

DesignCon 2007

Advanced Tools for High Speed Serial Data Measurements:

Equalizer Emulation
And
Virtual Probing™

Peter J. Pupalaikis, LeCroy Corporation
PeterP@LeCroy.com

Abstract

Two new tools that operate in LeCroy scopes – Equalizer Emulation and Virtual Probing – are provided for dealing with difficult test and measurement problems in measurements of equalized systems.

Author(s) Biography

Pete Pupalakis is Principal Technologist at LeCroy working on new technology development. Previously he was Product Marketing Manager for high-end scopes. He has nine patents in the design of measurement instruments. He holds a BSEE from Rutgers University and is a member of Tau Beta Pi, Eta Kappa Nu and the IEEE Signal Processing, Microwave and Communications Societies.

Introduction

Recent trends indicate that the next increases in serial data transmission speeds will not come from large improvements in the quality of the channel (defined as the media, being a cable or backplane) over which the data is transferred, but rather in increased sophistication built into the transceiver chips in the form of equalization. This means, in basic terms that channels will be non-ideal and that transceiver chips will increasingly employ equalization methods to compensate for these non-ideal channels. This trend leads to increased challenges in the area of test and measurement in equalized serial data systems. This paper explores these challenges and introduces some tools to confront these issues.

The Non-equalized Environment

The goal of this paper is to show how advanced tools are applied in the test and measurement of equalized serial data systems. That being said, it helps to explain the test environment in the non-equalized case to provide a foil for contrasting the two situations.

The role of test and measurement instruments with regard to serial data is to aid in the development of systems that transmit data at a high rate with low numbers of errors. The rate is specified in bits/s (or Gb/s) and the error rate is specified as the bit error ratio (BER) – the number of bits in error divided by the number of bits transacted. So, one goal of instruments is to measure BER in the fastest, most accurate manner possible. In compliance testing, another goal is to *guarantee* that a system, once produced, will operate properly and more importantly, that individual pieces of the system will interoperate. In other words, it is important to be able to budget and test to adherence of a budget the individual elements of the system such that when connected together they form an operating system. The ability to build systems in this manner forms the basis of our high-tech world. The ability to separately design and measure performance of elements in a standalone fashion is a level of abstraction necessary to allow engineers to design complex systems.

In the non-equalized world, the situation is such that isolated system components may be measured and analyzed to determine overall system performance. In this world, the measurement of connectors, backplanes, transmitters and receivers is performed separately and easily. For example, a backplane has been traditionally measured with time domain reflectometry (TDR) to determine whether the backplanes and connectors are “good” meaning they have low loss, the correct impedance, and few impedance discontinuities¹. Today, the vector network analyzer (VNA), traditionally a microwave engineer’s tool of choice, is being used to make these measurements in the form of scattering parameters² (s parameters). S parameter measurements and TDR measurements are roughly equivalent ways of generating the same information.

Measurements of backplanes and connectors using network analysis instruments are what can be thought of as *static* measurements, for the measurements described provide component measurements that essentially determine the behavior of these components in their entirety and for all time, albeit in a complex manner. In other words, the measurements provided of a channel component by a VNA allows one to predict or simulate anything about the behavior of that component in a system.

The measurement of transmitter behavior in the non-equalized world is essentially one of jitter analysis. The jitter measurement industry is filled with various methods of measuring jitter and ascribing the measurements to various sources including random and deterministic, with deterministic including periodic and data dependent jitter. Measurements also include measures of noise, including random noise, cross-talk, and distortion. Noise and its manifestation as jitter is of particular concern. Distortion can also be data dependent – it is this main effect that leads to equalized systems.

When analyzing jitter and noise effects, it is useful to examine the eye – the voltage waveforms in the vicinity of the decision point for each bit overlapped. Jitter has the effect of closing the eye horizontally and noise has the effect of closing the eye vertically. Eye closure leads to higher chance of error and therefore increasing of the bit error ratio.

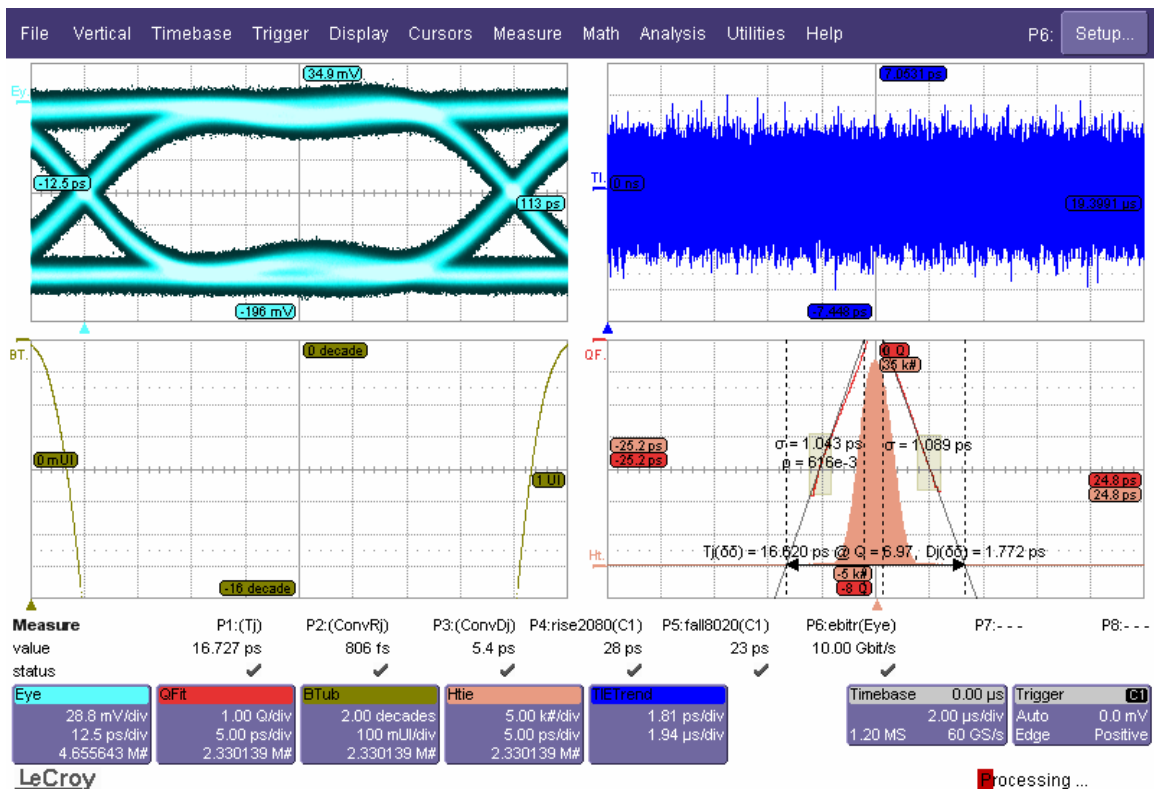


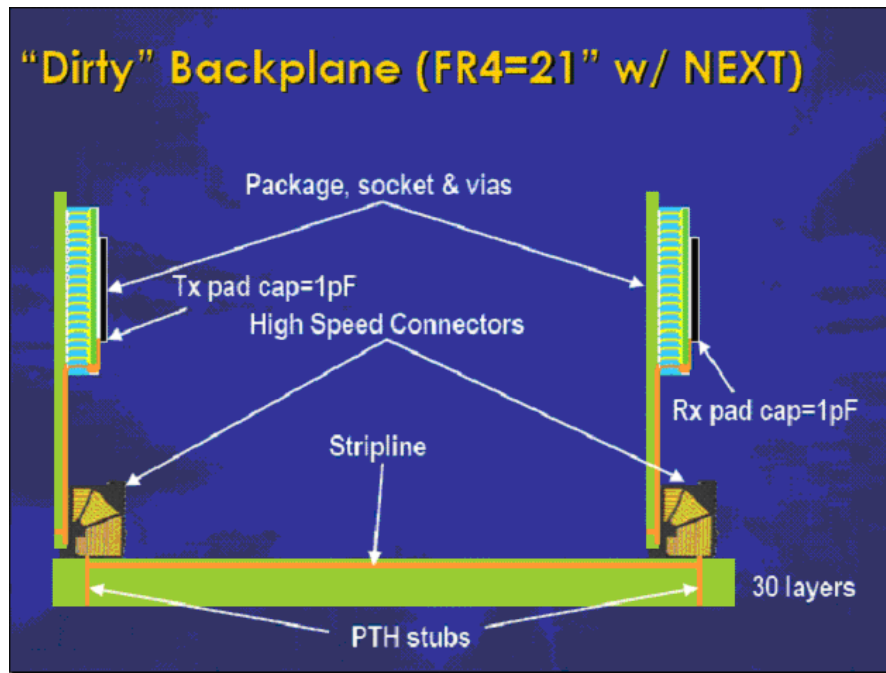
Figure 1 – Scope Based Eye Analysis

Figure 1 shows the elements of jitter analysis. It shows a 10 Gb/s signal transmitted in a near perfect environment. The analysis is performed using an SDA18000 18 GHz real-time scope. The total maximum eye-closure (total jitter) at a bit error ratio (BER) of 10^{-12} (1 trillion bits) is estimated at 16.7 ps. This has been estimated from 4.6 million unit intervals (UI). This means that if the estimate is correct, and 1 trillion bits were acquired in the eye diagram shown in the upper left hand corner, that the open area horizontally in the middle of the eye would be 83.3 ps wide.

