

TDR:

Reading the Tea Leaves



Time domain reflectometry (TDR) is an essential microwave measurement technique and also an essential technique for analyzing measured S parameters. However, extracting as much information as possible from a TDR trace can resemble the work of a fortune teller. There is a lot of information that escapes the casual observer but is readily apparent if one knows what to look for. This article reviews the basic techniques for interpreting TDRs and then adds a couple of simple techniques that are not as well known.

What Happens Where

The goal of time domain reflectometry (TDR) is to correlate the electrical behavior of a transmission line circuit with its physical features. Once one understands which physical features have the most significant effect on electrical behavior, it's a lot easier to improve the design.

TDR is almost always associated with measured data—either the TDR data is measured directly or else it's computed from measured S parameters. Either way, the

goal is to determine the magnitude and location of the transmission line discontinuities that are in the actual circuit, independent of what any model might predict. TDR data can also be generated for models, and that's useful for comparison to measured data.

Figure 1 and Figure 2 are examples of TDR data computed from measured S parameters. Each figure shows a time domain waveform which approximates the reflection coefficient as a function of time (and therefore distance). The two figures show TDR data looking from opposite ends of the same interconnect network. As will be explained in the following sections, the individual elements such as vias, traces, and connectors are clearly visible in the TDR trace. Some of these features are labeled in the figures.

How TDR Data is Measured

Time domain reflectometry measures the voltage step response at the input to the device under test. The source impedance is typically a standard reference impedance such as 50 ohms single ended or 100 ohms differential.

