SIEMENS

DIGITAL INDUSTRIES SOFTWARE

Signal integrity basics

John Golding

siemens.com/software

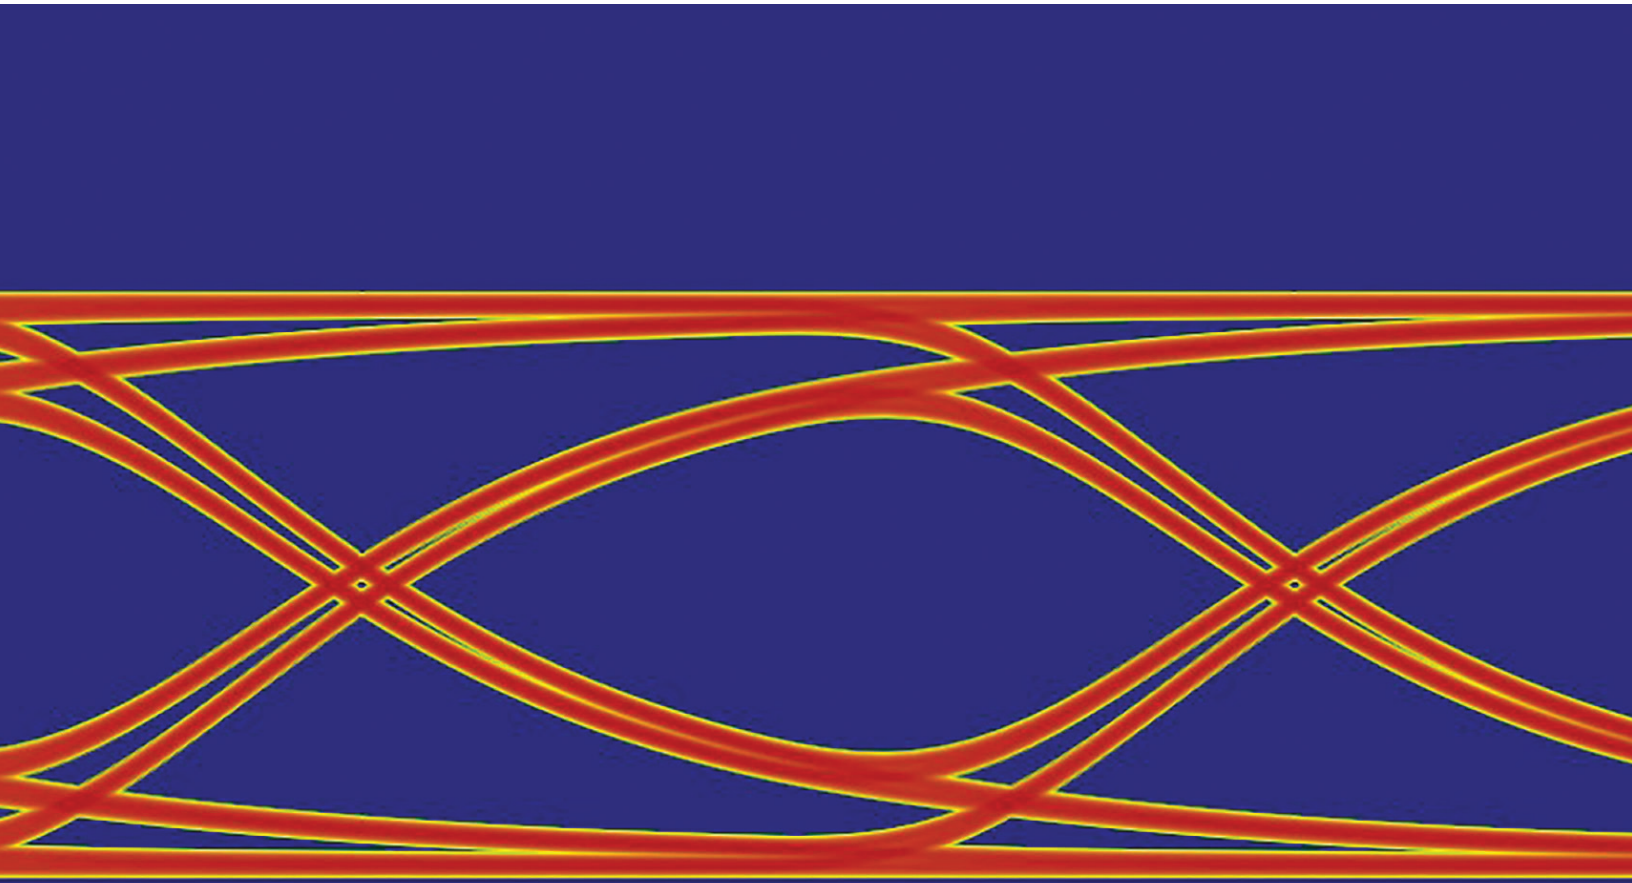

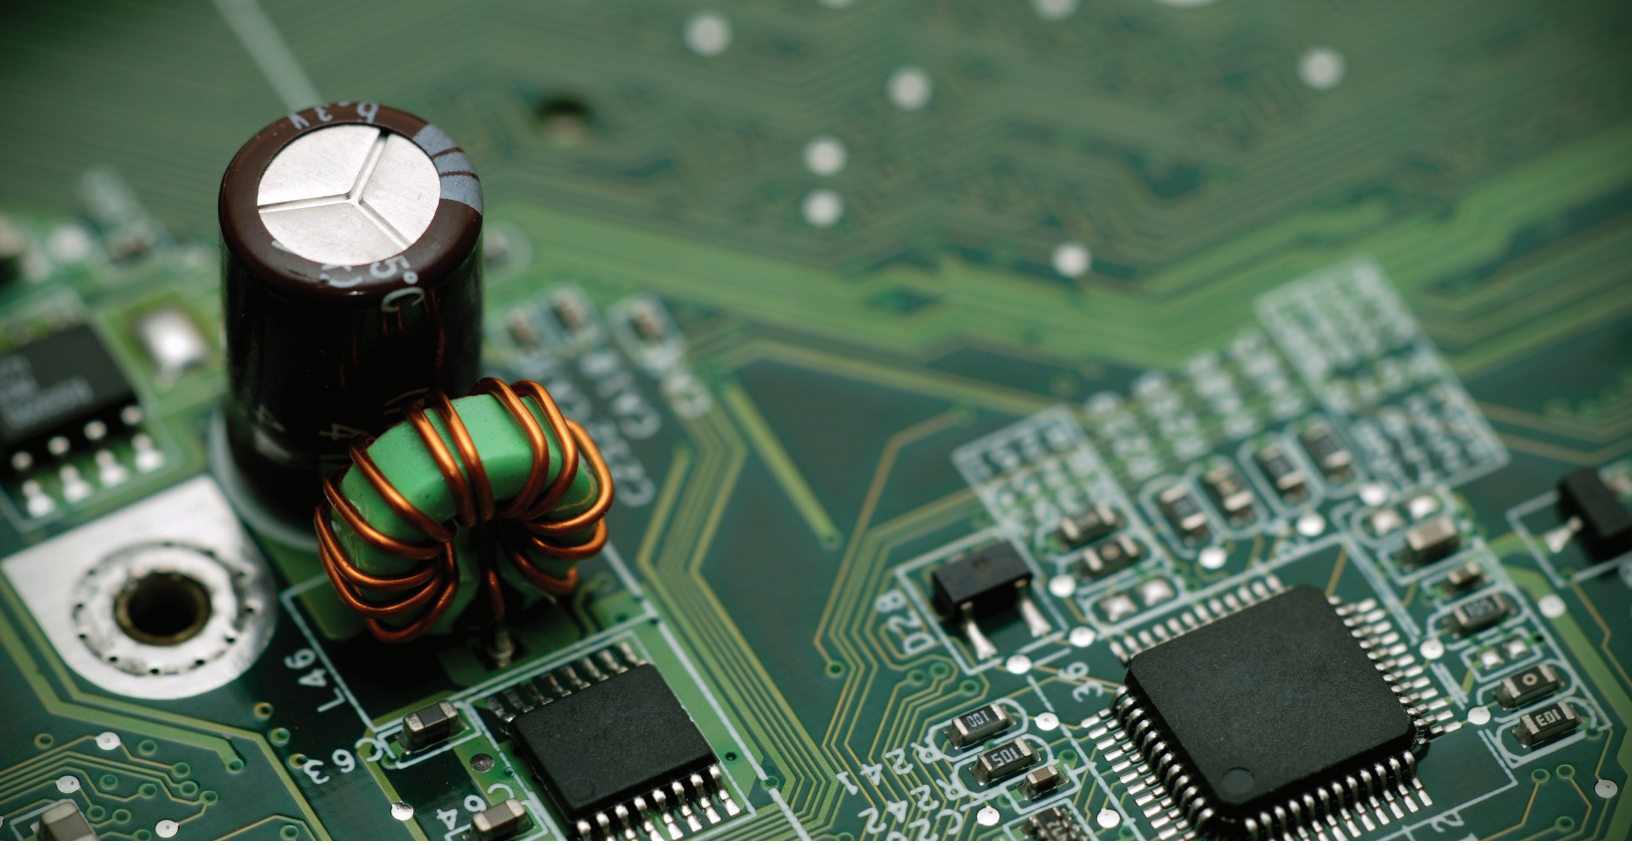

Table of contents

Introduction

In this orientation to signal integrity, we aim to introduce several important and fundamental concepts of signal integrity for the beginner. Most explanations are provided at a high level without a lot of depth and math, and examples are provided with a focus on comparison rather than detailed numerical results. Of course, background depth, math, and numerical details are very important and there are great resources available to help with those; however, the focus here is on orienting the reader to basic signal integrity concepts which can then be applied in practice and further study.

A simulator can be a powerful tool for learning about signal integrity and can help a user make design decisions and tradeoffs to meet performance requirements as well as other requirements such as cost and reliability for each use case. All simulations contained in this book were performed with the HyperLynx family of simulators and electromagnetic field solvers from Siemens EDA.

Many of the examples in this book are adapted from examples created by other engineers at Siemens EDA and have been previously used in presentations and workshops to teach signal integrity concepts and simulation using HyperLynx.

What is signal integrity?

When we talk about signal integrity, we're talking about the quality of a signal and how well it represents its intended shape.

The intended shape is defined by the receiving device. Though the transmitting device is critical to signal integrity, it already knows the message being transmitted. The receiving device, however, does not know what the message is ahead of time but needs to correctly receive it. To illustrate this concept, consider an analogy of a voice conversation. If I could speak perfect English but a listener doesn't understand English, we can't communicate. The message needs to be understood by the receiver of the message. In addition to language, there are other factors that could impact the success of voice communication. Some other possibilities are:

- Volume – If I speak too softly or too loudly, we may not communicate well. Too much volume could even damage your hearing. If I speak too loudly, I may also disrupt other conversations.
- Echo – If the sound bounces around, communication can be impaired.
- Noise – If there is excessive noise, the listener may not be able to hear or distinguish my voice. Noise could come from other people speaking or other sources.
- Timing – If I speak too quickly or slowly, or if my voice is delayed (on a phone call, for example), you recognize that the time at which the sound arrives affects the communication.

Of course, the speaking abilities of the speaker need to be compatible with what the listener requires. And, in a given environment, there may be some things the speaker could do to make communication more successful. Speaking louder or slower, for example, may make communication possible. There may also be something the listener could do to make communication possible or more successful, like using a hearing aid.

In some cases, however, you may need a better "connection." For example, you may need to move farther away from other conversations or into a room with better acoustic properties.

