being shifted back and forth to the two extremes of the square wave jitter excursion with virtually no 'dwell time' between. In the triangle cases, the random population can be viewed as linearly transitioning between the extremes. Clearly the square wave case effectively moves the <10E-12 probability events further away from the ideal edge point than does the triangle case. The result is that PJ (δ–δ) will be larger for a square wave jitter source than for a triangle wave jitter source. This indicates a greater likelihood that the PJ would cause a bit error.

Deterministic jitter (DJ): Similar to periodic jitter, DJ is also expressed with respect to a dual Dirac-delta model. The DJ (δ–δ) parameter indicates the jitter magnitude required to separate the RJ PDFs to match the dual Dirac-delta model with the TJ histogram population (see Page 24 for a detailed description of the DJ methodology).

Total jitter (TJ): The TJ value is interpreted as the total effective eye diagram closure in the time axis at a specified BER level (default is 10^{-12}). The closure is expressed in time or unit intervals. The probability that an edge will be misplaced beyond the TJ value is less than the user-defined probability threshold. Thus if the user defined closure threshold is 1E-9, and the TJ value provided by the 86100C is 50 ps, the likelihood that an edge will be more than 25 ps late or 25 ps early is less than 1 in 10^9. If for the same signal the threshold is set to a lower probability, such as 10^{-12}, the reported jitter value will increase significantly, as a much wider population of the distribution function is included in the assessment of TJ.

The case for jitter separation

In many communications systems, specifying jitter involves determining how much jitter can be on transmitted signals. Jitter is analyzed from the approach that for a system to operate with very low bit error ratios (BER)'s (1 error per trillion bits being common), it must be accurately characterized at corresponding levels of precision. This is facilitated through separating the underlying mechanisms of jitter into classes that represent root causes. Specifically, jitter is broken apart into its random and deterministic components. The deterministic elements are further broken down into a variety of subclasses (discussed later). With the constituent elements of jitter identified and quantified, the impact of jitter on BER is more clearly understood. This then leads to straightforward system budget allocations and subsequent device/component specifications.

The 3 Gb/s measurement barrier

There are a number of well established tools for jitter analysis. For many of them, the mechanism they use to acquire information allow them to be used for data rates up to 3 Gb/s and slightly beyond. Thus as systems are being developed requiring transmission rates of 4, 6, and 10 Gb/s, different measurement technologies are required. The wide-bandwidth oscilloscope has a bandwidth in excess of 80 GHz, thus these instruments' operating ranges are more than sufficient for general waveform analysis to 40 Gb/s and beyond. For the wide-bandwidth oscilloscope, rates that can be measured are dependent on channel bandwidth. Typically, the bandwidth should be at least double the data rate, with bandwidths ranging from 20 to 50, 70 and 80 GHz depending upon the oscilloscope plug-in configuration. As the oscilloscope is DC coupled, low data rates are also within its measurement range.

Historical limitations of legacy sampling scopes for jitter analysis

Where does the wide-bandwidth sampling oscilloscope fit into the jitter analysis picture? At a first glance, this instrument looks very attractive due to its very low jitter noise floor and the previously mentioned wide bandwidth. Historically, several barriers have existed for efficient and accurate measurements of jitter. Fortunately, the Agilent 86100 DCA has undergone significant architectural changes to overcome these issues. The new Agilent 86100C is referred to as the DCA-J, indicating its functionality as a digital communications analyzer with extensive jitter analysis capability.

The three key historical limitations of equivalent time sampling oscilloscopes, when they are used for jitter analysis are:

- Slow measurement speed due to a relatively slow data acquisition rate
- Requirement of a pattern trigger
- Measurement errors introduced when large ranges of timebase delay are used

While the bandwidth of these scopes is extremely high, the rate at which data is obtained is relatively slow, at much less than 1 Msample/s. This is in contrast with "realtime" oscilloscopes, that have bandwidths around 6 GHz, but which acquire data at rates of up to 20 Gsamples/s. The sampling oscilloscope then has special requirements for the data it can analyze. The signals must be repetitive, as it will take several passes to acquire enough samples to accurately reconstruct a waveform. (An exception to the repetitive signal requirement is the eye diagram. In this case, samples are acquired at essentially random, but synchronous locations in a bitstream. The common eye diagram is then a composite of samples acquired throughout the stream of data.)