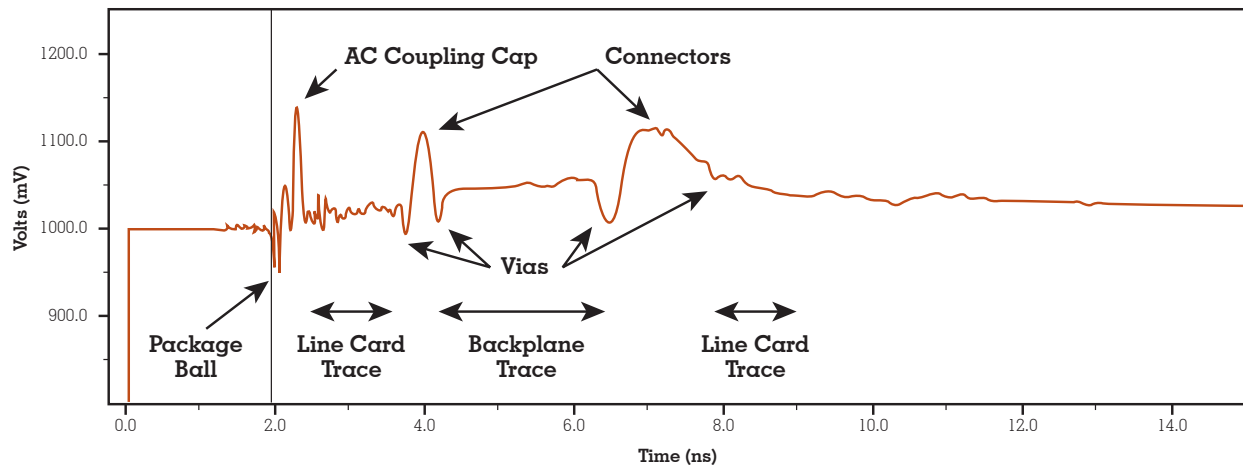


Figure 1: Example TDR derived from measured S parameter data

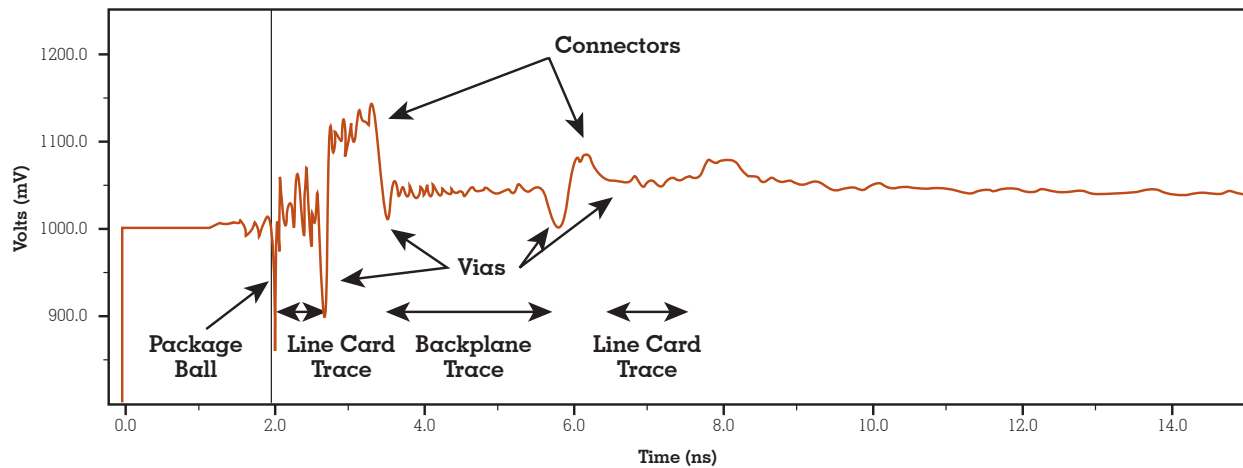


Figure 2: Example TDR looking at the same path from the opposite direction

For actual physical measurements, there is usually a length of transmission line needed to connect the device under test to the measurement circuit. As will be seen later, this length of transmission line also plays a role in calculating a TDR from measured S parameters.

When performing TDR on an interconnect network, the common practice is to terminate the network with the same impedance as the source. The single ended version of this setup is shown in Figure 3.

The idea is that when the rising edge of the voltage step stimulus encounters any change in the impedance of the transmission path, it generates a reflected edge that propagates back toward the stimulus source. When this reflected edge arrives at the input to the interconnect network, it gets added to the voltage at this node. There is a delay between the time the stimulus edge first passed the input to the interconnect network and the time that

the reflected edge arrived at the same circuit node; and this delay indicates the position of the discontinuity that generated the reflected edge.

To interpret TDR results with any degree of clarity, one must understand the three approximations on which the concept of TDR is based:

1. The impedance of the transmission path is approximately constant.
2. The transmission path is approximately lossless.
3. The propagation velocity is approximately constant.

In most applications, none of these approximations is very good. They are usually good at the beginning of the transmission path, but then gradually break down; and the rate at which these approximations break down should drive the way one interprets the results.




Figure 3: Time domain reflectometry measurement of interconnect network

Nonetheless, these approximations are good enough to obtain some useful information.

If one assumes that the three simplifying approximations are valid, then for a transmission path with varying impedance $Z(x)$, the reflection coefficient for a short length δ of that path is:

$$\rho(x) = \frac{Z(x + \delta) - Z(x)}{Z(x + \delta) + Z(x)}$$

Taking the limit as $\delta \rightarrow 0$ and applying the approximation of constant impedance,

$$\rho(x) \approx \frac{1}{2Z_0} \frac{dZ(x)}{dx}$$

Applying the second and third approximations (low loss and constant velocity) to the calculation of the impulse response at the driving end, and assuming that the propagation velocity is γ ,

$$h(t) \approx \rho \left(\frac{2x}{\gamma} \right) \approx \frac{1}{2Z_0} \frac{dZ}{dx} \left(\frac{2x}{\gamma} \right)$$

One can then get the step response (that which is actually measured) by integrating with respect to time. The result is:

$$s(t) \approx \frac{Z \left(\frac{2x}{\gamma} \right)}{2Z_0}$$

In other words, the amplitude of the step response at a particular time is approximately proportional to the impedance at some corresponding distance along the transmission path. This is the desired result.