In signal integrity, we are concerned with ensuring a signal can be transmitted, travel through the interconnect, and arrive at a receiver in a way that is able to be correctly received by the receiving device. As with the voice analogy and multiple conversations going on, we also need to consider all relevant signals and their possible interaction with each other. We begin by considering some example wave shapes.

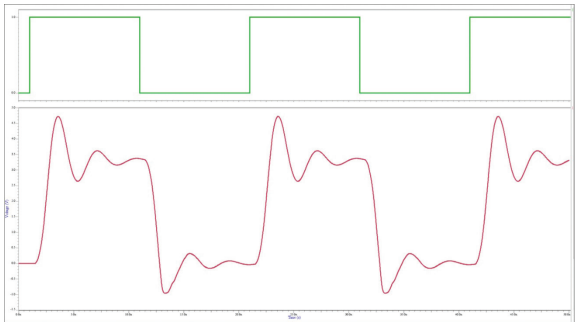

Figure 1. Digital and analog waveforms.

At the top of figure 1, we have a digital waveform that we intend to transmit. It would be great if we could deliver this waveform exactly to the receiver as is. On the bottom we have an actual waveform which is an analog waveform representing the digital pattern and it doesn't exactly match the digital signal. We'll compare them in more detail below.

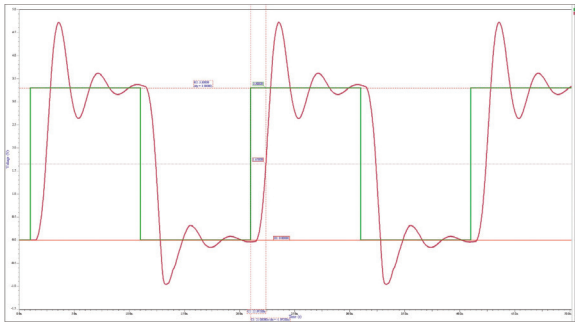

Figure 2. Comparing digital and analog waveforms.

In figure 2, we still have the ideal digital waveform (green) but plotted on a voltage scale (rather than just 0 and 1) for comparison to a more realistic signal. When we overlay and compare them, you'll notice some of the possible differences.

- Levels
  - The digital waveform has only discrete 0 and 1 levels (no voltage).
  - The actual waveform has continuous and varying high and low voltage levels (0V and 3.3V nominally).
  - The actual waveform has overshoot in the high and low transitions. Overshoot is a condition where the analog waveform goes above/below the nominal settling voltage.
  - The actual waveform doesn't immediately settle to the nominal high and low voltages but instead exhibits ringing.
- Timing
  - The digital waveform has perfect step changes between the 0 and 1 levels.
  - The actual waveform has finite rise and fall times when changing between voltage levels.
  - The actual waveform is shifted in time (delayed) a bit relative to the digital waveform.

Basic waveform issues

After recognizing basic differences between actual analog waveforms and ideal digital signals, we can further inspect waveform characteristics and possible issues. Real waveforms may (or may not) exhibit the characteristics shown in figure 3 in various combinations and to varying degrees.

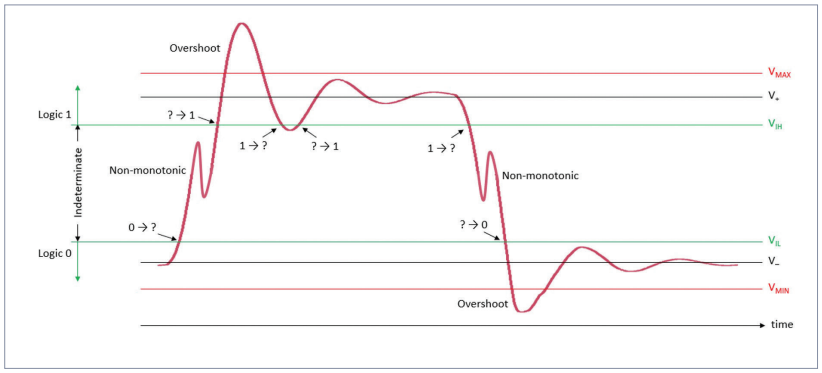

Figure 3. Basic waveform issues.

Overshoot occurs when a signal surpasses its intended voltage level before settling down to its final value. It often happens due to factors such as reflections from impedance mismatches which will be covered later. In digital systems, overshoot can lead to errors and even component damage if it exceeds the allowed limits. The overshoot is followed by ringing which may settle to the nominal voltage if there is time before the next transition.

A rising or falling edge that changes slope or direction during the transition is a non-monotonic edge. Non-monotonic behavior can also result in data transmission errors. Similar to overshoot, it has several root causes such as a reflection from an impedance mismatch.

Overshoot, ringing, and non-monotonic behavior are common issues encountered in signal integrity. Understanding their causes and implementing appropriate mitigation techniques is essential for ensuring the reliability and performance of electronic systems.

Intersymbol interference

Another possible waveform issue is intersymbol interference (ISI). This is a phenomenon that occurs in digital communication systems where symbols (or bits) transmitted over a channel interfere with each other, making it difficult for the receiver to correctly interpret the symbols. ISI can be caused by channel bandwidth limitations and other factors like reflections. We can easily see ISI caused by the channel's frequency response by comparing the waveforms in figures 4 through 9. We'll notice how it becomes more pronounced as the ratio of bit rate to channel bandwidth increases.

Bit rate = B1
Line length = L1
Bit pattern = PRBS7

Figure 4. Serial bit stream – bit rate B1.

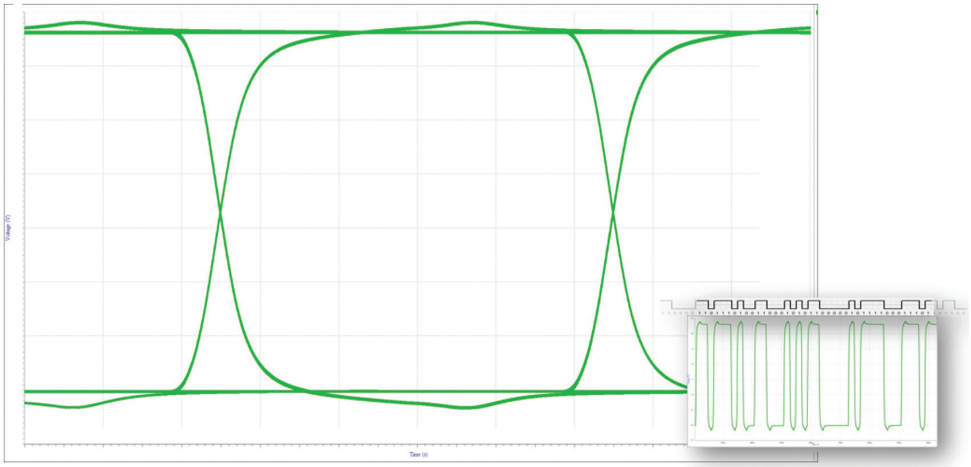

Figure 5. Eye diagram – bit rate B1.

In the image in figure 4 you can see a pretty good serial bit stream. The pattern is PRBS7, and you can see the digital representation above the analog waveform. The pattern starts before and ends after what is shown so we're somewhere in the middle of the bit stream. Notice the wider pulses and the narrower pulses reach the same high and low voltage levels.

The complete waveform, including bits from before and after the serial stream that we saw in the last figure is shown in figure 5. This is called an eye diagram because it looks a bit like an eye. That will be clearer in the next example. To create an eye diagram, we take all of the bits from the bit stream and overlay them on top of each other, so they're all represented in one bit time. In this eye diagram we have one full bit time in the center and then half a bit time to the left and half another to the right.

Figure 6 shows a waveform from the same interconnect and the same transmitter and receiver. The bit rate in this case is five times faster than in the last case. Notice now that not all the bits reach the same high and low levels. The narrower pulses are lower in amplitude than the wider ones.

Look at the long run of zeros just after the middle of the time axis. The voltage in that run of zero bits reaches the lowest voltage in this time capture. The following one bit does not reach the full high voltage level because it did not have time to reach

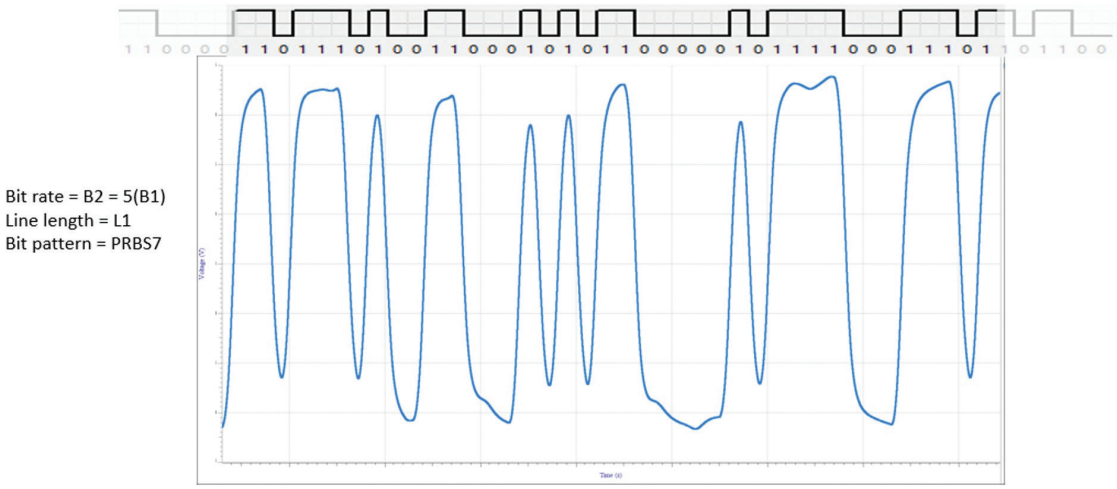

Figure 6. Serial bit stream – bit rate B2.

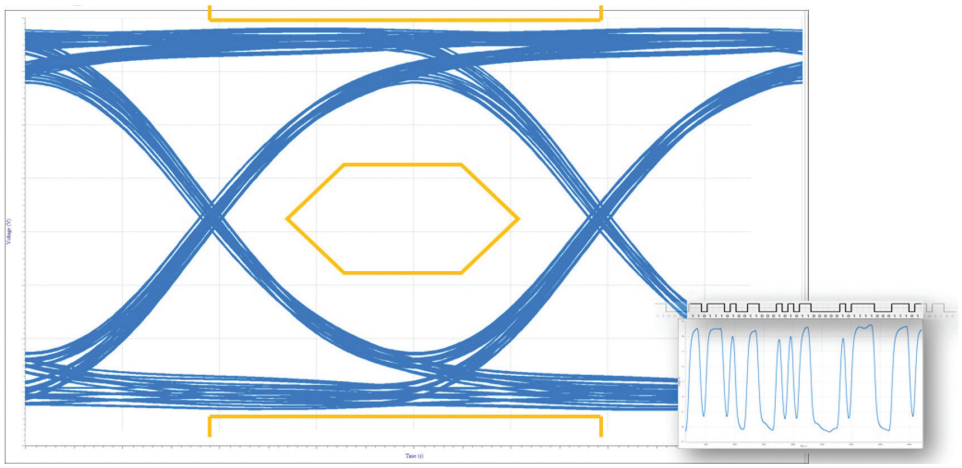

Figure 7. Eye diagram – bit rate B2.

that level before switching low again. Then, the following zero bit after that does not reach the same low voltage as the zero bit that is only two-bit times before it.