The simplest jitter analysis performed on sampling oscilloscopes consists of projecting a pixel wide slice of the eye diagram at the crossing point along the time axis (Fig. 1). This approximates the probability distribution function of the signal's jitter generation.

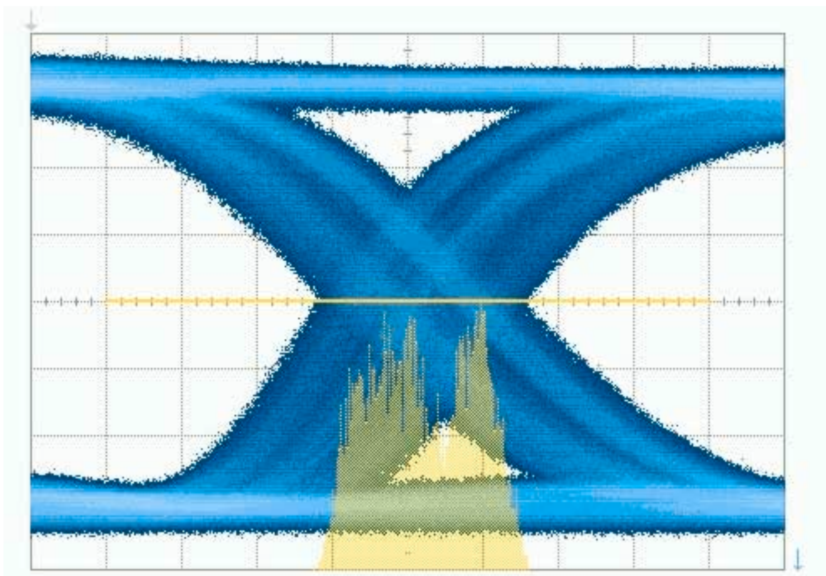

Figure 9. Jitter observed through the eye diagram crossing point

While an equivalent time sampling oscilloscope cannot capture a long contiguous data stream for analysis the way that a real-time oscilloscope can, the crossing point histogram can be post-processed to provide rudimentary jitter analysis: Random jitter (RJ, jitter due to random or stochastic processes) and deterministic jitter (DJ, jitter due to systematic repeatable mechanisms) can be separated by fitting Gaussian tails to the histogram and total jitter (TJ, the aggregate of all the jitter components) can be extrapolated. However, because of the low sampling rate and fluctuations in the histogram population over time, this simple analysis tends to be coarse.

The sampling scope still has the potential to do an accurate jitter measurement. The data dependent jitter (DDJ, jitter related directly to the sequence or pattern of bits) can be found by providing a pattern or frame trigger to the oscilloscope. The pattern trigger provides a timing edge that occurs at most once per repetition of the data pattern. Thus, a requirement for the DDJ measurement is a repeating data pattern. Random and deterministic jitter that is not correlated to the data pattern can be eliminated from the measurement by enabling trace averaging. The location of each edge can be compared to its ideal location through comparison to the clock edge position (if also displayed on the scope) or by comparison of its position relative to the subsequent position measurement of all the other bits in the pattern.

The RJ and uncorrelated periodic jitter (PJ, jitter that is periodic and therefore deterministic, but unrelated to the data pattern) can be measured by disabling any trace averaging and measuring a single data edge in the pattern. (Thus the pattern trigger requirement is still in place). Since all the samples are acquired on a single edge, there will be no DDJ component in this data. A data slice at the edge midpoint can be used to generate a time axis histogram that will contain the RJ and uncorrelated PJ. The histogram can then be analyzed and fitted to a curve to determine the critical statistics of these two jitter elements.

While this process is superior to a simple histogram of the eye crossing point in terms of precision and the separating out of the jitter elements, it still has some serious flaws. Perhaps the most important is the amount of time required to acquire a sufficient population for an accurate measurement. Consider that an oscilloscope trace of a data edge is composed of perhaps 1000 samples. The time between samples is several microseconds (25 microseconds for the Agilent 86100). Thus it can take 25 milliseconds to produce one edge from which a single jitter value can be obtained. Several jitter values are required to obtain sufficient data to produce a significant population from which the RJ and uncorrelated PJ values might be obtained. The time required for this analysis can easily approach several minutes.

