The Case of the Closing Eye

Ransom Stephens, Ph.D.

February, 2008

Abstract:

It would be nice if digital signals could forget their analog origin, but insidious processes like dispersion, skin effect, multiple reflections and random noise force us to think of a signal as a complicated waveform instead of a tidy stream of bits. A combination of time-domain and frequency-domain analysis distinguishes which problems can be fixed with simple circuit modifications, clever tricks at the transmitter or receiver, from those that require major investment.

The higher the data rate, the harder it is to work in the imaginary world of digital electronics. Signals with fast rise/fall times have steep edges that make waveforms look digital at the transmitter but not at the end of a trace. Once you get over a few hundred MHz, nothing but trouble comes from overlooking a digital signal's analog reality and, at a few GHz, you might as well have "microwave" stamped on your business card.

As long as we use FR-4 (Flame Retardant type-4) circuit boards and simple coaxial cables, we'll have to work with closed eyes.

In this paper, we look at the most common causes of eye closure and then, having accepted our closed-eye fate, we face the reality of microwave-circuit physics. Our approach to the microwave stuff is deeply conceptual; I guarantee that I can even keep you awake through *S*-parameters. Once we've faced reality, we do what engineers do best: fix it.

The Suspects

Figure 1 is a diagram of a serial data system with the primary components: Transmitter, transmission path, and receiver. The transmitter is our ally in keeping eyes open. A nice clean flexible transmitter providing the best rise/fall time with no Duty-Cycle Distortion (DCD), minimal random noise and random jitter, and the ability to emphasize different aspects of the waveform can fix a lot of problems. It is the receiver who is the victim of the closed eye – it has to identify ones and zeros at a Bit Error Ratio (BER) of at least 10^{-12}. In a well designed system, the transmission path is where the eye-closing perpetrators hide.

But you're an engineer. If you understand what's causing a problem, approximately 97.3% of the time you can solve it. If a component acts like an unwanted filter, then follow it with a filter that does the opposite; if it's lossy, stick an amplifier in there.

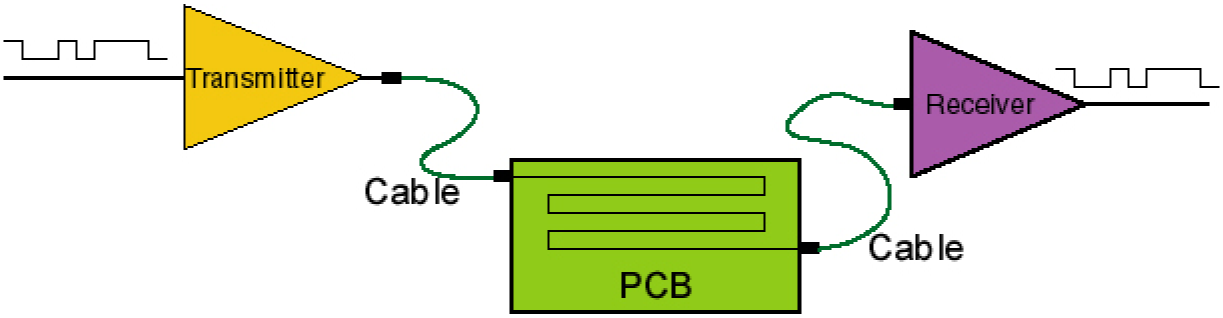

Figure 1: A straw diagram of a serial data system.

The culprits are reflections at connectors and vias, less obvious sources of return loss from sharp turns on a trace, any impedance mismatches – including the tiny but cumulative effect of varying trace width and thickness – and dispersion. That last one we usually reserve for optical systems but we can use it to shed light (get it?) on how even a perfectly straight trace on FR-4 filters a signal.

No Such Thing as Digital

The first thing we have to dispense with is the illusion that a signal is digital.