What we're seeing here is Intersymbol Interference. Pulses (or bits) spread into and interfere with the shape of other pulses (or bits). Again, this is caused by the limited bandwidth of the interconnect. In general, a typical interconnect acts like a low pass filter where higher frequencies are attenuated more than lower frequencies.

Figure 7 shows the eye diagram for the previous bit stream. Notice it is not as square as the previous one – and looks more like an eye with rounded corners.

We've also added an eye mask here which is the gold hexagon in the center. This is one way to measure the performance of a channel by plotting the eye diagram with a mask. If the waveform encroaches into the mask, the waveform fails the mask, and the receiver may experience bit errors. If the waveform stays away from the mask, then we have a high confidence we will have a low bit error rate.

The eye mask used here is an arbitrary eye mask. In general, these masks can be various shapes, such as hexagons, rectangles, diamonds, etc.

We'll look at one last example to further illustrate ISI by further increasing the data rate of the previous examples. In this case the bit rate is 10 times faster

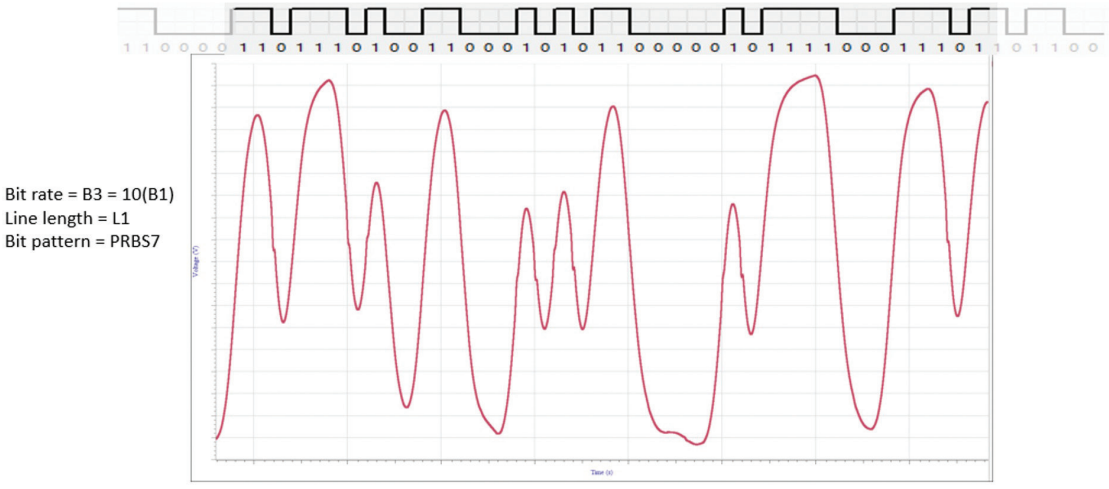

Figure 8. Serial bit stream – bit rate B3.

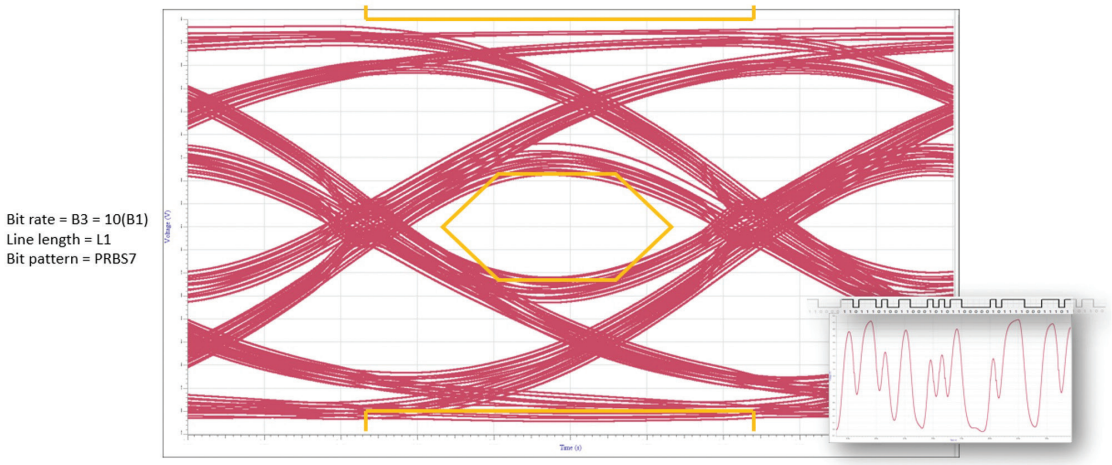

Figure 9. Eye diagram – bit rate B3.

than the original. Notice the pattern effects in figure 8.

We can easily see that the minimum and maximum voltages of the single isolated 0 and 1 bits depend on what came before them. The first isolated 1 bit reaches a higher voltage than the second isolated 1 bit because the waveform was biased to a higher voltage before the first isolated 1 bit. The second isolated 1 bit was preceded by multiple 0 bits, so it didn't reach as high a voltage. Also notice that the next to last isolated 0 bit reaches a lower voltage than the last isolated 0 bit. Again, this is ISI caused by channel bandwidth limitations. This case has a higher bit rate to channel bandwidth ratio, so the ISI effects are more pronounced than in the previous example.

Figure 9 shows the corresponding eye diagram. Notice now we're failing the mask (observe the signal encroaching on the center hexagon). This signal has a much higher chance of bit errors than the previous two examples.

The three previous cases are shown together in figure 10 to compare them. The interconnect, transmitter, and receiver are the same in all cases and we only changed the bit rate (increasing the bit rate from left to right in the figure). The horizontal dashed lines help give a comparison of the eye height of the three cases. The eye width is also different in each case. This can be observed by looking at the width of the eye opening where the low/high transitions cross.

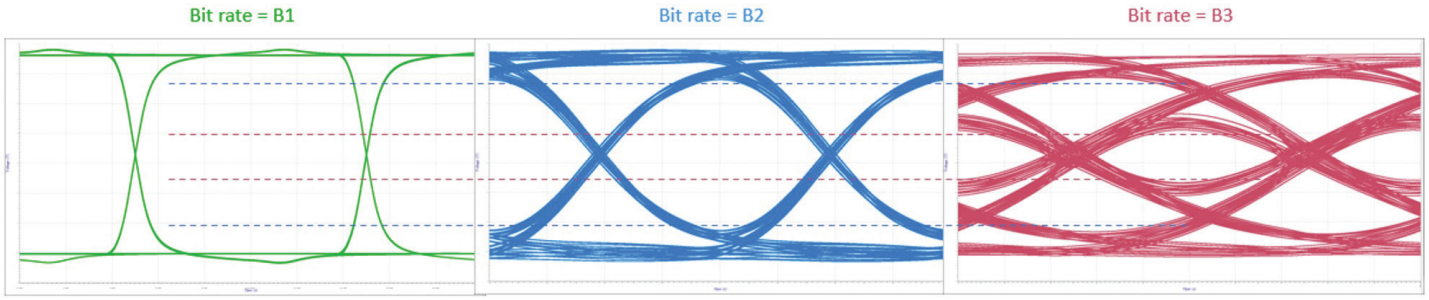

Figure 10. ISI – constant line length, varying bit rate.

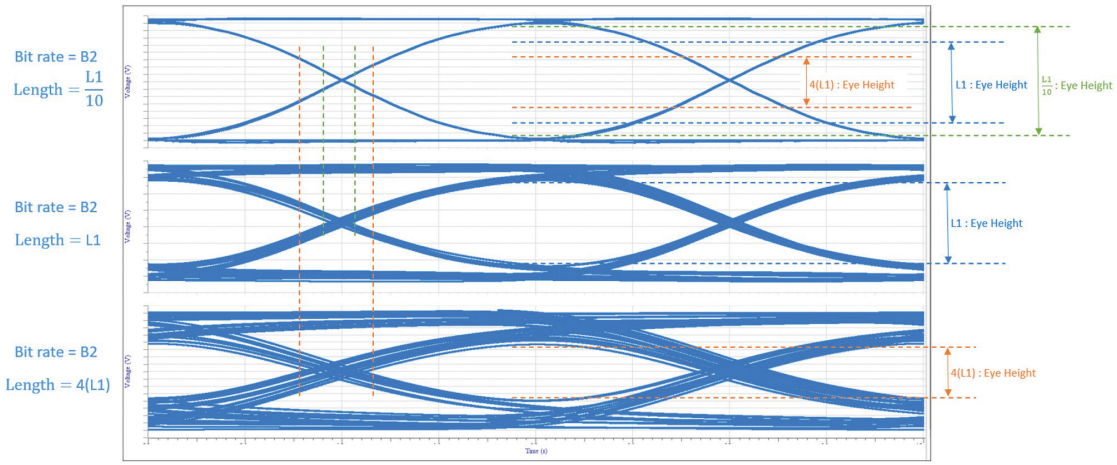

Figure 11. ISI – Constant bit rate, varying line length.

Now we'll look at one more example where we keep the bit rate the same but change the interconnect length. As mentioned before, the channel acts like a low pass filter and the longer it is, the more high-frequency signal attenuation we can expect. We choose bit rate two (B2) from the previous examples to test over the varying interconnect length. Again, the transmitter and receiver and bit rate are the same in all cases and we only changed the interconnect length. The results are shown in figure 11. The horizontal dashed lines help give a comparison of the eye height of the three cases. The vertical dashed lines give an indication of the eye width differences by comparing the width of the crossing of the three waveforms. Notice from these examples the relationship between the bit rate and channel bandwidth. If the bit rate is faster or the interconnect length is longer, the signal experiences more high frequency loss and the eye opening is reduced. We won't cover equalization here but note that equalization is a technique used to counteract this lowpass filter behavior.

With an awareness of some potential waveform issues, we can move on to discuss some of their causes and mitigation techniques.

Frequency

When a signal is static, it doesn't cause the waveform issues mentioned above. It simply remains at a constant voltage. When a signal is not static, its voltage is changing over time. A change over time implies there is some frequency associated with the rate of voltage change with respect to time. The waveform issues mentioned above arise when a signal is changing and are related to the frequency of change, so it's important to consider a signal's frequency content.