The situation is more severe when attempting to determine the DDJ, as each edge in the pattern must be analyzed. For an accurate measurement, only a few edges are acquired for a given waveform. To remove the uncorrelated jitter elements, trace averaging is used. Thus the number of points that must be acquired is multiplied by the averaging factor, perhaps 16. While the time required for a short pattern is manageable and might take a minute or two, patterns such as a 2^15-1 PRBS (32767 bits) can take many hours to measure.

A second problem with this sampling oscilloscope method is the requirement of a pattern trigger, that is a trigger coincident with the repetition of the pattern (sometimes also called a frame trigger). While this may be easy to obtain in a test system using an instrumentation pattern generator, many test scenarios will not have this signal readily available. Additionally, the pattern trigger from a pattern generator instrument will typically have some data dependent jitter associated with it. This can corrupt the measurement, particularly when measuring low levels of jitter.

A third limitation associated with legacy equivalent time sampling oscilloscope architectures is a time error that is introduced by using large ranges of timebase delay (required to observe events that occur significant amounts of time after the trigger event). In traditional applications of these scopes this has not typically been a problem. However, when applying these scopes to precision jitter analysis, this becomes a significant issue. In order to characterize jitter through a pattern (primarily for DDJ), large ranges of timebase delay are typically required. This can introduce significant error into the jitter measurement.

Thus, long test times, the requirement of a pattern trigger, and large ranges of timebase delay present significant obstacles for the sampling oscilloscope jitter solution.

Architectural Changes Yield Fast and Accurate Measurements

To overcome these historical limitations, fundamental changes to the sampling oscilloscope architecture are required. Once the architecture is improved, the door opens for new measurement algorithms, which can dramatically reduce the time required and significantly improve the accuracy of jitter measurements.

Deriving a pattern trigger within the oscilloscope
A pattern trigger can be derived from a system clock signal. One must know the length of the pattern and the ratio of the data rate to the system clock rate. For example, in the case of the full rate clock, if the pattern length is N, a strobe signal should be generated every time that N clock cycles have been counted. While the concept is very simple, the successful implementation is far from trivial. Consider that any timing imperfections or asynchronisms in generating the trigger are manifested directly in jitter measurement error. It is also important for the sampling oscilloscope to have control of the pattern trigger and where it fires in relationship to the pattern. This control enables the oscilloscope to target the sampling at specific locations in the data pattern and to force the data acquisition to be performed over a small range of the timebase. Thus, the internally generated pattern trigger architecture has eliminated the requirement for a pattern trigger and measurement errors induced by large ranges of timebase delay.

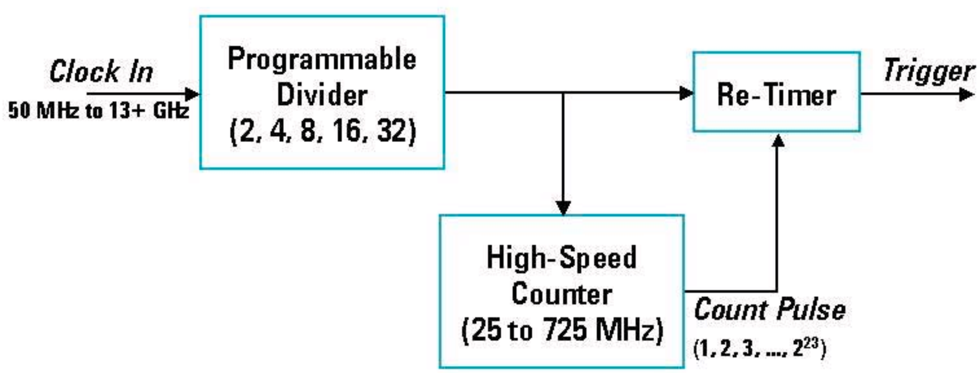

Figure10. 86100C Conceptual diagram for internally deriving a pattern trigger

With the internally generated and controlled pattern trigger, the oscilloscope can now systematically 'walk' through the pattern and determine the nominal location of each edge in the sequence. When executing this function, averaging is used to remove any uncorrelated jitter mechanisms in order to best determine the average location of each edge.