Consider a conducting trace on a piece of FR-4. The signal is a sum of electromagnetic waves that flow through the FR-4 dielectric medium guided by that conducting trace. The circuit is a complicated waveguide with no closed-form solution to Maxwell's Equations. The signal itself is a constantly changing sum of different frequency components. The highest frequencies are from the Fourier components of square waves. Remember, a true square wave is the sum of an infinite number of sinusoidal components, Figure 2.

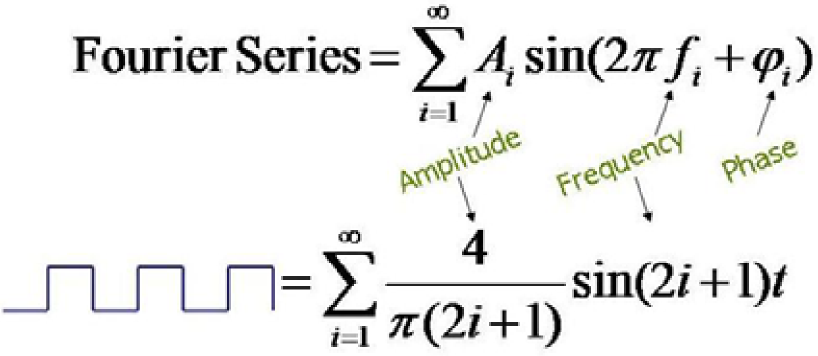

$$\text{Fourier Series} = \sum_{i=1}^{\infty} A_i \sin(2\pi f_i + \varphi_i)$$

$$= \sum_{i=1}^{\infty} \frac{4}{\pi(2i+1)} \sin(2i+1)t$$

Figure 2: Fourier components.

The frequency content presents us with competing effects. On the one hand, nice square waves look digital and convey the possibility of a big wide-open eye. On the other hand, the high frequencies necessary to produce those fast rise/fall times and tight square edges exacerbate the skin effect.

Aside: Skin effect

Recall from Ampere's Law that a DC current flowing in a straight line causes a magnetic field to flow in circles centered on that line. An AC current does the same thing except that, by virtue of Faraday's Law of Induction, as the current changes, eddy currents are induced in the conductor in the opposite direction of the change. The faster the current changes, that is, the higher the frequency, the greater the induced counter effect. In the limit of extremely high frequencies, the signal current plus the induced current is zero. At frequencies over 1 GHz, the currents cancel in the core of the conductor and the net current is restricted to a "skin depth" that gets thinner as frequency gets higher.

There are two important effects: high frequencies increase the effective resistance and reduce the inductance leaving the capacitance unchanged. Increased resistance absorbs power attenuating the signal. Changes in the impedance of different parts of the conductor change the phase relationship of the signal current and the induced currents; the net current is shifted in phase with respect to the transmitted signal. Remember, in building a square wave, the Fourier components are combined with both carefully chosen amplitudes and relative phases, Figure 2 – the skin effect messes up both.

Optimizing the desire for enough frequency content to keep the eye open but not so much that the skin effect takes over and ruins the signal requires a carefully tuned rise/fall time. The result is that, at high data rates, the highest frequency components of the signal are rarely more than a factor of three higher than the data rate.

The frequency content of a digital waveform changes constantly. Sections of the signal that have alternating ones and zeros, 01010, have a large contribution at a frequency of half the data rate (remember, it takes two bit periods to make a full 01 cycle). Sections of the signal with two consecutive identical bits, 0011, have a large dose of quarter-data-rate content, and so on.

Both the variety of frequency content and that it is constantly changing smears the signal as it propagates. The signal is guided by the trace but propagates in the dielectric and the dielectric "constant" of FR-4 is not constant, it varies with frequency,

$$\varepsilon = \varepsilon(f). \tag{1}$$

Equation (1) is called a dispersion equation and it means that different frequencies travel through the medium at different speeds. It's the same phenomenon that makes rainbows and blue skies. In this case, it makes the transmission path act like a poorly behaved filter and, along with the skin effect, is a major cause of Inter-Symbol Interference (ISI).