To summarize, in the non-equalized world, channels are measured for goodness using network analysis and transmitters are measured for goodness using jitter analysis. This is perhaps an oversimplification, but it basically describes the situation.

Transmission Speed Increases Without Improving Channels

As mentioned earlier, the latest increases in transmission speed are coming without commensurate channel improvements. This is not to say that channels are not being improved, just not at the same rate as the latest desired increases in speed. In some cases it is necessary to increase (perhaps double) the speed of transmission over existing, unchanged channel infrastructure.



Courtesy of Bryan Casper of Intel Labs – used with permission

Figure 2 – High Speed Serial Data Transmission Environment³

Figure 2 shows some sources of the problems in the high-speed transmission environment. In this environment, certain channel characteristics tend to dominate system performance. These are:

1. Channel loss, while insignificant at low speeds, becomes a huge effect at high speeds, causing signals to become smaller at the receiver.
2. Channel loss as a function of frequency causes serious distortion in the signal arriving at the receiver as higher frequency signal components are attenuated relative to lower frequency components.
3. Impedance discontinuities cause reflections that bounce around in the circuit causing remnants of earlier symbols to arrive later than the intended symbol (in many cases, much later) These are due to chip pad capacitance, connectors, and stubs formed by vias and connector pins.

There are other effects, including cross-talk, that are exacerbated in this situation.

These effects lead to data or pattern dependent signal distortion that in some cases distorts the signal beyond any capability for reliable data decoding.

Equalization

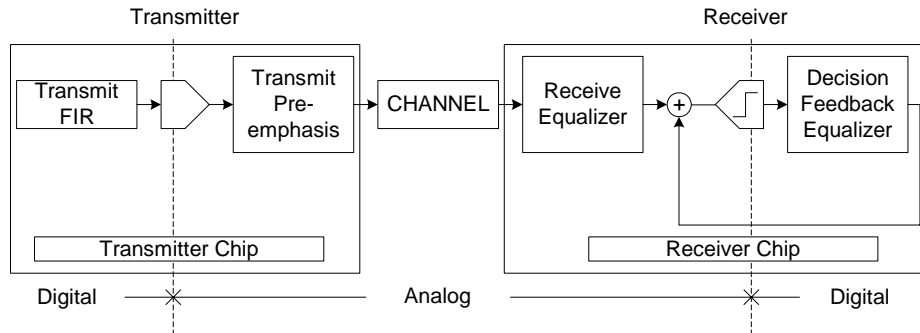


Figure 3 – Equalizer Topologies

Equalization is defined as the act of making equal. When channels have loss, specifically frequency dependent loss, frequency components of the transmitted signal arrive at the receiver with unequal magnitudes, causing distortion. Equalization, in this context, refers to the act of making these unequal magnitudes equal again, thus undoing the distortion. The same is true for phase. Frequency dependent phase variations (or group delay variations) cause frequency components of the signal to arrive at different times. So, equalization can be viewed simply as the act of making the components of a signal arrive equally in time and amplitude. Today, equalization refers to all of the techniques utilized to compensate for the adverse effects of the channel, usually confined to the data dependent effects. The various equalizer topologies are shown in Figure 3. The methods employed include transmit pre-emphasis, linear receive equalization, and decision feedback equalization.

Pre-emphasis (or de-emphasis, depending on how you want to look at it) is the pre-distortion of the transmitted signal to anticipate the channel distortion. Usually it involves increasing the magnitude of the high frequency components relative to low frequency components in anticipation of the channel loss. Transmit pre-emphasis is usually implemented either as a digital or analog finite impulse response (FIR) filter.

Linear receive equalizers (LEs) do the same as pre-emphasis, except at the end of the channel. Methods include transversal filters (also FIRs) that sum delayed versions of the received signal allowing equalization to be performed based on future and past symbols that arrive at the receiver (with a pipeline delay). These are sometimes called feed-forward equalizer (FFE).

Finally, Decision Feedback Equalization (DFE) is increasingly being employed. The DFE utilizes decisions made in the decoding of the bits at the receiver to increase the likelihood of future bit decisions being made correctly. Because there is a non-linear element in the system (the slicer, or bit decision maker), the decision feedback equalizer is non-linear.

Transmit pre-emphasis and the DFE are most commonly in use today. Usually multiple forms of equalization are utilized to leverage the pros and mitigate the cons of each equalizer type.

Figure 4 illustrates the equalization situation. It utilizes the tools described later to show the effects of the backplane, with characteristics shown in Figure 5 (ideally terminated) on serial data signal at various speeds, and the effects of equalization on those signals. It shows that at 2.5 Gb/s, equalization is not necessary, at 5 Gb/s the signal is degraded and a three tap DFE performs sufficiently, and at 10 Gb/s, the signal is unrecognizable at the receiver with the combination of a five tap linear equalizer and three tap DFE enabling proper data decoding. This is indicative of the equalization situation when data speeds are increased without making channel improvements.

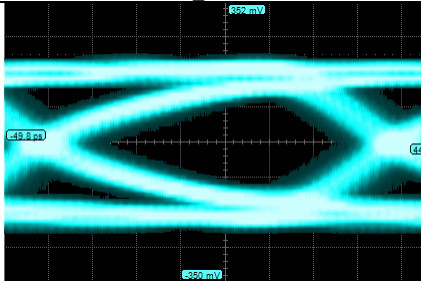
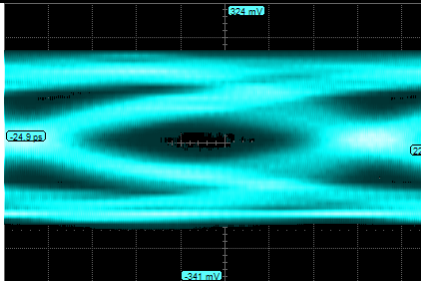
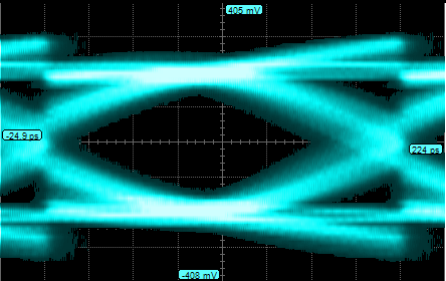
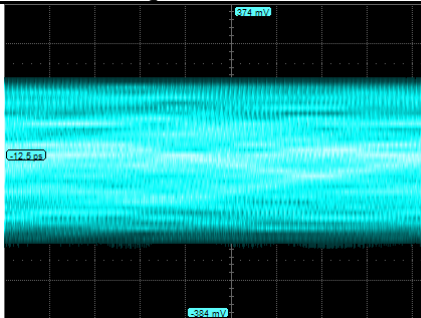
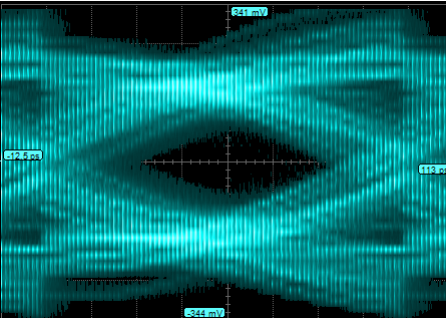
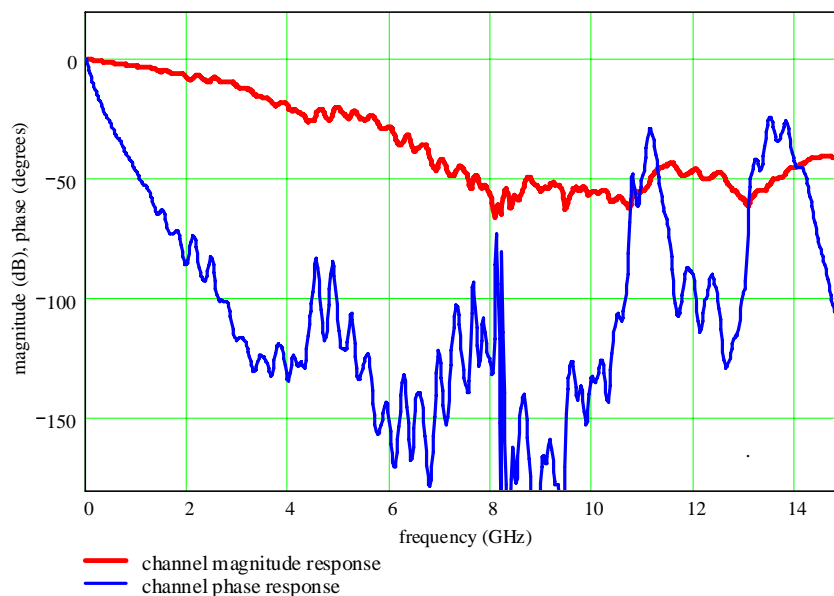
	Not Equalized	Equalized
2.5 Gb/s	 <p>Eye acceptable without equalization</p>	<p>Equalization Not Required Not Applicable</p>
5 Gb/s	 <p>Eye severely degraded without equalization</p>	 <p>3 taps DFE</p>
10 Gb/s	 <p>Data cannot be recovered without equalization</p>	 <p>5 taps FFE (2 pre + 2 post cursor) 50 ps tap delay 3 taps DFE</p>