The next significant change in the new approach to jitter measurement requires a methodology for dramatically increasing the measurement speed.

Increasing measurement speed through optimized sampling

Given that the sampling rates of equivalent time oscilloscopes are slow, improving the speed of jitter measurements requires that data be collected and processed in a fashion where as many samples as possible contribute directly to the jitter measurement. The main candidate for improvement stems from the fact that, historically, a single jitter value is derived from hundreds of waveform samples. This is because the oscilloscope is fundamentally an *amplitude* measuring instrument. To determine the position of an edge, many samples were required to reconstruct the waveform, and from that only a single jitter value was derived.

To improve the jitter measurement efficiency of the oscilloscope, two important changes are made to how samples are collected. First, rather than collecting samples from all portions of the waveform including where amplitude is somewhat constant (and hence there is no direct jitter information), samples are acquired only from edges. However, this alone would not solve the problem, as many samples from the edge would still be required to yield a single edge position and thus a single jitter measurement data point. The key enabler to measurement efficiency is to develop a transfer function that allows the amplitude of any single sample taken from an edge to be directly translated to a jitter measurement value. This is achieved by generating models of the edges.

Consider the edge of figure 11. If the time of the sampling is set up to take place at the middle of the ideal edge (this is enabled by the internally generated and controlled pattern trigger), an early arriving signal will have an amplitude above the middle, while a late arriving edge will have an amplitude below the middle level. To determine the amount of jitter on the edge from which the sample was taken, one must know the amplitude versus time shape of the edge. This is effectively an amplitude-to-jitter transfer function – an edge model. Once the edge is modeled, *every* sample that is taken along the edge yields a jitter measurement.

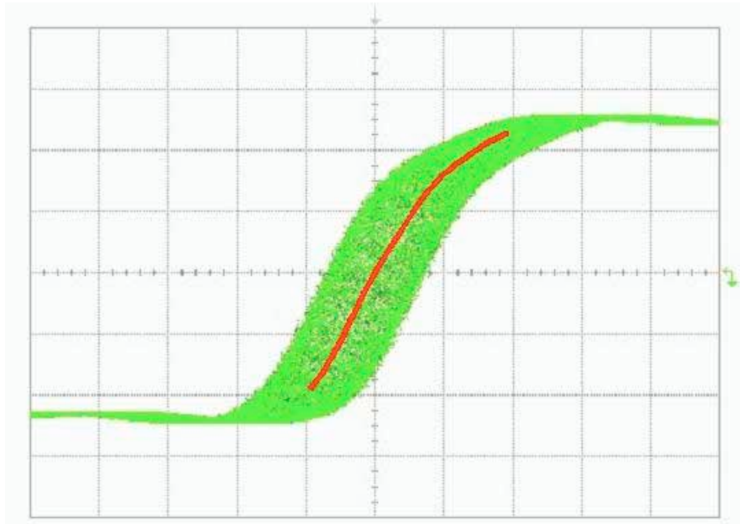

Figure 11. The edge model yields the amplitude to jitter transfer function

Two types of edge models are presented – single-edge models and composite-edge models. (The single edge models are used in uncorrelated jitter measurements while the composite edge models are used for data dependent jitter measurements. This is discussed later). A single-edge model is constructed by taking 1024 samples across the entire span of one edge. A mathematical function is constructed that delivers the best fit, in a least squares sense, to the sampled data. A composite-edge model is very similar, except the samples used to construct the model are taken from multiple edges.

A single-edge model is used to describe the amplitude-to-jitter characteristics of a specific edge of a pattern. Composite-edge models are composed of 4096 points and are used to describe the amplitude-to-jitter characteristics of a class of edges. The shape of an edge can be dependent upon the preceding bits. The 'memory' generally lasts for three or four bits, thus there are several classes or groups of edge shapes possible. A class is defined by two factors – rising versus falling and the four preceding bits. For example, '00001' is one class of rising edge, and '00101' is another. Similarly, '11110' is one class of falling edge, and '11010' is another. Consequently, 16 edge classes are defined – 8 rising and 8 falling.