Frequency can be used to refer to the number of times a digital signal repeats or completes a full cycle in a unit of time (like a clock frequency, for example). We'll call that the cycle rate for this discussion. But this doesn't fully describe the frequency content of a digital signal. In signal integrity, we also must consider the rise and fall times, or the rising and falling edge rates, which we'll often simply refer to as edge rate. That is important, as we look into the waveform issues we saw previously, with overshoot and ringing, etc. It is also important for the ISI issue. Consider the two waveforms in figure 12. The green waveform is repeating at a rate of 10 MHz. The blue waveform is

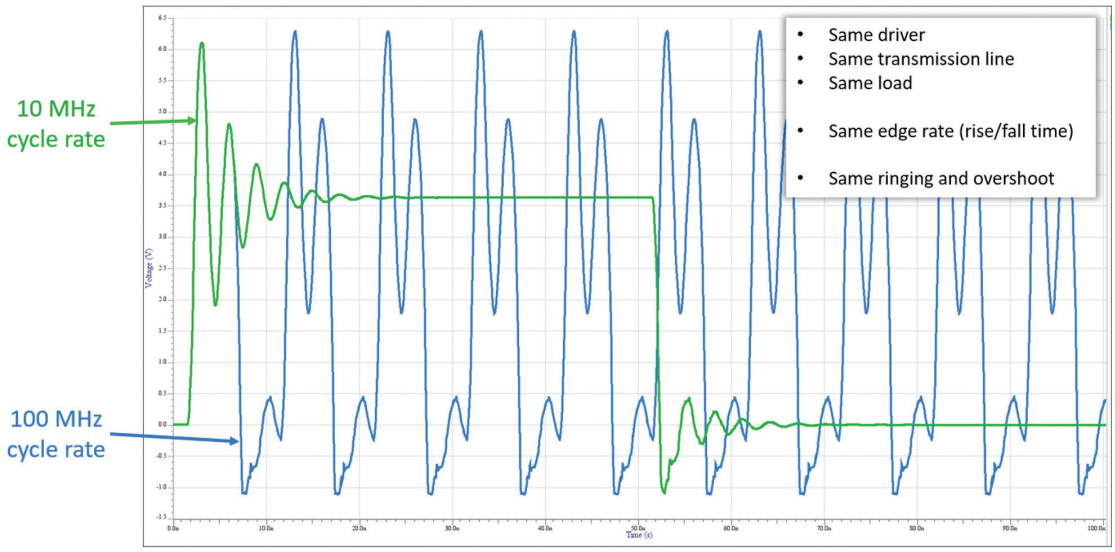

Figure 12. Frequency – cycle rate vs. edge rate.

repeating at a rate of 100 MHz, which is 10 times faster than the green waveform. Both waveforms have the same amount of overshoot and ringing. In this case we have a waveform that is repeating (cycle rate) 10 times slower, but the edge rate is just as fast and therefore we have the same ringing and overshoot.

When we think about phrases like high-speed or high-frequency signals, or fast or slow signals, we need to distinguish between the cycle rate and the edge rate. The high-frequency content of digital signals is in the signal edges as illustrated in the figure 13. On the left we have signals with different frequencies. These are sinusoidal signals and each has only a single frequency. One has a frequency that three times higher than the other.

On the right we have two digital signals, both cycling at the same rate but one has faster edge rates. The digital waveform on the top can almost be represented by the single sine wave. The digital waveform on the bottom is significantly different from the sine wave. For both digital signals we need higher frequency content to make up those edges. The bottom waveform needs more high-frequency content than the top one to achieve that edge rate. The flat parts of the signal where it is static are DC and, while we care about that part of the signal, that is not what causes signal integrity issues.

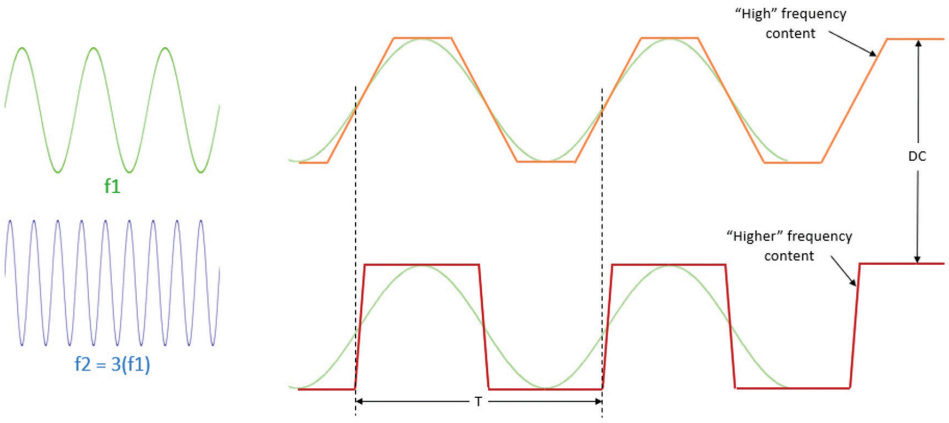

Figure 13. Frequency – cycle rate vs. edge rate (2).

Transmission lines

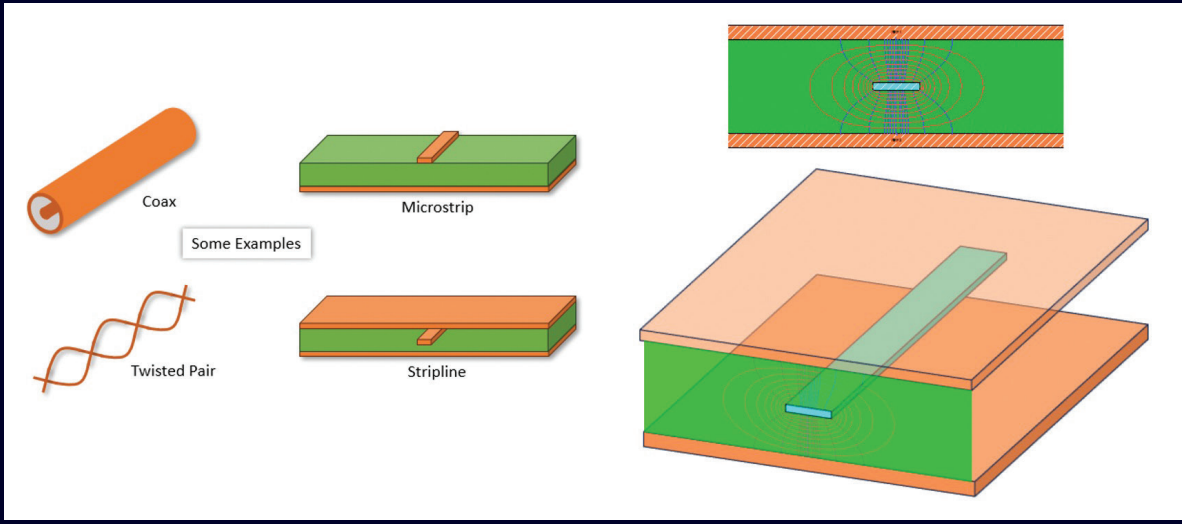

Figure 14. Transmission line examples.

In signal integrity design and analysis, one of the fundamental elements we're concerned with are the transmission lines carrying those signals. A transmission line is essentially a set of conductors separated by dielectric used to carry or guide electromagnetic energy from one place to another. Some examples are shown in figure 14.

The waveforms we've looked at thus far are voltage waveforms representing the signal and, of course, we have voltages and currents on transmission lines. But we also have electric and magnetic fields in the dielectric media, propagating from one place to another. Although we consider voltages and currents, we don't want to lose sight of the fact that we have electromagnetic energy moving around. This helps us think about concepts like propagation velocity and coupling to other structures – such as other signals or references. In a transmission line, we consider at least one conductor as return or reference, and another as a signal. Sometimes the intended return is called ground (or even a power net) on a schematic but it's important to remember that the return is a real, physical conductor, and not an ideal equipotential node. Figure 15 shows some example transmission lines designating return and signal conductors.

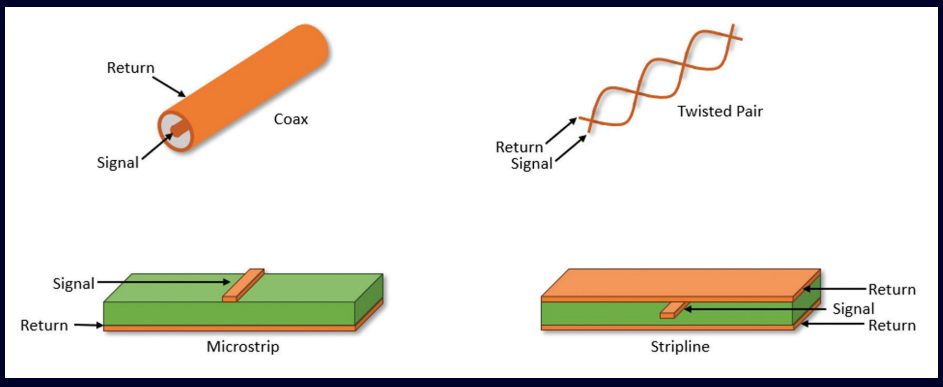

Figure 15. Transmission line examples (2).

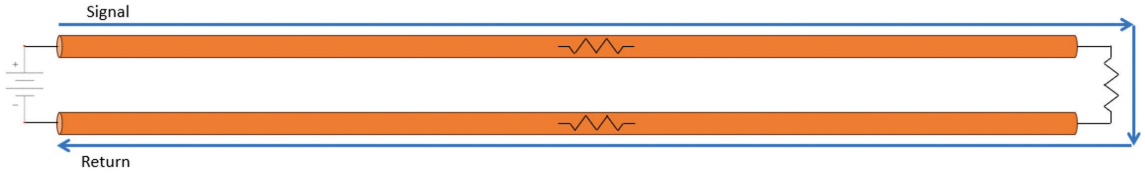

Figure 16. Transmission line at DC.

In a typical schematic, the interconnections between components are logical connections without physical definition. Even in basic circuit simulations, the connections between components are usually considered ideal. However, the connections between components on a printed circuit board (PCB) have physical material properties and dimensions. This means those connections are not ideal and they exhibit delay, loss, and coupling. This impacts the signal on a given connection and potentially other signals. We need to understand that impact when managing the integrity of those signals.

To get a basic understanding of these interconnects, consider first an interconnect carrying DC current. Figure 16 shows an arbitrary transmission line with a load resistor at DC. The DC current flows down the line through one conductor, through the load resistor, and back to the source along the return conductor. The conductors also have resistance associated with them.

Now we'll consider the same transmission line at AC (switching between voltage levels). Figure 17 shows the same arbitrary transmission line with a rising edge injected onto the transmission line at one end by quickly changing the voltage there from low to high. As mentioned previously, this rising edge has "high" frequency content. The AC current flows as displacement current through the distributed capacitance between the conductors and returns to the source. It starts at the end of the transmission line where the signal is injected and continues with the changing voltage propagating down the line and reaching the end after some time. The time it takes for the signal to travel the full length of the transmission line is the length of the line divided by the signal's propagation velocity. The propagation velocity is the speed of light in the dielectric media surrounding the conductors.

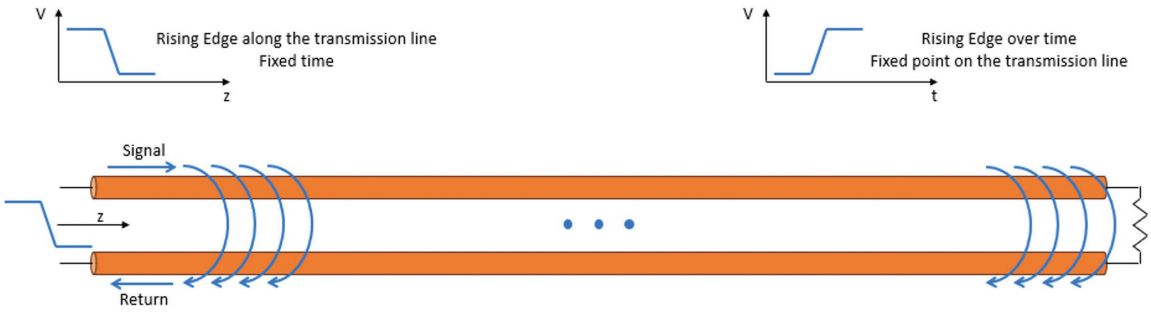

Figure 17. Transmission line at AC.