The other dominant causes of ISI are reflections and multi-path interference. Since the signal propagates through the dielectric, it's possible for fractions of the signal to take different routes through the circuit board from transmitter to receiver. The sum of the multiple reflections and signals traveling different paths at the receiver can look pretty nasty, like a black eye.

Okay. Now that we understand ISI, we can fix it. If we know (1) where major reflections occur and their return loss, (2) the different possible paths and the fractions of the waveform that leaks onto them, and (3) the filtering and attenuation properties of the transmission path, then we ought to be able to figure out a way to either transmit a signal so that it looks good at the receiver or a way to receive the signal so that we can discriminate ones and zeros even if the eye is closed.

In other words, if we know the geometry of the circuit, the medium of which it is composed, and the shape of the signal at the transmitter, then we can predict the eye-closure due to ISI. If we think of the received signal as the transmitted signal with this other information encoded on it – plus some random noise/jitter – then we ought to be able to decode the signal at the receiver. Decoding these predictable causes of eye-closure is called *equalization*.

Of course, before we can open the eye, we need to measure the attenuation, frequency response, and reflections that the signal must endure between the transmitter and receiver.

Time Domain Reflectometry/Transmission and *S*-parameters

Time Domain Reflectometry/Transmission (TDR/T) is the technique of sending a sharp voltage step into a circuit and then analyzing the reflected and transmitted responses. Gross circuit problems can be quickly and conclusively identified from a TDR/T trace, Figure 3, shorts, shunts, loads, inductive and capacitive mismatches each has a recognizable signature.

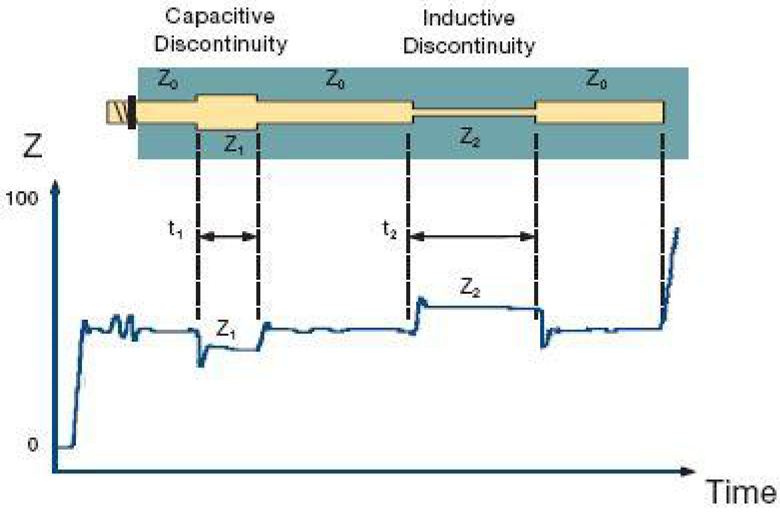

Figure 3: Example of TDR measurement.

