


Signal Integrity Analysis Series

Part 1: Single-Port TDR, TDR/TDT, and 2-Port TDR

Introduction



Interconnect Analysis is Simplified with Physical Layer Test Tools

The Time domain Reflectometer (TDR) has come a long way since the early days when it was used to locate faults in cables. Time Domain Reflectometry can be used for more than 40 characterization, modeling, and emulation applications, many of which are illustrated in this application note series.

If your applications involve signals with rise times shorter than one nanosecond, transmission line properties of the interconnects are important. TDR is a versatile tool to provide a window into the performance of your interconnects to quickly and routinely answer the three important questions: does my interconnect meet specifications, will it work in my application, and where do I look to improve its performance?

The TDR is not just a simple radar station for transmission lines, sending pulses down the line and looking at the reflections from impedance discontinuities. It is also an instrument that can directly provide first order topology models, S parameter behavioral models, and with up to four channels, characterize rise time degradation, interconnect bandwidth, near and far end cross talk, odd mode, even mode, differential and common impedance, mode conversion, and the complete differential channel characterization.

To provide a little order to the wide variety of applications we explore in this application note series, we divide the series into three parts covering four general areas. Part 1: Those which use a single-port TDR, those which use TDR/TDT, and those which use 2-port TDR. Part 2: Those which use 4-port TDR or 4-port Vector Network Analyzer (VNA) with Physical Layer Test System (PLTS). Part 3: Those which use advanced signal integrity measurements and calibration. The principles of TDR and VNA operation are detailed in other application notes and references listed in the bibliography. We concentrate this application note series on the valuable information we can quickly obtain with simple techniques that can be used to help us get the design right the first time.

Table of Contents

1.0 Single-Port TDR	5
1.1 Overview	5
1.2 Measuring characteristic impedance and uniformity of a transmission line.....	5
1.3 Measuring time delay of a transmission line.....	13
1.4 Accurate measurement of signal speed in a transmission line.....	16
1.5 Extracting bulk dielectric constant of the laminate	17
1.6 Building a model of a discontinuity such as a corner, test pad, gap in the return path, SMA launch, terminating resistor	19
1.6.1 Extracting a model for capacitive discontinuities	19
1.6.2 Extracting a model for inductive discontinuities.....	24
1.6.3 Modeling termination resistors	26
1.7 Building a high bandwidth model of a component.....	30
1.8 Directly emulating the impact on a signal with the system rise time from a discontinuity.....	32
2.0 2-Port TDR/TDT	34
2.1 Overview	34
2.2 Introduction to TDR/TDT	34
2.3 Measuring insertion loss and return loss.....	35
2.4 Interconnect modeling to extract interconnect properties.....	39
2.5 Identifying design features that contribute to excessive loss.....	41
3.0 2-Port TDR/Cross Talk	45
3.1 Overview	45
3.2 Measuring NEXT	46
3.3 Measuring FEXT.....	47
3.4 Emulating FEXT for different system rise times.....	48
3.5 Identifying design features that contribute to NEXT	49
3.6 Exploring the impact of terminations on NEXT and FEXT	52
3.7 Measuring ground bounce.....	53
3.8 Identifying design features that contribute to ground bounce	55
3.9 Emulating ground bounce noise for different system rise times.....	56

4.0 2-Port Differential TDR (DTDR)	60
4.1 Overview	60
4.2 Measuring each of the five impedances associated with a differential pair	62
4.3 Measuring the degree of coupling between lines in a differential pair	66
4.4 Measuring the differential impedance of a twisted pair cable	67
4.5 Measuring the reflected noise of a differential signal crossing a gap.....	69
4.6 Measuring the mode conversion in a differential pair	71
4.7 Identifying specific physical features that contribute to mode conversion in a differential pair	71
References	74
Web Resources	74

1.0 Single-Port TDR

1.1 Overview

In this section we will look at the seven most important applications of a 1-port TDR. The first two refer to the complete characterization of a uniform transmission line, extracting the characteristic impedance and time delay.

But, we can get more than this if we use specially designed test structures. We can also get a fundamental, intrinsic property of the transmission line, the velocity of a signal, and from this, the intrinsic bulk dielectric constant of the laminate.

When the line is not uniform, but has discontinuities, we can build first order, topology based models right from the front screen. If this isn't high enough bandwidth, we can bring the measured data into a simulation tool such as ADS, and build very high bandwidth models which can then be used in simulations to evaluate whether this interconnect might be acceptable in a specific application.

Finally, we can emulate the final application system's rise time with the TDR to directly measure the reflection noise generated by physical structures in the interconnect and whether they might provide a potential problem, or equally of value, might be ignored.

1.2 Measuring characteristic impedance and uniformity of a transmission line

Historically, the most common use of the TDR has been to characterize the electrical properties of a transmission line. For an ideal, lossless transmission line, there are only two parameters that fully characterize the interconnect: its characteristic impedance and its time delay. This is the easiest and most common application for the TDR.

The TDR sends a calibrated step edge of roughly 200 mV into the device under test (DUT). Any changes in the instantaneous impedance the edge encounters along its path will cause some of this signal to reflect back, depending on the change in impedance it sees. The constant incident voltage of 200 mV, plus any reflected voltage, is what is displayed on the screen of the TDR.



Figure 1. Measured TDR response from a microstrip transmission line. Blue trace is the reflection from the end of the cable; yellow trace is the reflected signal from the DUT.

In Figure 1, the yellow line is the measured TDR response when the DUT is the microstrip trace shown. The first two inches of the transmission line have a characteristic impedance of roughly 50 Ohms, while the next four inches of the transmission line has a characteristic impedance of roughly 40 Ohms. The far end of the line is open.

The voltage displayed on the screen is the total voltage: the incident, constant 200 mV, plus the reflected voltage. Note on the bottom of the screen, the vertical voltage scale is 100 mV/div. The blue data is the TDR response for the cable not connected to the transmission line. This defines the beginning of the cable, which is an open. On the yellow line, at this instant of time, is the small reflected voltage from the SMA launch, followed by the roughly 50 Ohm section of the line, and about one division later, the small drop in voltage from the lower impedance, second half of the transmission line.

Contained in this reflected signal is the information about the impedance profile of the transmission line. We could read the voltages off the front screen and use pencil and paper to back out the impedance of the line, or we can take advantage of some of the built-in features of this TDR.

We can use the two markers which will automatically perform the calculations to back out the instantaneous impedance from the measured data. There are clearly two regions of relatively uniform impedance on this transmission line. We move the markers so that one is in each region, as shown in Figure 2, and then we can read the impedance of each region from the screen.

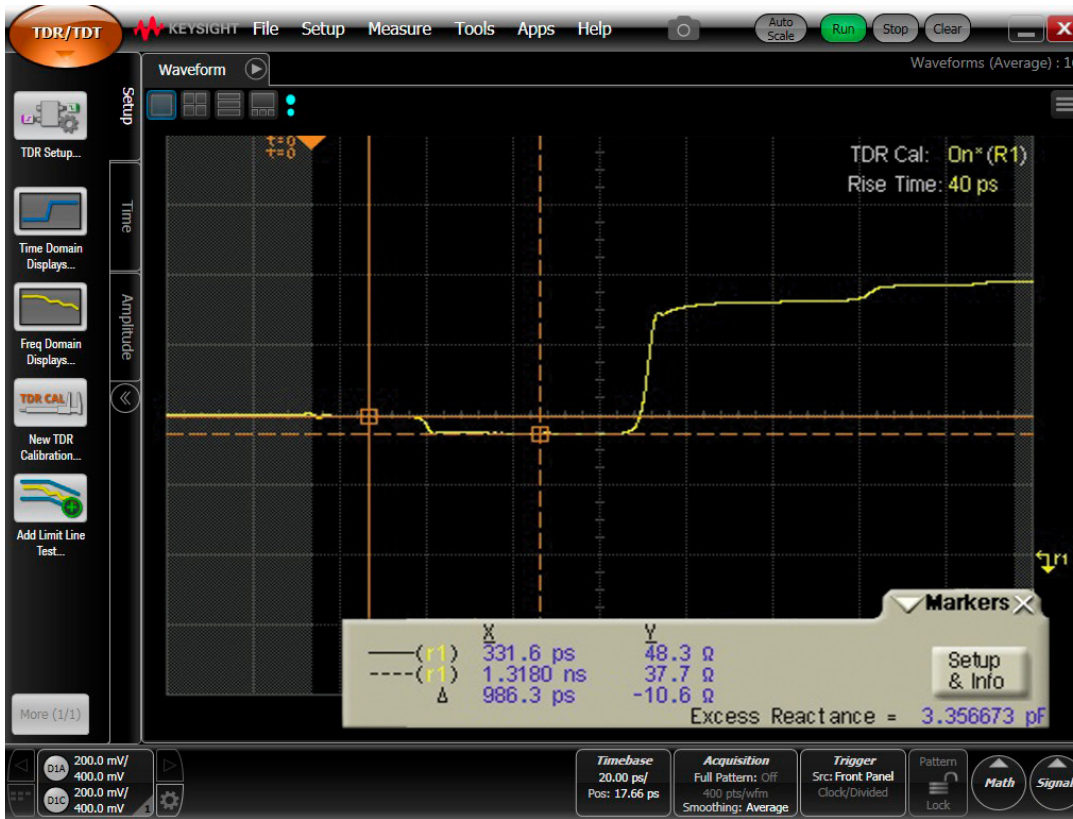


Figure 2. Using markers to measure the characteristic impedance of a transmission line.

The impedance of the first region is read from the solid line marker as 48.3 Ohms. The impedance in the second region, as read by the dotted line marker is 37.7 Ohms. The nominal design impedances were 50 Ohms and 40 Ohms, so we see that the actual, as fabricated impedances, are off by about 3.5 and 6 percent.

The one caveat when using the markers is to watch out for masking effects. The impedance read by the marker can be interpreted as the instantaneous impedance of the transmission line at the location of the marker, as long as it is the first interface, or there have been only small impedance discontinuities up to the location of the marker. This feature makes extracting the instantaneous impedance of a uniform transmission line almost trivial. In addition, we can see that the impedance in each region is relatively uniform, as there is little deviation in the reflected voltage up and down the line segments.

In addition to using the marker to identify the specific instantaneous impedance of the transmission line, we can also convert the vertical voltage scale into an impedance scale.



Figure 3. Using the advanced settings function to adjust the vertical scale to display the impedance directly.

By selecting time domain displays, then Ω , then new chart, we can choose to display T11 in Ω to see the impedance profile vertical scale 1, 100 mV/div: 200.0 mV and then clicking Advanced, we can access the scale settings command functions, as shown in Figure 3. When we select the Ohms scale, the TDR will convert every point of the reflected voltage into an equivalent instantaneous impedance.

Effectively, the TDR takes each measured voltage point, subtracts 200 mV to get the reflected voltage, then takes the ratio of this voltage to the 200 mV of incident voltage to get the reflection coefficient, and from the reflection coefficient, uses the simple relationship: $Z = 50 \times (1 + \rho) / (1 - \rho)$ to calculate the instantaneous impedance of each point. Finally, this extracted instantaneous impedance is plotted on the screen. The offset and scale settings, now calibrated in Ohms, can be used to adjust the scale for our application.

Figure 4 is the same TDR data for this two segment transmission line, but now displaying the instantaneous impedance directly on the vertical scale. In this case, the scale is 10 Ohms/div with the center location set to 50 Ohms. On this scale we can literally read off the screen the impedance of the first section as about 48 Ohms and the impedance of the second section as about 38 Ohms.




Figure 4. The same transmission line displayed on the impedance scale at 10 Ohms/div, with 50 Ohms in the center.

This scale setting allows a direct and effortless graphical display of the impedance profile of a transmission line, with the one caveat that we are assuming all the measured voltage coming back to the TDR is due to reflections from impedance discontinuities. This is a good assumption as long as the impedance changes up to each point are small.

It looks like, for this transmission line, the impedance of the first section is decreasing slightly down the line, while the impedance profile of the second section is mostly constant. We can use this technique to evaluate how uniform the impedance of a transmission line is.




Figure 5. High resolution TDR profile of a nominally uniform transmission line, at 2 Ohms/div and 50 Ohms in the center of the screen.

In Figure 5, we have the measured TDR response of a nominally uniform transmission line, on an expanded scale of 2 Ohms/div. The impedance at the center of the screen is set at 50 Ohms. This scale information can be read next to the channel 1 button of the screen.

The large peak at the beginning of the line is the inductive discontinuity of the SMA launch, which on this high resolution scale, looks huge. At 2 Ohms/div on the vertical scale it looks like this uniform transmission line is not so uniform. It appears to have a variation of as much as 1 Ohm from the beginning to the end of the line. This is roughly 2 percent.

Is this variation real, or could it be some sort of artifact? There are two important artifacts that might give rise to this sort of behavior. It could be there is rise time degradation in the incident signal. It may not be perfectly flat, like an ideal Gaussian step edge. After all, the reflected signal displayed on the TDR is really the reflection of the incident signal. If the incident signal has a long tail, we will see this long tail in the TDR response and may mistakenly interpret this as an impedance profile variation. One way around this problem is to use the calibrated response feature of the DCA 86100C TDR, which we are doing in this case.

Another source of artifact is if there is either distributed series resistance in the trace or distributed shunt conductance in the trace due to the lossy nature of the line. The series resistance will cause the reflected voltage to increase as we move down the line, while the shunt conductance will cause the reflected TDR response to decrease as we move down the line, as in this case.

One way to evaluate whether an impedance profile is really showing a variation in the instantaneous impedance of the transmission line or an artifact, is to measure the TDR response of the line from both ends. If it is real, we should see the slope of the response change, depending on which end of the line we launch from. If it is one of the two artifacts, the response will look the same on the screen, independent of which end we launch from.




Figure 6. High resolution TDR response from each end of the same uniform transmission line, verifying the impedance variation is real.

In Figure 6, we show the measured TDR response launching from each end of the line. The yellow trace is the TDR response from the left edge of the trace while the blue line is the TDR response when launching from the right edge of the trace. In both cases, the scale is 2 Ohms/div.

The yellow line shows the left side of the line is the higher impedance. The blue line, the TDR response looking from the right side, also confirms that the impedance of the trace is higher on the left side. This variation in the instantaneous impedance is confirmed to be real, and is not due to the series resistance or shunt conductance or the nonideal step edge. Using the technique of comparing the launches from both ends, we can unambiguously identify real, nonuniform impedance effects in a transmission line.

In this example, the microstrip is showing a variation of about half a division, or 1 Ohm out of 50 Ohms, or two percent, from one end to the other. This could be due to a variation in the laminate thickness, the slight drift in the alignment of the trace width over a fiber bundle in the glass weave, or a variation in the etching of the line due to photo resist developer variation across the board.

By measuring the variation in other lines across the board, or inspecting the dimensions of the board, the root cause might be identified and the process stabilized.

1.3 Measuring time delay of a transmission line

The second important property of a transmission line is the time delay from one end to the other. This can also be measured directly from the screen of the TDR using markers. However, to get an accurate measure of the time delay, we need to know the starting point of the transmission line.

By removing the DUT and recording the TDR response from the open end of the cable, we can use this as a reference to define the beginning of the line. This is the blue line in Figure 7. When we reconnect the DUT and record the TDR response, we see the reflection from the open at the far end of the transmission line, just visible at the far right edge of the screen.




Figure 7. TDR response of a uniform 6 inch transmission line open at the far end.

The total round trip time delay is the time interval from the beginning of the reflection from the open end of the cable to the reflection from the open far end of the DUT. To increase the accuracy, we use the time from the mid-point of the two open responses. This can be measured simply and easily using the vertical markers, directly from the screen.

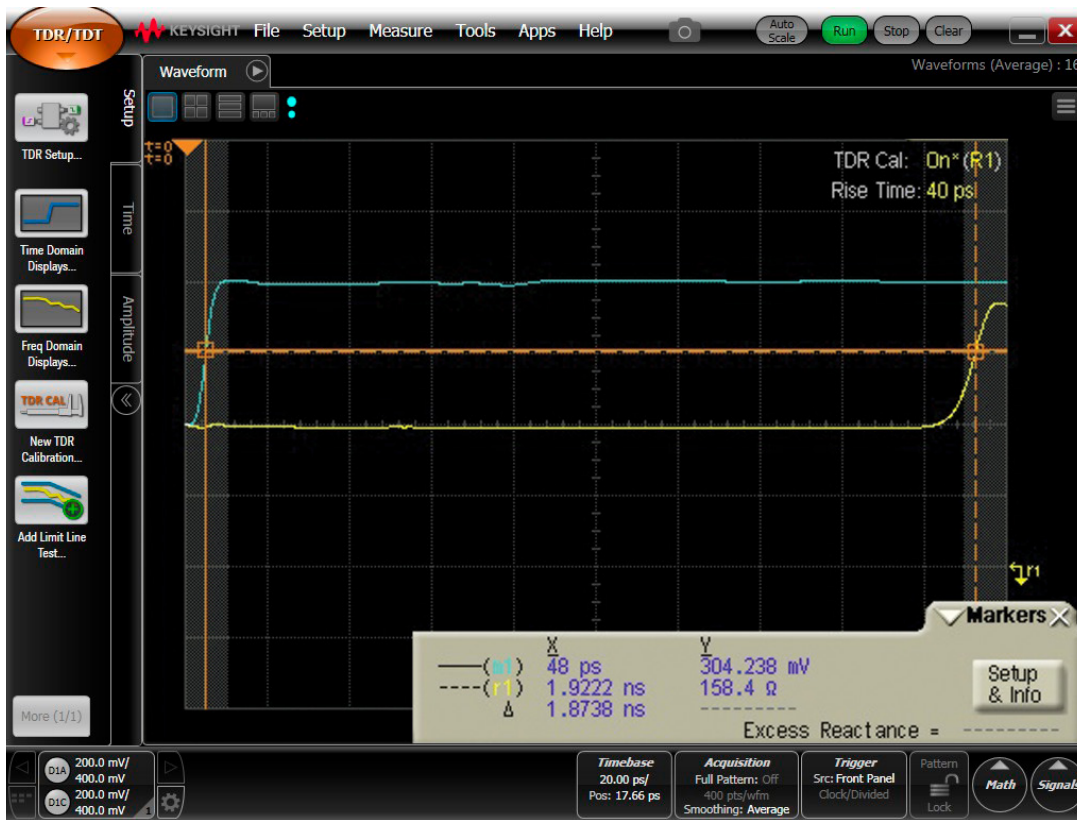


Figure 8. TDR response of the reference open and uniform six inch transmission line, with markers showing the beginning and end of the traces.

Using the marker buttons below the screen, we can position the markers in Figure 8, so that they define the midpoint to midpoint distances. We can read off the screen that the total time delay is 1.87 ns. This is the round trip time delay. The one way time delay is half of this, or 0.935 ns. This is the time delay (TD) of the transmission line.