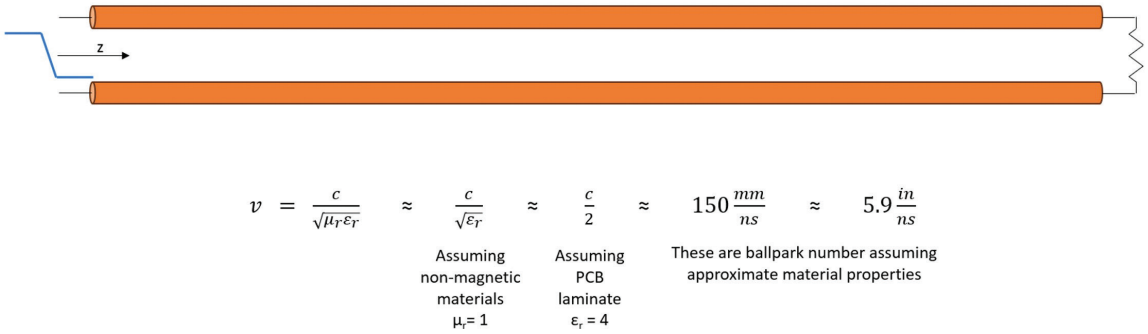

Figure 18. Propagation of velocity in a PCB.

As shown in figure 18, the electromagnetic fields propagating in a PCB travel at the speed of light in vacuum divided by square root of the relative permeability times the relative permittivity of the dielectric materials surrounding the transmission line. For non-magnetic materials, as in a typical PCB, the relative permeability is 1. The relative permittivity (which we also call dielectric constant) of PCB materials can have a range of values. But, if we assume a common approximate value of four, we see that the signal propagates in a typical PCB at about half the speed of light in a vacuum. This is about 150 mm/ns or about 6 in/ns which are convenient approximate numbers to remember.

When we think about interconnect length and signal integrity, we want to consider it in relation to the signal's rise and fall times. A signal's rise/fall time is usually 20%-80% or 10%-90% of the time it takes to transition from one voltage level to another. For a given interconnect length, we know it takes some time for a signal to propagate from one end to another as we just saw. If that time is short, relative to the rise/fall time, the voltage is almost constant over the length. To get a feel for this, consider figure 19. On the left side of the figure, are two identical signal transitions at time t=t0. On the right side of the figure are two transmission lines. The top transmission line is "short" and the bottom transmission line is "long." Now imagine sliding the signal transitions over each transmission line at the same speed and notice the voltages at the start and end of the transmission lines at time t=t1 in figure 20. On the short transmission line, the voltage at the start and end of the line is approximately the same. However, on the long transmission line, at time t=t1, we have a full voltage change at the start of the line, while the voltage at the end of the line has not changed at all because the transition hasn't had enough time to propagate there yet.

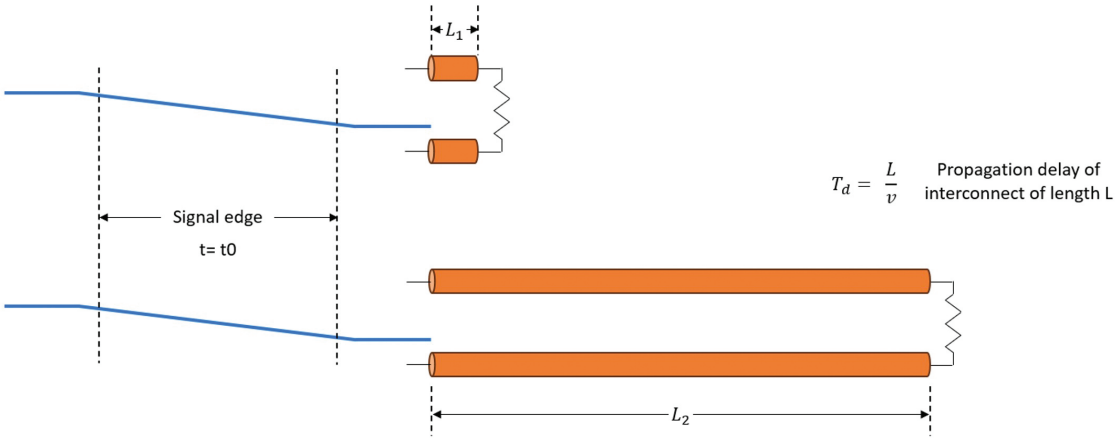

Figure 19. Signal edge relative to interconnect length.

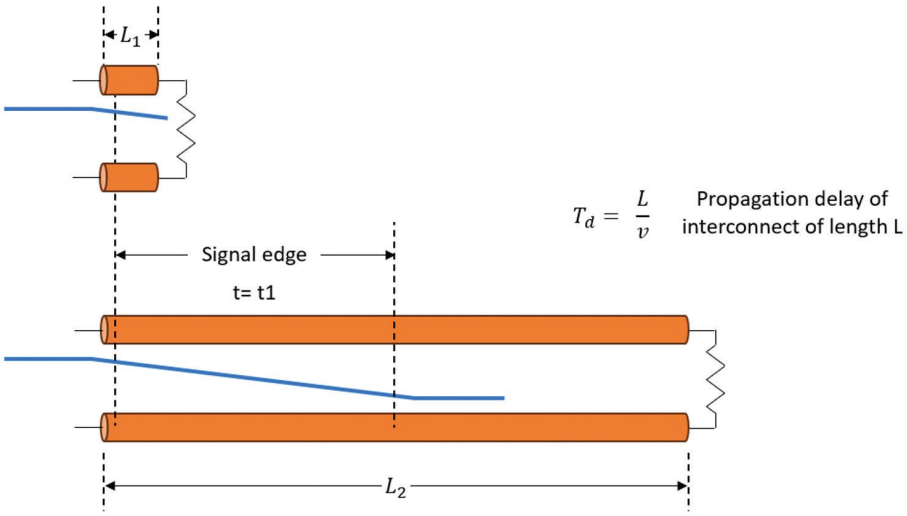

Figure 20. Signal edge relative to interconnect length (2).

Said differently, if the rise/fall time is very slow compared to the time it takes for the signal to propagate down the line, then the initial change will arrive at the far end of the line while the starting point of the is approximately the same. When the propagation delay (T_d) is much less than the rise/fall time (T_{rf}), we have this condition. When Td is an appreciable fraction of T_{rf} then the voltage varies along the length of the line (the voltage at the end is not the same as the beginning of the line). People use different benchmarks for when the interconnect, or transmission line, T_d/T_{rf} ratios begin to exhibit effects that can cause signal integrity problems. Conceptually though, if the propagation delay is not short relative to the rise/fall times, those effects can arise. For a transmission line whose delay is not very short relative to the rise/fall time, then it can be helpful to think of interconnect as distributed network as shown in figure 21.

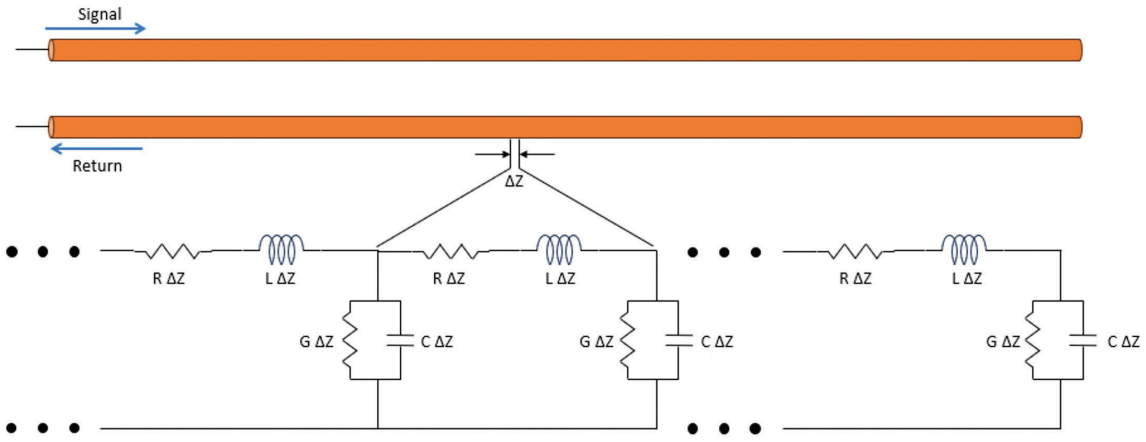

Figure 21. Transmission line distributed network model.

The figure shows an arbitrary transmission line with signal current flowing to the right and return current flowing to the left. It also shows a helpful representation of the transmission line as a distributed network of circuit elements. The values of these per-unit-length elements are determined by the geometry and material properties making up the transmission line. Each Δz is very short (in reality, it's a continuous structure), and we can think about our signal propagating through this line charging up each element as it goes. In relation to the previous example, consider short Δz as short delay relative to the rise/fall time. L and C are the per-unit-length inductance and capacitance, respectively, associated with the conductors. This model also reminds us of losses associated with a transmission line. R is the per-unit-length resistance representing conductive losses. G is the per-unit-length conductance representing dielectric losses. The resistance and inductance are associated with the signal and the return conductors but have been grouped together in the figure for simplicity.

In figure 22, we have another arbitrary transmission line. It has a uniform cross-section, meaning the geometry and material properties do not change along the length. As the signal begins to charge the line, current begins to flow. The current is determined by the voltage and the impedance of the transmission line. This impedance is known as the characteristic impedance, and it is a very important parameter and concept in signal integrity. We identify it as Z0 or Z_0. As the signal travels farther, and at each point, we have the same impedance because the structure is constant. If we change the cross section of the transmission line (materials and/or geometry), we can change the characteristic impedance and the delay.

Illustrated in figure 23 and figure 24 are some transmission line examples showing their corresponding delays and characteristic impedances. We start with a baseline transmission line in the left column and then change one variable at a time to see the effects. Characteristic impedance consists of more than just capacitance but we can quickly anticipate characteristic impedance increases or decreases by thinking of a parallel plate capacitor. We do this by remembering the capacitance $C=\varepsilon_r\varepsilon_0 A/d$ and the impedance $Z_c=1/j\omega C$, where ϵ_r is the dielectric constant, A is the area of the plates and d is their separation. This will be illustrated in the following examples.

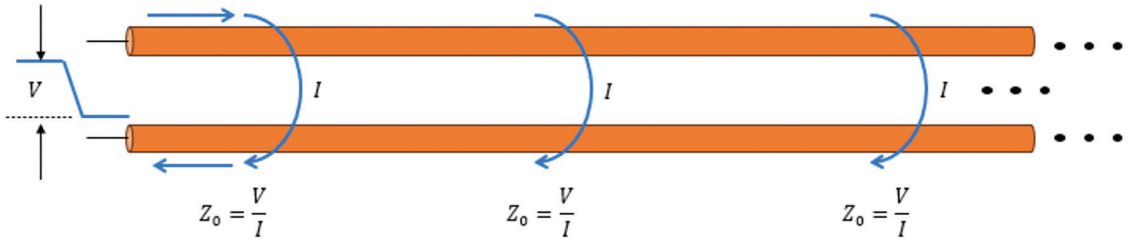

Transmission line with uniform cross section

Figure 22. Characteristic impedance.