Figure 4 – Equalization Scenarios for Various Data Rates (all data taken using backplane as shown in Figure 5)

The Challenges of Equalization from a Measurement Standpoint

The channel in an equalized system cannot be considered ideal. This immediately presents a challenge. This is because previously, the channel was measured using, for example, a VNA, and the response of the channel was examined to verify that it did not have appreciable loss at the frequencies of interest – let's say a few dB, for example. Furthermore, it was verified that it did not have lots of bumps, wiggles and resonant cavities indicating that the impedance is homogenous and there are not big effects from vias or stubs. Finally, the return loss was examined to verify that the line presented a good match and that appreciable amounts of signal are not reflected back.

This is not easy in the equalized system environment. Depending on the load being carried by the equalizers, channels may have huge amounts of loss, huge suck-outs in the response due to stubs and very bumpy magnitude and phase responses. Figure 5 is such an example. It represents one of six cases studied by the 802.3 ap Backplane Ethernet task force. The goal is to transmit 10 Gb/s over this backplane.



Tyco Backplane Example #6, Vancouver 802.3ap meeting, March 2004
Measurements performed by University of New Hampshire Interoperability Labs

Figure 5 – Channel Example⁴

Although efforts are being made to set specification limits using data as shown in Figure 5⁵, it is difficult to find correlation between backplane VNA measurements and overall system performance. In general, the measurements simply confirm the badness of the situation (without equalization). Furthermore, while it may be possible to determine whether a particular channel will work, there is still wide variation in channel performance and no good recipe that can fully determine the ability of the channel to interoperate. Also, the equalization employed is tunable or sometimes adaptive, and therefore able to accommodate large variations in channel performance that are not readily understandable simply by looking at the s parameters.

Channels themselves do not operate in isolation – they are terminated on one side at the receiver and on the other at the transmitter. As such, the channel interacts in a very complex manner with the transmitter and receiver. The differential nature of the system only adds to the complexity. Mixed mode s parameters⁶ help in simplifying things, but depending on how they are used, can simplify away significant effects. For example, in a differential system, even when differential signals are transmitted, there tends to be mode conversion (conversion of signals between differential and common modes, and vice versa) at the terminations which cause unintended components to reflect around the system.⁷ Viewing the system as a simplified differential system is often inadequate.

All of this leads to a simple fact:

Channels and Transceivers cannot be easily measured independently to guarantee interoperability with each other.

To compound the situation, the connectors, cables and channels are being designed to specifications that are measurable with network analyzers, measurements which are usually readily available, but the important chip specifications are not readily available. This is not because the chip manufacturers are hiding something, but because there is no language to easily describe the conditions of interoperability. Sometimes, a chip manufacturer will specify channels that the chips will operate with – but these tend to be overconstraining on the channel design (i.e. the channel can be much worse and still work fine), or they specify channel conditions that will not work – but these tend to be underconstraining (i.e. there will be situations where the channel specifications are met and the system still doesn't work). Even when things do work, there is not a good idea of what the margin of error might be – although some insight is gained by examining how close various equalizer tap coefficients are to the edges.

The bottom line is that there are challenges summarized as follows:

1. Engineers lack the language to define interoperability conditions between channels and transceivers.
2. Compliance tests dictate tests that in the end, don't meet the goal of guaranteeing interoperability.
3. When systems don't interoperate, it is not possible to determine which portion is not working – and when they are working, it is not possible to determine margins.

It is generally not possible to get access to the equalized signal. Referring to Figure 3, this means that probing points are usually limited to the point immediately at the end of the channel (which is the input to the receiver chip). The point at the input to the slicer is inside the chip and is not considered accessible under normal circumstances.

Possible Solutions

There are several possible solutions to the above problems. One already exists – the plugfest. A plugfest is where system elements are interchanged in a manner that tests large numbers of permutations. This has the benefit of gaining confidence that systems will interoperate and serves to provide some ability to weed out the culprits. When devices do or don't interoperate,

we are still left with the problem of determining why. So, while this may serve as a Monte-Carlo method for compliance testing, it does not help much with debugging problems. Another, similar method is the development of “golden” or reference implementations. These include multiple corner-case channels, and reference transmitters and receivers. The idea is that if a transceiver pair can interoperate over the corner-case channels (assumed to represent the worst possible situations), then they are deemed interoperable. Similarly, if a receiver can receive data from a reference transmitter transmitting over the corner-case channels, then it is deemed good. There are still problems resulting from these methods:

1. The transmitters and receivers must still be tuned to the corner-case channels.
2. The corner-case channels must be truly representative of the worst-case situations.
3. The channels must be somehow compared to the corner-case channels to determine if they are truly better.
4. The corner-case channels and reference transceivers must exist in large quantities for testing and must all be similar in characteristics.
5. The tests are time-consuming and expensive (i.e. there is lots of connecting and reconnecting of reference or corner-case components).
6. The situation still lacks debugging capability.

These solutions are a possible start, but do not solve all of the problems.

Another potential solution that will not be explored here is instrumentation built into receiver chips themselves that allow on-board bit error rate testers (BERTs) and BERT-scan capability.

Simulation

Because of the intractability of the problems involved in analyzing equalized systems, there has been an increased reliance on simulation. This trend is increasing and is present in many other areas of engineering and has spawned legions of the often maligned “Spice-Jockey”. Although an over reliance on simulation, at least a reliance that supplants analysis and understanding with massive amounts of simulated data cannot be condoned, it is increasingly the solution of choice. That being said, the real problem with simulation is that oftentimes the model being simulated does not match reality, either because it is oversimplified and leaves something out or that oftentimes simulation is just too slow. It should be mentioned that there are simulation tools, such as StatEye, that are specifically designed for simulation in the equalized environment.⁸

What if:

1. corner case channels (backplanes, connectors, cables, via models etc.) could be represented in software?
2. reference equalizers or equalized receivers could be represented and optimized in software?
3. what if simulation techniques could be leveraged without sacrificing correlation with reality through oversimplification?
4. what if simulation techniques were fast?

Advanced Tools for Serial Data Measurements

LeCroy has produced two tools that aid in the measurement of equalized serial data systems and address many of the problems outlined. These tools are: Equalizer Emulation and Virtual Probing.

The equalizer emulation feature of Eye Doctor allows you to emulate ideal equalizers.

The Ideal Equalizer Emulation feature enables the emulation of various equalizer configurations including:

1. Linear transversal, tapped delay line, linear feed-forward equalizers.
2. Clock recovery.
3. Decision Feedback equalization.

The equalizer emulation feature allows you to see the recovered data, clock and equalized waveform. The equalizers can be specified by tap delay, number of taps, and tap coefficients, or the equalizers can be automatically trained to optimize settings utilizing blind adaptation.

The Virtual Probing feature enables a variety of advantages in signal probing situations including:

1. The ability to compensate for probe loading effects by allowing you to see waveforms that occur in a circuit as they would with and without the probe connected to it.
2. The ability to acquire waveforms that occur in locations other than the probing point.
3. The ability to acquire waveforms that would occur in circuit configurations different than the configuration actually used for the measurement.

Virtual probing works by providing:

1. An accurate description of the circuit configuration both in the probing configuration and the configuration in which you want to see the virtual waveforms.
2. Accurate measurements, models or assumptions about the behavior of the various circuit components in the system, including the probing elements and the scope.

Having this information, you simply provide this information to the scope in two key formats:

1. A *system description file* – a file that describes the various circuit descriptions to the scope.
2. S parameter files – Touchstone®⁹ format files that describe, through S parameters, the behavior of the circuit components.

Equalizer Emulation

The equalizer emulation components consist of processing components configured in LeCroy's processing web. The main component is a full, ideal equalized receiver.