From the physical length of the transmission line, six inches, and the time delay, 0.935 ns, we can also calculate the speed of the signal down this transmission line. The speed is $6 \text{ in}/0.935 \text{ ns} = 6.42 \text{ in/ns}$. This is an intrinsic property of the transmission line, and would be true for any transmission line of the same width built on this layer of the board, independent of the length of the line.

One of the artifacts in this measurement is the uncertainty of how much of the total time delay is due to the connector at the front of the line. Is the open reference really the beginning of the line, or is there some contribution to the launch into the transmission line of the circuit board? We can take advantage of a simple test feature to get around this artifact and extract a more accurate value for the speed of a signal on the trace

This trick is useful only if we have the option of designing the test line to aid in the characterization of the circuit board and each particular layer. The secret is to add small imperfections to the line, such as reference pads at two locations with a known separation.



Figure 9. TDR response of a uniform transmission line with two small reference pads, located on four inch centers.

Figure 9 shows an example of a six inch long transmission line with two reference pads (in close up), located with a center spacing of four inches. These pads can be easily detected with the TDR. The TDR response, the yellow line, is displayed on a scale with 2 Ohms/div. This is a high sensitivity scale. On the far left, it shows the beginning of the line with a few ripples from the SMA launch. About two divisions from the beginning is the dip from the first pad, which acts as a small capacitive load, a lower impedance. Some time later, the TDR signal shows the response from the second reference pad.

1.4 Accurate measurement of signal speed in a transmission line

The time difference between these two negative dips is the round trip time difference between the pads separated by four inches. By measuring this time delay from the screen, we can get an accurate measure of the speed of the signal, independent of the nature of the launch into the transmission line.

We can measure the time delay between the dips using the on screen markers. By aligning each marker with the center of the dip, we can measure this location within a few picoseconds accuracy. We can see from the screen in Figure 10 that the round trip time delay is 1.238 ns. From this round trip delay, we can calculate the one way time delay as half this or 0.619 ns.




Figure 10. TDR response from a microstrip with two reference pads, using markers to measure the round trip time delay.

Given the physical distance between the two reference pads as four inches, the speed of the signal down the microstrip can be calculated as $4 \text{ in}/0.619 \text{ ns} = 6.46 \text{ in/ns}$. This is very close to the 6.42 in/ns calculated as the speed of the signal using the end to end method.

Using this value of the speed of the signal, we can extract the laminate's dielectric properties.

1.5 Extracting bulk dielectric constant of the laminate

The speed of the signal down a transmission line is directly related to the dielectric constant the signal sees. In a stripline structure, such as shown in Figure 11, the signal sees a uniform, homogeneous material, with a composite dielectric constant that is made up of a combination of the resin dielectric constant and the glass weave dielectric constant. Small variations in the local relative combination can affect the local dielectric constant which is an important source of skew between adjacent lines in a differential pair.



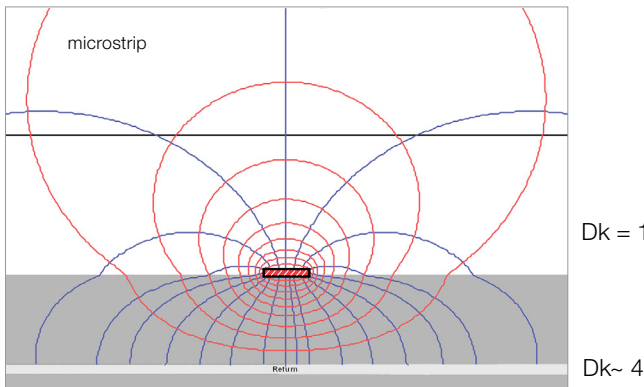
$$Dk = \left(\frac{11.803}{\text{velocity}} \right)^2$$

Figure 11. A stripline construction and extracting the bulk dielectric constant.

From a measurement of the speed of a signal down a stripline transmission line, the effective dielectric constant the signal sees, Dk , can be extracted using the simple relationship shown. The 11.803 number is the speed of light in air, in the units of inches/ns.

However, in a microstrip, the effective dielectric constant the signal sees is not the bulk value of the laminate.

In a microstrip, some of the electric field lines are in the bulk laminate, and see the laminate composite dielectric constant, but some of the field lines, as shown in Figure 12, are in the air, with a dielectric constant of one. The signal sees a composite of these two materials, which creates an effective dielectric constant, Dk_{eff} . It is this value which affects the signal speed and can be extracted from the measured speed of the signal.



$$Dk_{\text{eff}} = \left(\frac{11.803}{\text{velocity}} \right)^2$$

Figure 12. Effective dielectric constant in microstrip.

In this example, the speed is 6.46 in/ns. The extracted effective dielectric constant would be 3.34. This is unfortunately not a very useful number. It is not the bulk dielectric constant of the laminate. We cannot use this value of the effective dielectric constant in a field solver or approximation to help us calculate the impedance of any other geometries, for example. We really need to convert the effective dielectric constant into the actual bulk dielectric constant.

This conversion is related to the precise nature of the electric field lines, and what fraction is in the air and the bulk laminate. It also depends very much on the cross section geometry of the microstrip. The only way of converting the extracted, effective dielectric constant into the bulk laminate dielectric constant is using a 2D field solver.

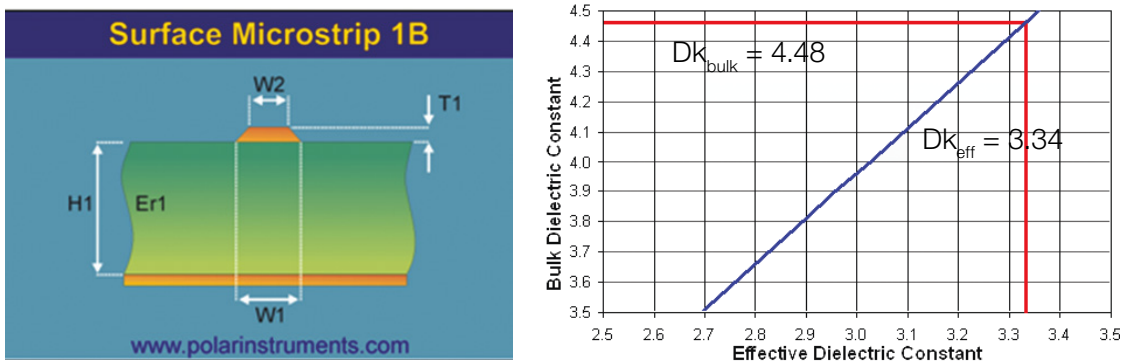


Figure 13. Using a field solver to back out the bulk dielectric constant form the effective dielectric constant.

In this example, we use a 2D field solver to calculate the effective dielectric constant for different bulk values, using the same geometry as the trace that is measured. We set up the field solver with the cross section information about the specific microstrip that was measured, and use the field solver to calculate the effective dielectric constant for different bulk dielectric constant values.

When we plot up the bulk dielectric constant verses the effective dielectric constant, we get a relatively straight line, as shown in Figure 13. We use this curve to back out the bulk dielectric constant, given the effective value we measured of 3.34. This analysis gives a bulk dielectric constant for this laminate of 4.48.

The TDR enables the measurement of the effective dielectric constant while the 2D field solver enables the conversion of the effective dielectric constant into the bulk dielectric constant.

1.6 Building a model of a discontinuity such as a corner, test pad, gap in the return path, SMA launch, terminating resistor

1.6.1 Extracting a model for capacitive discontinuities

Not all interconnect structures are uniform transmission lines. As much as we might try to eliminate them, there will often be discontinuities that are unavoidable. For example, test pads, component leads, 90 degree corners, gaps in the return path or even engineering change wires, will all create discontinuities. These structures, by their nature, are nonuniform and often times difficult to calculate, other than with a 3D field solver. Sometimes, the quickest way to evaluate their impedance is to build a structure, and measure it.

From the measured response, we can empirically evaluate the impact on the signal if we match the TDR's rise time to the rise time of the application. We could then directly measure off the screen of the TDR, the amount of reflected voltage noise we might see in the system. Alternatively, we could use the TDR to extract a simple, first order model for the structure and use this model in a system level simulation to evaluate the impact of the discontinuity. Finally, if we need more accuracy or a higher bandwidth model than what we can get directly from the screen, we can take the measured data from the TDR and bring it into a modeling and simulation tool, such as SPICE or ADS to fit a more accurate model. These processes are illustrated in this section.

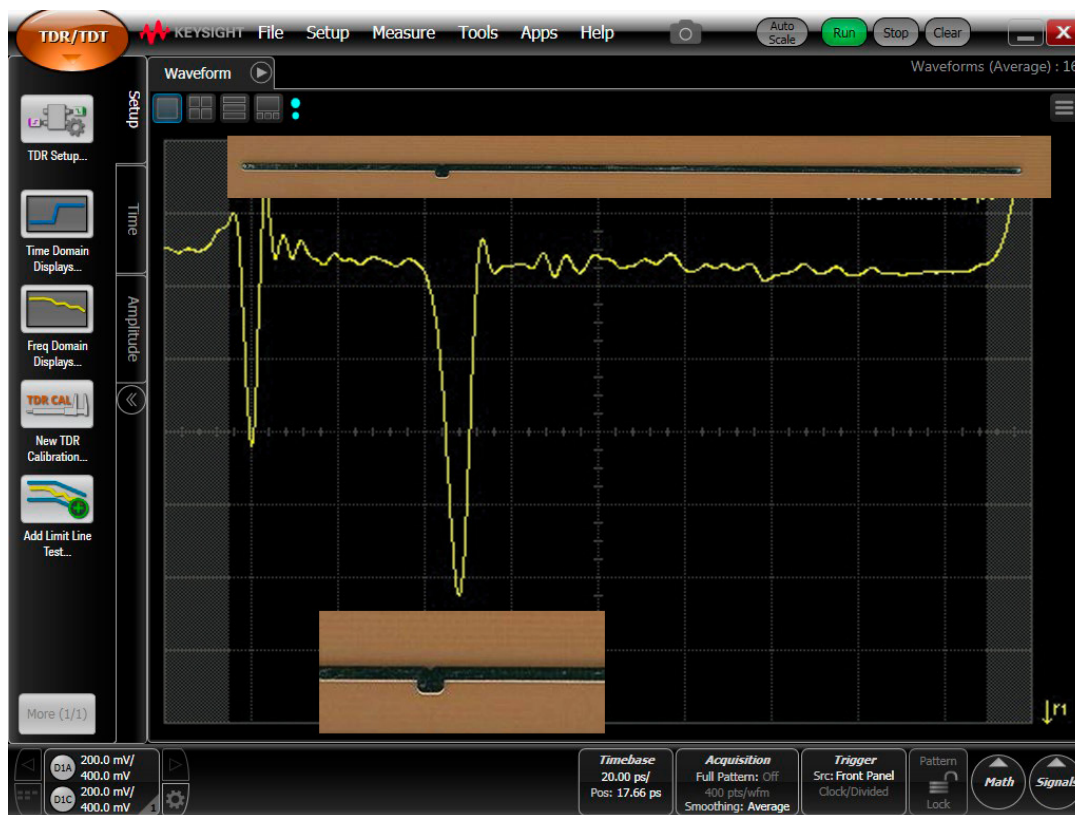


Figure 14. TDR Response from a uniform transmission line having a small test pad.

Let's start with a simple test pad on an otherwise uniform line, as shown in Figure 14. The TDR response is shown as the yellow line on the screen, displayed in an Ohms scale, with 2 Ohms/div. The small dip near the beginning of the line is due to the SMA launch. The large dip about three divisions from the left edge is from the test pad.

On this scale, the reflected signal from the small test pad looks huge, but is a discontinuity of only 4.5 divisions or about nine Ohms. This can be interpreted as the instantaneous impedance a signal would see, if it had the rise time of the TDR, in this case, about 40 ps. Since this test pad is not a uniform transmission line, the instantaneous impedance is not related to a characteristic impedance, and the impedance a signal would see is going to depend on the rise time of the signal. We can use the TDR to directly emulate any rise time from as fast as 20 ps up to longer than one nanosecond, to directly evaluate the impact of the discontinuity on the system rise time.

Using the built-in calibration feature of the DCA 86100C, we can change the effective rise time of the stimulus and directly display the response from this small discontinuity. The structure is the same, and the scale is the same for each of the four rise times of 40, 100, 200, and 500 ps.




Figure 15. TDR response for a uniform transmission line with a small test pad, at four different rise times of 40, 100, 200, and 500 ps.

From Figure 15, we see clearly that the instantaneous impedance a signal would see encountering this test pad is strongly dependent on the rise time of the signal. If the rise time were 40 ps, the signal would see an impedance discontinuity of about nine Ohms. At 100 ps, this is only about five Ohms, at 200 ps, it is 2.5 Ohms, and at 500 ps, it is less than one Ohm, hardly noticeable to the signal. Based on the noise budget allocated for discontinuities, we could determine what is the shortest rise time at which this discontinuity would begin to cause problems, or could be ignored.

For example, if three Ohm discontinuities were allowable, this particular test pad could be used for rise times as short as 250 ps. Much below this, and the impact might be felt. The way to know for sure would be to build a model for the discontinuity and use it in a simulation.

By inspection, the simplest model for this discontinuity is a single lumped capacitor. At 40 ps rise time, the TDR response is close to that from an ideal lumped capacitor. We can use the built in “excess reactance” feature to build a model and extract the parameter values directly from the screen using markers.

The excess reactance feature built into the DCA 86100C, will model the DUT as a uniform transmission line having a single discontinuity, either a lumped inductor or lumped capacitor. The software will use the position of the two vertical markers to define the region of the response where the capacitance or inductance will be extracted.

To use this feature, position the markers on either side of the discontinuity and read the amount of capacitance or inductance from the “excess reactance” value on the screen. One hint in using this feature is to position the markers so that they have roughly the same impedance value on either side of the discontinuity. It doesn't matter what the vertical scale is when the excess reactance function is used.

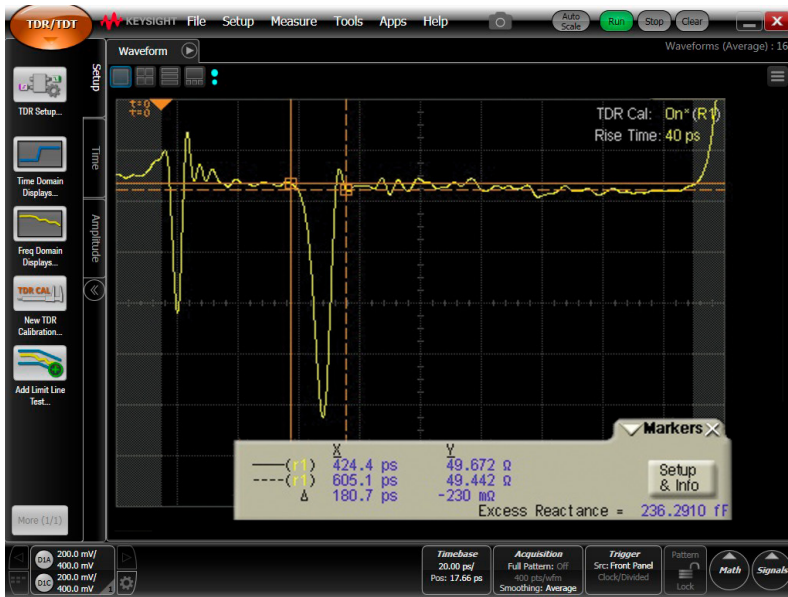


Figure 16. Using the excess reactance feature to extract the capacitance of a test pad.

In Figure 16, the markers are used to extract the capacitance of the test pad. The model we are assuming is a single lumped capacitor. The value of this capacitance, we can read off the screen as 236 fF. This capacitance, plus the impedance of the uniform part of the line, 49 Ohms, provides a complete model for this transmission line structure.

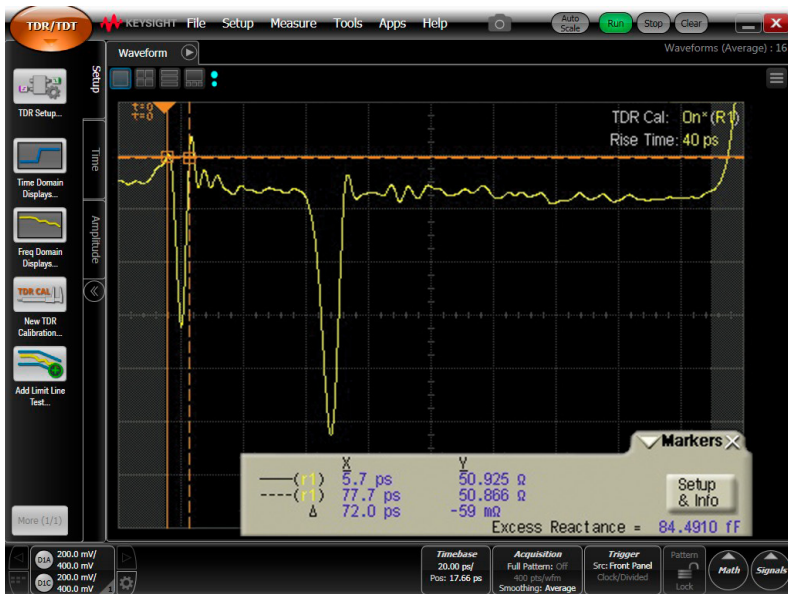


Figure 17. Using markers to extract the excess capacitance of the SMA launch.

While we are at it, we can also extract the capacitance associated with the pads used on the end of the transmission line for the SMA launch at the beginning of the line. Using the markers in Figure 17, we get 84 fF of capacitance. It's clear that the TDR has a very high sensitivity for extracting discontinuity values. On this scale, 84 fF of capacitance is a very large and easily measurable effect.