In the microstrip examples of figure 23, notice first the delay for the baseline example and remember the ballpark 6 in/ns propagation velocity discussed previously. Here we have a 3 inch long transmission line and the delay is approximately 0.5 ns, but slightly less. Because this is microstrip, not all of the electromagnetic field is contained in the dielectric having a dielectric constant of four. Instead, some of the electromagnetic field (above and next to the signal trace) is in the solder mask and the air, which have a lower dielectric constant. Therefore, the propagation velocity is a little faster than the estimate that was based on a dielectric constant of four. In the second column, the height of the dielectric under the signal trace is decreased which causes a lower characteristic impedance. If we think of moving a capacitor's conductors closer together, the capacitance increases and its impedance decreases. In the third column we keep the thinner dielectric and make the signal trace narrower which increases the characteristic impedance. Again, we can expect this by thinking of a parallel plate capacitor. By making capacitor plates have less area, the capacitance decreases and the impedance increases. In both the second and third columns, the delay also changed slightly, even though the dielectric material properties did not change. This is because changing the cross-section geometry, also changed the amount of electromagnetic field contained in the dielectric under the signal trace, as well as in the solder mask and the air above and next to the signal trace. Finally, in the fourth column, we increase the dielectric constant and keep everything else the same. This results in a decrease of the characteristic impedance, as we expect when thinking of a capacitor. We also see a larger change in the delay in this case and it increased. This is because the larger dielectric constant causes a slower propagation velocity.

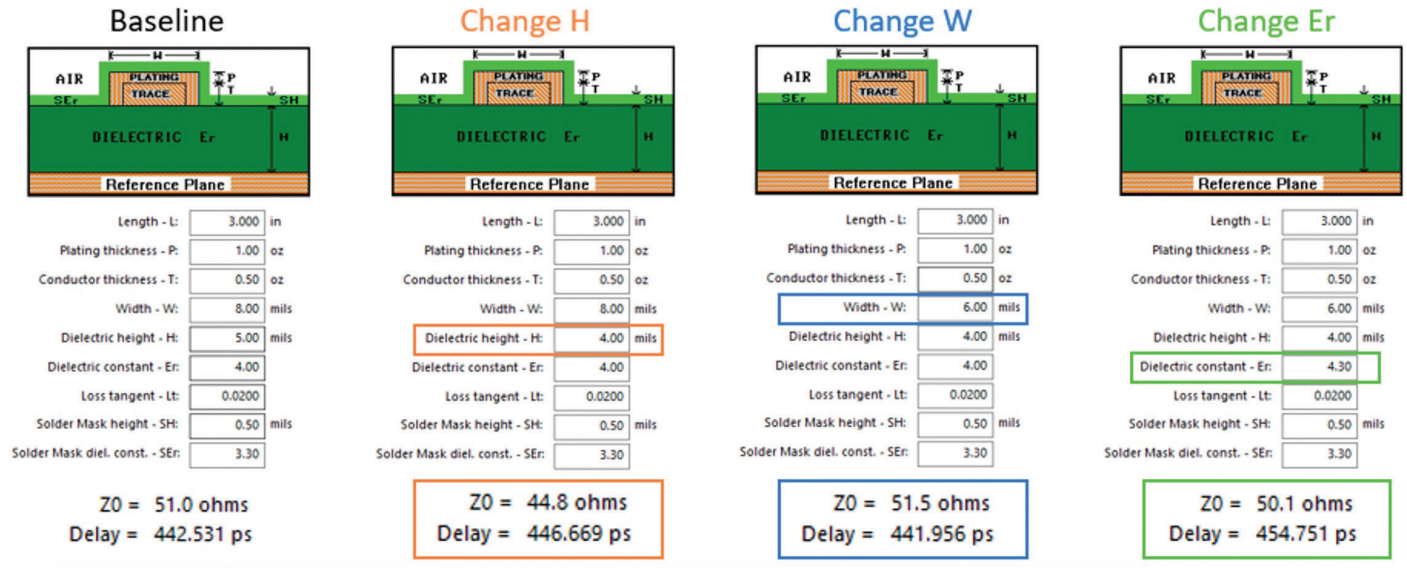

Figure 23. PCB microstrip examples.

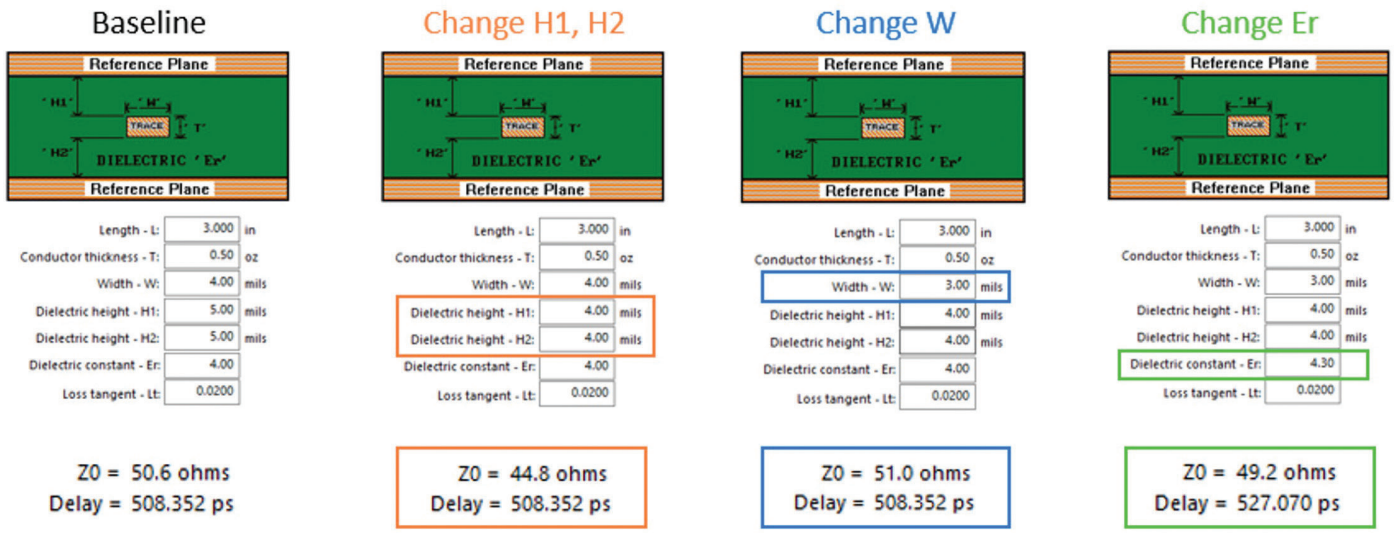

Figure 24. PCB stripline examples.

Now let's look at some stripline examples in figure 24. Notice the delay for the baseline and remember the ballpark 6 in/ns propagation velocity discussed previously. Here we have a 3 inch long transmission line, and the delay is approximately 0.5 ns, as we would expect. Notice the delays for all the stripline examples are longer than the microstrip examples. This is because some of the electromagnetic field is contained in materials with lower dielectric constant for the microstrip cases and therefore signals propagate faster than in the stripline cases.

In the second column, the heights of the dielectrics are decreased which causes a reduced characteristic impedance. Again, if we think of moving a capacitor's conductors closer together, the capacitance increases and its impedance decreases. In the third column we keep the thinner dielectrics and make the signal trace narrower which increases the characteristic impedance. In both the second and third columns, the delay stayed exactly the same. This is because no change was made to the dielectric constant and all of the electromagnetic field in the stripline case is contained in this dielectric. Finally, in the fourth column, we increase the dielectric constant and keep everything else the same. This results in a decrease of the characteristic impedance. We also see an increase in the delay as we expect because the larger dielectric constant causes a slower propagation velocity.

Thus far, we've only been considering a dielectric constant to illustrate concepts. However, we should note that real dielectric material properties have some frequency dependence which will be briefly addressed later.

We have also been considering only transmission lines with uniform cross section and one characteristic impedance until now. But if the characteristic impedance changes, this has an impact on the signal, and we will consider that now. Let's say we do change the cross section of a transmission line somewhere along its length. We will then have sections with different characteristic impedances touching each other. This is called an impedance discontinuity. In the microstrip and stripline examples we saw that if we change a trace width, we get a change in characteristic impedance. We also got a change in characteristic impedance by changing the separation between signal and return conductors and by changing the material properties. These geometry changes could happen if a signal trace changes layers in a PCB or changes width on a layer. We will see a few other possible causes of impedance discontinuities later.

In figure 25, we have a representation of a transmission line with two sections having different characteristic impedances which we denote as Z_{01} and Z_{02}. Note in lower right corner of the figure that we have a change in representation of a transmission line. Rather than explicitly showing both the signal and return conductors, we will use the coax-like cylinder symbol to represent a transmission line to simplify the diagrams. We need to remember though, that the transmission line has both signal and return conductors. Both are important to signal integrity, and both are represented in the simplified representation.

When a signal propagating on a transmission line encounters an impedance discontinuity, some of the signal will be reflected back toward the source and some will be transmitted through. Qualitatively, we can think of this by remembering that the characteristic impedance is determined by the transmission line cross section (materials and geometry) and that impedance is the ratio of voltage to current. Now consider a signal transition (rising or falling edge) propagating down the line on section 1 and arriving at the impedance discontinuity where section 1 meets section 2. The characteristic impedances in the sections are different so the voltage/current ratios cannot be the same in both sections. The signal must therefore change at the point of discontinuity to have voltage and current continuity there.

The signal changes by establishing a reflected voltage and current that propagate back towards the source in section 1. Quantitatively, we can calculate the reflected voltage using the equation in figure 25, which shows the reflection coefficient equation. The reflected voltage then is the incident voltage multiplied by the reflection coefficient. The total voltage at the boundary, and in section 1, is the sum of the incident and reflected voltages. The transmitted voltage moving into section 2 is the sum of the incident and reflected voltages. Keep in mind that the reflected signal experiences the same transmission line in section 1 as the incident signal does, including delay. Therefore, it will take time for the reflected signal to propagate from the discontinuity back down the transmission line toward the source and we will see an example of that.

Examining the reflection coefficient equation, we can see that if the characteristic impedance of section 2 is larger than the characteristic impedance of section 1, then we get a positive reflection so the voltage at the discontinuity increases. If the characteristic impedance of section 2 is smaller than the characteristic impedance of section 1, then we get a negative reflection and the voltage at the discontinuity decreases. Of course, if the characteristic impedance of section 2 is the same as section 1 then the reflection coefficient is zero and we get no reflection as we would expect.

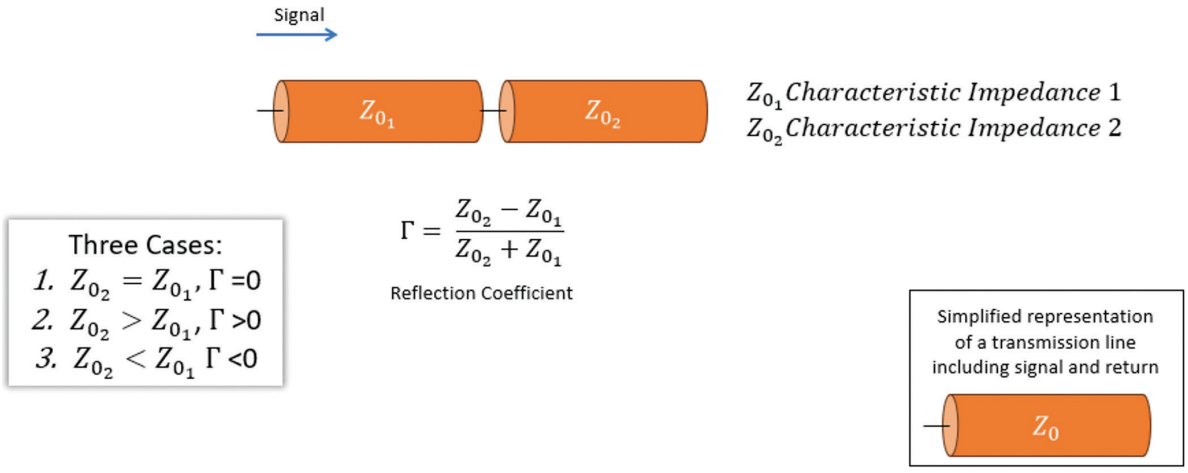

Figure 25. Reflection coefficient equation.