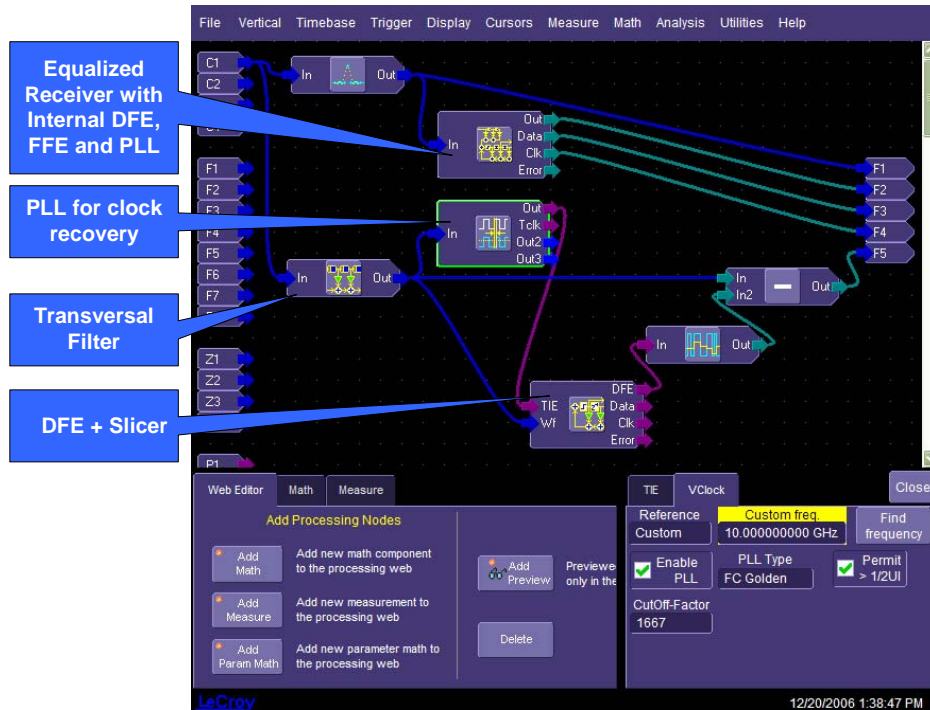


Figure 6 – LeCroy Processing Web with Equalizer Emulation Components

The input pin to the equalizer receiver component is an analog serial data waveform – ideally the waveform that appears at the receiver in a system.

The four output pins are:

1. Out – the equalized waveform that appears at the input to the slicer in the emulated receiver.
2. Data – a data waveform representing the waveform at the output of the slicer in the emulated receiver.
3. Clock – the clock waveform representing the clock recovered by the receiver.
4. Error – Provide an output representing the error between the ideal data waveform level at the slicer output and the equalized waveform at the slicer input to allow users to use this data to implement their own optimization algorithms.

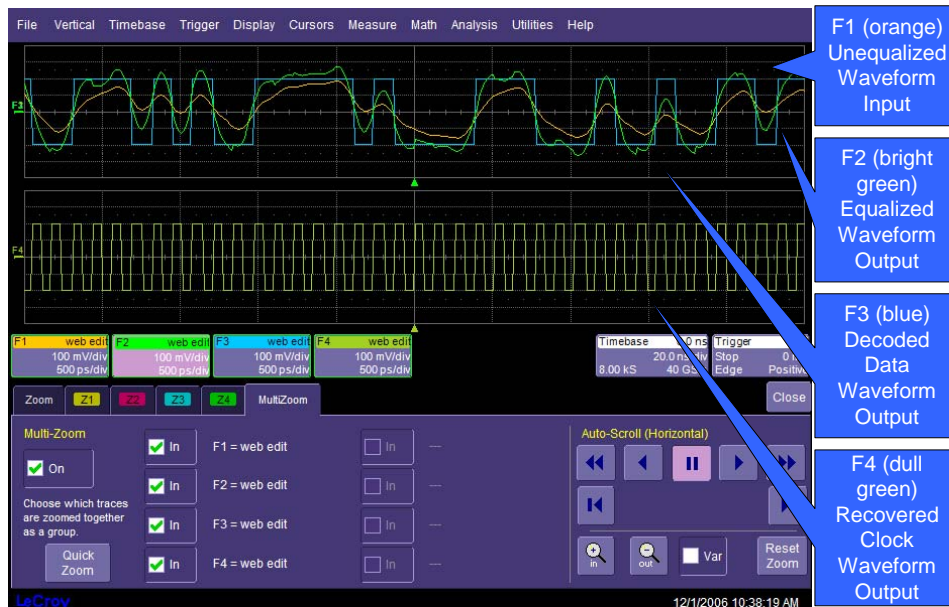


Figure 7 – Equalized Receiver Emulation Component Output Waveforms

The equalized receiver is made up of several internal processors. These processors are all available for use independently and include:

1. An equalizer optimizer component – determines optimum FFE and DFE coefficients from the input waveform.
2. A tapped delay line equalizer – provides linear, feed-forward equalization.
3. Time Interval Error component – acts as a PLL that extracts the clock from the tapped delay line equalizer component output.
4. A decision feedback equalizer/receiver component – provides decision feedback equalization and serves to determine the decoded data and extracted clock.



Figure 8 – Equalized Receiver Configuration Dialog



Figure 9 – Equalized Receiver Dialog Showing Manual Filter Tap Settings

Let's explore some of the problems that the equalized receiver solves.

The equalized receiver is a simple component, conceptually, that allows for specification of a reference equalized receiver. For example, if a new standard specifies that a receiver must be able to decode data at a particular error rate based on a three tap DFE, then signals can be probed at a reference termination and applied to the oscilloscope, an equalizer specified as a three tap DFE can be configured and optimized (either automatically or through user or standards defined methods) and the equalized output of this component can be analyzed through the traditional, real-time scope based methods of jitter analysis to qualify the transmitter and channel together. In other words, this component can be used to qualify a backplane and transmitter combination. If the reference receiver can receive the data and the real receiver cannot, then the receiver fails the test. If the reference receiver cannot receive the data, then the real receiver cannot be expected to and the transmitter and/or backplane fails the test. The receiver analysis has therefore been separated from the backplane analysis. The equalized receiver therefore forms the basis for a software defined reference equalized receiver.

While the equalization emulation is ideal and does not fully emulate hardware, it can be used to determine hardware characteristics that are necessary for proper data reception. For example, if the data could not be received properly with three optimized DFE taps, but could be received

properly with the addition of three linear tapped delay line equalizer taps, then an argument could be made for either channel improvements or more specifically, the addition of some linear receive equalization or transmit pre-emphasis.

If the transmitter used in this testing was a reference transmitter, then these methods can be utilized to completely budget the backplane and equalizer development. It breaks the problem into making backplanes that can transmit data that can be equalized with a three tap DFE, for example, and designing three tap DFE equalized receivers that emulate the ideal equalized receiver.

Interestingly enough, we are no longer talking about channel characteristics in a VNA measurement context – which is what will be needed for compliance testing in the equalized system environment.

Equalizer emulation bridges the gap between the waveform that appears at the receiver chip (interesting, but not useful for measurements) and the waveform that appears inside the receiver chip at the input to the slicer, which is the truly important waveform.

Virtual Probing

Equalizer emulation provides some solutions to a number of problems in the equalized system, but not all – and its use is limited by some other problems that creep in.

First, if you probe at the receiver at the end of a very lossy channel, the signal received has undergone severe distortion due to the channel, especially high frequency attenuation. Speaking simplistically, an effect of (at least linear) equalization is to boost the high frequency components, thus undoing the distortion. In real-time scope based measurements, this can be a problem since any processing that boosts frequencies will boost the noise floor along with it. Depending on the amount of boost needed, this can present a measurement challenge. Figure 10 illustrates this situation with four simulations. The top row is a 5 Gb/s simulation and the bottom row is a 10 Gb/s simulation. The left column is without virtual probing while the right column shows the noise benefit of virtual probing. In the case where virtual probing is not employed, the waveform at the receiver is acquired with a signal-to-noise ratio (SNR) of 30 dB (approximating the SNR of an 18 GHz scope). A linear receive equalizer is employed and optimized, causing high-frequency boost, as expected. In the case where virtual probing is employed, the waveform is acquired at the transmitter with the same 30 dB SNR. The virtual probing component generates the waveform at the receiver from the acquired waveforms, which causes large amounts of high frequency attenuation of both the signal and the noise injected by the scope. When the linear receiver equalizer is added, it boosts the high frequency components, but the net result is less noise generated by the measurement system, yielding less noisy waveforms. This is more realistic, since in most cases, most of the noise comes from the measurement instrument.