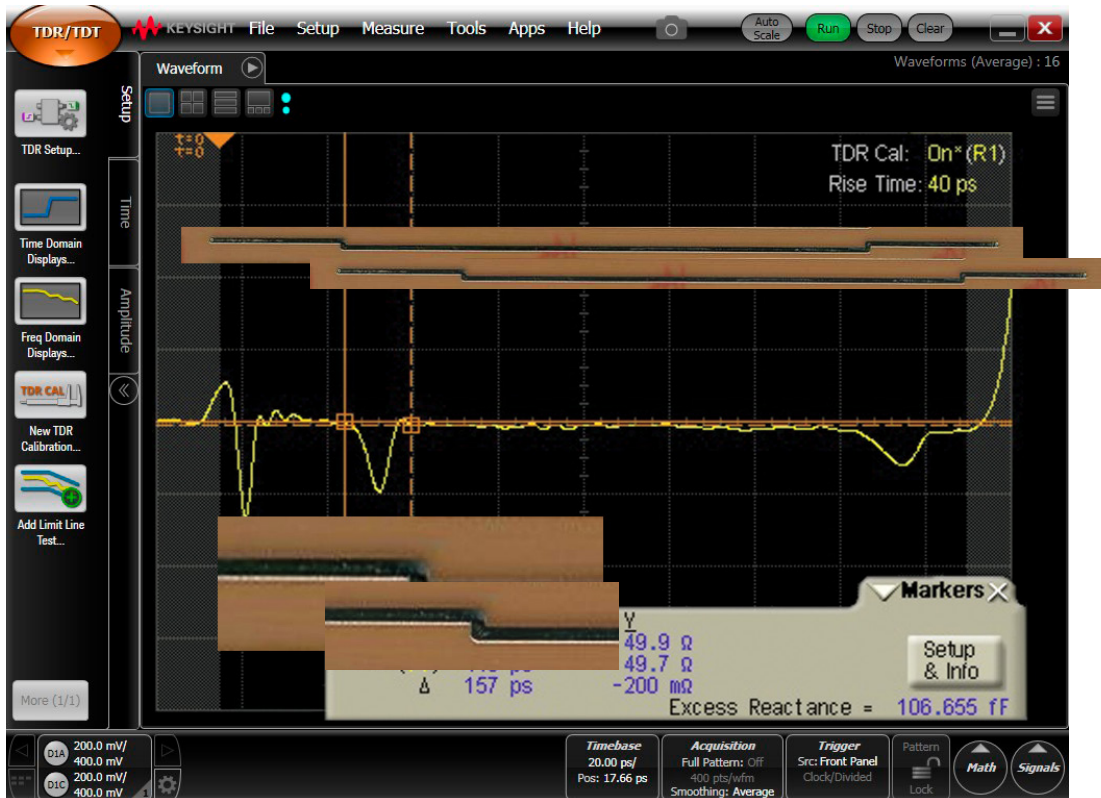


Figure 18. Using markers to extract the excess capacitance of two corners.

We can apply this technique to measure the capacitance associated with a corner. Corners, or 90 degree bends on traces, have stimulated a lot of discussion and concern over the years to signal integrity engineers. In Figure 18, we see a simple, uniform transmission line structure built as a microstrip, that has two small jags in it. Each jag is a combination of two, 90 degree bends.

On this scale of 10 mV/div, it is absolutely clear that a corner causes an impedance discontinuity that is easily measured. On this scale, each division is a reflection coefficient of 10 mV out of 200 mV incident signal, or five percent. The two corners in each jag, together create a reflected signal on the order of four percent, so each corner creates about two percent reflection at a rise time of 40 ps. This is in a 50 Ohm line, with a line width of about 60 mils.

In many analog, RF circuits where a flat response over a narrow frequency range is important, or a return loss below -40 dB is sometimes required, a corner can introduce big problems. Historically, many RF and microwave circuits were built on thick ceramic substrates where line widths were 100 mils or wider. This would almost double the impact from a corner. This is one of the reasons why corners have developed a reputation as a potential problem and should be avoided.

Using the TDR measurement, we can build a model for a corner and use this model in a system simulation to evaluate whether a corner might pose a potential problem, or can be ignored. Clearly, from the TDR response, we can see that the impact of the two corners in this jag looks like the response from a single lumped capacitor. Using the two markers, we can measure the excess capacitance from the two corners as 107 fF. Since this is from two corners, this corresponds to about 53 fF of capacitance per corner. This value can be put into a circuit simulation tool, such as SPICE or ADS to simulate the impact from a 53 fF capacitor.

It is also important to note that the amount of capacitance in a corner will scale with the width of the line, for the same impedance traces. If a 60 mil wide line has a capacitance in one corner of roughly 60 fF, then a five mil wide line will have a capacitance of roughly 5 fF. This is a good rule of thumb to remember: the capacitance of a corner is about 1 fF per mil of line width for a 50 Ohm line.

1.6.2 Extracting a model for inductive discontinuities

The second type of discontinuity is an inductive discontinuity. These arise, for example, when the line width of a trace necks down, as when passing through a via field, or if the return path is disturbed, such as when the trace crosses a gap, or when there is an engineering change wire. An inductive discontinuity will look like a higher impedance and give a peak reflection as the TDR response.

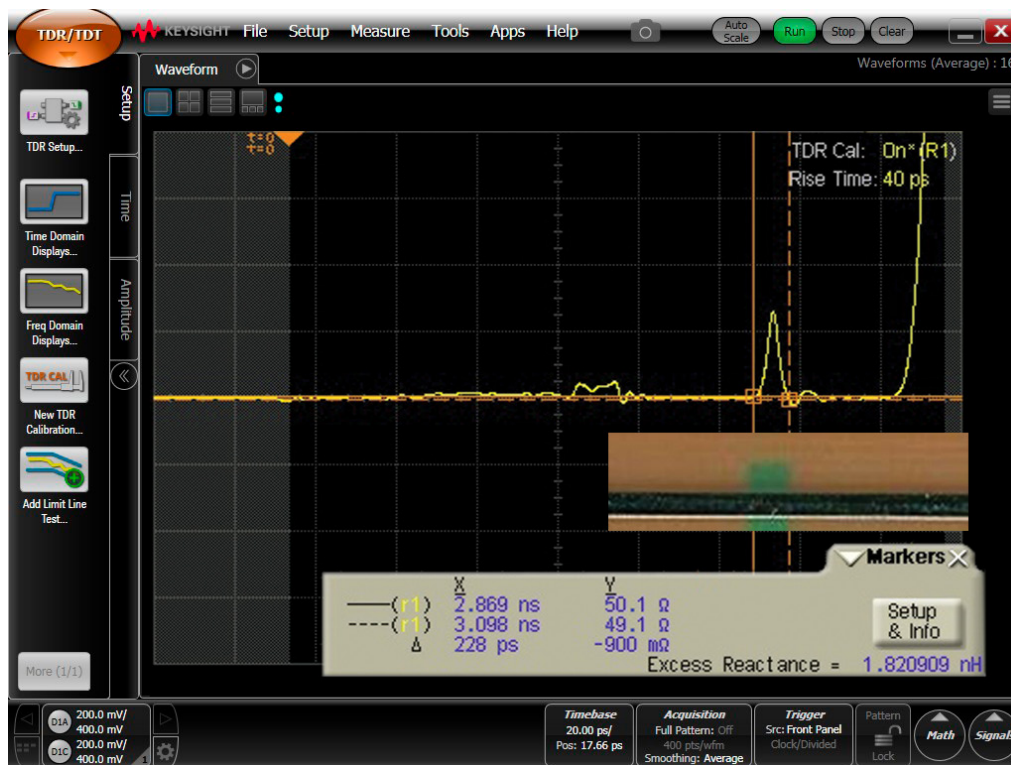


Figure 19. Using markers to extract the excess inductance of a short gap in the return path.

Figure 19 is an example of the TDR response from a uniform transmission line with a signal line that passes over a very short gap in the return path. This commonly occurs when passing through a via field where the clearance or antipads are large enough or on tight enough centers to overlap, inadvertently creating a gap.

The TDR response is a positive peak, just as we would expect from a lumped inductor. We can position the markers on either side of the discontinuity and read the lumped inductance right off the screen as 1.8 nH. For this gap, roughly 100 mils long, the loop inductance created is about 1.8 nH.

If the gap length were increased, the inductive discontinuity created would increase as well. Figure 20 is an example of the TDR response from two different return path gap structures. The blue line is the response described previously, from a 100 mil long gap. The inductance was about 1.8 nH.



Figure 20. Using markers to extract the excess inductance of a large gap in the return path.

The yellow trace in Figure 20 is the TDR response from a longer gap of 500 mils length. Using the markers, we can measure this lumped inductive discontinuity as 6.3 nH. If we were concerned about the impact from these features, we could easily use the model of the uniform transmission line with these lumped series inductors in a circuit simulation to determine if the additional noise or impact on timing was sufficient to warrant attention.

1.6.3 Modeling termination resistors

It is not just interconnects that can be characterized and modeled with a TDR. We can also use a one port TDR to build a model for discrete components such as termination resistors. Figure 21 is an example of the TDR response from a 50 Ohm axial lead termination resistor connected to the end of a 50 Ohm transmission line.

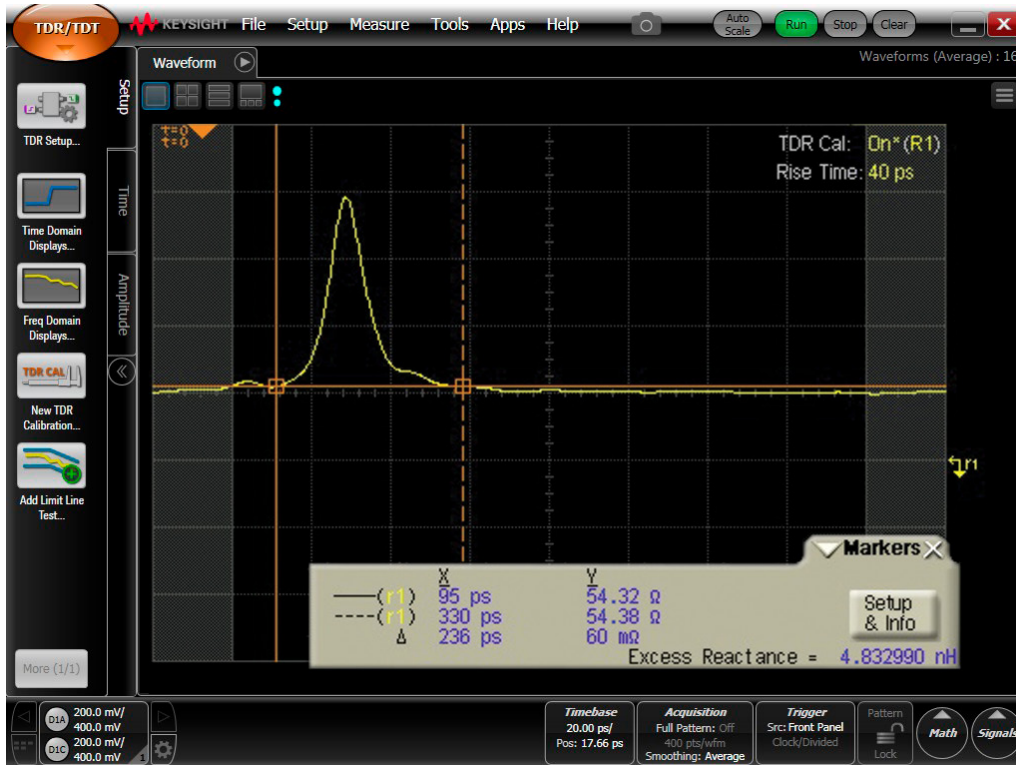


Figure 21. Using markers to extract the excess inductance of an axial lead termination resistor.

On the left side of the peak is the transmission line going to the resistor, which on this scale of 50 Ohms per division, we can read off the screen, has an impedance of about 50 Ohms. The resistor itself also has an impedance on the order of 50 Ohms. This is seen on the right side of the peak. It's just that it also seems to have some lumped inductance associated with it. This series inductance arises from the long body of the resistor and the leads connecting the signal to the return path.

Using the excess reactance function, we can read the excess lumped inductance of this resistor by positioning the markers on either side of the discontinuity and reading the series loop inductance off the screen as about 4.8 nH. The equivalent circuit model we are assuming is a uniform transmission line with an ideal 4.8 nH series inductor, followed by an ideal resistor of 50 Ohms.

With a rise time of 40 ps, the signal sees a peak impedance of about 200 Ohms. This is the 150 Ohms discontinuity in addition to the 50 Ohms of the line. Of course, as we saw earlier, the impedance a signal would see when it interacts with this inductance will depend on the rise time of the signal. A longer rise time will see a lower impedance, however, the excess inductance of this resistor will not change with the rise time. It is only a function of the geometry of the device.

This is a huge amount of inductance and would probably limit the operation of any circuit it was used in to rise times greater than about 1 ns. At 1 ns rise time, the roughly 5 nH inductance would create a noise level of about 10 percent.

Of course, for high speed circuits, axial lead resistors are out of the question. Surface mount technology (SMT) resistors are required which are physically smaller and can be mounted with much less equivalent series inductance. In Figure 22 is an example of an 0603 SMT resistor soldered between the signal and return path on a signal integrity test board, available from BeTheSignal.com. On the top side of the board is an SMT SMA connector which is connected to the TDR.

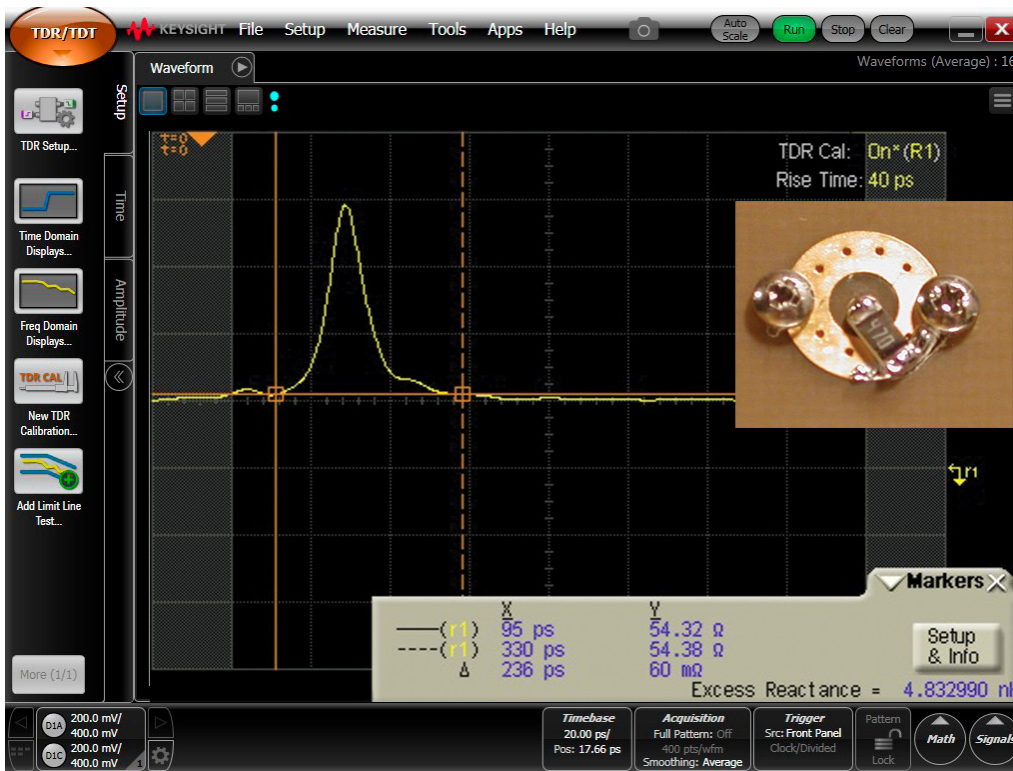


Figure 22. Using markers to extract the excess inductance of a SMT termination resistor mounted to the MCW620 test board.

The TDR response from this component is also shown in Figure 22, using a scale of 10 Ohms per division, much more sensitive than for the axial lead example.

On the left side of the peak is the transmission line and connector going to the resistor, which we can read from the screen, has an impedance of about 50 Ohms. The resistor also has an impedance on the order of 50 Ohms, within about one percent. This is seen on the right side of the peak. The small peak is the reflection from the series inductance arising from the resistor body, the surface trace and the vias going to the top layer. The design for the attach of this particular SMT resistor has been optimized for low mounting inductance.

Using the marker function, we can read the excess lumped inductance of this resistor as about 480 pH. This is about an order of magnitude lower series loop inductance than an axial lead resistor, and is typical of what can be obtained with an optimized mounting design for a body size of 0603. This applies to resistor components as well as capacitor components.

With a rise time of 50 ps, it looks like it has a series impedance of about 11 Ohms. Of course, as we saw earlier, the impedance a signal would see when it interacts with this inductance will depend on the rise time of the signal. A shorter rise time will see a higher impedance. However, the excess inductance of this resistor will not change with the rise time. It is only a function of the geometry of the device.

Using the marker function to read the excess inductance off the front screen assumes a simple model for the DUT. In the case of this terminating resistor, we assume the model is a single, series lumped inductor. The excess inductance we read off the screen is then the series loop inductance of this resistor. However, we don't have any clear sense from looking at the screen how high the bandwidth is for this simple model, nor can we build more sophisticated models of components easily from just the front screen. This is a case where switching to the frequency domain can get us to the answer faster.

In TDR measurements, a time domain, fast rising step edge is sent into the DUT and the reflected signal measured. In addition, we can look at the signal that is transmitted through the DUT. This is the time domain transmitted signal, TDT.

If we look at the signal incident into the DUT, it can be thought of as being composed of a series of sine waves, each with a different frequency, amplitude and phase. Each sine wave component will interact with the DUT independently. When a sine wave reflects from the DUT, the amplitude and phase may change a different amount for each frequency. This variation gives rise to the particular reflected pattern. Likewise, the transmitted signal will have each incident frequency component with a different amplitude and phase.

There is no difference in the information content between the time domain view of the TDR or TDT signal, or the frequency domain view. Using Fourier Transform techniques, the time domain response can be mathematically transformed into the frequency domain response and back again without changing or losing any information. These two domains tell the same story, they just emphasize different parts of the story.

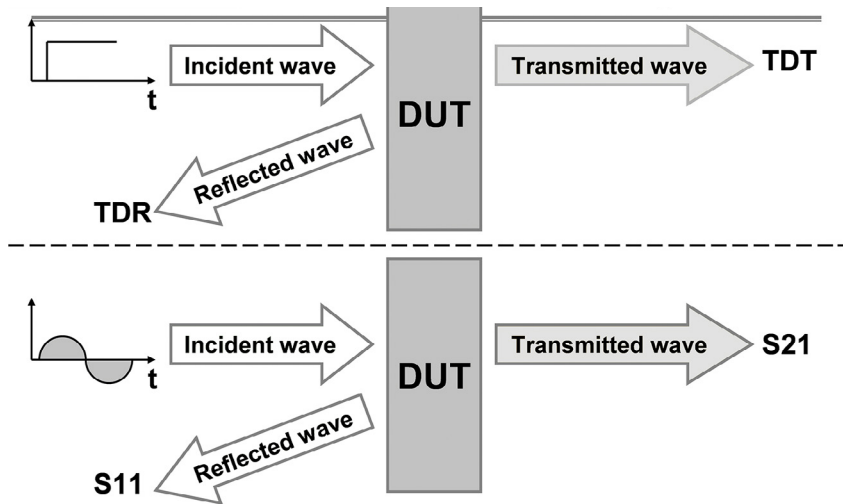


Figure 23. TDR and VNA techniques.