It can be helpful to consider the extreme cases of impedance discontinuities, and these are shown in figure 26. First, consider the case of an open circuit, or infinite impedance, at the discontinuity. The reflection coefficient is 1 which means we have 100% voltage reflection. We can recognize that no current flows off the end of the line, so the 100% reflection means the same voltage as the incident signal appears but with opposite propagation direction and opposite current flow direction. Again, the total voltage at the discontinuity is the sum of the incident and reflected voltages which, in this case, is two times the incident voltage. Now consider the other extreme case of a short circuit, or zero impedance, at the discontinuity. The reflection coefficient is -1 which means we have -100% voltage reflection. Here we can recognize that there is no voltage across a short circuit so the -100% reflection means a voltage appears that is opposite the incident voltage and propagating in the opposite direction. The total voltage at the discontinuity is zero because if we add two equal and opposite voltages, we get zero. All other cases are somewhere between these extremes. The extremes are convenient to keep in mind because they approximate frequently encountered conditions like low output-impedance and high input-impedance buffers, for example.

Reflections can cause signal integrity problems so it's important to recognize potential sources of impedance discontinuities and avoid them where possible. It is not usually possible to avoid all discontinuities so in those cases we need to minimize and/or manage them. So far, we've considered changes in transmission line cross-section as a source of discontinuity but there are other sources such as vias, connectors, and component pins and packages. Remember too that if the transmission line has a branch or stub or if the return path is interrupted, all these effectively change the transmission line's cross section and therefore can result in impedance discontinuities.

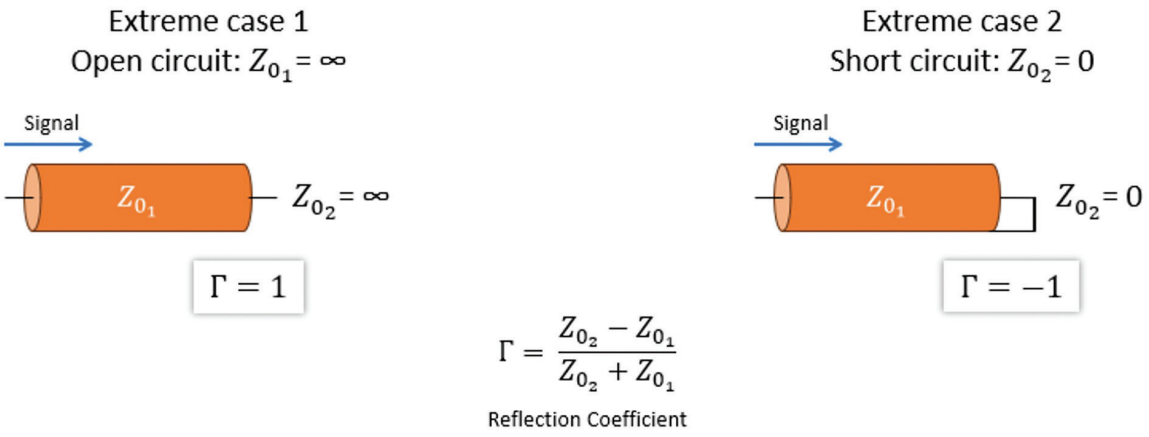

Figure 26. Reflections – extremes.

One way to mitigate the effect of reflections is to use a short interconnect relative to the rise/fall time such that reflections are effectively masked by the signal's edges. This effect can be seen in figure 27 where the only change causing the three different waveforms is the length of the interconnect. For the "long" interconnect (red) we see significant overshoot and ringing. For the "short" interconnect (green), we see minimal overshoot and ringing. In practice, we often can't achieve a total interconnect length relative to the rise/fall time that is short enough to manage reflections. This is because physical layout constraints and device technologies usually mean we can't make the length short enough and/or the rise/fall times long enough. Though not always practical, this technique is mentioned to reinforce the important signal integrity concept of the relationship between interconnect delay (length) and signal rise/fall times (frequency content).

Another technique that is frequently used to avoid or minimize reflections is keeping a consistent (or controlled) impedance throughout the interconnect. We can also try to avoid vias or design vias to minimize signal impact (more on this later). There are some discontinuities we may not be able to completely avoid – like connectors for example (if we need connectors) – so we have to minimize and manage those effects. Note that even if we could avoid impedance changes along the line, we also have to consider the ends of the lines.

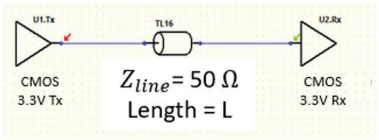

Keep the transmission line or line segments "short" (it's not always possible in practice)

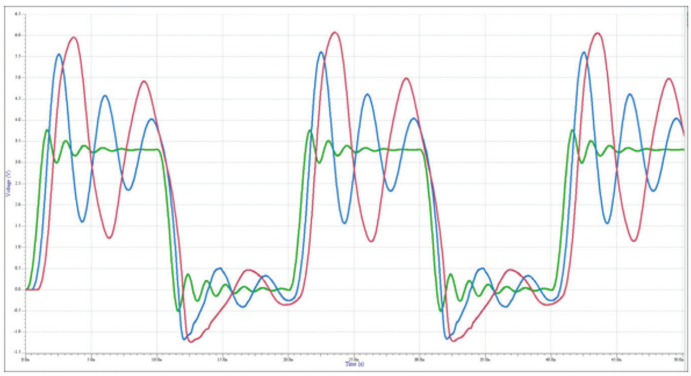

Figure 27. Reflections and length.

Consider the example as shown in figure 28. We have a 3.3 V CMOS driver and receiver at the ends of the line, and a 50 Ω transmission line with a 10 ns delay from driver to receiver. The CMOS driver has a low output impedance (approximately 4 Ω). The CMOS receiver has a very high input impedance (significantly higher than 50 Ω). In this case, we only have one transmission line with a consistent characteristic impedance, so we don't expect any reflections from the interconnect between the driver and receiver. However, we can still get reflections at the ends of the transmission line when the signal encounters the impedance of the buffers. In the figure, the red waveform is at the driver and the green waveform is at the receiver. When the driver launches the rising edge into the transmission line, it rises to approximately 3 V. It doesn't rise all the way to the full 3.3 V because there is some small output impedance associated with the driver. This signal propagates toward the receiver on the 50 Ω transmission line and eventually encounters a very high impedance when it arrives at the receiver 10 ns later. The reflection coefficient there is approximately 1, so we get a reflected voltage of approximately 3 V and the voltage at the receiver rises to 6 V (3 V + 3 V). The reflected 3 V signal then propagates 10 ns back toward the driver on the 50 Ω transmission line when it eventually encounters a 4 Ω impedance at the driver. The reflection coefficient there is -0.85 so we get a reflected voltage of -85% of the 3 V signal that was reflected off the receiver. The voltage at the transmitter approximately 20 ns after the initial signal was launched, is the sum of the original launched voltage (3 V) plus the reflected voltage from the receiver (3 V) plus the new reflected voltage at the transmitter (-2.5 V). This is approximately 3.5 V (3 V + 3 V – 2.5 V). The -2.5 V reflection from the transmitter then propagates toward the receiver where there will be another positive reflection and so on. This is how the signal integrity issue of overshoot is caused followed by ringing as the reflected signal bounces back and forth along the transmission line.

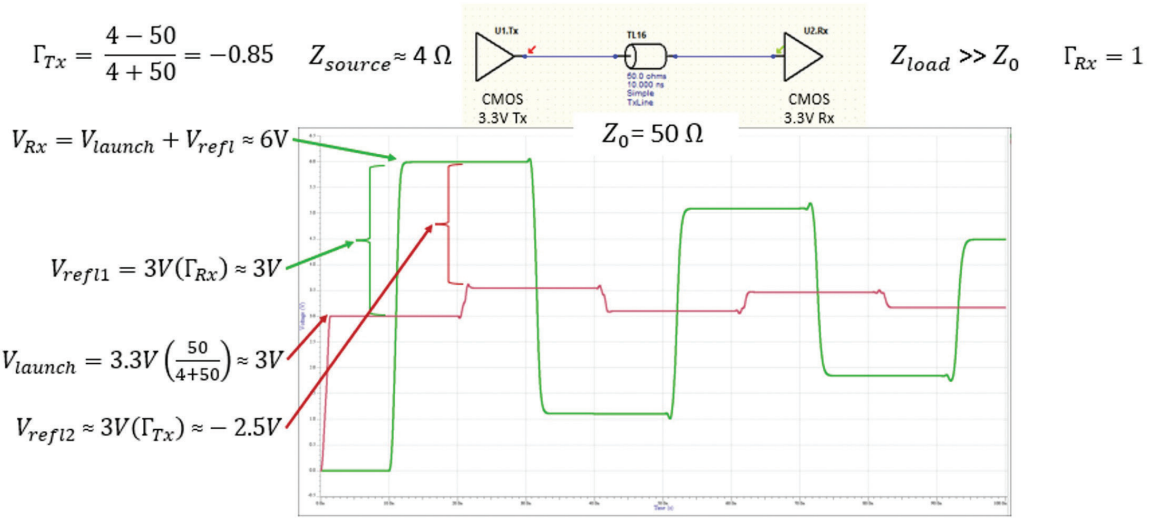

Figure 28. Reflections example.

Termination

We're now going to continue with the discussion of managing reflections using terminations. The specific examples are used to demonstrate concepts, but in general, a simulator can be used to determine the correct values and topologies for a given situation. We begin by revisiting the previous reflection example with the CMOS buffers and a 50 Ω transmission line between them. The example is shown again in figure 29 but this time the output impedance (4 Ω) of the transmitting buffer is shown explicitly.

Now, we modify our example by adding a resistor, R_s, close to the transmitter, as shown in figure 30. If we keep resistor R_s very close to the transmitter, we can effectively combine it with the output impedance of the buffer to get a new effective transmit impedance, Z_{TX}, which is the sum of R_s and the buffer's output impedance (Z_{source}). This technique is called series termination and with this approach we are taking advantage of the large reflection from the receiver. The idea is to allow that first reflection from the receiver end but then prevent any further reflections by not allowing any reflections at the transmitter. If we don't want any reflections from the transmitter then the reflection coefficient there should be zero and we use this to determine the value of R_s. Setting the reflection coefficient to zero means setting the equation's numerator to zero or $Z_{TX} - Z_0 = (Z_{source} + R_s) - Z_0 = 0$. This then requires $Z_{source} + R_s = Z_0$, and $R_s = Z_0 - Z_{source}$.