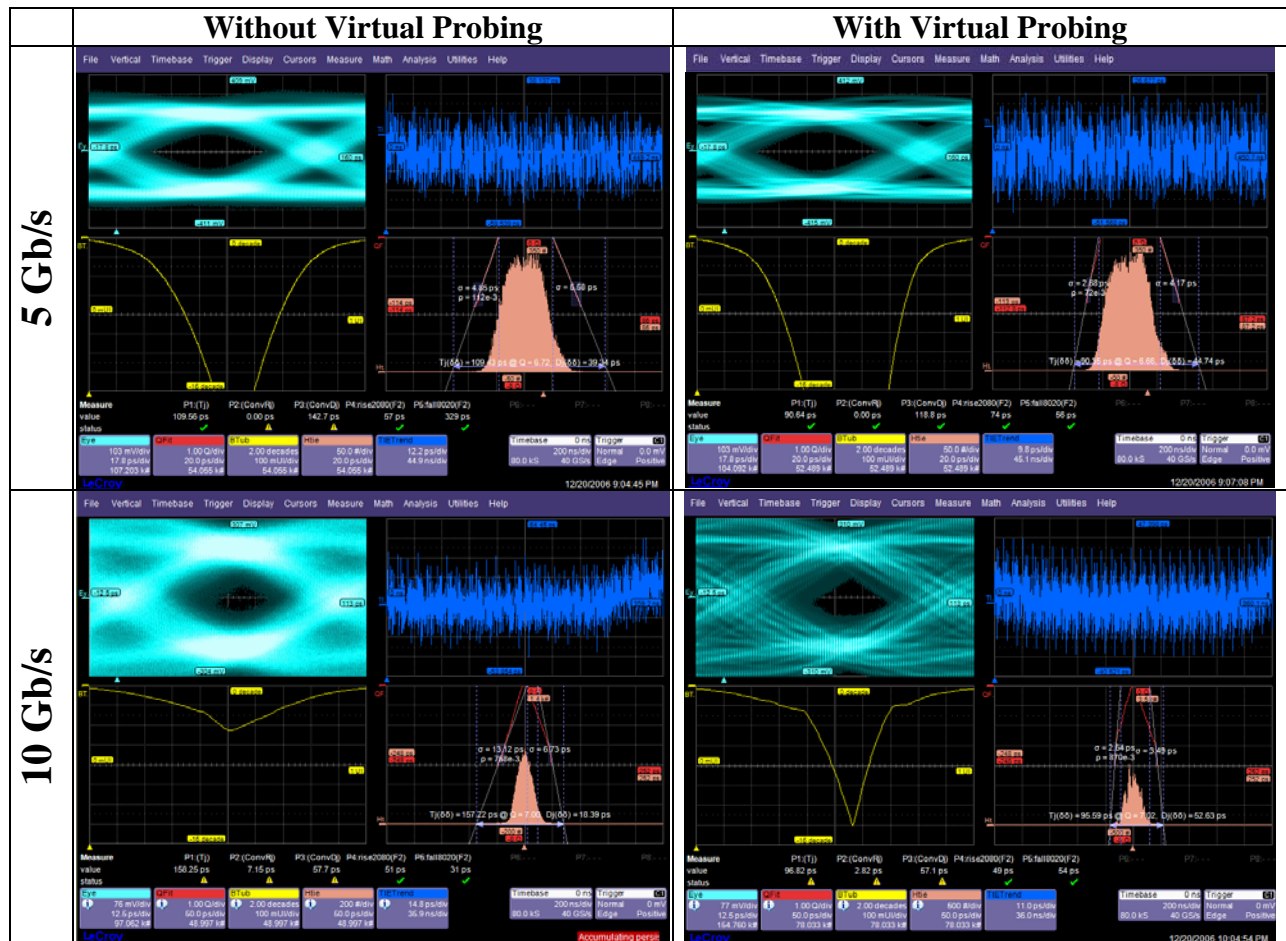


Figure 10 – Noise Comparisons

In addition to noise injected by the measurement instrument, probes, cables, and scope inputs have other issues. Most probing methods especially at high frequencies, interact with the circuit under test to create measurement errors. There are ways of mitigating some problems (for example through the use of inline attenuators) but not all. There are also ways to mitigate the severe loss of the channel through the use of inline coaxial equalizers¹⁰, but the effects of these equalizers must be undone in order to properly show the waveform present. The goal is usually to obtain the waveform at the receiver chip *in the configuration that it will be used*, not the configuration it is measured (i.e. with the measurement instrumentation de-embedded).

Finally, we still have the problem of needing actual backplanes to make all measurements at the receiver with the solutions proposed up to now.

These are all measurement problems. The virtual probing component provides the solution to these problems.

Virtual probing works by utilizing s parameter data that describes system circuit elements in conjunction with a circuit topology description that describes the interconnection of the system circuit elements to produce filters capable of converting waveforms measured at one portion of

the circuit into waveforms present simultaneously at other circuit locations or under different circuit configurations. Some think of this as de-embedding, but it is actually more than that and does not come with some of the difficulties of implementation. The system description defines how you will measure a signal and the system and location where you want to see the signal.

Figure 11 shows one possibility. In this configuration, the virtual probing component would produce waveforms that represent VRP and VRM (plus and minus voltages at the receiver) as a function of VMD measured on a scope channel with a probe connected at the transmitter. These waveforms are the waveforms that are present with the probe connected.

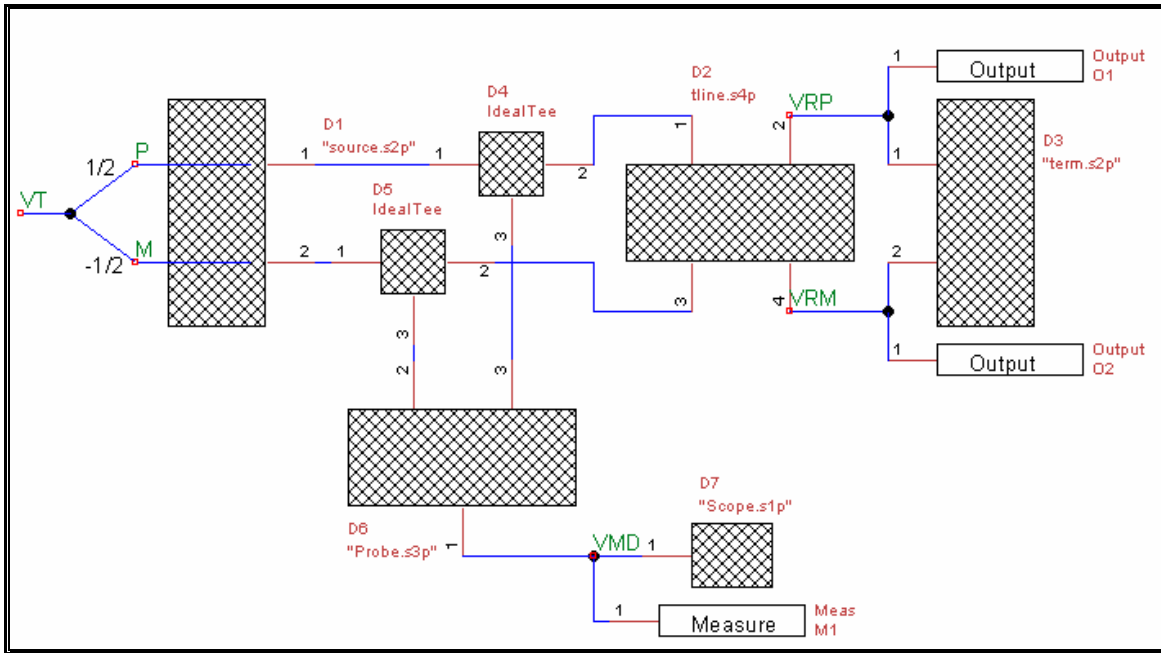


Figure 11 – Virtual Probing System Diagram Example #1

Figure 12 is a similar example, except that VRP and VRM become waveforms that would be present without the probe connected – but they are waveforms based on measured waveforms at VMD with the probe connected. Figure 12 is an example of a measurement where the measurement system is de-embedded. Again, it should be stressed that the user doesn't have to do any calculation to do the de-embedding – just provide accurate s parameter data and an accurate description of the system configuration with and without the measurement instrumentation connected to the system.

