Fundamentals of Signal Integrity

ARTICLE #6

By Dr. Howard Johnson
Signal Consulting, Inc.
March 15, 2009.

Summary

LeCroy Corporation, in association with Signal Consulting, Inc., has prepared an eight-part series on the fundamentals of signal integrity. Authored by the world's foremost authority on signal integrity, Dr. Howard Johnson, the series is a "must read" for engineers who need a clear understanding of issues essential to high-speed performance.

Other papers in the series include Confirm The Diagnosis, Adequate Bandwidth, and Step Response Test. To read other parts in the series, please visit: http://www.lecroy.com
:

Jitter Characterization

Today's topic concerns the characterization of jitter. I wish I could begin by stating the definition of jitter. Wouldn't it be great if there was only one definition? Unfortunately, the subject isn't that simple. Here's a sampling of definitions from various sources:

Jitter, "Abrupt, spurious variations in… [time, amplitude, frequency, or phase]." IEEE Standard Dictionary of Electrical and Electronics Terms.

This IEEE definition seems to me very general and somewhat obtuse, like it was written by a committee. Not surprising. The IEEE Standard Dictionary is an unedited collection of definitions from all the IEEE standards developed over the years. In an open standard, anybody can join the process, and anybody can propose a definition. No matter how poorly worded a definition may be, if it doesn't break anything it will get voted into the standard, the theory being, "Why fix something that isn't broken?"

The next is a definition more relevant to my cause today:

Jitter, "Unwanted pulse position modulation." Patrick Trischitta and Eve Varma, *Jitter (in Digital Transmission Systems),* Artech House, 1989. Their book is extremely relevant to SONET and similar systems with multiple stages of signal propagation. It is very mathematical.

I prefer the term "unwanted" rather than "abrupt", because "unwanted" implies that the definition of jitter depends on you and what you expect from your circuit, which is precisely the case. The term "abrupt" implies an absolute delineation between "fast" and "slow" events. Such a delineation does not exist. What's fast in one application is slow in another, and vice-versa.

OK, now let's get down to a definition written by people who have actually made *lots* of high-speed digital system measurements: the Fibre Channel committee. That group was among the first to write an explicit procedure for measuring jitter.

Jitter, "Deviation from the ideal timing of an event." Methodology of Jitter Specification, NCITS (National Committee for Information Technology Standards), T11.2. This document is informally referred to as the "MJS".

The MJS report was produced as part of the Fibre Channel effort. In that process numerous companies involved in producing, using, and measuring serial data communications hardware battled out the details of how jitter should be defined and characterized. If you haven't read this piece of work, I highly recommend it. Steve Joiner of HP did a beautiful job of technical editing for the document. It forms the basis of most high-speed jitter analysis used today.

I like this definition because it highlights the relative nature of the problem. The definition suggests a comparison between actual events and ideal events. That's not quite what we measure in real life, though. In actual practice jitter is always measured as deviations between one signal and another. The other is never ideal.

Random Processes

Imagine in front of you a jittery data stream. Think of jitter as a second signal, separate from the data stream, that encodes only information about timing variations in the data. When a data edge arrives early, the jitter swings negative. When the data is late, the jitter goes positive. In a system with slowly-wobbling timing variations, the jitter signal may undulate from one extreme to the other over a period of thousands or millions of bits. In a system with sudden, quick variations the jitter signal might jump around like a drop of water sizzling on a hot stove. Either way, the jitter signal constitutes a *random process*.

Many random processes, jitter included, can be characterized by two types of statistics: an auto-correlation function and an amplitude histogram. These two dimensions of specification are independent. Both types of information must be stipulated to completely represent a random process.

The auto-correlation function speaks to the time-domain predictability of a waveform, while the histogram speaks only to the distribution of amplitude values, independent of time. Both specifications play an important role in jitter analysis.

For example, consider the familiar term "white Gaussian noise" (WGN). When used to describe a random process, the word "white" implies no auto-correlation between the value at any one point of time and the value at any other time. In other words, the signal is not *predictable*. In frequency-domain terms, the power in a white signal distributes itself evenly throughout the spectrum.

The hiss of escaping steam, or the sound you hear on an FM radio when tuned between stations, are both white audio processes.

The term "Gaussian" refers to the shape of the amplitude histogram, independent of time. A Gaussian histogram is shaped like a bell. It has one central blob surrounded on either side by long lingering tails. Any random process built from a large number of tiny, independent, superimposed events tends to have a Gaussian histogram.

The distribution of SAT math scores forms a Gaussian curve (truncated at 0 and 800).

No physical random process can ever be truly white, because a truly white process implies spectral content (and thus power) extending from DC to infinite frequency. That's not possible. All physical processes that we know of have limited bandwidth. In practical terms a random processes is "white" (i.e., the power density spectrum is flat) only over a limited range of frequencies.

Similarly, no physical random process can ever be truly Gaussian, because a truly Gaussian histogram includes values extending to plus and minus infinity. All physical processes that we know have upper limits to their extreme values. What we call Gaussian is usually a distribution that looks Gaussian over the range of several standard deviations to either side of the mean. Beyond that limit the distribution falters.

Let's look at an example of jitter using three different tools: amplitude histogram, spectral plot, and time-domain jitter waveform.

Amplitude histogram

Figure 1 illustrates a decidedly non-Gaussian histogram obtained using the Serial Data Analysis feature on my LeCroy SDA6020 scope.