There is some measurement time overhead associated with generating edge models for the several edge classes. However, consider that once the models have been generated, from that point forward almost every sample obtained provides a jitter value. The time required to generate the models is on the order of one or two seconds. This small investment results in orders of magnitude improvement in measurement efficiency. In some recent trials, complete jitter measurements performed on patterns greater than 7600 bits in length were performed in less than 15 seconds. This measurement took over 6 hours with the conventional equivalent time sampling oscilloscope.

Using the New Architecture to Separate Jitter

The process to determine TJ, RJ and DJ is now described along with the process for determining the subcomponents of DJ - DDJ (including inter-symbol interference (ISI), duty-cycle distortion (DCD), uncorrelated periodic jitter (PJ), and subrate jitter (SRJ)).

As described earlier, the oscilloscope will examine the signal to automatically determine the triggering clock rate, bit rate and pattern length (alternatively, the values can be entered manually). Once these values are determined, hardware within the instrument will generate a pattern trigger. The pattern trigger is manipulated to execute the edge modeling process within the pattern. Samples are acquired only on the edges and converted directly to jitter values.

The jitter separation methodology approaches the task by independently targeting the jitter that is correlated to the data pattern and the jitter that is uncorrelated from the data pattern. The correlated jitter is by definition the data dependent jitter (DDJ). The uncorrelated jitter is made up of random jitter (RJ) and uncorrelated periodic jitter (PJ).

Correlated Jitter

Averaging is used to isolate the DDJ. Averaging eliminates the uncorrelated elements. The edge deviation effects – jitter - that remain are those that are correlated to the data pattern – the DDJ. Composite-edge modeling is used in order to maintain maximum sampling efficiency, as each edge in the pattern must be measured in order to characterize the DDJ. The pattern trigger is "walked" through the pattern and samples are taken from every edge. The composite-edge model associated with each edge's class is used to translate each amplitude measurement into a jitter measurement.

The jitter on edges is segregated for the rising edges and the falling edges. A probability distribution function (PDF) histogram is created for both the jitter of the rising edges and the jitter of the falling edges as well as a jitter histogram of all edges. The DDJ is given by the peak-to-peak spread of the histogram of all edges. It is the arrival time difference between the earliest arriving edge and the latest arriving edge. The jitter induced by inter-symbol interference (ISI) is given by the peak-to-peak spread of the rising edges or the falling edges, whichever is greater. Duty cycle distortion (DCD) is given by the difference between the mean of the rising edge positions and the mean of the falling edge positions.