TDR measurements are more sensitive to the instantaneous impedance profile of the interconnect, while frequency domain measurements are more sensitive to the total input impedance looking into the DUT. To distinguish these two domains, we also use different words to refer to them. Time domain measurements are TDR and TDT responses, while frequency domain reflection and transmission responses are referred to as S (scattering) parameters. S11 is the reflected signal, S21 is the transmitted signal. They are also often called the return loss and the insertion loss. This is illustrated in Figure 23.

Depending on the question asked, the answer may be obtained faster from one domain or the other. If the question is, what is the characteristic impedance of the uniform transmission line, the time domain display of the information will get us the answer faster. If the question is what is the bandwidth of the model, the frequency domain display of the information will get us the answer faster.

In the DCA 86100C, the frequency domain response of the TDR or TDT waveform can be instantly displayed by pulling down the S-Param window, located in the upper right hand corner of the screen.

1.7 Building a high bandwidth model of a component

To evaluate the bandwidth of the model for this SMT resistor, we can bring the measured S11 data of the DUT into a simulation tool like the Keysight Technologies, Inc. ADS, and perform more sophisticated modeling. The corresponding frequency domain response is S11, the return loss. This is one of the S parameter matrix elements.

By selecting the S-Param tab on the upper right corner of the screen, we can bring down the converted time domain response as a frequency domain response, S11. This is still the reflection coefficient, but now it is displayed in the frequency domain, and is related to reflections, not from instantaneous impedances, but from the total, integrated impedance of the entire DUT, looking into its input.

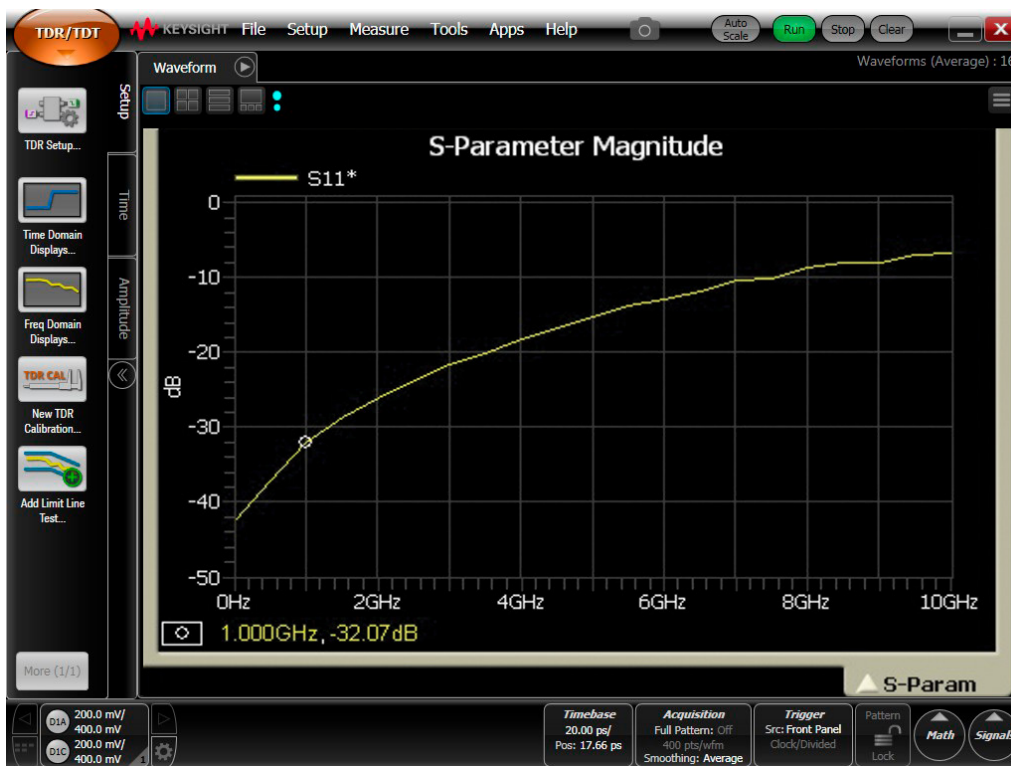


Figure 24. Converted S11 of the SMT termination resistor.

Displayed in Figure 24 is the measured S11 of this terminating resistor component, up to about 10 GHz. At low frequency, the magnitude of S11 is very low; not much of the signal comes back because the impedance the TDR sees is a pretty good match to 50 Ohms. As we go up in frequency, the impedance of the terminating resistor increases due to the higher impedance of the series inductance, causing more signal to reflect, which we see as a decrease in the negative dB values.

We can save this data in a .s1p Touchstone formatted file and bring it into ADS for further analysis.

ADS is a powerful analysis and modeling tool. In this example, we are interested in how well this simple model of the SMT resistor, being a series resistor and inductor, fits the actual, measured response, and up to what bandwidth it is accurate.

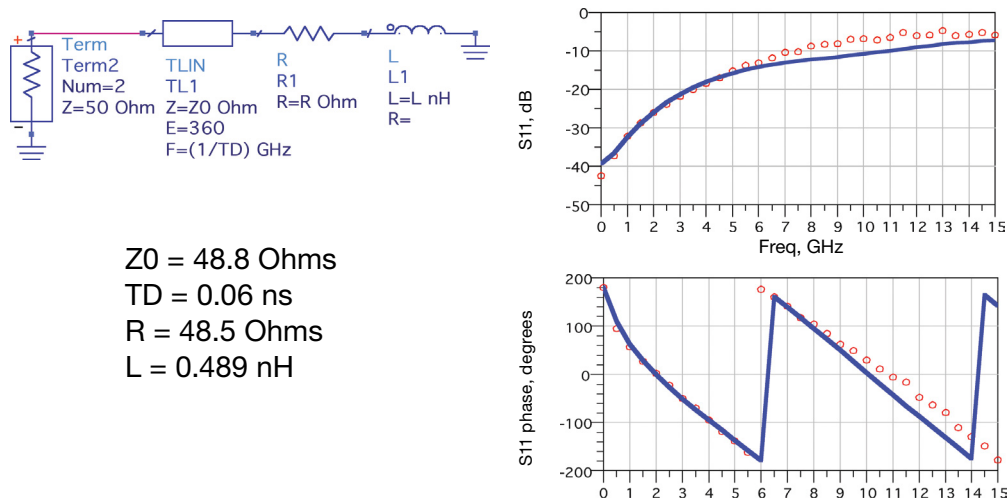


Figure 25. ADS model of resistor and the measured and simulated S parameters.

In ADS, we build a circuit topology to match what we think this DUT is. This simple circuit is shown in Figure 25. We start with an ideal, lossless, uniform transmission line with an ideal resistor in series with an ideal inductor, which terminate the end of the line. There are four parameters that fully characterize this circuit, the characteristic impedance and time delay of the line, the resistance and the inductance. We just don't know what their parameter values are.

However, we can take advantage of the built in, powerful optimization features in ADS to find the best set of parameter values that gives the closest agreement between the actual measured S11 values from the TDR to the simulated S11 values of this circuit model. We perform the optimization for the measured data at 4 GHz and below and find the best set of parameter values is a characteristic impedance of 48.8 Ohms, a time delay of 0.06 ns, a resistance of 48.5 Ohms and an inductance of 0.489 nH.

Also shown in Figure 25 is the comparison of the measured return loss, as red circles, and the simulated return loss, as solid blue lines, up to 15 GHz. We see that the simulated return loss of this simple model matches the actual, measured return loss of the component very accurately up to about 7 GHz. The bandwidth of this model is 7 GHz.

We also find that the extracted value of the equivalent series inductance of this SMT resistor is 0.489 nH. Using the markers on the screen, we estimated it to be 0.481 nH. This estimate from the excess reactance is within two percent of what we get using higher bandwidth modeling techniques.

It is also important to note that the actual frequency values that are measured by the TDR are rather sparse. In the time domain response, the scale was 200 ps per division. This scale gave a comfortable time resolution to see the inductive spike, using the 50 ps rise time. With 10 horizontal divisions full scale, the entire time sweep was 2 ns. When we convert this measured data from the time domain to the frequency domain, the 2 ns window converts to a first harmonic frequency value of $1/(2 \text{ ns})$ or 500 MHz. This is the step size used for all the frequency values.

We see in the display of the measured data in Figure 25 that there is a circle at every 500 MHz. This is the frequency resolution of the TDR. If we wanted finer frequency resolution, we would have had to use a longer time window, and a larger number of picoseconds per division. For example, if we use 1 ns per division, for a total of 10 ns full scale, the frequency resolution would have been $1/(10 \text{ ns})$ or 100 MHz.

1.8 Directly emulating the impact on a signal with the system rise time from a discontinuity

A short discontinuity will look like either a lumped capacitor or inductor. The impact on the reflected signal will depend on the rise time of the incident signal. A shorter rise time will create a larger reflected signal. This means that the instantaneous impedance of a discontinuity, as measured off the screen of a TDR is rise time dependent. It really doesn't mean anything to describe the impedance of a discontinuity, unless we also specify at what rise time. Even then, we can't do much with this information.

One way of evaluating the impact of this discontinuity in a particular application is to build a model for the structure, using excess reactance, and bring it into a circuit simulator. We could then use the driver models to simulate signals at the system's rise time and calculate the amount of reflection noise based on this interconnect model.

Another way of evaluating the expected reflection noise is to use the TDR to emulate the system rise time. After the TDR stimulus is calibrated, we can change the rise time of the stimulus to match any rise time from as low as 20 ps to well above 1 ns. We can then directly measure the reflected noise for different rise times and measure the impact on the signal a specific interconnect discontinuity would have.

In Figure 26 is an example of a uniform 50 Ohm transmission line which has a region about 200 mils long which necks down. This is what typically occurs when a line has to pass through a via field associated with a connector foot print. The impedance of this line goes from 50 Ohms up to about 70 Ohms and then back to 50 Ohms.

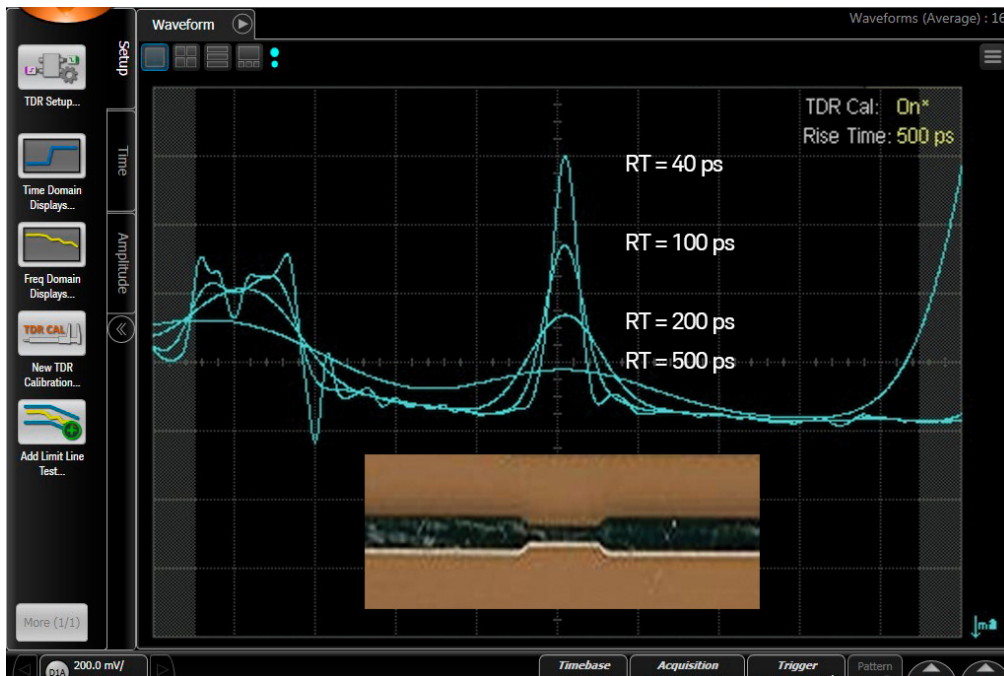


Figure 26. Emulating system rise time responses for a 200 mil long neck down region with RT = 40, 100, 200, and 500 ps.

By changing the rise time of the stimulus, we can directly measure the reflected noise at rise times of 40, 100, 200, and 500 ps. The vertical scale is 10 mV/div. In the region of the discontinuity, the peak reflected noise is about 37, 22, 11, and 5 mV, respectively. This is with an incident signal of 200 mV, or a reflection coefficient of 18, 11, 5.5, and 2.5 percent. The decrease in reflected amplitude does not scale directly with rise time because of the complication of the finite size of the discontinuity and the rise time of the signal. This complexity is automatically taken into account in the TDR measurement.

For example, if a 500 ps signal were to encounter this 200 mil long neck down, it would see a reflection, but it would only be 2.5 percent, which might be acceptable. This would demonstrate the advantage of necking down the line to get through the via field is a reasonable compromise, compared with possibly adding more layers in the board to keep the line width uniform.

Alternatively, if the system rise time were 100 ps, we would see that the reflection noise of 11 percent might exceed the typical 5 percent noise budget allocated to reflection noise, and it might not be acceptable to neck down the trace, but may require a rerouting around the connector field.

Using the built in adjustment of the rise time of the stimulus, we can emulate the actual system's rise time for a specific application and directly measure the performance of an interconnect in the specific application without having to first build a model and run a simulation. This can save a lot of time and help us to get to an acceptable answer faster.

2.0 2-Port TDR/TDT

2.1 Overview

As we saw in the previous section, a TDR generates a stimulus source that interacts with an interconnect. With one port, we were able to measure the response from one connection to the interconnect. This limited us to just looking at the signal that reflects right back into the source. From this type of measurement, we got information about the impedance profile and properties of the interconnect and extracted parameter values for uniform transmission lines, with discrete discontinuities.

By adding a second port to the TDR, we can dramatically expand the sort of measurements possible and the information we can extract about an interconnect. There are three important new measurements that can be performed with an additional port: the transmitted signal, coupled noise, and the differential or common signal response of a differential pair. The most important applications that can be addressed with these techniques and examples of each, are illustrated in this chapter.

2.2 Introduction to TDR/TDT

When the second port is connected to the far end of the same transmission line and is a receiver, we call this time domain transmission, or TDT. A schematic of this configuration is shown in Figure 27. The combination of measuring the TDR response and TDT response of an interconnect allows accurate characterization of the impedance profile of the interconnect, the speed of the signal, the attenuation of the signal, the dielectric constant, the dissipation factor of the laminate material and the bandwidth of the interconnect.



Figure 27. Configuration for TDR/TDT measurements.

The TDR can be set up for TDR/TDT operation by selecting TDR setup, choosing single-ended stimulus mode, selecting change DUT type and choosing a 2-port DUT. You can assign any available channel to Port 2 or click auto-connect. This is shown in Figure 28.

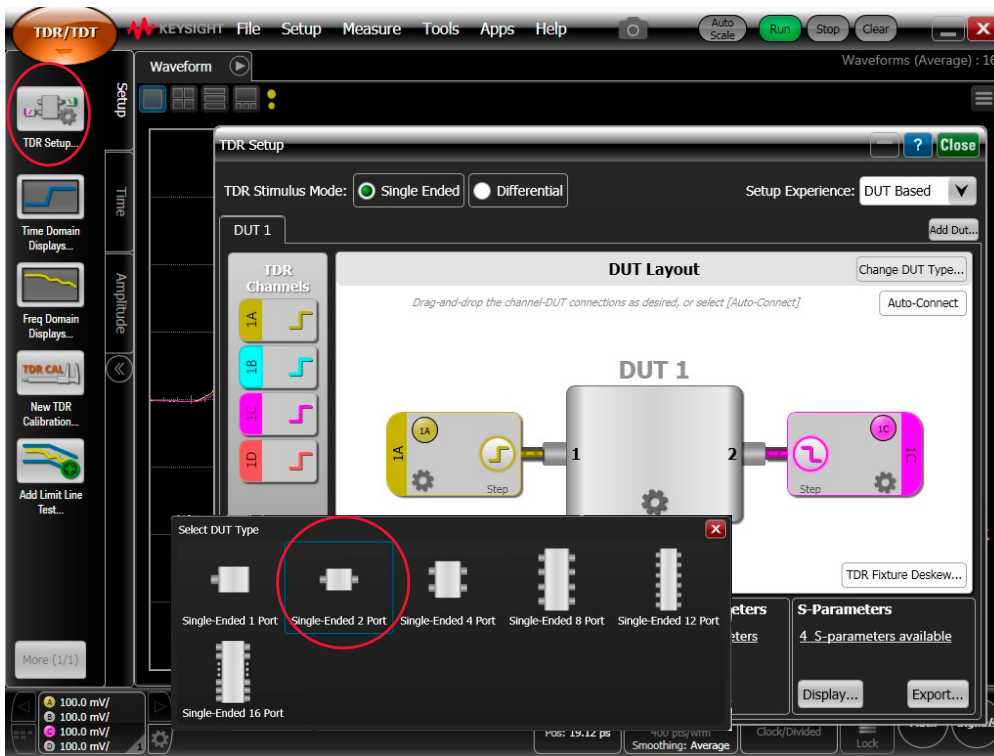


Figure 28. TDR set up screen for TDR/TDT operation.

2.3 Measuring insertion loss and return loss

In the simplest application, the ports of the TDR are connected to each end of the single ended transmission line. Port 1 is the TDR response we are familiar with, while channel 2 is the transmitted signal. In the TDR response of a uniform, 8 inch microstrip transmission line, as shown in Figure 29, the end of the line is seen as having an impedance of 50 Ohms. This is the cable connected to the end of the DUT, and then ultimately, the source termination inside the TDR's second channel.