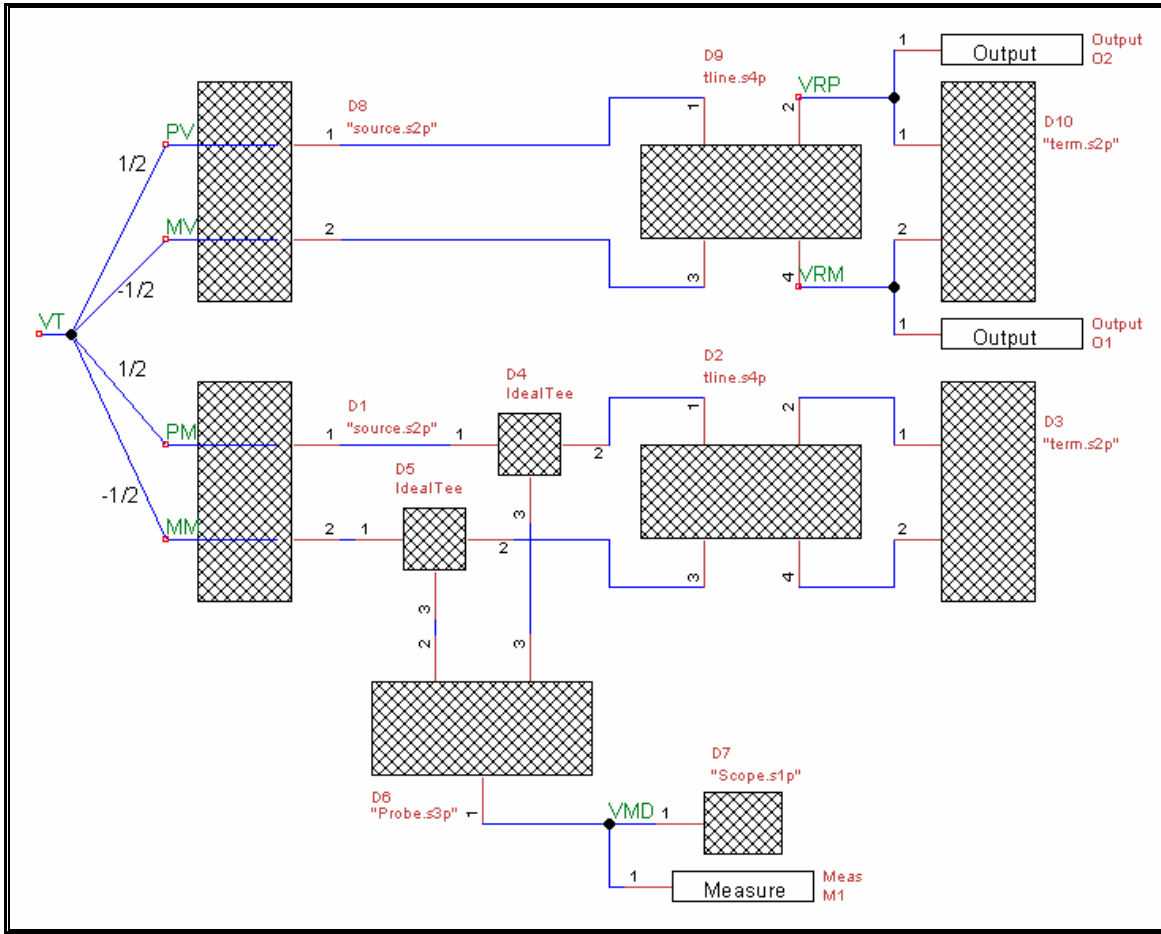


Figure 12 - Probing System Diagram Example #2

It should be noted that in this particular measurement, where VMD is a single differential measurement, the user does need to impose an extra constraint on the system. In this case, the constraint of a balanced differential transmitter has been imposed – but it must be emphasized that this constraint does not mean that the signal components in the system are all differential – this constraint still handles common mode components in the system with a balanced transmitter. This constraint can be removed by adding another probed measurement, or by probing with two single-ended probes.

Traditionally, as in Figure 11 and Figure 12, the definition of where you want to see the signal will be at the end of a channel or backplane and the measurement location will be signal points closer to the transmitter. Virtual probing therefore allows the receiver signal to be boosted through equalization without incurring the noise penalty typically incurred by probing at the receiver.

Because the system is described in software using s parameters, it is easy to switch backplanes or channels in your measurements by simply specifying different s parameter files. This solves the problem of having multiple physical versions of backplanes for compliance testing that must have the same characteristics – tests share the same s parameter files and therefore have exactly the same characteristics. Since the s parameter files are in the universally accepted Touchstone

format, the files can be generated from all VNA or time-domain network analyzers (TDNA) instruments. They can even be generated from RF/microwave simulators such as electromagnetic simulators and therefore the circuits need never physically exist.

Virtual probing, like equalizer emulation, is utilized in the LeCroy processing web as shown in Figure 13. Here it is shown measuring two single ended signals on channels 1 and 2 and producing a differential voltage that is applied to the equalizer emulation component. In this configuration, we are measuring waveforms at the transmitter, producing the differential waveform at the receiver along with the equalized waveforms.

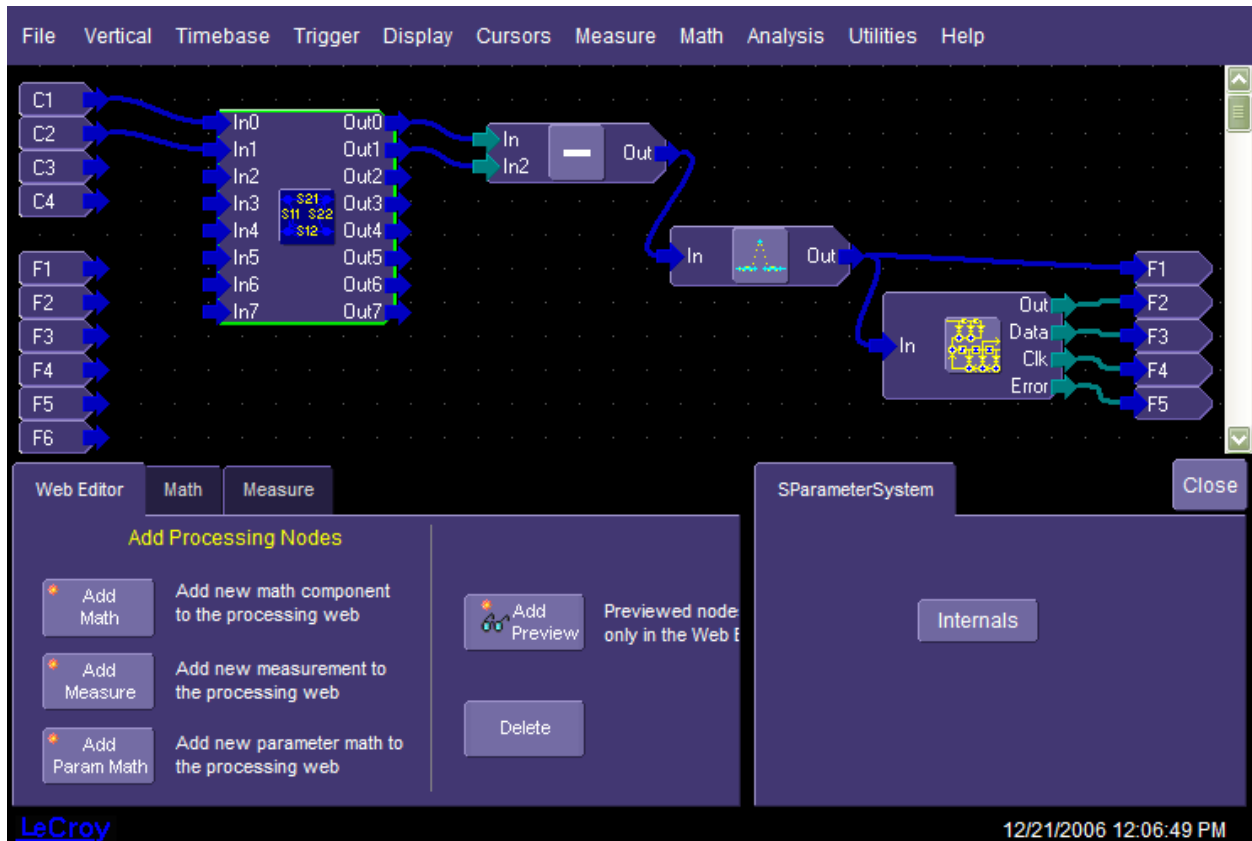


Figure 13 – Virtual Probing Component in Processing Web

Virtual Probing Component with up to eight inputs and outputs

Log shows System Description File compilation results

Component configured by specifying System Description File, and other measurement characteristics

Compiled System dictates component input and output pin assignments

System Description File references touchstone format s parameter files that describe network behavior

System Description File describes devices (networks), interconnects and measurement and output waveform nodes

FREQ	S11 MAG	S11 ANG	S21 MAG	S21 ANG	S12 MAG	S12 ANG	S22 MAG	S22 ANG
10.000	0.006417	-76.862383	0.990105	-15.961459	1.014052	-13.956389	0.017676	17.902539
20.000	0.004534	14.453200	0.989363	-31.838911	0.990916	-31.329403	0.009879	-48.903844
30.000	0.006822	-21.295701	0.993067	-47.963954	0.991417	-47.364315	0.004552	-61.501254
40.000	0.007014	-45.449837	0.989357	-63.870409	0.991044	-63.519627	0.004078	-29.067818
50.000	0.005194	-71.086289	0.989579	-79.794395	0.989965	-79.571359	0.006262	-55.059562
60.000	0.004312	-100.698489	0.988497	-95.829078	0.989736	-95.526974	0.007932	-104.200769
70.000	0.003954	-133.068589	0.987480	-111.78966	0.989137	-111.478889	0.007063	-141.525301
80.000	0.003382	-176.346025	0.987239	-127.373812	0.988419	-127.373812	0.005932	-177.082431
90.000	0.003855	-133.9817	0.987239	-127.373812	0.988419	-127.373812	0.004700	-157.7111

Figure 14 – Virtual Probing Component Configuration

Figure 14 illustrates the configuration of the virtual probing component. Once placed in the processing web, a system description file is specified. The system description file contains a

netlist that describes the system. It is a text description of the diagrams like those shown in Figure 11 and Figure 12. The system description file describes the system, referring to s parameter files for device behavior. The file is compiled from within the component either automatically when waveforms arrive at the component input pins, or by pressing a button on the dialog. A log shows the progress of compilation and helps in the debugging of problems with the system descriptions. Once compiled the log will show no errors, and the dialog will contain node names for both the measured nodes and the output waveform nodes along with the component pin assignments. The component is then wired into the processing web and when measured waveforms arrive at the input, the component produces the desired output waveforms.