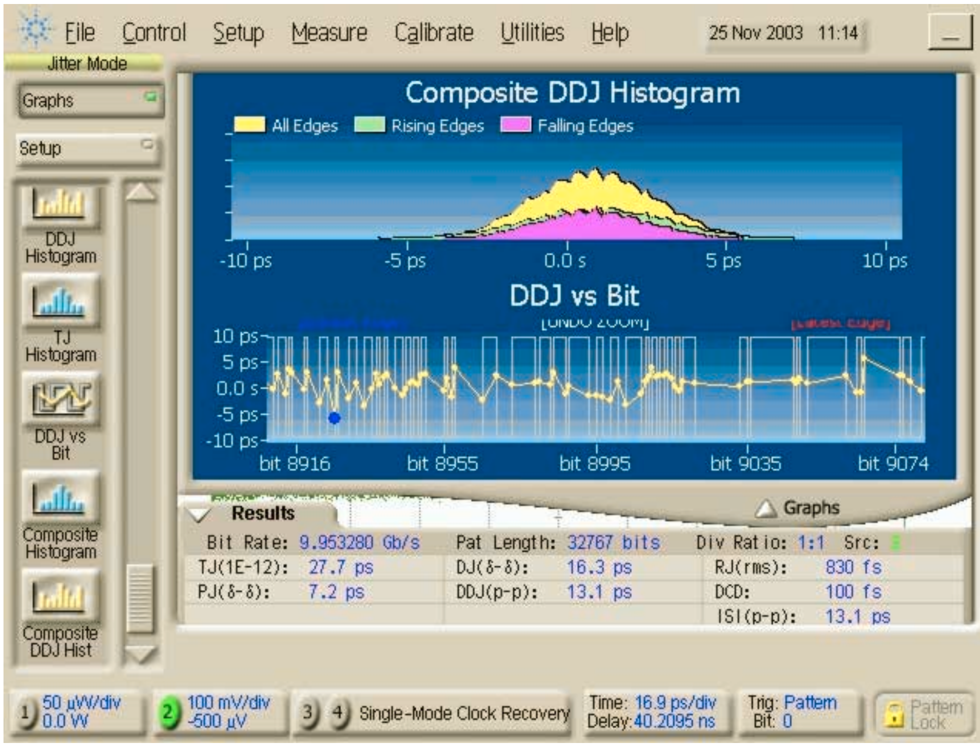

Figure 12. DDJ histogram and DDJ versus bit – split display

Sub-rate jitter (SRJ) is a new term for jitter which is periodic, correlated to the data, and whose frequency is an integer sub-rate of the data rate. This form of jitter is often associated with multiplexing data structures. For example, if one leg of a parallel data structure systematically results in an edge to be late or early, the jitter will be periodic in nature. It is correlated to the data stream, but may not be seen systematically on the same bits in a given pattern depending on the relationship between the pattern length and MUX size. Another common source of SRJ is coupling of a sub-rate reference clock onto the full-rate data stream. Extraction of SRJ is somewhat complex, as SRJ can show up as DDJ when the pattern length is a multiple of any clock divisor used to trigger the oscilloscope. Otherwise it shows up as PJ. This fact is then used to segregate the SRJ component and display it as a unique quantity.

Uncorrelated Jitter

The mechanism for determining the uncorrelated jitter (RJ and PJ) takes advantage of the fact that any edge taken in isolation has no knowledge of the pattern dependent elements that effect it. Its deviation about its mean position is totally dictated by the uncorrelated elements. For each data acquisition cycle targeted at uncorrelated jitter, the oscilloscope will acquire all of its samples on a specific edge within the pattern, thus all pattern dependent jitter is removed from this data. A single-edge model is used, and samples are taken specifically at the edge. The edge model technique is used to efficiently convert amplitude samples to jitter values. Subsequent acquisitions are taken from other edges in the pattern, but each individual acquisition record uses data from a single edge. Data from all acquisitions is accumulated in a histogram.

The internally generated pattern trigger is controlled such that the sampling interval is very precise and consistent sample-to-sample. The result is that the samples are taken in a highly periodic fashion. This allows the jitter values to be transformed into the frequency domain using a fast Fourier transform (FFT). This results in the spectrum of the jitter that is uncorrelated from the data pattern, which includes the RJ and uncorrelated PJ. The RJ makes up the noise floor of the spectrum and the PJ shows up as discrete frequency components or line spectra. The RJ is obtained by integrating the noise floor of this spectrum. Prior to the integration, the PJ line spectra are removed, and interpolation is used to fill the 'gaps' left behind by the missing lines. The remaining spectrum is due to the random components of jitter. This is integrated to determine the root mean squared (RMS) RJ, that is, the standard deviation of the random jitter distribution.

The line spectra are not used to determine the PJ. The maximum periodic sampling rate of the 86100C is 40 KHz. Any jitter spectral content that is above 20 KHz will be aliased. This limits the analysis of the jitter spectrum to that of the noise floor described above and will not allow the periodic elements to be determined accurately. The PJ is determined by returning to the accumulated histogram of the jitter values obtained for the targeted edges.

A dual-Dirac delta model is used to determine the PJ. The standard deviation of the RJ distribution is given by the measured RMS RJ described above. A dual-Dirac delta model is constructed with two identical Gaussian distributions each defined by the measured RMS RJ value. The separation of the two Gaussian distributions is adjusted in order to match the model to the histogram. The match is made where the peak-to-peak separation representing 99.8% of the volume of the model matches the corresponding width containing 99.8% of the volume of the measured histogram. (See figure 4.) The PJ value is given by the resultant separation between the means of the two Gaussian distributions.