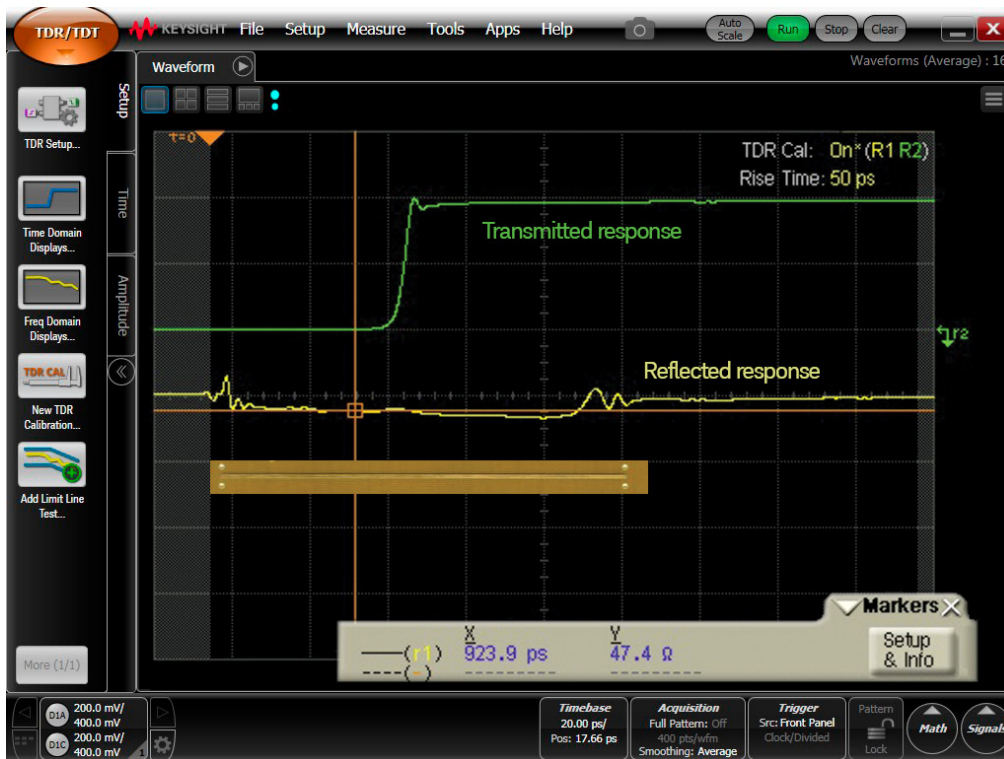


Figure 29. Example of TDR/TDT response from eight inch long microstrip transmission line on 20 mV/div and 500 ps/div scales.

The time base in this application is 500 ps/div, with the vertical scale at 20 mV/div. The marker is being used to extract the impedance of the line as 47.4 Ohms. Note that the green trace, the signal transmitted through the interconnect, on 100 mV/div scale, shows the signal coming out exactly halfway between the time the signal goes into the front of the line, reflects off the back end, and is received at the source.

The TDR signal looks at the round trip time of the flight down the interconnect, then back again to the front, while the TDT signal sees the one way path through the interconnect. In the time domain display, we can see the impedance discontinuities of the SMA launches on the two ends of the line, and see that the line is not a perfectly uniform transmission line. On this scale of 20 mV/div or a reflection coefficient of 10%/div, the variation in impedance is about 1 Ohm down the line.

The transmitted signal is a relatively fast edge, but it is difficult to get much information off the screen from this received signal. Though we could measure the 10-90 or 20-80 rise time directly off the screen, its not clear what we would do with this information, as the interconnect distorts the edge into a not really Gaussian edge. This is a case where we can take the same information content, but change how it is displayed, to interpret it more quickly and easily.

Figure 30, shows the same measured response as shown in the time domain, but now transformed into the frequency domain. This screen is accessed by clicking on the S-Param tab in the upper right hand corner of the TDR response screen. In the frequency domain, we call the TDR signal, S11 and the TDT signal, we call S21. These are two of the S parameters that describe scattered waveforms in the frequency domain. S11 is also called return loss and S21 is called insertion loss. The vertical scale is the magnitude of the S parameter, in dB.

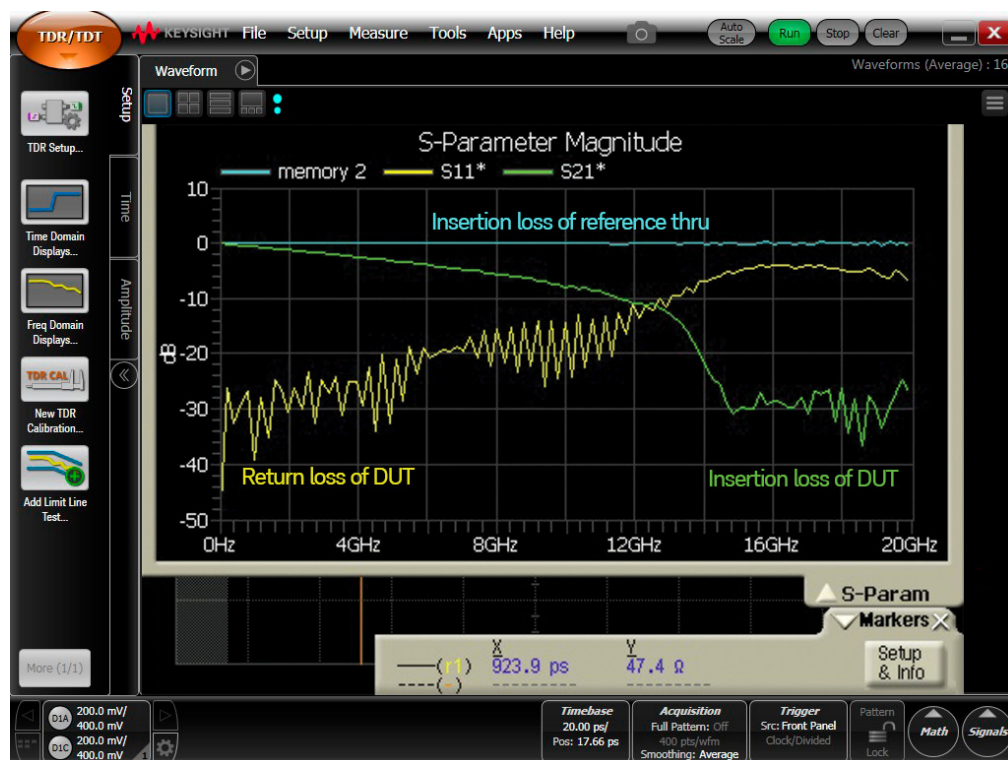


Figure 30. TDR/TDT response converted into frequency domain: return loss/insertion loss.

The blue trace is the insertion loss for a reference thru. Of course, if we have a perfect thru, every frequency component will be transmitted with no attenuation, and the amplitude of the received signal is the same as the incident signal. The magnitude of the insertion loss is always 1, and in dB, is 0 dB. This is flat across the entire 20 GHz frequency range.

The yellow trace, starting at about -30 dB at low frequency, is the return loss for this same transmission line, which is really S11 in the frequency domain. The green trace is the insertion loss of this transmission line, or S21. On this display, we are only showing the magnitude of the S parameters, the phase information is there, it is just less important to display.

The return loss starts at relatively low values, near -30 dB, and then creeps up, eventually reaching the -10 dB range, above about 12 GHz. This is a measure of the impedance mismatch of this transmission line and the 50 Ohm connections on either end.

The insertion loss has immediately useful information. In a high speed serial link, the transmitters and receivers work together to enable high bit rate signals to be transmitted and then received. In simple CMOS drivers, an insertion loss of -3 dB might be acceptable, before a significant bit error rate. With simple SerDes chips, an insertion loss of -10 dB might be acceptable, while for state of the art, high end SerDes chips, -20 dB might be acceptable. If we know the acceptable insertion loss for a particular SerDes technology, we can directly measure off the screen the maximum bit rate an interconnect is capable of.

As a rough rule of thumb, if the bit rate in Gbps is BR, the bandwidth of the signal is BW, the highest sine wave frequency component is roughly $BW = 0.5 \times BR$, or $BR = 2 \times BW$. The BW is defined by the highest frequency signal that can be transmitted through the interconnect, and still have less attenuation than the SerDes can compensate for. Using low end SerDes, the acceptable insertion loss might be -10 dB, and the bandwidth for this eight inch long microstrip, we can read right off the screen in Figure 30, would be about 12 GHz. This would allow operation well above 20 Gbps bit rates. But, this is for a wide conductor, only eight inches long. In a longer backplane or motherboard, with connectors, daughter cards, and vias, the transmission properties are not as clean.




Figure 31. Return and insertion loss of a 24 inch interconnect on a motherboard with two daughtercards.

The TDR/TDT response of a 24 inch long stripline interconnect in a typical motherboard, is shown in Figure 31. In this example, an SMA launch connects the TDR cable into a small card, through a connector, a via field, back through a connector, and into the second channel of the TDR.

The green trace is the insertion loss displayed as S21. For this interconnect, the -10 dB insertion loss bandwidth is 2.7 GHz. For this interconnect, the maximum transmitted bit rate would be about 5 Gbps, using low end SerDes drivers and receivers.

2.4 Interconnect modeling to extract interconnect properties

The ability to take the measured data and display it as either time domain responses or frequency domain responses, means we can easily extract more information than if we were limited to just one domain. Further, by exporting the frequency domain insertion and return loss measurements as a Touchstone formatted file, we can use sophisticated modeling tools, such as Keysight’s ADS, to extract more information than we could get off the screen.

In this example, we will look at the uniform, eight inch long microstrip and how we can use modeling and simulation tools to extract material properties. The simplest model to describe this physical interconnect is an ideal transmission line. We can use the built in multilayer interconnect library (MIL) of ADS to build a physical model of this very microstrip, with the material properties parameterized, and extract their values from the measurement.

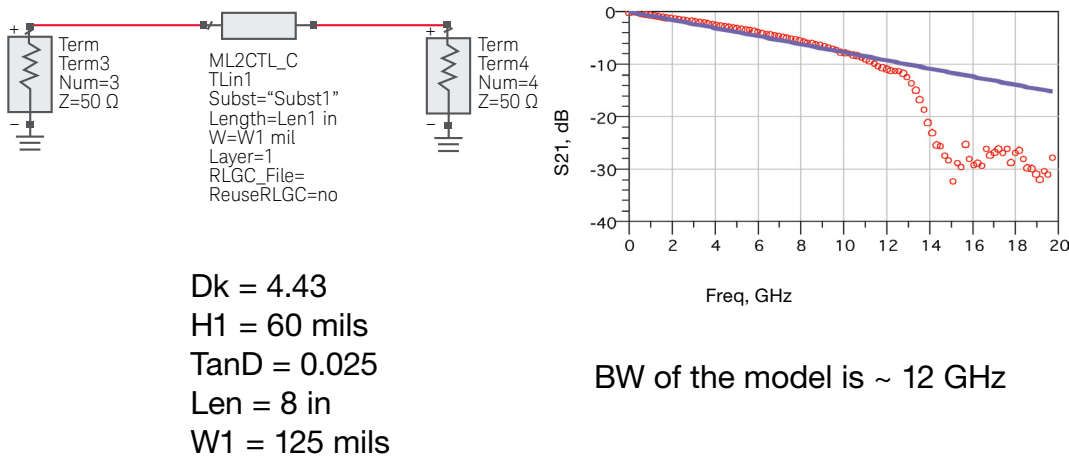


Figure 32. ADS modeling of a uniform eight inch long microstrip, showing the bandwidth of the simple model to be ~12 GHz.

Figure 32 shows the simplest model to describe this transmission line, as a single trace on a substrate, with a length of eight inches, a dielectric thickness of 60 mils and a line width of 125 mils. These are parameters measured directly from the physical interconnect. What we don't know initially are the laminate's bulk dielectric constant and its bulk dissipation factor. However, we have the measured insertion loss. Figure 32 displays the measured insertion loss of the interconnect as red circles. This is exactly the same data as displayed previously, from the screen of the TDR. The phase response is also used in the analysis, it's just not displayed in this figure.

Given this simple model, with the two unknown parameters, the dielectric constant, and dissipation factor, we use the built in optimizer in ADS to search all parameter space for the best fit values of these two terms to match the measured insertion loss response to the simulated insertion loss response. The blue line, in Figure 32 is the final value of the simulated insertion loss using a value of 4.43 as the dielectric constant and 0.025 as the dissipation factor. We can see from this display, the agreement between the measured and simulated insertion loss is excellent, up to about 12 GHz. This is the bandwidth of the model. There is even better agreement in the phase, not shown in this figure.

By building a simple model and fitting parameter values to the model, and taking advantage of the built-in 2D boundary element field solver and optimization tools of ADS, we are able to extract very accurate values for the material properties of the laminate from the TDR/TDT measurements. We are also able to convince ourselves that, in fact, this interconnect is very well behaved. There are no unusual, unexplained properties of this transmission line. There are no surprises, at least up to 12 GHz.

2.5 Identifying design features that contribute to excessive loss

Being able to quickly and easily bring measured TDR/TDT data from the TDR instrument directly into a modeling tool can sometimes cut the debug time from days to minutes by helping us unravel the root cause of surprising or anomalous behavior. Figure 33 is an example of the measured TDT response from three structures. The top horizontal line is the insertion loss measured from a reference thru, showing the very flat response when the interconnect is basically transparent. This is a direct measurement of the capabilities of the instrument.




Figure 33. Measured insertion loss of a reference thru, a uniform line (DUT-1) and a uniform line that is part of a differential pair (DUT-2).

The second line from the top is the insertion loss of the eight inch, single ended microstrip we saw before. The third line is the measured insertion loss of another nine inch long, uniform microstrip transmission line. However, this transmission line has a very large dip in the insertion loss at about 6 GHz. This dip would dramatically limit the usable bandwidth of this interconnect. The -10 dB bandwidth of the first transmission line is about 12 GHz, while the -10 dB bandwidth of the second line is about 4 GHz. This is a reduction of a factor of three in usable bandwidth. Understanding the origin of this dip would be the first step in optimizing the design of the interconnect. What could cause this very large dip?

In this second transmission line, there are no vias. It is a uniform microstrip. The SMA launch is identical as in the first transmission line. It happens that, though this is a single ended measurement, there is another transmission line physically adjacent and parallel to this measured transmission line, with a spacing about equal to the line width. However, this adjacent line is also terminated with 50 Ohm resistors on its ends. Is it possible that somehow, the proximity of this other trace could cause this large dip? If so, what feature of this other line influences the dip frequency?

One way we can answer this question is by building a parameterized model for the physical structure of the two coupled lines, verifying its simulated insertion loss matches the measured insertion loss and then tweaking terms in the model to explore design space.

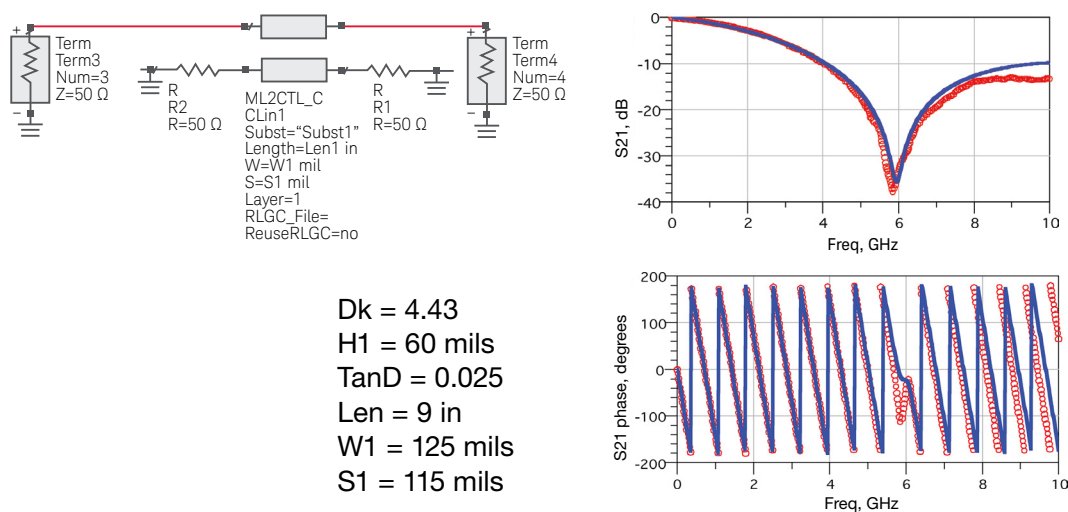


Figure 34. ADS model of the nine inch long trace, modeling the coupling to the adjacent, quiet line, showing the bandwidth of the model to be ~8 GHz.

Figure 34 shows the simple model of two coupled transmission lines using the MIL structures in ADS. All the physical and material properties are parameterized so that we can vary them later. We assume a simple model of two uniform, equal width lines, with a spacing, length, dielectric thickness, dielectric constant, and dissipation factor. We use all the geometry terms as measured with a micrometer from the structure and use the same dielectric constant and dissipation factor as measured from the uniform transmission line.

The integrated 2D field solver in ADS will automatically take these geometry values and calculate the complex impedance and transmission properties of the line and simulate the frequency domain insertion and return loss performance, configured exactly as in the actual measurement.

We bring the measured insertion loss data in Touchstone format from the TDR into ADS and compare the measured response and the simulated response. Shown in Figure 34 is the magnitude of the insertion loss in dB and the phase of the insertion loss. The red circles are the measured data, same as that displayed on the screen of the TDR instrument. The blue lines are the simulated response based on this simple model, with no parameter fitting.

The agreement is astonishingly good up to about 8 GHz. This says, there is nothing anomalous going on. There is nothing that is not expected from the normal behavior of two coupled, lossy lines. In this case, the second line, which is not being driven, is terminated by 50 Ohm resistors at the ends. The model is set up to match this same behavior. That we see this anomalous dip in the insertion loss in a single line when it is part of a pair of lines, but not when a similar line is isolated, and that we confirm this with the field solver, suggests there is something about the proximity of the adjacent line that causes this dip.

The effect that gives rise to this disastrous behavior is not anomalous, it is just very subtle. We could spend weeks spinning new boards to test for one effect after another, searching for the knob that influences this behavior. For example, we could vary the coupled length, the line width, the spacing, the dielectric thickness, even the dielectric constant, and dissipation factor, looking for what influences the resonance frequency. Or, we could perform these same experiments as virtual experiments using a simulation tool like ADS. It's only after we have confidence the tool accurately predicts this behavior that we can use it to explore design space.

Similar coupled stripline
 $w = 125$ mils, $s = 115$ mils




Figure 35. Changing separation between the two transmission lines showing the impact on the insertion loss dip.

One obvious virtual experiment to try is varying the line spacing. What happens to the resonant absorption dip in the insertion loss of one line as the traces are moved closer or farther apart? Figure 35 is the simulated insertion loss of one line in the simple two coupled line model, when we use separations of 50, 75, 100, 125, and 150 mils. The red circles are the measured insertion loss for the single ended trace. Each line is the simulated response of the insertion loss with a different separation. The trace with the lowest frequency resonance has a separation of 50 mils, followed by the 75 mil, and finally 150 mil.

As the separation distance increases, the resonance frequency increases. This is almost counterintuitive. Most resonance effects decrease in frequency as a dimension is increased. Yet, in this effect, the resonance frequency increases, as the dimension- the spacing- increases. If we did not have the confirmation of the very close agreement between the simulation and measurement in the previous figure, we would possibly begin to doubt the results from the simulation.

The explanation of the dip is clearly not a resonant effect. Its origin is very subtle, but is intimately related to far end cross talk. In the frequency domain, when the sine wave enters the front of the first line, it will couple into the second line. As it propagates along, there is a frequency where all of the energy couples from the first line into the adjacent line, leaving none in the first line, and hence, the large dip.