As each of the approximations breaks down, however, the step response exhibits additional artifacts.

1. Impedance Variations

a. As the impedance variations become more and more extreme, the impedance scale in the TDR waveform becomes distorted. For example, an open circuit (essentially infinite impedance) does not result in an

infinite step response amplitude; it results in a doubling of the step response amplitude.

Therefore the amplitude scale of a TDR waveform is really incident wave plus a reflection coefficient $\left(1 + \frac{Z - Z_0}{Z + Z_0}\right)$ between 0 and 2 rather than a linear impedance scale.

b. If there are multiple discontinuities in the transmission path, then the step response edge reflected by a discontinuity which is further from the measurement point will be reflected by the nearer discontinuities and then reflected again by the further discontinuity, resulting in multiple reflections arriving at the measurement point.

2. Transmission Loss

As the step response edge propagates along the transmission path, it encounters loss, especially at higher frequencies. This loss smooths out the edge that is incident on discontinuities which are further from the measurement point. The step response edge that is reflected also gets smoothed out on its way back to the measurement point. Therefore features which are further from the measurement point become less distinct.

For example, the features on the right hand side of Figure 1 and Figure 2 are much less distinct than those on the left hand side, even though the features on the left hand side of one figure are exactly those of the right hand side in the other figure.

3. Propagation Velocity Variations

If the propagation velocity changes, for example because the transmission medium changes from PC board to cable, then the translation from time to distance will be distorted accordingly.

Interpreting the TDR Trace

Figure 4 shows a view of the example TDR in Figure 1 that has been plotted on an impedance scale and expanded to show some detail at the beginning of the TDR trace.

Measurement/Calculation Setup

The first two nanoseconds of the TDR trace show some details about how the TDR was calculated, or how the equivalent trace would be measured in the lab.

Since the original S parameter data was supplied to




Figure 4: Example TDR expanded and shown using impedance scale

20GHz, there is an initial delay of 17pS corresponding to the rise time of the measurement stimulus. This 17pS is also the rise time/resolution quoted in the data sheet for at least one 20GHz bandwidth sampling oscilloscope that is used to measure TDR.

Also, the measurement setup shown in Figure 3 includes a transmission line inserted in the measurement circuit between the measurement point and the device under test. For the calculations shown here, the delay of that transmission is 1nS, thus injecting a 2nS round trip delay in the measurement.

This calculation technique was suggested to me by Donald Telian. Note that although the reflections due to an ideal transmission line should be zero, there appear to be nonzero reflections in the first 2nS of the TDR trace. These apparent reflections are actually numerical artifacts due to the fact that the S parameter data was limited to 20GHz whereas the real circuit has a nonzero response beyond 20GHz. If these numerical artifacts are not included in the integration of the TDR as a step response, the TDR is offset vertically by the integral of these numerical artifacts.

Identifying Features and Trace Lengths

There is an abrupt increase in impedance approximately 280pS after the beginning of the circuit under test. Since 280pS is the round trip delay, this discontinuity occurs

after the stimulus has traveled 140pS on the line card. Assuming that the line card dielectric is a relatively high quality dielectric with a dielectric constant of 3.2, the speed of propagation is 150 pS/inch. Therefore, this discontinuity occurs at approximately 0.92" from the beginning of the measured trace. At this particular position on this particular trace, there happens to be an AC coupling cap. Therefore, the discontinuity is most probably due to the AC coupling cap.

Similarly, after the AC coupling cap there is 1.37nS of delay to the next significant discontinuity, an abrupt reduction in the apparent impedance. The position of this discontinuity is approximately $1.37\text{nS}/2/150\text{pS} = 4.57\text{"}$ trace length from the AC coupling cap. This corresponds to the connector via at the edge of the board. Therefore, the structure on the right hand side of Figure 4 is most probably the connector with its vias. To produce the labels in Figure 1, this process was continued along the entire TDR trace.