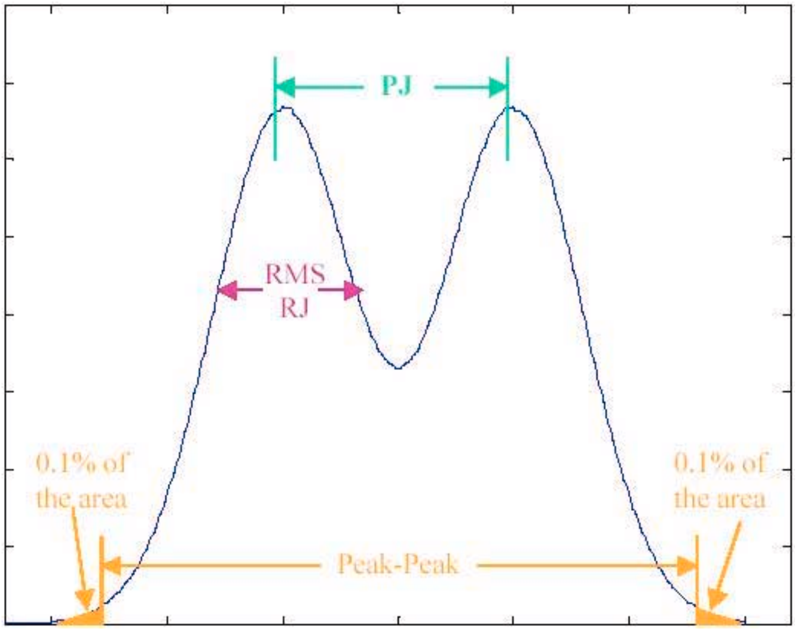

Figure 13. The dual-Dirac delta jitter model

Aggregate Deterministic Jitter (DJ)

The above analysis has resulted in specific values for RJ, PJ, ISI, DCD, SRJ, and DDJ. These contain all the elements of jitter. The task now is to accurately combine these elements to produce an aggregate deterministic jitter (DJ) value and finally a total jitter (TJ) value. So far, probability density functions (PDF) have been determined for the uncorrelated jitter (RJ and PJ) and for the correlated jitter (DDJ). The DJ value is composed of both the DDJ and the PJ, but it is not a simple sum of the values, as each is defined by a statistical distribution. To determine the aggregate DJ, a methodology similar to that used to determine the PJ from the RJ, PJ PDF is used. The RJ, PJ and DDJ PDF's are convolved together. The aggregate histogram is called the total jitter histogram, as it is the PDF of all of the measured jitter – both correlated and uncorrelated – combined in a single histogram. The dual-Dirac delta model methodology described above for extracting PJ from the RJ, PJ histogram is applied to the total jitter histogram. The same measured RMS RJ value describes the two Gaussian distributions. The same fitting technique is used. However, as the total jitter histogram is the PDF of all jitter elements, the resultant separation is the aggregate DJ.

Aggregate Total Jitter (TJ)

The ultimate measure of jitter generation performance of a device under test (DUT) is total jitter (TJ). The aggregate TJ histogram must then be further analyzed to provide a numerical value that can be used to assess the quality of the DUT. Note that the tails of the aggregate PDF extend indefinitely, so a peak-to-peak value has meaning only when associated with a specific probability. A typical approach is to determine the jitter level such that the probability of exceeding it is less than 10^{-12}. The dual-Dirac delta model associated with the total jitter PDF is used to determine this value. TJ is determined by extending the dual-Dirac delta model down to the point where the probability of occurrence is less than 1 part per trillion of the whole. The width of the model at this threshold is the total jitter. TJ values at other levels of probability can be obtained in a similar fashion.