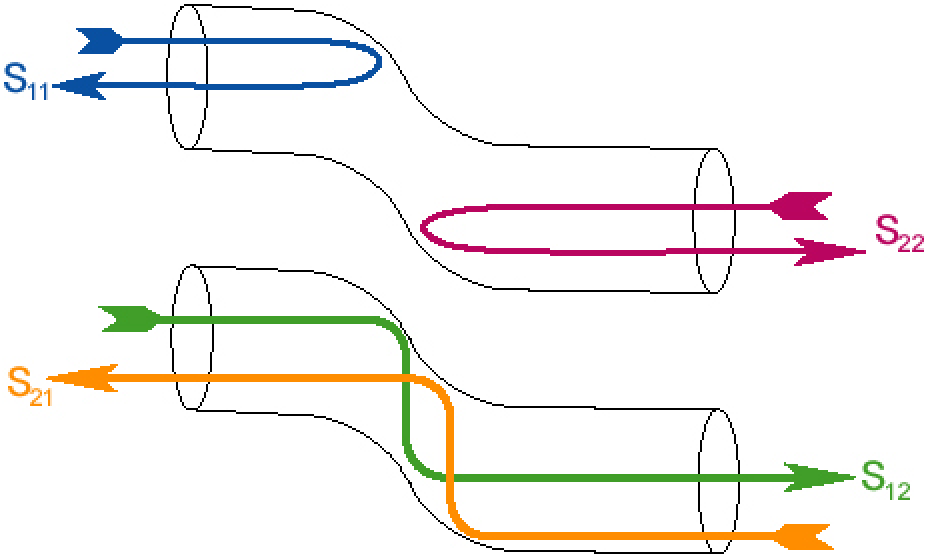

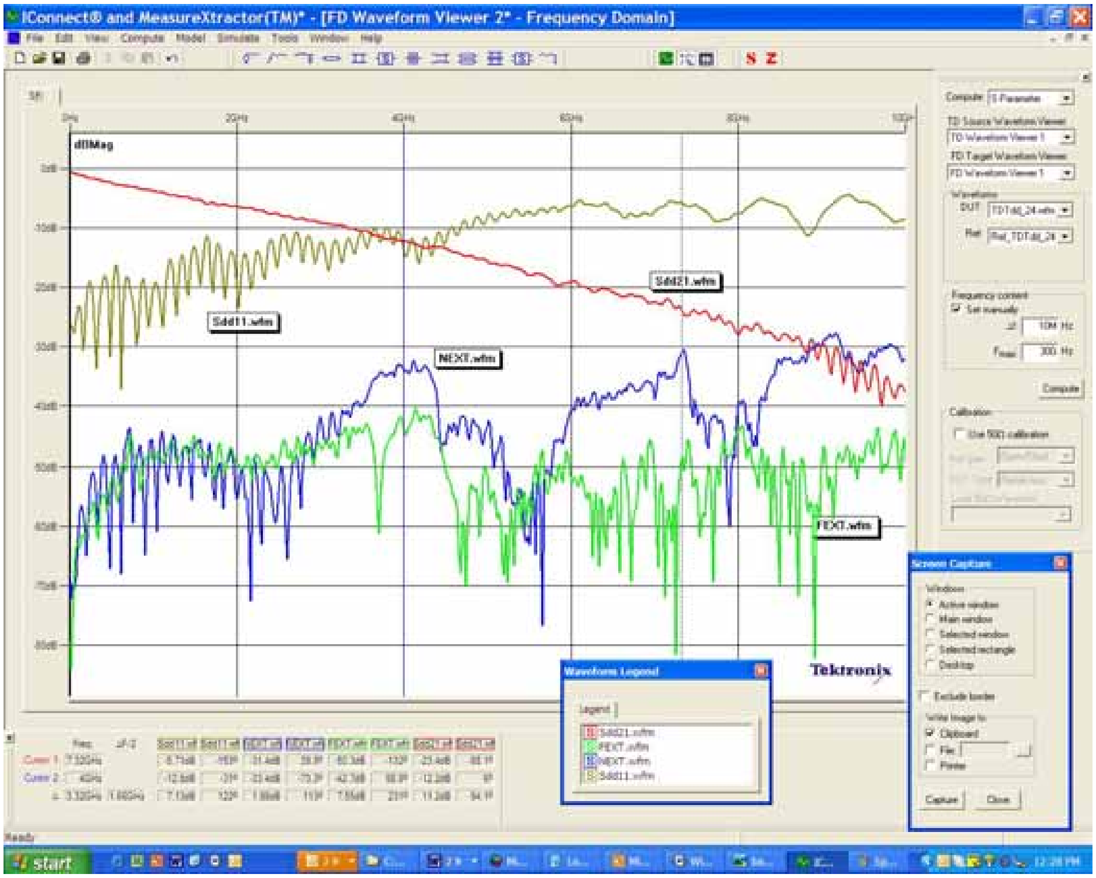

Figure 4: ***S*-parameters measure the frequency response of transmission and reflections from both ends of a device.**

TDR/T can find big problems, but the subtle genius is in the *S*-parameters.