As we increase frequency more, the energy will couple back to the first line. This process will repeat. This is a fundamental property of modes and tightly coupled systems. It is ultimately related to the fact that the two modes, which propagate down the pair of lines, the odd and even modes, travel at different speeds in microstrip. If this were the true explanation, and if the two coupled lines were constructed in stripline, where the even and odd modes travel at the same speed, there would be no dip.

Also shown in Figure 35 is the simulated insertion loss of a single stripline transmission line, with the same line width, having an adjacent, terminated trace with a spacing of 115 mils. There is no dip at 6 GHz and the insertion loss is smoothly decreasing with frequency, all due to the dielectric loss of the laminate.

This suggests an important design rule: To get the absolute highest bandwidth in a single ended transmission line, you want to avoid having a closely spaced adjacent line, however terminated it may be.

3.0 2-Port TDR/Cross Talk

3.1 Overview

So far, we've evaluated the electrical performance of single transmission lines. When an adjacent transmission line is present, some of the energy from one line can couple into the second line, creating noise in the second line. To distinguish the two lines, we sometimes call the driven line the active line or the aggressor line. The second line is called the quiet line or victim line. This is illustrated in Figure 36.

One end of the active line is driven by the TDR stimulus. We get the TDR response of the active line for free. If we connect the second port to the far end of the active line, we can measure the TDT response. If we connect the second port to the end of the quiet line adjacent to the stimulus, we can measure the noise induced on the quiet line. To distinguish the two ends of the quiet line, we refer to the end near the stimulus as the near end, and the end far from the stimulus as the far end.

The ratio of the voltage noise measured on the near end of the quiet line to the incident stimulus voltage going into the active line is defined as the Near End Cross Talk (NEXT). The ratio of the far end voltage noise on the quiet line to the incident stimulus voltage going into the active line is defined as the Far End Cross Talk (FEXT). These two terms are figures of merit in describing the amount of cross talk between two parallel, uniform transmission lines. They can be measured directly by a 2-port TDR.




Figure 36. Configuration for two-port TDR measurements.

3.3 Measuring FEXT

By connecting the second channel of the TDR to the far end of the quiet line, the far end noise can be measured. At the same time, a 50 Ohm terminating resistor is added to the near end of the quiet line. The measured far end and near end noise on the quiet line is shown in Figure 38. Both are on the same scale of 20 mV/div. This corresponds to 10 percent cross talk per division. The white line is the near end noise, while the blue line is the measured far end noise.

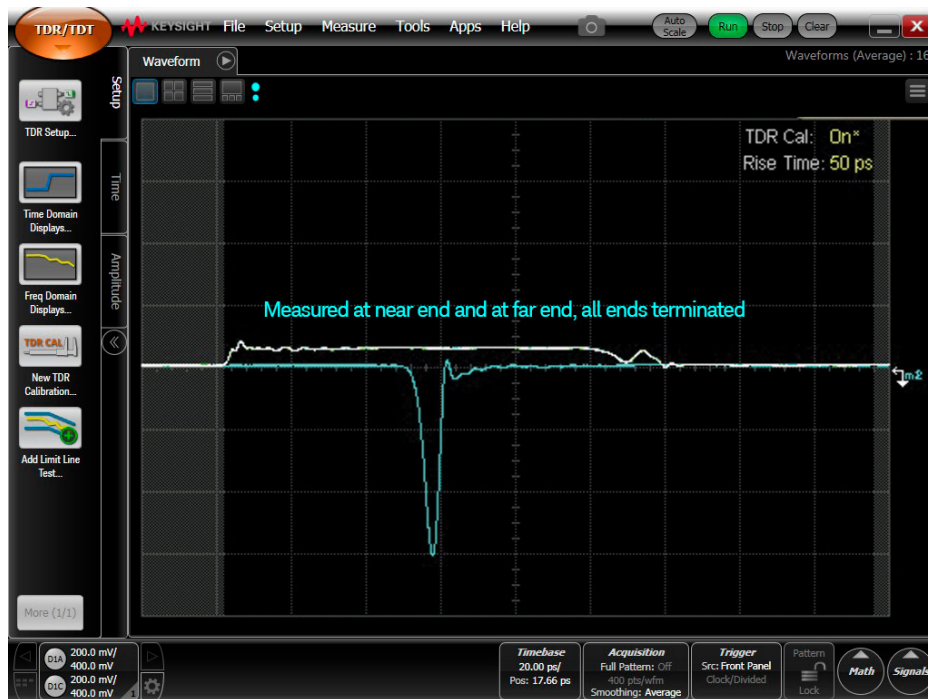


Figure 38. Measuring both the NEXT and FEXT with the second channel in the TDR.

In this example, while the NEXT is only 2.6 percent, the FEXT is seen to be about 30 percent, a huge amount. It appears to be coming out of the quiet line at a time equal to half the round trip time of flight, which is the one way time of flight, the time it takes the signal to propagate from the input to the output.

The width of the far end noise is the rise time of the signal. In fact, the shape of the far end noise is roughly the derivative of the rising edge of the signal.

These values of NEXT and FEXT are defined for the special case of all ends of the two lines terminated, so there are no reflections of the signals or noise. Normally, performing these measurements of NEXT and FEXT requires connecting and disconnecting the second port of the TDR to each of the two ends of the quiet line, while connecting a termination to the unused end. By taking advantage of reflections, we can actually perform both measurements of near and far end noise from one end only. We just need to understand that changing the terminations will change the noise voltages picked up.

3.4 Emulating FEXT for different system rise times

If we measure the noise voltage coming out of the near end, we will obviously pick up the near end voltage. However, if we keep the far end of the quiet line open terminated, the far end noise propagating down the quiet line to the far end will reflect at the open and head back to the near end, where it can be measured from the near end side of the quiet line.

In addition, if we keep the far end of the active line terminated in an open, the signal will also reflect. As it heads back down to the source end of the active line, it will be generating additional far end noise in the quiet line, but it will be heading in the direction back to the near end. This will increase the amount of far end noise picked up in the quiet line, at the near end.




Figure 39. Emulating the FEXT with different system rise time responses with RT = 100, 200, 500 ps, and 1 ns.

In Figure 39 is the measured noise at the near end of the quiet line, showing the initial near end noise, followed, one round trip time of flight later, with the reflected far end noise.

The signature of the far end noise is the derivative of the rising edge of the signal. This means that as the rise time changes, the peak value of the far end noise will change. However, the area under the curve of the far end noise will always be constant, as long as the signal voltage transition levels are the same.

In Figure 39 is an example of the combination of near and far end noise, measured at the near end, as we change the rise time of the signal. We see two important features. The magnitude of the near end noise is independent of the rise time. Second, the peak value of the far end noise decreases as the rise time increases, but the area under the far end noise voltage is constant. As the rise time decreases, the peak decreases, but the far end noise spreads out, since the derivative of the rising edge spreads out with longer rise time.

Understanding the impact of terminations on the measured noise at one end or the other of a quiet line can often help resolve the interpretation of the noise picked up.

3.5 Identifying design features that contribute to NEXT

For this example, we will look at the measured noise between two parallel interconnects in a mother board with two plug in cards. In this case, the lines are stripline, but the dielectric distribution is not uniform, so there will be some far end noise. The total parallel run length is about 24 inches, including the path on the motherboard.

The TDR response of one of the single ended lines is shown in Figure 40. The scale is 10 Ohms/div. We see the very first peak at the SMA launch. The constant impedance region, with an impedance value we can read off the first, solid marker as 56.8 Ohms, is the interconnect on the daughter card.

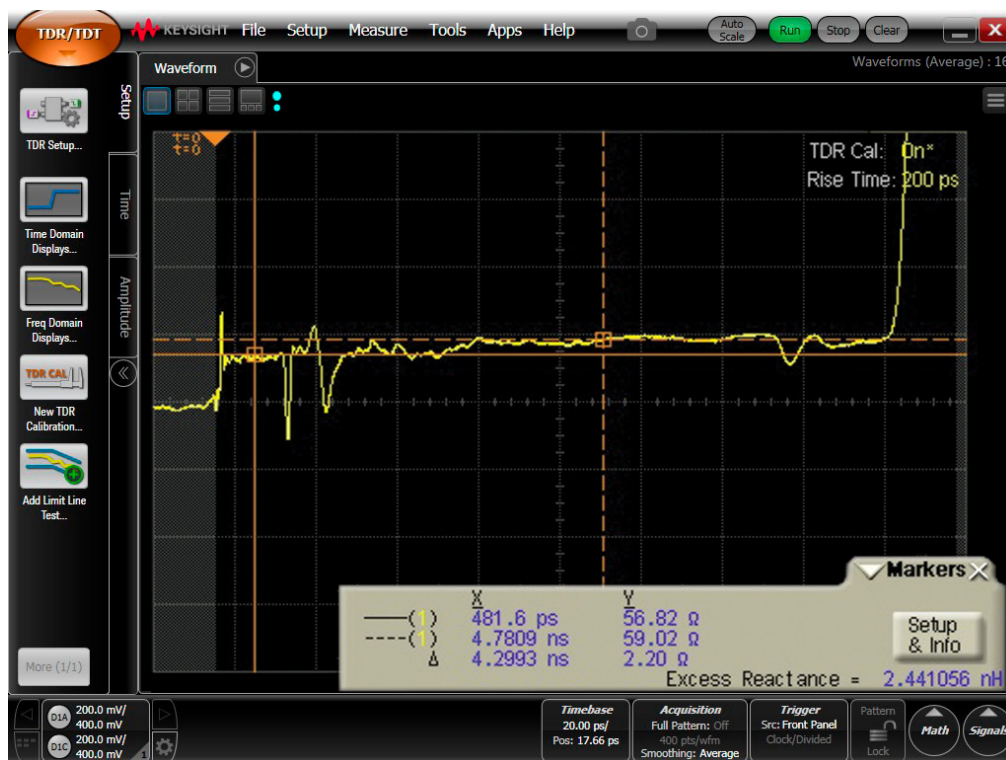


Figure 40. Measured TDR response of 24 inch long trace in a mother board using markers to measure the impedance in the daughter card and mother board.

The first dip is the via field on the daughter card where the connector is. This short uniform region is inside the connector, followed by the second dip of the via field in the motherboard. The long region of uniform impedance, with a value read from the second, dashed marker, about 59 Ohms, is the interconnect in the motherboard. At the end of the motherboard trace is the second daughter card. Because of the lossy interconnects, the initial rise time of the TDR stimulus has increased by the time it gets to the far end of the active line, and the spatial resolution has decreased.

This incident signal from the TDR will be used as the active signal, while we measure the near and far end noise on the adjacent quiet line.

The near end noise on the quiet line is shown in Figure 41 as the blue trace, on a scale of 10 mV/div. This is with a signal magnitude of 200 mV, so the scale is really 5 percent per division. We see that in this case, the near end noise has a peak of about 11 percent. With a typical cross talk noise allocation in the noise budget of about 5 percent, the 11 percent near end noise can be considered a lot. In order to even consider how to reduce it, the first step is to identify where it is coming from.




Figure 41. Measured NEXT and FEXT in a 24 inch long trace on a mother board with all ends terminated.

By comparing the location of the near end noise with the TDR response, we can quickly identify where in the interconnect path this noise is created. As the active signal propagates down the active line, the TDR response picks up reflections from impedance changes. The time at which these reflections are picked up at the near end of the active line is the round trip time from the source to the discontinuity.

Likewise, the time value when we pick up the near end noise at the near end of the quiet line corresponds to the time it takes for the signal to hit that region of the interconnect, plus the time it takes for the generated noise to propagate back to the near end of the quiet line. This is the round trip time of flight. This means that by comparing the time response of the near end noise to the TDR response, we can identify which specific interconnect features in the active line might have generated the near end noise.

Looking at the blue line in Figure 41, and comparing it to the yellow line of the TDR response, we can see that the near end noise in the daughter card trace is small, only about four percent. The very large peak corresponds to the connector between the daughter card and mother board. We see the double peak corresponding to the via field in the two boards, and a large contribution from the connector itself.

The near end noise from the traces in the motherboard is also high, on the order of five percent, followed by a peak in near end noise from the connector on the other end of the interconnect. This says, to minimize the near end noise in this motherboard application, the place to look is in the design or selection of the connector. For single ended signal applications, we would want to use a connector with lower coupling than this particular one. The coupling between signal lines in the motherboard interconnects is high, but probably would meet a typical noise budget.

Also shown, as the green line in Figure 41, is the measured far end noise on the quiet line. We see on this scale, the far end noise is only about four percent. Even though the interconnects are stripline, any inhomogeneities in the dielectric distribution will generate far end noise. This is low enough to not really cause a problem.

3.6 Exploring the impact of terminations on NEXT and FEXT

Given these noise levels on the quiet line when all the ends of the lines are terminated, we can also evaluate what would be the impact if the ends were not terminated, as would be the case if the receivers were tri-stated. Because of the high impedance, any voltages hitting these open ends will reflect. There is one situation in particular that can cause significant problems.

If the far end of the active line is tri-stated and open, and the far end of the quiet line is a receiver, while the near end of the quiet line is tri-stated and open, as illustrated in Figure 42, the noise picked up at the far end of the quiet line can exceed 15 percent, a very large amount. In this configuration, the signal propagating down the active line will generate near end noise in the quiet line, which will propagate back to the source end of the quiet line. But, if this end of the quiet line is tri-state and open, the near end noise will be reflected and head back to the far end of the quiet line.

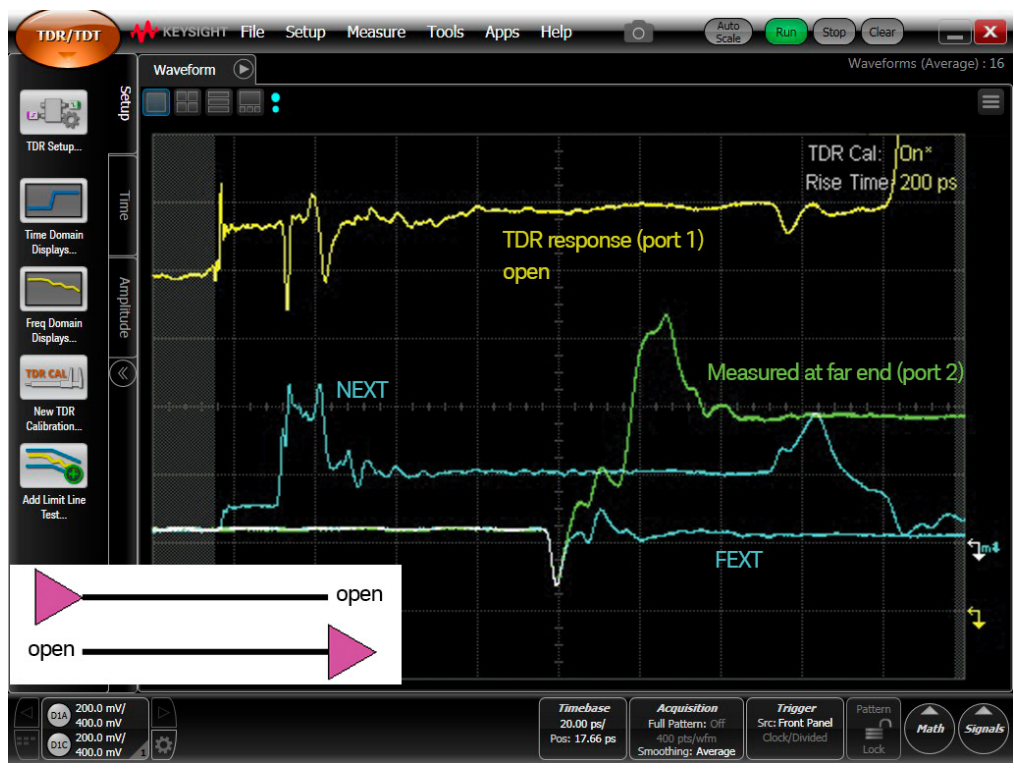


Figure 42. Measured cross talk in quiet line with worst case termination.

The front edge of the reflected, near end noise will be coincident with the front edge of the signal heading down to the receiver on the active line. When this signal hits the far end of the active line and the receiver there is still tri-stated, the signal will see an open and reflect. This reflected signal will be heading back to the driver and generate in the quiet line, near end noise that is heading to the far end receiver.

The near end noise generated on the quiet line that reflected, hits the far end of the quiet line at the same time the active signal hits the far end and reflects, generating another round of near end noise. This means there will be almost twice the near end noise picked up by the receiver on the far end of the quiet line.

In Figure 42, the green line is the measured noise picked up on the far end of the quiet line when the far end of the active line is open and the near end of the quiet line is open. This is a noise value with a peak of 15 percent. It is not the 20 percent double of the NEXT, mostly because of the smearing out of the near end noise across the short length of the connector, compared to the rise time of the signal.

This example points out that the actual measured noise is related to not just the coupling between the lines, but also to how the ends of the lines are terminated. When measuring cross talk, care should be taken to consider the termination configuration for worst case coupling.

3.7 Measuring ground bounce

Cross talk is created by capacitive and inductive coupling. In uniform, adjacent transmission lines, the capacitive and inductive coupling is uniformly distributed along their length. In addition, coupling can be localized. In particular, any discontinuities in the return path under signal conductors, can dramatically increase the inductive coupling. We call the noise generated by inductive coupling, switching noise. A special case of switching noise is when the return currents of multiple signal paths share the same return paths. We call the noise generated in this case, ground bounce.

Ground bounce occurs mostly in connectors and packages, where multiple signal lines typically share the same return pin. Ground bounce can also occur in board level interconnects if there are any discontinuities in the return path that force return currents to overlap. This happens if there is a gap in the return plane and return currents are confined to narrow paths.

Figure 43 is an example of two gaps in the return path underneath two coupled microstrip lines in FR4. Each transmission line is about 50 Ohms, and the spacing is about equal to the line width. The dark green color adjacent to the copper traces is the circuit board color when there is a solid ground plane underneath. In the two regions marked by arrows, the copper plane has been removed, making the area a light green in color.




Figure 43. Tightly coupled pair of transmission lines with small gaps in the return path that will generate ground bounce.

In this configuration, the near end noise is about 2.5 percent, an acceptable level in most applications. However, where the gap is under the traces, the return currents will have to meander around the gap. This will increase the loop inductance in this section of the transmission line and increase the mutual inductance between the two lines. The noise generated on the quiet line due to the higher mutual inductance in this region is called ground bounce and can be measured at the near end of the quiet line.

One of the lines in the pair is used as the aggressor. The TDR response of this line is shown in Figure 44. We see an initial peak from the SMA launch, another reflection peak from the first gap, a uniform region and then another peak from the second gap, a uniform region and then the open at the far end.



Figure 44. TDR of a single ended transmission line crossing gaps in the return path, showing the inductive discontinuities.