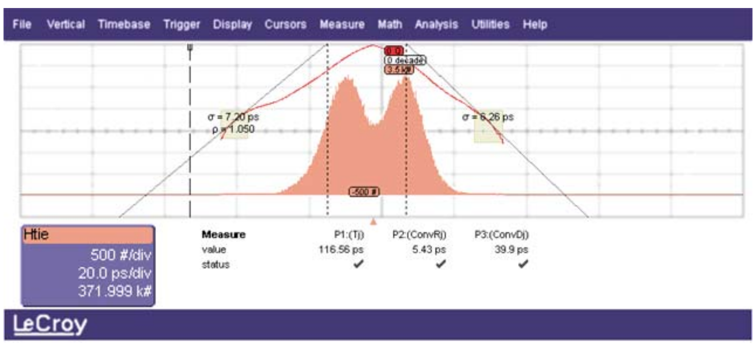

***Figure 1** - This double-humped histogram indicates significant amounts of deterministic jitter.*

The signal under test is a 1.25 Gb/s data signal, transmitting a repeating 10-bit JK pattern. The display shows the histogram of received jitter. The histogram is the pink blob in the center of the screen. The horizontal axis represents the range of possible times-of-arrival of data edges on a scale of 20 ps/div. The vertical axis represents how often each time-of-arrival happened to be detected.

The histogram reveals a double-humped distribution of jitter values. It looks like a superposition of two Gaussian blobs at different positions. The width of each blob is related to random Gaussian effects, while the separation between blobs is, apparently, related to some deterministic effect that causes some data edges to consistently appear 20 ps earlier than others.

A jitter histogram tells you *how much* jitter exists and *how often* it occurs. It may also give clues as to the nature of deterministic effects.

A spectral plot, on the other hand, delivers completely different information. The spectral plot relates only to the frequency concentration of the jitter signal.

Spectral plot

Figure 2 plots the frequency content of jitter within the same system as before. The vertical axis is 10 dB per division, the horizontal axis is 125 KHz per division. This plot was computed using an FFT function. Several FFT's were power-averaged to produce this plot. The averaging process suppresses random variations, clearly exposing the underlying power density spectrum.

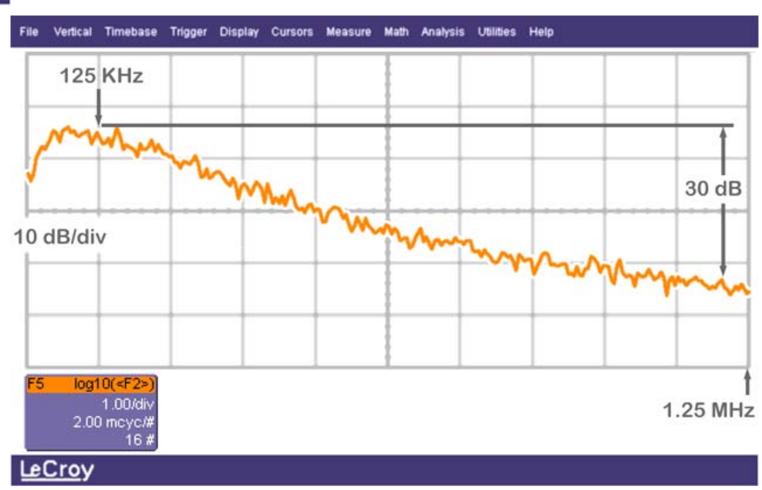

***Figure 2** - The power density spectrum is concentrated below 250 KHZ.*

The density of jitter power in this signal appears 30 dB higher at 125 KHz than it is at the right side of the picture (at 1.25 MHz). Apparently, some process operating above 125 KHz acts to attenuate the jitter in this signal. The steepness of the spectral curve indicates the efficacy of the jitter-filtering mechanism.

A spectral plot is used for some types of phase-noise compliance tests, it is useful for testing the performance of jitter-smoothing PLLs, and it can identify concentrated peaks of periodic jitter emanating from switching power supplies and other periodic sources of interference.

Time-domain waveform

The jitter spectrum in Figure 2 appears concentrated at low frequencies. In the time-domain, that implies that the jitter signal may be somewhat slow-moving.

Since the maximum frequency content (6-dB point) of the yellow jitter spectrum is about 250 KHz, you might expect in the time-domain to see rising/falling undulations on a scale of time equal to approximately 2 usec. I arrive at this number as the ratio of one-half divided by the upper (6-dB) bandwidth limit.

Figure 3 shows the captured jitter on a scale of 2 usec/div. The jitter variations appear in line with the prediction we just made from the frequency-domain chart. Smoothness in the time domain translates to an attenuated degree of high-frequency power in the spectral plot. A completely random jitter signal (totally uncorrelated noise, different at every data edge) possesses a flat jitter spectrum.

The time-domain waveform is particularly useful when diagnosing phase-slippage and other PLL locking issues. Combined with advanced triggering concepts it also serves to identify rare and erratic events.

Figure 3 - *Each undulation has a rise/fall time of approximately 2 usec.*

Conclusions

The three tools I have introduced present different aspects of jitter. The more conversant you become with these tools, and the overall language of jitter, the more quickly you will be able to diagnose and pinpoint jitter problems.

I hope to complete soon a series of articles about jitter and jitter analysis tools that will explore more thoroughly the topics I raised briefly today.

Best Regards,
Dr. Howard Johnson