Internal to the component, it has generated filters. The filters are FIR filters as shown in Figure 15. Figure 15 shows four impulse responses to exemplify the situation where two voltages are measured single-ended at a transmitter and two output voltages are generated at the receiver (with the idea that these outputs would be subtracted to form a differential voltage).

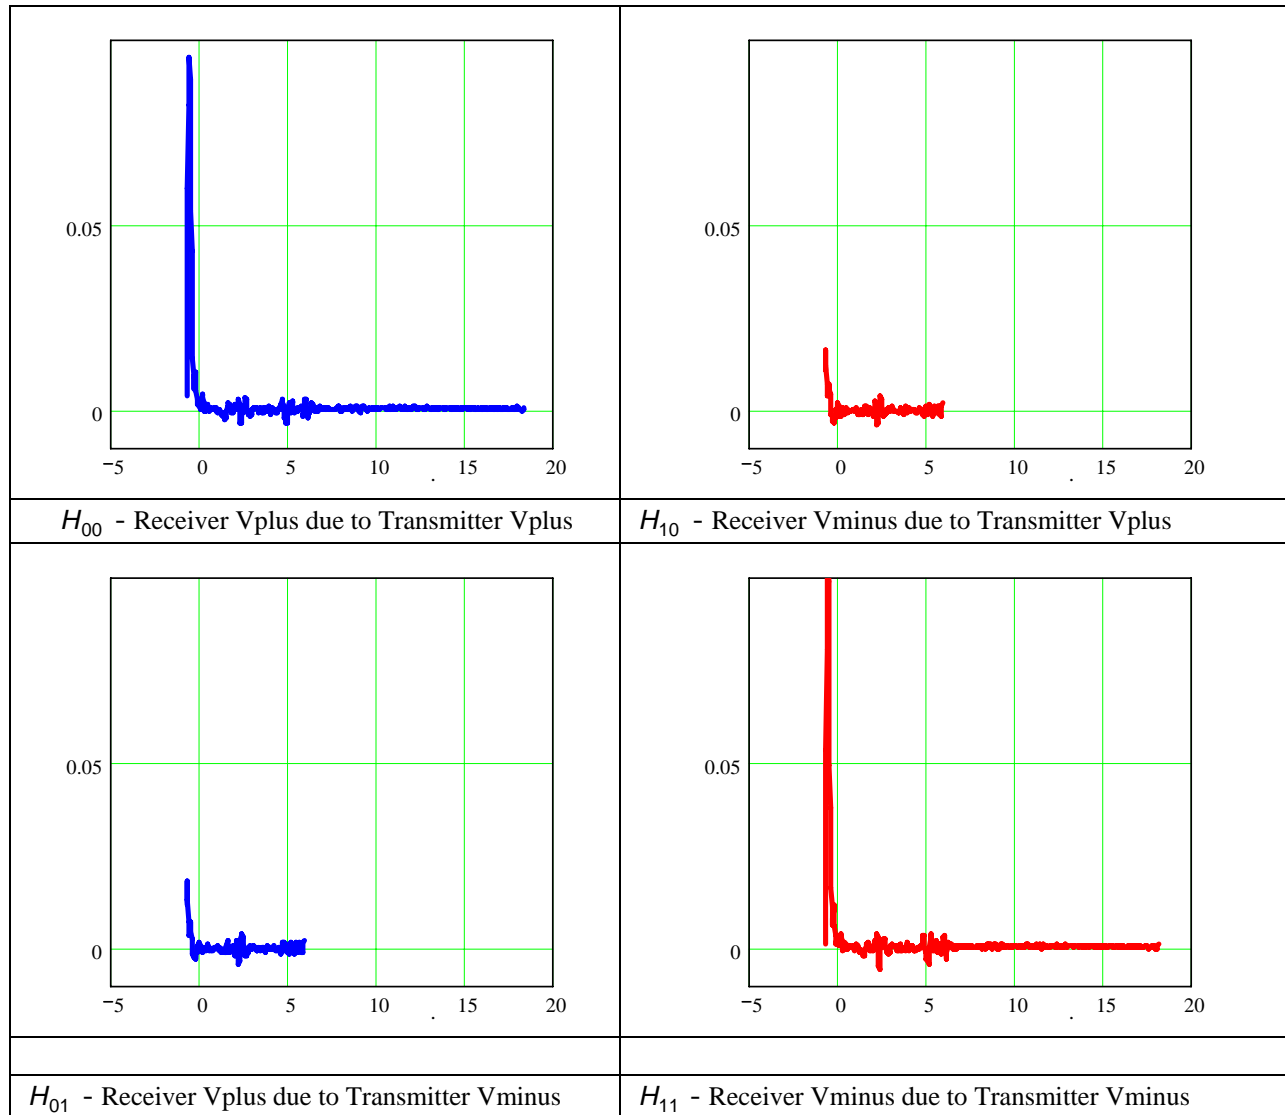


Figure 15 – Filter Impulse Responses

The filters in Figure 15 are illustrative of several facts. First, although the effect of the minus voltage at the receiver on the plus voltage at the transmitter (and vice versa) is small, it is not negligible (and can be bigger in some cases). The virtual probing component handles this situation. Second, The impulse responses occur before time zero. This is because the cabling to the scope was longer than the backplane. The scope handles this properly and produces all waveforms in a manner that properly reflects the time relationships. Thirdly, the impulse responses vary in size. This is because the system optimizes filter length based on the significance of the response on the output measurement.

The filter topology used by the system is always one of summed versions of multiple filtered waveforms as shown in Figure 16. This is due to the nature of the math involved in solving the system.





Figure 16 – Example Two Input – Two Output Filter Topology

Contrast between Virtual Probing and Other Methods

In the last section, I showed how virtual probing solved measurement problems, but there are other problems as well. Mostly, these are what I term simulation or modeling problems. It is important to see how virtual probing addresses these problems, as well.

Figure 17 shows an oversimplified view of the channel. This view models the channel as the differential through response of the channel and assumes that this is the transfer function of the system. Methods of converting four-port sets of s parameters to so called mixed-mode are well known¹¹. This widely held view models the system as the differential mode thru response¹².




$$H = sdd_{21}$$

$$V_o = H \cdot V_i$$

Figure 17 – Over-simplistic Channel View

The view of a system shown in Figure 17 is correct for the differential response if the differential and common mode terminations are perfectly matched to the channel. It neglects common mode components which, as I said before, will usually be present even if the transmitter is perfectly balanced.

Another view of a system is shown in Figure 18. This is the simulator-jockey view. The circuitry or model in the black box can be anything and can be made to arbitrarily accurately match reality. Not shown is the transmitter model in this simulation. Some think that circuit simulation is practically useless¹³ but no one will argue that it tends to be very slow. The output voltage is usually whatever the simulator puts out, but you can infer a transfer function from input to output.



$$H = \frac{V_o}{V_i}$$

Figure 18 – Simulated Channel View

The virtual probing view of a simple system is shown in Figure 19, which can be represented as Figure 20.




Figure 19 – Virtual Probing System View

Figure 20 is a signal flow diagram representation¹⁴ of a two port (differential) source, labeled S, connected to a four port channel, labeled T, terminated in a two port load, labeled L. The diagram shows the flow of power waves in a system. The system is assumed linear.