The *S* in *S*-parameters comes from “scattering.” They are a complex (both literally and mathematically) representation of how a signal is scattered as it propagates through a circuit. By “scattered,” I mean the way that it interacts with the medium. At every step through a circuit, the signal moves into a different environment, whether the change is an infinitesimal variation in impedance from one point in a trace to the next or a big bouncing reflection at a connector with a cold-solder joint.

Since a circuit is basically one dimensional – we only really care about what goes in and what comes out of each end – there are four *S*-parameters: S11, S12, S21, S22,

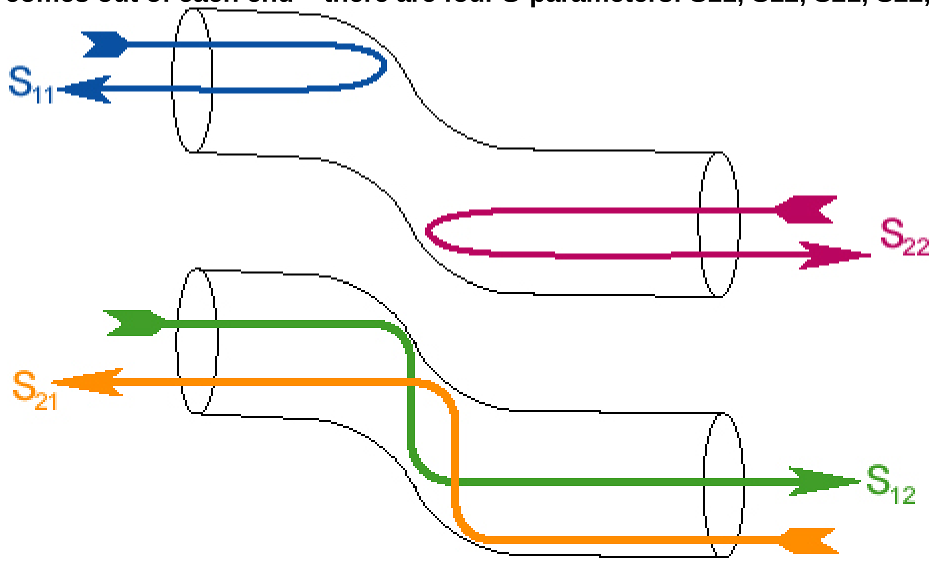

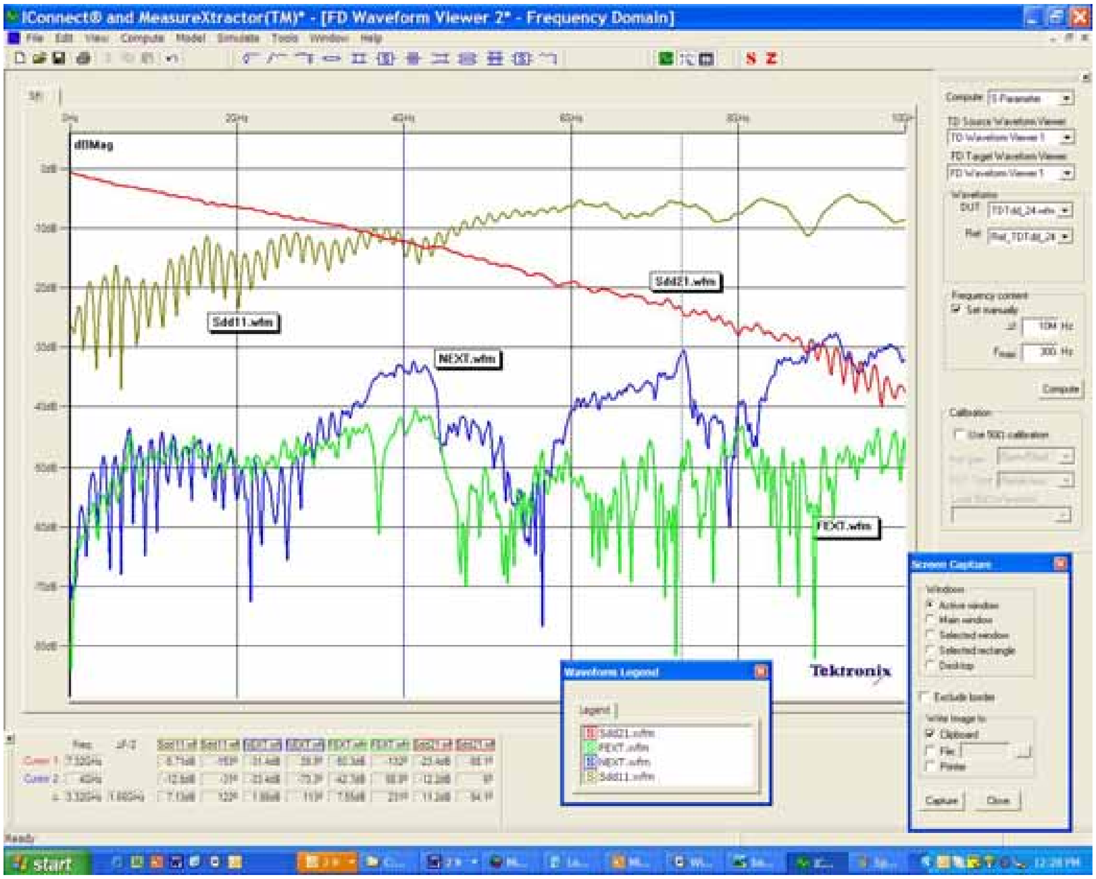

Figure 4. S11 is the reflectivity of a signal injected from the left and S12 is the transmissivity of that signal; S22 is the reflectivity of a signal injected from the right and S21 is the transmissivity of that signal. The situation is more complicated in differential systems – a complete set of *S*-parameters for differential and common-mode stimulus and response result in a total of 16 parameters – but the idea is the same.

S-parameters can be measured in either the frequency domain or time domain. In the frequency domain *S*-parameters are measured by analyzing the reflectivity and transmissivity in narrow frequency bands across a large frequency spectrum. In the time domain, the *S*-parameters can be measured with TDR/T by analyzing the Fourier components, Figure 2, of the reflected and transmitted remnants of the incident step. The faster the rise-time of the step, the more frequency components are included, and equivalently, the better that resolution with which close together reflectors can be distinguished.

Impulse Response

Regardless of how they are measured, *S*-parameters describe the frequency response and loss of the circuit. In fact, they embody the *impulse response*. The impulse response is what a circuit does to an extremely narrow pulse. Think of an infinitely narrow pulse with extremely high amplitude, Figure 5a, and remember that the inverse relationship between the time and frequency domains implies that the narrower the pulse, the greater its frequency content. Now think of the circuit as a conglomeration of tiny capacitors, inductors, dielectrics and resistors. As the impulse works its way through the circuit, the different frequency components arrive at slightly different times, with slightly different loss, and with slightly different amounts reflecting back and transmitting forward – unless there are big impedance mismatches, in which case we get big reflections that cause images of the pulse to bounce back and forth about the circuit. The dielectric, FR-4, is woven fiberglass; not only do different frequencies propagate at slightly different speeds, but the propagation speed also depends on the direction of the trace relative to the weave and whether a segment of trace is on a more or less dense portion of the weave. The result is a smeared pulse at the output, Figure 5b.