3.8 Identifying design features that contribute to ground bounce

While the signal is propagating down the active line, noise is being generated on the quiet line from the capacitive and inductive coupling. In Figure 45, we show the measured noise on the near end of the quiet line, with the far end of the quiet line open terminated. The near end noise is on a scale of 20 mV/div or 10 percent noise per division.




Figure 45. Measured ground bounce on the quiet line from gaps in the return path.

Initially, the small near end noise of roughly 2.5 percent is from the uniform section of the transmission lines. The first peak in the near end noise of roughly 11 percent is due to the small gap in the return path. This is a direct measure of the ground bounce voltage generated across the gap that is picked up in the adjacent, quiet line.

On the other side of this gap is the 2.5 percent near end noise from the uniform region, followed by another 11 percent of near end noise from the second gap. Then comes the near end noise from the uniform section and finally, we see the reflected far end noise at the end of the quiet line. The ground bounce also contributes to an increase in the far end noise.

3.9 Emulating ground bounce noise for different system rise times

All switching noise is generated because of a switching current through some mutual inductance. The peak value of the switching noise is related to the mutual inductance times the di/dt . This magnitude will depend on the rise time of the signal. If we can slow the edge down, we can reduce the magnitude of the switching noise. If the rise time can be increased enough using slew rate control in the driver and not impact timing, the switching noise might be reduced below a problem level.

Using the rise time control feature of the DCA 86100C, we can change the rise time of the stimulus and measure the resulting ground bounce when the rise time is longer. Figure 46 shows the TDR and near end cross talk response for these two coupled lines with the two gaps in the return path. In blue is the response we saw previously, with a rise time of 100 ps. In the yellow is the TDR response for a rise time of 500 ps, while in green, is the noise measured on the near end with a rise time of 500 ps.




Figure 46. Emulating impact of rise time on the ground bounce noise in a pair of coupled lines with a rise time of 500 ps.

We see that in each case, the magnitude of the noise peak has been dramatically reduced. The ground bounce has been spread out over a larger area, to a level that could be perfectly acceptable. The far end noise has been significantly reduced by the increase in system rise time, but is still large, approximately 15 percent. However, we are measuring the far end noise from the near end side. In this configuration, the magnitude of the measured far end noise is actually twice what would appear at the far end if there were a receiver present, so in fact, the far end noise may also be acceptable with a 500 ps rise time. If we know the final application system's rise time, we could emulate the system's signal and empirically determine if the switching noise generated in the interconnect was acceptable or if it had to be reduced.

In the previous example, the gap in the return path was very slight, and the ground bounce noise generated was small. We could probably find this level of ground bounce acceptable. But often, the gap in the return path is large. In the next example, the gap has been increased to be a large, wide slot. Figure 47 shows the top view of a pair of 50 Ohm microstrip transmission lines with a solid plane as the dark tan colored region. In the middle of the board, the copper return plane has been removed in a region about an inch long and an inch wide. The region with no copper plane is a yellow color.

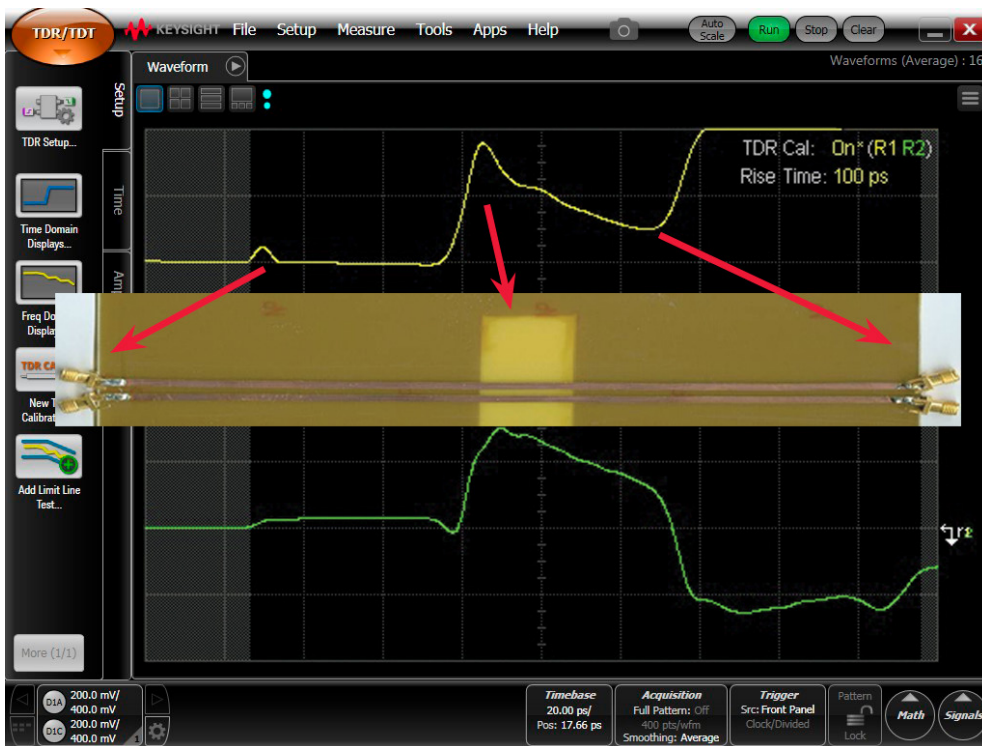


Figure 47. Measured TDR response of a single line crossing a large gap in the return path and the ground bounce noise in the quiet line.

The TDR stimulus launches into one of the lines and the second channel is used to measure the noise on the near end of the trace. The two far ends of the transmission lines are left terminated in an open. As the TDR signal propagates down the line, it is sensitive to impedance discontinuities. The measured TDR response, the yellow trace on a 50 mV/div scale in Figure 47, shows the small inductive peak at the beginning of the line from the SMA launch. The very large peak midway down the line is from the gap in the return path. The gap dramatically increases the loop inductance of the signal path, as the return current must make a large detour around the gap to reach the source. This extra path length increases the series loop inductance of the signal path. Just before the reflected voltage settles down, the open end of the line is reached.

At the same time the TDR stimulus is propagating down the active line, the second channel is measuring the near end noise on the quiet line, the green trace, also on a 50 mV/div scale. We see the very slight near end noise initially due to the tight coupling between the uniform transmission line segments. On this scale it is barely at the detectable level. However, as soon as the TDR signal hits the inductive discontinuity and generates the ground bounce voltage across the two regions of the circuit board, this voltage is picked up in the quiet line. In fact, we see that the ground bounce voltage in the quiet line is just about the same magnitude as the reflected voltage in the active line. All the reflected voltage was really ground bounce voltage, shared by the quiet line.

This amount of noise in the quiet line, about 75 mV out of 200 mV, or 37 percent of the incident signal, is far higher than any reasonable noise budget and would be a disaster. In fact, every trace in a bus that shared this return path, meandering around the gap, would see the same ground bounce. The more lines that switched simultaneously, the more di/dt ground bounce would be generated, and the larger would be the switching noise on the quiet lines.

One way to identify switching noise is to look for narrow, isolated regions where the near end noise dramatically increases. The TDR response of the active line can be used to guide us to the physical location where the near end noise is being generated.

While increasing the rise time will decrease the magnitude of the switching noise, sometimes, it can still be too large. Figure 48 is an example of comparing a 100 ps rise time and a 1 ns rise time as the system rise time for this large gap.

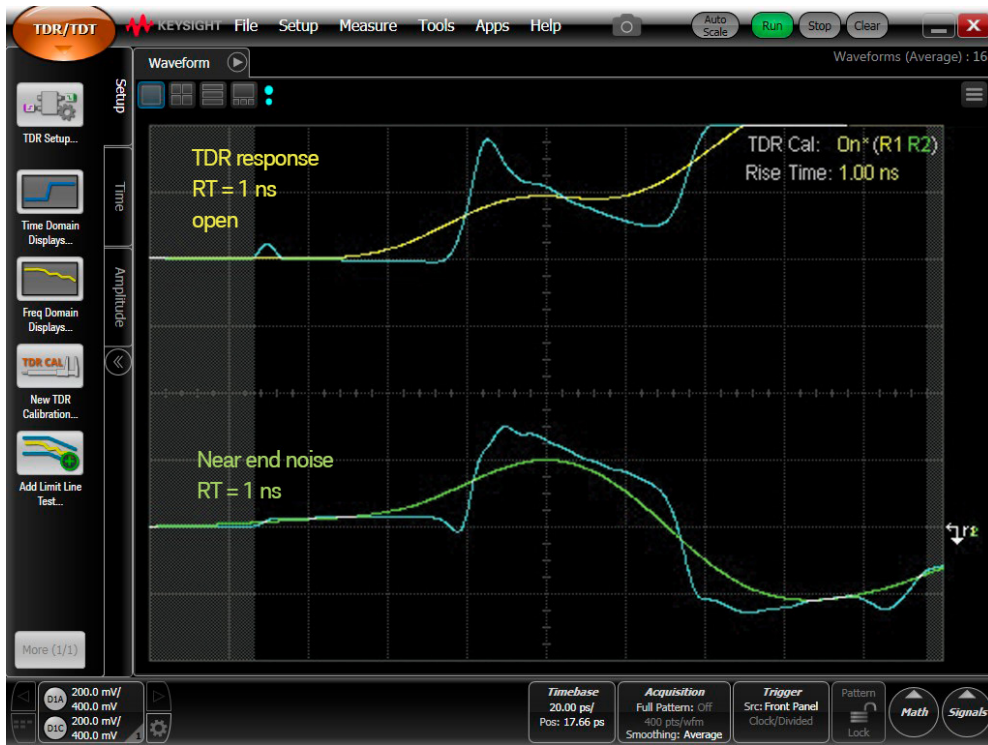


Figure 48. Emulating ground bounce noise from large gap at rise times of 100 ps and 1 ns.

Though the ground bounce voltage does decrease a little when we increase the rise time by an order of magnitude, the amount of ground bounce is still too large. This is an example of using the TDR to emulate the system rise time and evaluate the impact of a discontinuity on the amount of cross talk generated. This suggests that no amount of slew rate control would have a hope of getting around this problem. Instead, it would be necessary to identify the source of the ground bounce and either remove the gap or route the signals around the gap, rather than letting them cross it.

4.0 2-Port Differential TDR (DTDR)

4.1 Overview

Previously, we explored two single ended lines with coupling. Each line had its properties of an impedance profile and time delay and there was near and far end noise on one line from the signal on the other line. This is one way of describing these two individual lines.

An equivalent way of describing these same two lines is as a single differential pair. Two types of signals can propagate on a differential pair: a differential signal and a common signal. In a differential signal, the voltage on one line is the negative of the other. The differential signal component on a differential pair is the difference in voltage between the two lines. This means in a differential signal, the voltage on one line, measured with respect to the return plane, is the negative of the other. Most high speed serial links use a differential signal to transmit information. Because of the nature of the receivers, a differential signal can have much better signal to noise ratio and noise immunity than a single ended signal.

The common signal component is the average of the two signals on each line of a differential pair. This means a common signal is really a measure of how much voltage a pair of lines have in common. While a common signal is rarely used to carry information, it can sometimes cause complications if it becomes so large as to saturate the differential receivers, or if it were to get out of the product on external cables, as it would contribute to EMI.

When a differential signal propagates on an interconnect, it drives the odd mode of the differential pair and the differential signal sees the differential impedance of the interconnect. When the common signal propagates on an interconnect, it drives the even mode of the differential pair and the common signal sees the common impedance of the differential pair.

To characterize a differential pair, the TDR must drive either a differential signal or a common signal, and measure the response as the reflected differential signal or common signal. This requires two channels to be connected to the same end of the diff pair, and have the equivalent of two, simultaneous stimulus- either launching a differential signal or launching a common signal into the device under test. This is done with a differential TDR (DTDR).

When set for differential stimulus, as shown in Figure 49, the stimulus from the two channels is exactly opposite, while, when it is set for common stimulus, the output voltages are exactly the same.




Figure 49. Configuration for differential pair characterization.

In application, the DTDR is set up for one operating mode or the other. To adjust the DTDR for the differential stimulus operating mode, the TDR setup window is opened by clicking TDR setup and then select differential for TDR stimulus mode.

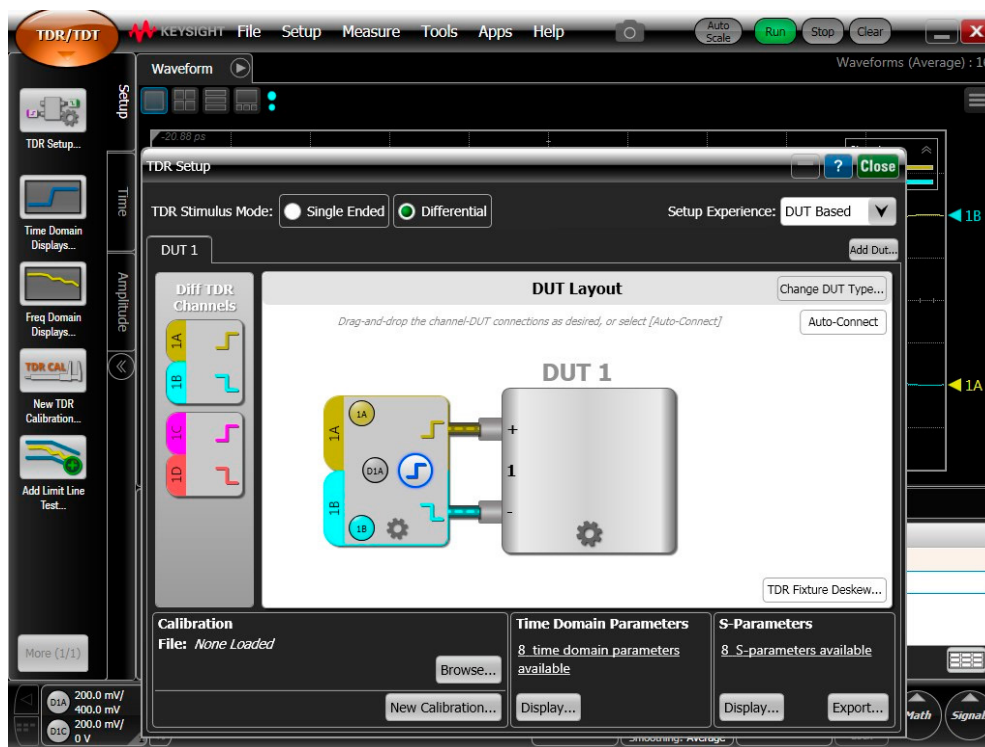


Figure 50. DTDR set up screen for differential measurements.

Figure 50 shows the set up screen when the operating mode is adjusted for differential operating mode. As a side note, though it is sometimes confusing, don't mix up the common mode of operation with the even mode in which the differential pair can be driven. The "mode" in the screen label with common refers to the "mode of operation," not, a mode in which the differential pair is driven.

4.2 Measuring each of the five impedances associated with a differential pair

If we have a single ended transmission line that is part of a differential pair, it really has three different impedances that characterize it. It has a single ended impedance, its instantaneous impedance when the other line in the pair has a constant voltage on it; an odd mode impedance, the instantaneous impedance of the line when the pair is driven in the odd mode; and an even mode impedance, the instantaneous impedance of the line when the pair is driven in the even mode.

In Figure 51 is an example of the measured TDR response from a single line in a differential pair. The lines are nine inches long, roughly 50 Ohm microstrip traces, with a spacing about equal to their line width. On this scale of 20 mV/div, one division corresponds to a reflected voltage of 10 percent.

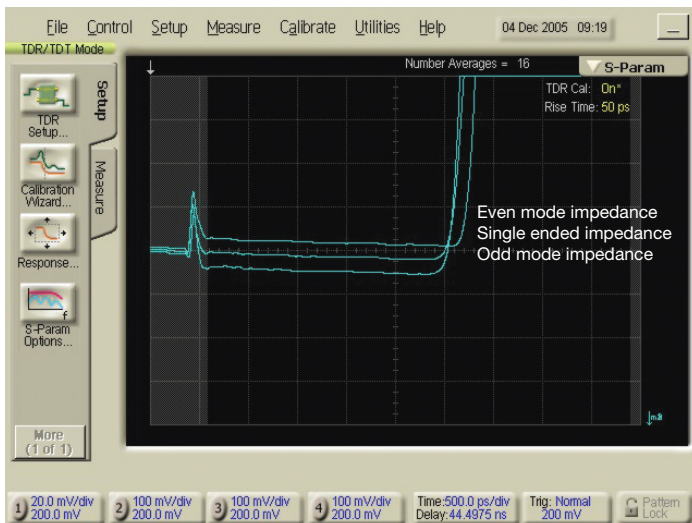


Figure 51. Measured TDR response of a single transmission line configured for the even mode.

When the DTDR is set up as single ended, the response is the single ended impedance. When it is set up as differential, the TDR response from each channel is the odd mode impedance of the line, and when the DTDR is set up in the common mode of operation, the TDR response from either channel is the even mode impedance of the line. While we could take the measured reflected voltages and calculate the corresponding impedances, it is much easier to let the TDR do it for us.

The vertical scale can be changed to Ohms so that the first order impedance is directly displayed. Figure 52 is the same measured response as previously, but the reflected voltage has been converted into the instantaneous impedance. The scale has been expanded to 2 Ohms/div with 50 Ohms right at the center.

We can now read the three impedances of this line directly off the screen. In each case, the impedances start high on one end and drop about 1 to 2 Ohms by the other end. The even mode drops 2 Ohms, while the odd mode impedance changes by only 1 Ohm. This suggests it is probably a dielectric thickness variation that causes the small change in impedance across the length of the board, as the even mode impedance is more sensitive to dielectric thickness than the odd mode impedance.




Figure 52. The three impedances of a single line displayed directly on an impedance scale.

Everything we ever wanted to know about the impedance properties of the differential pair is contained in these three impedance values of each line.




Figure 53. Measured odd mode impedance of each line in a differential pair, displayed directly on an impedance scale.

In DTDR measurements on a differential pair, the odd mode impedances of both lines are measured simultaneously. Previously, we displayed the impedance of just one line. Figure 53 is the odd mode response of both lines on a 2 Ohms per division scale. This is the individual response from each channel. In this case, we see the differential impedances of the two lines are matched to within a small fraction of an Ohm.




Figure 54. Measured even mode impedance of each line in a differential pair, displayed directly on an impedance scale.

Likewise, when the stimulus is set for the common mode of operation, the even mode of each line can be measured as the responses from the two channels. Again, we see from Figure 54, that for this differential pair, the even mode impedances are matched to within a small fraction of an Ohm.