Lumped Elements

Short, clearly defined discontinuities such as the AC coupling cap in Figure 4 can be approximated as lumped circuit elements: either series inductors when the discontinuity impedance is higher or shunt capacitors when the discontinuity impedance is lower.

Consider, for example, that the AC coupling capacitor




Figure 5: Example TDR with series resistance shown

increases the apparent transmission line impedance by about 14Ω for a round trip delay of 77pS in the TDR trace, or 38.5pS on the PC board. This is enough information to estimate the series inductance of the AC coupling cap.

From [1], the impedance of a transmission line as a function of its propagation velocity v , its inductance per unit length l and its capacitance per unit length c is

$$Z_0 = \sqrt{\frac{l}{c}} = \frac{1}{vc} = vl$$

Consider also that the time delay due to a propagation across a distance d is

$$t = \frac{d}{v}$$

The total excess inductance is therefore

$$L = dl = tZ_0$$

For the AC coupling cap in Figure 4, this comes out to a series inductance of 0.54nH , which is about right for a surface mount capacitor.

For capacitive discontinuities, the calculation must be performed using the reciprocal of the characteristic impedance. That is,

$$C = \frac{t}{Z_0}$$

For example, the via in Figure 4 increases the conductance of the transmission line from $1/52.5$ Semens to $1/49.6$ Semens, for an increase of 0.00114 Semens. One half of the round trip delay of 72pS in the TDR trace is 36pS , so the excess capacitance of the via is approximately 41fF .

Steady State Values

It is often observed that the steady state value of a TDR trace is higher than the initial stimulus amplitude. This is apparent in Figure 1 and Figure 2, and is illustrated again using an impedance scale in Figure 5.

This feature of the TDR trace is not a numerical artifact; it's a fundamental behavior of the interconnect network and a useful piece of information. Consider that when the TDR trace (as a step response) settles to a long term value, the equivalent circuit is




Figure 6: Steady state measurement of interconnect network

That is, at steady state, the only relevant circuit elements are the resistances in the circuit. For most interconnect networks (those with negligible DC leakage), the

resistances are the source resistance, the load resistance, and the series resistance of the interconnect network. Therefore from Figure 5, the series resistance of that particular interconnect network is approximately 3Ω .

For another view of the measurement of interconnect network behavior at low frequencies, see [2].

References

[1] Matthei, Young and Jones, Microwave Filters, Impedance Matching Networks, and Coupling Structures, section 5.02, equation 5.02-5, page 164, McGraw-Hill, copyright 1964.

[2] <http://realtimewith.com/pages/rtwvprofile.cgi?rtwvcid=13&rtwvid=1691>

About the Author

Michael Steinberger, Ph.D., is responsible for leading SiSoft's ongoing tool development effort for the design and analysis of serial links in the 5-30 Gbps range. Dr. Steinberger has over 30 years experience in the design and analysis of very high speed electronic circuits. Dr. Steinberger began his career at Hughes Aircraft designing microwave circuits. He then moved to Bell Labs, where he designed microwave systems that helped AT&T move from analog to digital long-distance transmission. He was instrumental in the development of high speed digital backplanes used throughout Lucent's transmission product line. Prior to joining SiSoft, Dr. Steinberger led a group of over 20 design engineers at Cray Inc. responsible for SerDes design, high speed channel analysis, PCB design and custom RAM design. ■

Simulating Large Systems with Thousands of Serial Links



Run 10's of 1,000's of
Simulation Cases to Spot
System-Level Trends.

Learn how a virtual system model supports **48 hour turnaround**
from PCB layout changes to updated results
versus weeks for a physical prototype.

Visit www.sisoft.com to learn more.

 **SiSoft™**
We Are Signal Integrity