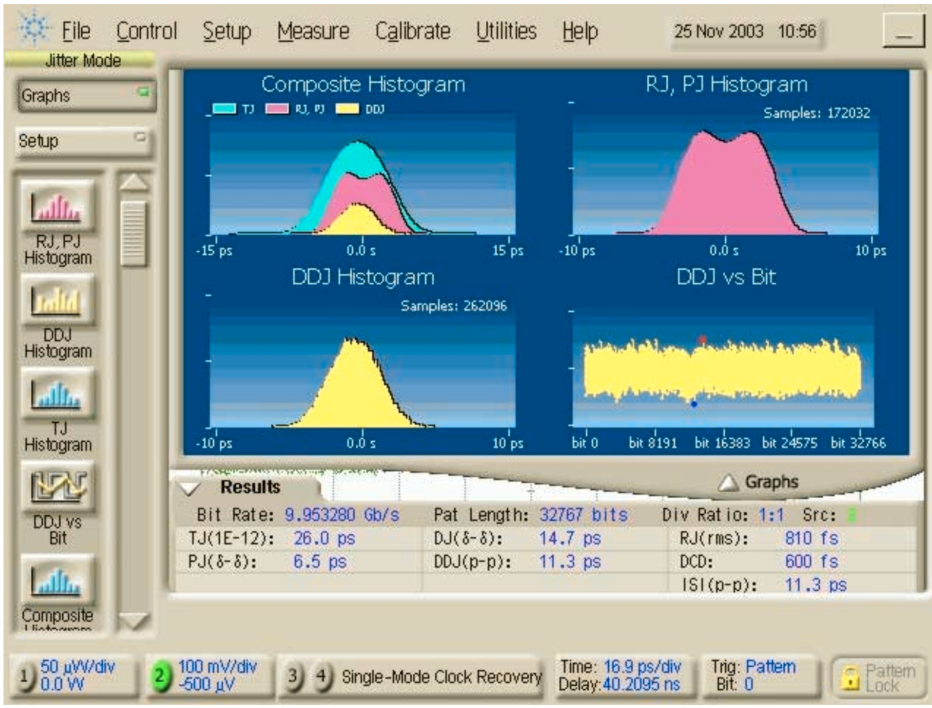

Figure 14. Combined histograms and tabular results

Measurement requirements and limitations of the DCA-J in Jitter Mode

The basic requirements for the 86100C to perform jitter measurements are specific to the signal used to trigger the instrument and the data sequence of the signal being measured. The trigger signal must be synchronous to the data. If there is no trigger signal, the measurement cannot be made. Allowable divide ratios (the ratio of the data bit rate to the trigger clock rate) are 1, 5,10, 20, 25 and 2^N, for virtually any integer value of N. If no clock is available, a clock recovery system can be used to provide a clock.

The data pattern must be continuously repeating, with no idle states. The pattern length must not exceed 2^{15} (32268 bits). The instrument will automatically determine the pattern length if it is available in its lookup table. The lookup table includes 2^N and 2^N-1 for N=1 through 23, as well as the common jitter test patterns (K28.5,CJPAT, CRPAT, SCPAT, JTPAT, SPAT). Other pattern lengths can be added to the lookup table. Once added, the instrument will then be able to automatically detect and lock to that sequence. (The lookup table can be edited by selecting the “Pattern Lock Setup” tab of the trigger setup dialog, pressing the “Select From List” button next to the Pattern Length, and choosing “Edit List” from the displayed selections.)

The instrument itself will contribute to the jitter. However, this can be as low as 150 femtoseconds when using the 86107A precision timebase module, thus the effective jitter noise floor or measurement sensitivity level is extremely low.

As the jitter values are obtained through the edge model transfer function, this places some limits on the magnitudes of jitter that can be measured. As edge speeds become very fast relative to the bit period, the magnitude of jitter that can be accurately measured is reduced. The risetime or falltime of the signal should exceed the sum of the peak to peak PJ and double the RMS RJ. The DDJ magnitude is not restricted by the edgespeed provided the edgespeed exceeds the peak PJ and double the RMS RJ. A good guideline is:

$$\text{Risetime or Falltime} > (2\ \text{RJ rms} + \text{PJ pp})$$

Thus the RJ and PJ jitter magnitudes that can be observed must be less than half of a bit period.

Although the magnitude of the periodic jitter is provided, the frequency values are not. This is due to the relatively low sampling rate which in turn yields an aliased spectrum for signals of frequencies beyond 20 KHz.

Conclusions

Increased rates of data transmission are a fact of life. Jitter, also a fact of life, will continue to be a significant obstacle to increasing transmission speeds. Thus measurement tools with the ability to make accurate assessments of jitter to 10 Gb/s rates and beyond represent a significant technology breakthrough. The complex nature of jitter requires complicated analysis. However, well-designed instrumentation can simplify the task. Jitter measurements can now be performed in seconds, with instruments that have effectively reduced the measurement process to a single key press.

Printed in USA December 9, 2003

Agilent Technologies