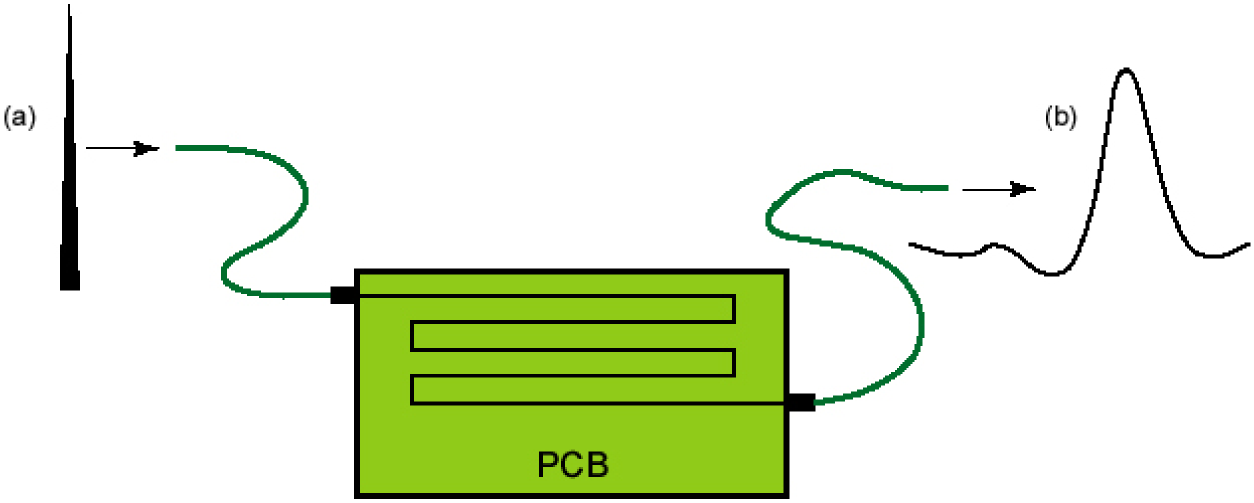

Figure 5: Impulse Response, (a) the ingoing impulse and (b) the outgoing impulse.

We will discuss the impulse response in the next installment of this series, *Acting on an Impulse: Equalization and Emphasis*, but suffice it here to state that the power of *S*-parameters is that we can calculate what any signal will look like at the receiver. In other words we can try all kinds of transmitted waveforms until we find one which appears at the transmitter with a wide-open eye! This is the heart of de-emphasis (a.k.a., emphasis or pre-emphasis).

Okay, there are a couple of caveats. First the true impulse response requires the whole infinite bandwidth; measured *S*-parameters, of course, are band-limited, so be careful to make sure that your *S*-parameters are measured with an accuracy equivalent to a bandwidth of at least three times the data rate and, even then, continue to be suspicious.

Transmitter De-emphasis

A de-emphasized signal is one, as shown in Figure 6, where a signal is given a voltage boost whenever there is a logic transition. Think of the transmission path as a filter. The high frequency components, Figure 2, are filtered more than the low frequency components. By giving the signal a voltage boost just prior to a logic transition, the high-frequency amplitudes are enhanced. If the boost is tuned just right, then the filtering effect reduces the now enhanced amplitudes to the not-enhanced levels and a signal that would otherwise fail to cross the receiver's logic-decision threshold, resulting in a bit error, is properly identified.

In other words, by transmitting a signal that might look funny at the transmitter, we can get a signal that looks good at the receiver. In principle we could design a transmitter that would emit a signal that encodes the inverse properties of the transmission path. Just out of the transmitter, the eye would be closed, but at the receiver it would be open. It turns out that it's cheaper and easier to design a receiver that can undo the adverse effects of the transmission path. Fixing the signal at the receiver is called equalization, though rarely called equalization, de-emphasis at the transmitter is too. Stay tuned for the next seminar in this series to learn about equalization.