The odd and even mode impedances are only part of the story. Though each line may have an odd mode impedance when a differential signal propagates down the differential pair, the differential signal itself sees a differential impedance. It is numerically equal to the sum of the odd mode impedances of both lines. When the odd mode impedances of the two lines are the same, the differential impedance of the pair is just twice the odd mode impedance of either line.




Figure 55. Measured differential impedance of a pair of microstrip traces, displayed directly on an impedance scale.

The DTDR can simply and easily display the differential impedance profile of the pair of lines. With the stimulus set for differential mode of operation, the differential impedance is selected in the Response 2 setting. As shown in Figure 55, the differential impedance profile can be plotted directly from the screen. In this case, it is on a five Ohms per division scale with 100 Ohms at the very center. The marker can be used to read the differential impedance as about 91 Ohms.




Figure 56. Measured common impedance of a pair of microstrip traces, displayed directly on an impedance scale.

In the same way, the common impedance profile can be displayed directly on the screen. The stimulus mode is set for common mode of operation and the common impedance is selected for Response 2. Figure 56 shows the common impedance profile on a scale of 2.00 Ω /div. The marker is set to read the common impedance off the front screen as about 25.5 Ohms. In this way, we can extract the complete impedance profile characterization of either line or both lines in a differential pair.

4.3 Measuring the degree of coupling between lines in a differential pair

The differential impedance of a pair of lines is twice the odd mode impedance. When there is very little coupling, the single ended impedance of one line is the same as the odd mode impedance of that line, and the differential impedance is really twice the single ended impedance of the line.

However, if there is any coupling, the single ended impedance is not the same as the odd mode impedance. The odd mode impedance of that line will be reduced by the coupling. We cannot easily and accurately measure the odd mode impedance or differential impedance of a pair of lines unless we drive the pair in the odd mode with a differential signal.

The difference between the odd mode impedance and the single ended impedance for a typical trace on a motherboard is shown in Figure 57. The blue line is the single ended TDR response of one line on the motherboard. On the scale of 5 Ohms/div, the single ended impedance of the trace on the daughterboard and motherboard is seen to be about 58 Ohms.

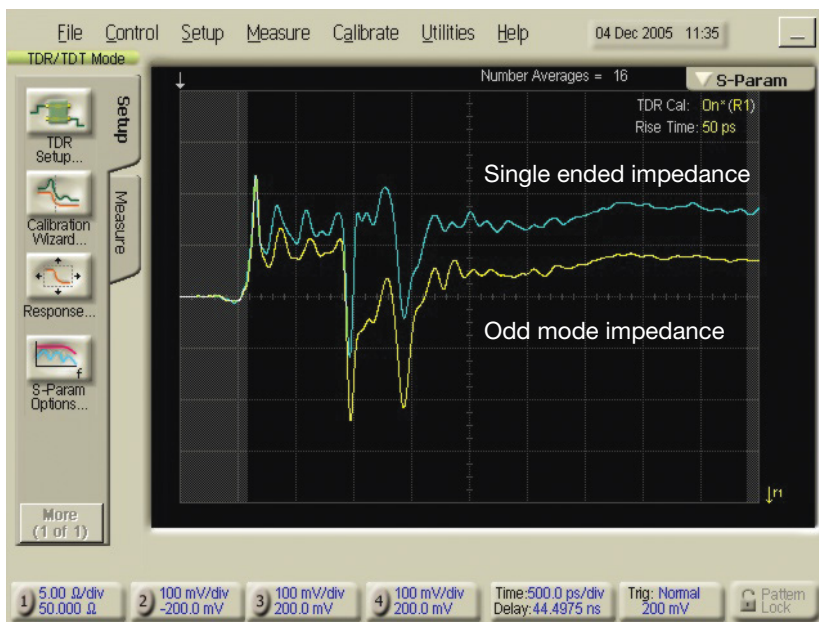


Figure 57. Comparison of the measured single ended impedance and odd mode impedance of a single line in a long motherboard trace.

When the pair is driven with a differential signal, the odd mode impedance of the same line is seen to have dropped in some cases by as much as 5 Ohms. The daughter card trace is still a little high, at about 55 Ohms, while the connector is seen to be very close to 50 Ohms, though still with large capacitive dips from the via field. The long line on the motherboard is about 53 Ohms.

If the odd mode impedance is 53 Ohms, the differential impedance of the pair would be 106 Ohms; close to the target value of 100 Ohms.

4.4 Measuring the differential impedance of a twisted pair cable

Many applications require a differential signal to be transported from one board to another through twisted pair cables. The DTDR can be used to measure not only the differential impedance of a differential pair on a circuit board, but also the differential impedance of a twisted pair cable.

In this case, there is no return plane in proximity. However, as long as the coupling between the two lines in the pair is much tighter than either line to an adjacent plane, the return currents for each line, when driven with a differential signal, will exactly overlap in the adjacent plane and the presence of the plane will be irrelevant. In a twisted pair, any plane, or literally the ground that could act as a return path, carries no current, and will not play a role in determining the differential impedance of the pair, or in a measurement of the differential impedance.

To measure the differential impedance of the twisted pair, we have to connect each of the lines in the twisted pair to the signal lines in the cable. This establishes a 100 Ohm launch into the twisted pair. In this example, we look at the measured differential impedance of two different types of twisted pair cable. The first case is a two foot length of twisted pair cable taken from a low cost, POTS telephone cable. The second case is a two foot length of twisted pair, taken from a CatV ethernet cable.

In Figure 58, the DTDR response from these two different twisted pair are shown. The blue is the DTDR response from a twisted pair of wires as found in a low cost telephone hook up cable. On this scale of 20 Ohms per division, the differential impedance of the cable can be seen to be relatively constant, but on the order of 125 Ohms. This impedance is related to the precise wire diameter and dielectric thickness of the insulation. This cable is typically specified for 120 Ohms, and is not rated for high bit rate. As we can see, it is a relatively controlled impedance.

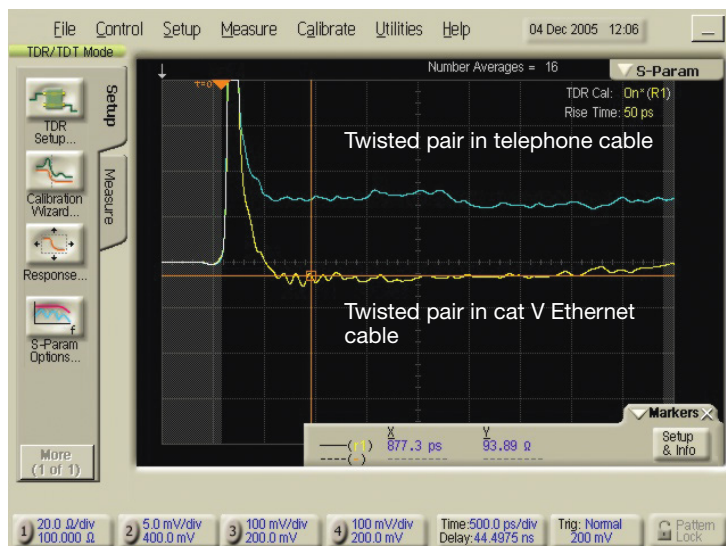


Figure 58. Measured differential impedance of two different twisted pair cables connected to a coax launch.

The yellow line is the measured differential impedance of the CatV twisted pair. It is specified for 100 Ohms, and we can read its impedance with the marker as 94 Ohms. It is also seen to be a very constant, controlled impedance.

The large peak at the beginning of the DTDR response is due to the poor launch into the twisted pair. In both cases, the wires were pulled apart in order to solder two, separate SMA connectors which connected to the coax cables from the DTDR. Part of the twisted pair connector design is optimizing this launch to minimize the discontinuity.

In this example, we are measuring the reflected differential signal. Once it gets through the connector, the differential impedance of the twisted pair cable is very close to the differential impedance of the two coax cables. What about the common impedance? While the signal is in the two coax cables, the common impedance is half the even mode impedance of either cable, which is about 25 Ohms. What is the common impedance of the twisted pair?

The common impedance is the impedance between the two signal lines, with respect to the return path, which in the case of a twisted pair, is literally the floor. As we might imagine, when the return path is far away, this common impedance can be pretty high, easily a few hundred Ohms. To the TDR, it will look like an open.

In addition to measuring the differential impedance profile of the twisted pair, we can also measure the common impedance profile as the common signal travels from the coax cable to the twisted pair. In Figure 59, the measured differential impedance profile on a scale of 25 Ohms/div is shown as the blue trace.

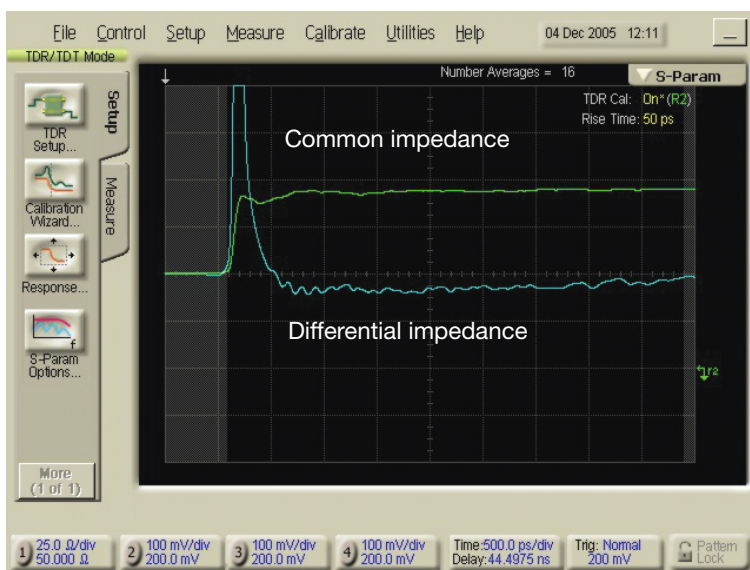


Figure 59. Measured reflected common signal from a coax to twisted pair transition with an incident common signal.

The DTDR was set up to use a common signal as the stimulus and then the reflected common signal is measured. The impedance was so high, we changed the scale format to voltage scale, and recorded the reflected voltage of the common signal on a scale of 100 mV/div. The incident common signal is 200 mV. We can see in this plot that the reflected common signal is almost 200 mV. In the transitions from the coax cable to the twisted pair, other than the discontinuity of the connector, the differential signal is able to transition to the twisted pair and propagate down the twisted pair, but, virtually all of the common signal is reflected, due to the very high common impedance of the twisted pair. The little bit of common signal that does get out on the cable will contribute to radiated emissions. This is why it is important in the design of twisted pair connections, to make the impedance the common signal sees as absolutely high as possible, so there is little common signal on the external cable to radiate.

4.5 Measuring the reflected noise of a differential signal crossing a gap

When a single ended signal encounters a large gap in the return path, it will see a large inductive discontinuity and generate ground bounce in the plane which will be picked up by any adjacent signal traces. Crossing a gap in a single ended transmission line can be a disaster.

However, the same gap can be crossed with less of a problem by a differential signal. Figure 60 shows an example of a microstrip differential pair crossing a gap in the return path. The gap is the light yellow color in the board where the copper plane has been removed. In blue is the single ended TDR response for the signal on one of the lines crossing the gap. There is a huge reflected signal, which affects the reflected signal for the duration of the time delay down the line.

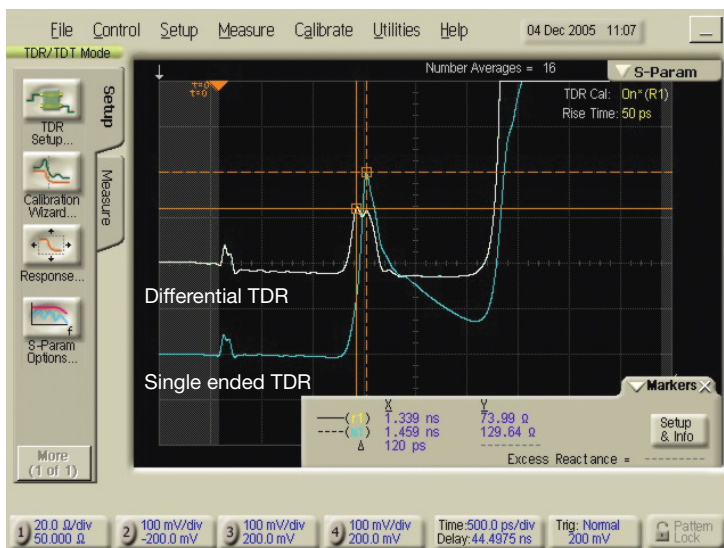


Figure 60. Measured differential impedance profile of a differential pair crossing a wide gap in the return path.

The white trace is the DTDR response for a differential signal on this same pair of traces, crossing the large gap. The differential signal sees a differential impedance of about 100 Ohms in the region where the plane is continuous. In the region where the plane is removed, the differential impedance of the pair is about 130 Ohms, as read by the dotted marker. Where the plane is removed, the differential impedance is uniform, it is just high. This 130 Ohm discontinuity lasts for the time of flight of the gap, and then the differential impedance the signal sees comes back down to roughly 100 Ohms.

Like all discontinuities, if we keep the length of the discontinuity short compared to the rise time of the signal, the impact of the discontinuity can be reduced.

The impact of this gap on the system's rise time can be emulated by changing the DTDR rise time. For this same differential pair with the one inch long gap, the DTDR response was measured for four different rise times of 100, 200, and 500 ps and 1 ns. Figure 61 shows the DTDR response of this discontinuity for these rise times, on a scale of 10 Ohms/div.




Figure 61. Emulating the differential impedance profile of a differential signal crossing a large gap at four different rise times.

The longer the rise time, the lower the effective impedance the gap appears as. When the rise time is 1 ns, the impact of the gap has almost disappeared. This suggests an important design rule: if a signal must cross a gap, keep the length of the gap short, the rise time of the signal as long as possible, and use a tightly coupled differential pair to cross the gap.

Using the DTDR, adjusted to the system rise time, would allow a quick and simple evaluation of the impact on the signal's reflected noise from this gap.

4.6 Measuring the mode conversion in a differential pair

In addition to the impedance profiles outlined so far, there is another problem a DTDR can assist in debugging. When a differential signal enters a differential pair, some of the differential signal can reflect back to the source, due to discontinuities in the differential impedance of the interconnect. Of course, these reflected differential voltages are detected by the receiver in the TDR and we use this received differential voltage to extract information about the differential impedance profile of the interconnect. Under some situations, not only will differential signals reflect and head back to the source, but the incident differential signal can be converted into a common signal and head back to the source as well.

The generation of a common signal, sent back to the source, when a pure differential signal is incident, is called mode conversion. In mode conversion, some of the differential signal is converted into common signal. The presence of the converted common signal is only an issue if the common signal level is large enough to saturate a receiver or if any of the common signal gets out on an external twisted pair, where it can radiate and cause the product to fail FCC EMC certification testing.

Otherwise, it is not the converted common signal that causes a problem, but the distortion of the differential signal, because of the mode conversion. After some of the differential signal is converted into a common signal, what's left of the differential signal will have a distorted rise time which can cause intersymbol interference, deterministic jitter, and collapse of the eye diagram. All are factors which will limit the maximum bit rate through the interconnect.

4.7 Identifying specific physical features that contribute to mode conversion in a differential pair

Identifying the physical sources that cause mode conversion will be the first step in eliminating them and enabling higher bit rates and lower radiated emissions. The fundamental cause of mode conversion is an asymmetry between either the individual signal launches into each line of the differential pair or an asymmetry between the two lines that make up the pair. When the differential signal source is the DTDR stimulus, the asymmetries in the signals can be reduced to below -40 dB. Using the Keysight DCA 86100C DTDR, we are mostly sensitive to asymmetries in the interconnect.

By using the timing information of when the converted common signal returns, we can identify physically where, down the line, the asymmetry might be located. As an example, Figure 62 shows in yellow, the differential TDR response from a symmetric microstrip differential pair, with a differential impedance of about 90 Ohms. We see a small peak at the beginning of the measurement, corresponding to the SMA launch into the differential pair, and a small dip about midway down the line and then the open of the line. In this example, the signal was a $+200$ mV and -200 mV differential voltage launched into the differential pair.




Figure 62. Measured mode conversion from differential to common signal due to an asymmetry on one line in a pair.

The dip in the middle was caused by adding a small capacitive load to one of the lines. This caused the differential impedance of the pair to decrease a small amount and reflect some of the differential signal. In DTDR operation, the receivers are sensitive to the reflected differential signal. In addition, while the stimulus is set to the differential mode of operation, we can adjust the receivers to measure the common signal, by selecting the Response 2, and set it for the common mode of operation, so it measures the common signal which reflects back.

The detected, received common signal voltage is displayed in Figure 62 as the green line. If there were no common signal reflecting back, the green trace would have registered as zero common voltage. The scale for the common signal is 5 mV per division.

Instead, we see that near the beginning of the line, there is a small common signal generated, where the SMA launches are, and very little common signal except in the middle, where the asymmetric capacitive load is. Finally, we see an additional common signal detected, coincident with when the signal has hit the end of the line and reflected back. This last peak in the common signal is the common signal generated by the asymmetry, moving in the forward direction that hit the end of the differential pair, where the common signal saw an open and reflected back to the source.

The sign of the reflected converted common signal depends on whether the discontinuity occurred on the + or the - line of the pair. If we move the asymmetry to the other line, we change the sign of the converted common signal.

In Figure 63 is the measured common signal at the receivers, for the same capacitive discontinuity, first on line 1 and then taken off line 1 and placed on line 2. We see that the time at which the converted common signal is detected is the same which means the physical location of the discontinuity is the same.




Figure 63. Measured mode conversion on a differential pair when the capacitive asymmetry is moved from one line to the other.

At the discontinuity, the common signal is converted and scatters both in the backward direction, back to the source and the detectors in the DTDR, and in the forward direction. The forward traveling common signal propagates down the differential pair and hits the open at the far end, where it reflects with a reflection coefficient of 1.

This reflected wave heads back to the receiver, where it is detected as a common signal, one round trip time of flight later. We see that both the backward and forward scattered common signals have the same sign when we place a capacitive discontinuity alternatively on one line and then the other.

References

1. Many of the principles described in this application note are introduced in detail in the book, Signal Integrity-Simplified by Dr. Eric Bogatin, published by Prentice Hall, 2003
2. Signal Integrity Solutions - Brochure, literature number 5988-5405EN, Aug. 29, 2005
3. Limitations and Accuracies of Time and Frequency Domain Analysis of Physical Layer Devices - White Paper, literature number 5989-2421EN, Nov. 1, 2005

Web Resources

For additional product information and application literature, visit our Web sites:

Physical Layer Test Systems:

www.keysight.com/find/plts

Learn more at: www.keysight.com

For more information on Keysight Technologies' products, applications or services, please contact your local Keysight office. The complete list is available at:

www.keysight.com/find/contactus