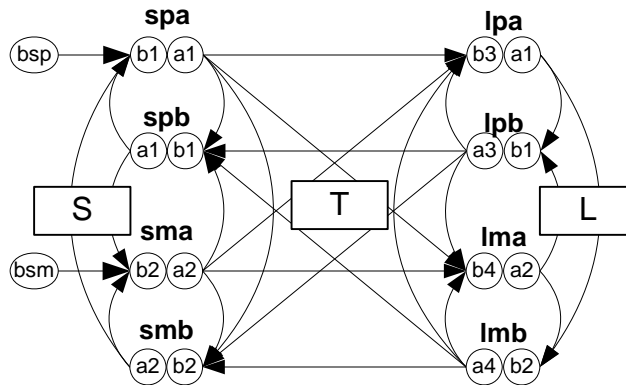


Figure 20 – Signal Flow Diagram

Figure 20 serves to illustrate the complex flow of signals in the system. It also serves to illustrate that the source and load terminations may send waves around the system in a complex manner.

The virtual probing component is provided with the s parameters of the devices and the device interconnects, and is capable of solving the sets of equations implied. Equation 1 through Equation 4 show the solution of this system. The details of how it solves them are left out and are beyond the scope of this paper.

$$S_i = \begin{bmatrix} 1 & 0 & -sS_{11} & -sS_{12} & 0 & 0 & 0 & 0 \\ 0 & 1 & -sS_{21} & -sS_{22} & 0 & 0 & 0 & 0 \\ -sT_{11} & -sT_{12} & 1 & 0 & 0 & 0 & -sT_{13} & -sT_{14} \\ -sT_{21} & -sT_{22} & 0 & 1 & 0 & 0 & -sT_{23} & -sT_{24} \\ -sT_{31} & -sT_{32} & 0 & 0 & 1 & 0 & -sT_{33} & -sT_{34} \\ -sT_{41} & -sT_{42} & 0 & 0 & 0 & 1 & -sT_{43} & -sT_{44} \\ 0 & 0 & 0 & 0 & -sL_{11} & -sL_{12} & 1 & 0 \\ 0 & 0 & 0 & 0 & -sL_{21} & -sL_{22} & 0 & 1 \end{bmatrix}^{-1}$$

Equation 1

$$VS = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} Si_{0,0} & Si_{0,1} \\ Si_{1,0} & Si_{1,1} \\ Si_{2,0} & Si_{2,1} \\ Si_{3,0} & Si_{3,1} \\ Si_{4,0} & Si_{4,1} \\ Si_{5,0} & Si_{5,1} \\ Si_{6,0} & Si_{6,1} \\ Si_{7,0} & Si_{7,1} \end{bmatrix}$$

Equation 2

$$H = \begin{bmatrix} VS_{2,0} & VS_{2,1} \\ VS_{3,0} & VS_{3,1} \end{bmatrix} \cdot \begin{bmatrix} VS_{0,0} & VS_{0,1} \\ VS_{1,0} & VS_{1,1} \end{bmatrix}^{-1}$$

Equation 3

$$Vdo = [1 \quad -1] \cdot H \cdot \begin{bmatrix} Vsp \\ Vsm \end{bmatrix} \quad Vco = \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix} \cdot H \cdot \begin{bmatrix} Vsp \\ Vsm \end{bmatrix}$$

Equation 4

Equation 3 shows how the virtual probing component solves for the transfer function of the system and Equation 4 shows how the transfer function is utilized to generate the differential and common mode signal present at the receiver, based on voltage measurements at the transmitter.

What is shown here is illustrative of several key concepts, that I shall point out:

1. The virtual probing component allows solutions that do not require over-simplifying assumptions of the system component interactions.
2. The solution allows for output of differential mode, common mode, or any voltages present at nodes in the system.
3. The component will utilize filters (from these transfer functions) that may be combinations of multiple measured voltages.
4. This component is not a simulator – transfer functions are generated that convert measured voltages into output voltages. In essence, the portion of the system (the channel) which is most easily modeled as a linear system is utilized in conjunction with measured voltages to form output voltages.
5. While a simple example was shown with a somewhat complicated solution, the component hides this complexity from you. Making a bigger system with more devices and more nodes, while conceptually difficult, is easily handled by the virtual probing component – again, all you provide are the s parameters for the networks and the network interconnections.
6. Again, while a simple example was shown, a bigger example would have probes, cables and scope inputs in the model, allowing the component to undo the effects of these sources of measurement error. All that is needed (beyond reasonably good signal fidelity and appropriate bandwidth of measurement components) are accurate models of these components – with the models always described via s parameters.
7. Since all of the components utilize s parameters, there is a common interface to virtually all other measurement and modeling systems. For example, if spice topological models

are used for a network description, most simulators are capable of converting the model to s parameters. If they can't directly, you can measure the s parameters yourself by measuring voltages and currents at various frequencies in your simulations.

8. While the conversion of circuits to s parameters via simulation and the measurement of s parameters using a VNA is a time-consuming process, once you have this data, the solution of the system diagram using the virtual probing component is fairly fast, and the use of the filters generated is lightning fast. A reasonable estimate of the order of magnitude of the times involved are about 10 seconds for system model compilation and about 1 Mpoint/s system throughput once the filters have been generated. This enables large amounts of serial data to be analyzed quickly.
9. Although the example shown is a properly constrained system, it can also be used to solve overconstrained cases (i.e. the case where the matrix to be inverted in Equation 3 is not square)¹⁵.
10. Finally, the tool runs inside your scope. So if you've been comfortable with scope measurement tools in the past, like those provided by the LeCroy Serial Data Analyzers, this tool allows you to bridge the gap in equalized systems and to continue using these tools.

Co-simulation

Virtual probing seems a lot like simulation, but a better word for this type of measurement system is co-simulation. Co-simulation is defined as a system that consists of elements that are simulated (usually using a digital computer) and actual hardware elements. In the cases outlined so far, the transmitter is generally real, the backplane channel is either simulated or real, and the receiver is generally simulated in an ideal sense. Measured waveforms are applied to the oscilloscope channels live. Earlier, I stated that channel measurements using network analyzers are static measurements with results that completely define the system. Since the channels are generally linear, passive elements, it seems logical to utilize these measurements in a simulation environment – in fact, they usually are used this way. The transmitter, on the other hand, is much more complicated and hard to describe. Therefore, it is more dynamic and varying in behavior and not necessarily well suited for simulation. The co-simulation environment offered by virtual probing is a good utilization of dynamic and static system elements and, in the end, produces measurements that are better matched to reality and generated much faster than data provided only through system simulation.

Summary

Equalization is increasingly being relied upon for speed increases in serial data transmission. The use of equalization present many challenges in compliance testing and test and measurement in general. LeCroy is providing two key tools to help: Virtual Probing and Equalizer Emulation. These tools have been shown to alleviate some of the problems involved in testing equalized systems.

References

- ¹ D. A. Smolyansky and S. D. Corey, *Characterization of differential interconnects from time domain reflectometry measurements*, Microwave Journal. Vol. 43, no. 3, (2000), pp. 68, 70, 72, 74, 76+
- ² K. Kurokawa, *Power Waves and the Scattering Matrix* IEEE Trans. Microwave TheoryTech., vol. MTT-13, pp. 194-202, Mar. 1965
- ³ B. Casper, *Tradeoffs of Receive and Transmit Equalization Architectures*, ISSCC2006
- ⁴ J. D'Ambrosia and J. Abler, *The Need For a Normative Channel Model Approach*, IEEE Vancouver Interim Meeting, January 24th, 2005
- ⁵ http://grouper.ieee.org/groups/802/3/ap/public/jan05/dambrosia_01_0105.pdf
- ⁶ D. E. Bockelman and W. R. Eisenstadt, *Combined Differential and Common-mode Scattering Parameters: Theory and Simulation*, IEEE Trans. Microwave Theory Tech., vol. 43 (Jul. 1995), pp. 1530-1539
- ⁷ H. Johnson, *Differential Termination*, Electronic Design News, June 5, 2000
- ⁸ A. Sanders, M. Resso and J. D'Ambrosia, *Channel Compliance Testing Utilizing Novel Statistical Eye Methodology*, DesignCon 2004 Proceedings.
- ⁹ *Touchstone® File Format Specification Rev 1.1*, EIA IBIS Open Forum, <http://www.vhdl.org/pub/ibis>
- ¹⁰ G. Foster, *Measurement Note: Using the XFI EQ Filter to Equalize a Channel to Enable Eye Measurements*, www.bertscope.com
- ¹¹ *Three and four port s-parameters: calibrations and mixed mode parameters*, Anritsu Application Note, 11410-00279, Nov. 2001
- ¹² *Modeling a Differential High-Speed Backplane in Simulink*, www.mathworks.com
- ¹³ B. Pease, *What's all this Spicey stuff, anyhow? (Part IV)*, Electronic Design, Sep 29, 2003
- ¹⁴ *S-Parameter Design*, Hewlett Packard Application Note 154, March, 1990, (5952-1087).
- ¹⁵ G. Strang, *Introduction to Applied Mathematics*, Wellesley-Cambridge Press, 1986, pg. 38