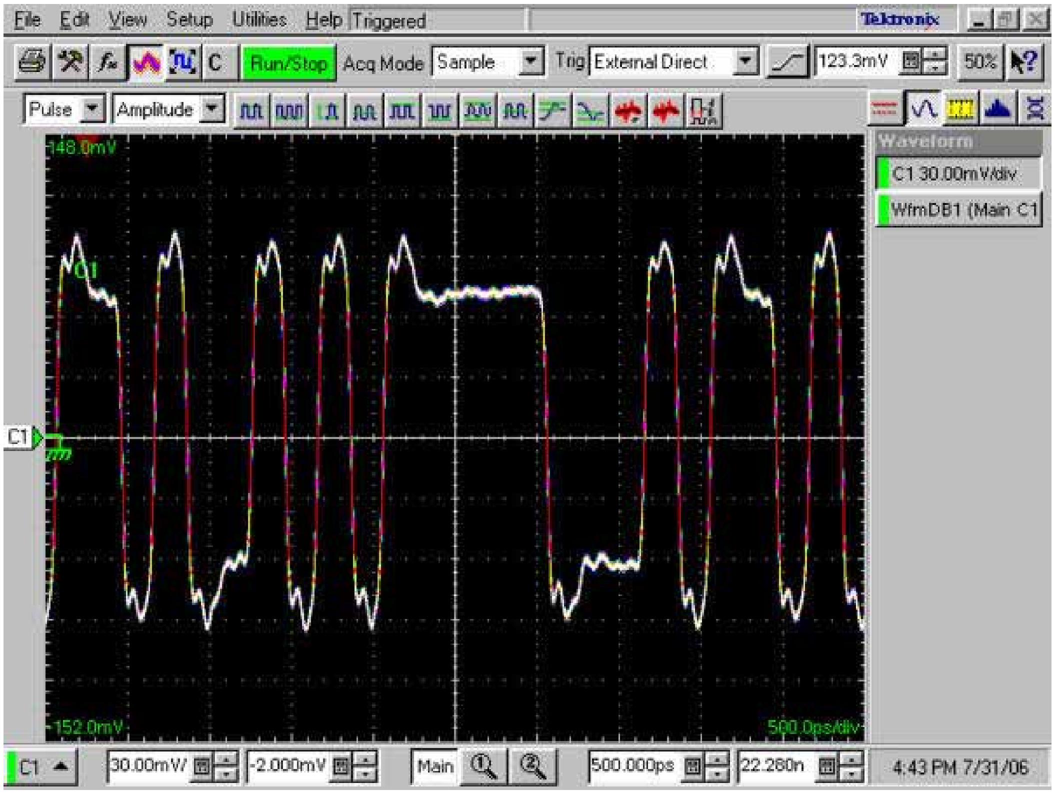

Acq Mode Sample
Trig External Direct
123.3mV
50%
C1 30.00mV/div
WfmDB1 (Main C1
C1 30.00mV/ -2.000mV
500.000ps 22.280n

Figure 6: A signal with de-emphasis. Notice that there is an extra voltage boost at every logic transition.

Conclusion

At high data rates we can no longer ignore the analog reality of digital signals. The result of intentional analog ignorance is a closed eye and a high bit error ratio.

Facing analog reality reveals the perpetrators that close eyes diagrams. Time Domain Reflectormetry/Transmission (TDR/T) reveals discrete problems, like reflection at a mismatched connector, that can be fixed by a simple redesign. Analysis of the scattering properties of the transmission path, as represented by *S*-parameters, reveals subtler problems that, in some cases, can be fixed by a re-design like more careful trace layout, but in others may be too expensive to fix, such as when the FR-4 circuit board simply can't maintain an open eye.

In summary, there are four approaches to solving the case of the closing eye:

1. Circuit redesign – e.g., fix mismatched connectors, reroute traces, remove vias, et cetera.
2. Transmit a signal of a shape that evolves as it propagates through the circuit resulting in an open eye at the receiver – e.g., transmitter de-emphasis.
3. Implement an equalization scheme at the receiver – the subject of the next installment of this series.
4. Surrender – replace FR-4 circuit board with a less dispersive medium, use fiber-optics instead of electrical signals, give up altogether and vow to never work at data rates above 300 baud.