The effective transmit impedance (Z_{TX}) and the characteristic impedance of the transmission line (Z_0) create a voltage divider. When they are equal, the voltage launched into the transmission line is half of the source voltage (seen in the red waveform in figure 30). This voltage then propagates on the transmission line with characteristic impedance Z_0 to the receiver where it encounters a very large impedance. The reflection coefficient there is approximately 1. With approximately 100% reflection, the voltage at the receiver is approximately two times the arriving voltage. The incident voltage was half of the source voltage so when adding the incident and reflected voltages, we get the full source voltage at the receiver which is our goal. This is seen in the green waveform in the figure. Now the reflected voltage propagates back toward the transmitter along the transmission line with characteristic impedance Z_0. When arriving at the transmitter, the signal encounters the effective output impedance Z_{TX} which equals Z_0 so there is no reflection at the transmitter which also is our goal. In summary, we have managed reflections with a series resistor (series termination) by allowing the

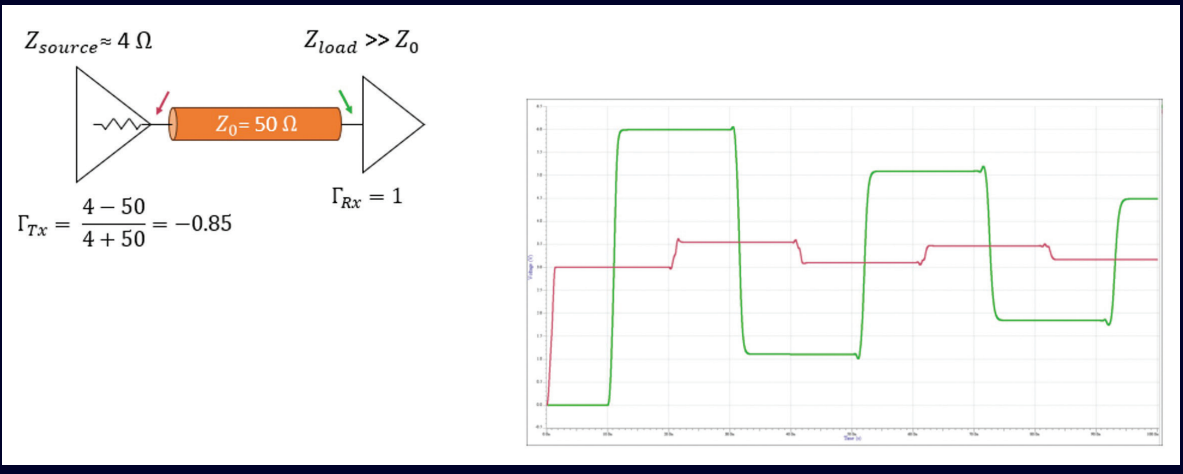

Figure 29. Reflections example (2).

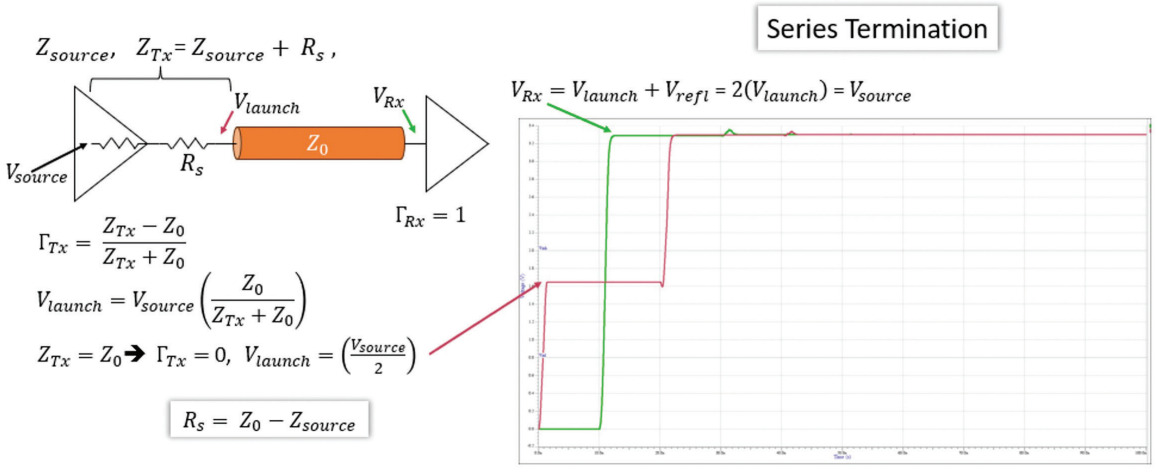

Figure 30. Managing reflections – series termination.

reflection at the receiver and terminating the reflected signal at the transmitter to eliminate further reflections. We must keep R_s close to the transmitter. If we don't, then the reflected wave doesn't encounter a match to the transmission line characteristic impedance. Instead, it encounters R_s and then some other impedance of the connection between R_s and the transmitter before encountering Z_{source}.

In the example shown in figure 31 we again have CMOS buffers and a transmission line with characteristic impedance of approximately 50 Ω. We also have a series termination resistor whose value is varied from 0 Ω to 50 Ω in 5 Ω steps. In the smaller picture in the lower left, we can observe all the resulting waveforms measured at the receiver overlaid. In the larger picture on the right, we see three of those waveforms with Rs = 0 Ω, 25 Ω, and

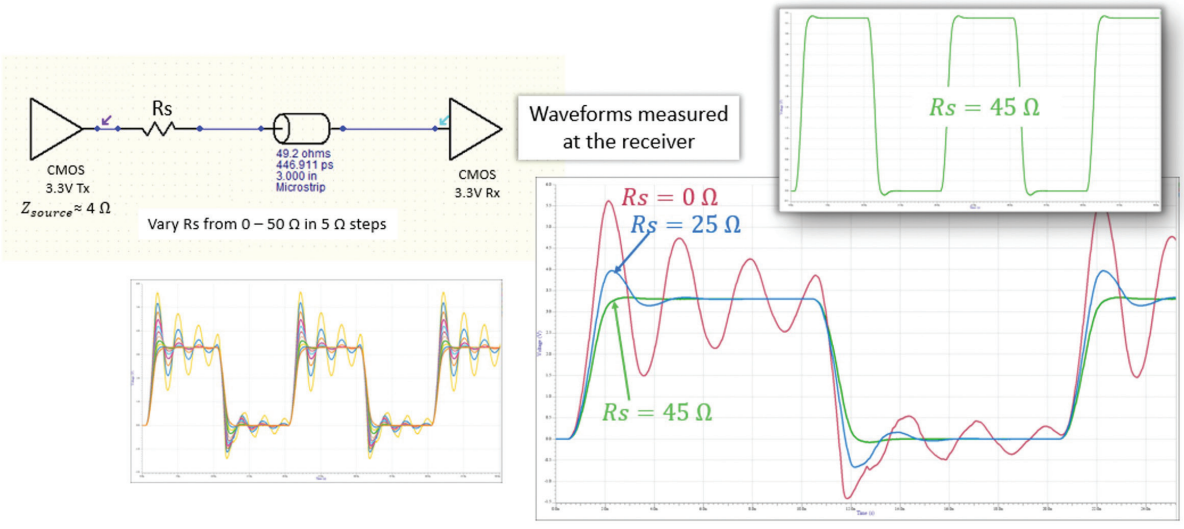

Figure 31. Series termination example.

45 Ω. The top right waveform shows the receiver waveform for just the case of R_s = 45 Ω. Notice the dramatic improvement in overshoot and ringing made with this resistor compared to the case with Rs = 0 (which is essentially no termination). It makes sense that a series termination value of 45 Ω would do a good job in this case because adding this to the output impedance of the buffer (4 Ω) is a good match for the transmission line characteristic impedance (49 Ω). We should also observe that increasing the series termination value has the effect of slowing down the edge rates.

Series termination is an effective technique, particularly for point-to-point signals, but it is not the only termination technique. Figure 32 shows some other termination possibilities. All the examples in the figure use the same transmitter and receiver and the same transmission line. A series termination example is also included for comparison (green waveform). All of the other examples make use of a termination at the receiving end rather than at the transmitting end of the transmission line. The general intent of these approaches is to match the transmission line impedance at the receiving end and prevent reflections there. Signals arriving at the receiver have been traveling down the characteristic impedance of the transmission line and then encounter an impedance of the same value at the receiver so there are no reflections. The impedance at the receiving device is the parallel combination of the termination impedance and the input impedance of the receiving device. Since the buffer input impedance in this case is very large, the effective impedance at the receiver is approximately just that of the termination. Take a moment though to consider that the devices and their packages can also have a significant impact on signal integrity. Signals arriving at the pins of a receiving component for example, may still have to continue through additional interconnect of the IC package before arriving at the actual input buffer. Some devices also provide terminations on the die which puts the termination as close as possible to the receiver. A signal measured in the lab at the pins of a device outside its package may look different than the signal on the die. For now, however, we're focusing on concepts and assume the IC package has minimal effect.

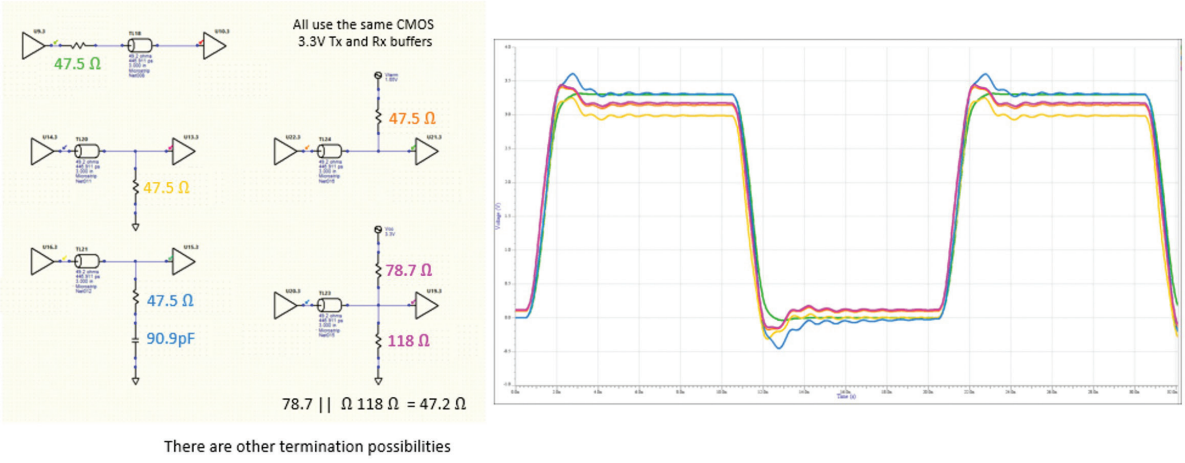

Figure 32. Additional (parallel) termination examples.

Let's briefly examine the terminations in figure 32. For all of them, we notice that if the termination impedance is a good match to the transmission line characteristic impedance, the signal at the receiver has minimal overshoot and ringing. The yellow waveform corresponds to the parallel resistor termination to ground and the orange waveform corresponds to the parallel resistor to some other DC voltage. These are DC parallel terminations. Compared to the series termination, these terminations result in higher power dissipation, lower signal swing, and potentially a different DC bias, because of the DC load resistor. They may also require an additional DC termination voltage to connect the resistor to. The blue waveform corresponds to the AC parallel termination with a resistor and capacitor. At high frequency (during the rising and falling edges), the capacitor is a low impedance, and the termination looks like a matching termination resistor. In addition, the capacitor blocks DC current so this termination doesn't dissipate DC power. The signal swing is larger compared to the parallel DC terminations because there is no DC load. This does require two components, however. Finally, the pink waveform corresponds to the two-resistor termination, and this is called a Thevenin termination. The resistors are chosen such that their parallel combination matches the characteristic impedance of the transmission line and such that the voltage divider sets the desired DC bias. This gives two conditions (equations) to choose the two resistor values. This termination provides DC bias without requiring an additional supply, but it also requires two components and dissipates DC power.

There are other termination and interconnect topology possibilities. We'll first consider a multi-load example as shown in figure 33 with no termination. The waveform at the load farthest from the driver (green) is perhaps the best-looking waveform. Moving back towards the transmitter, the waveform at the next receiver (yellow) looks similar to the last

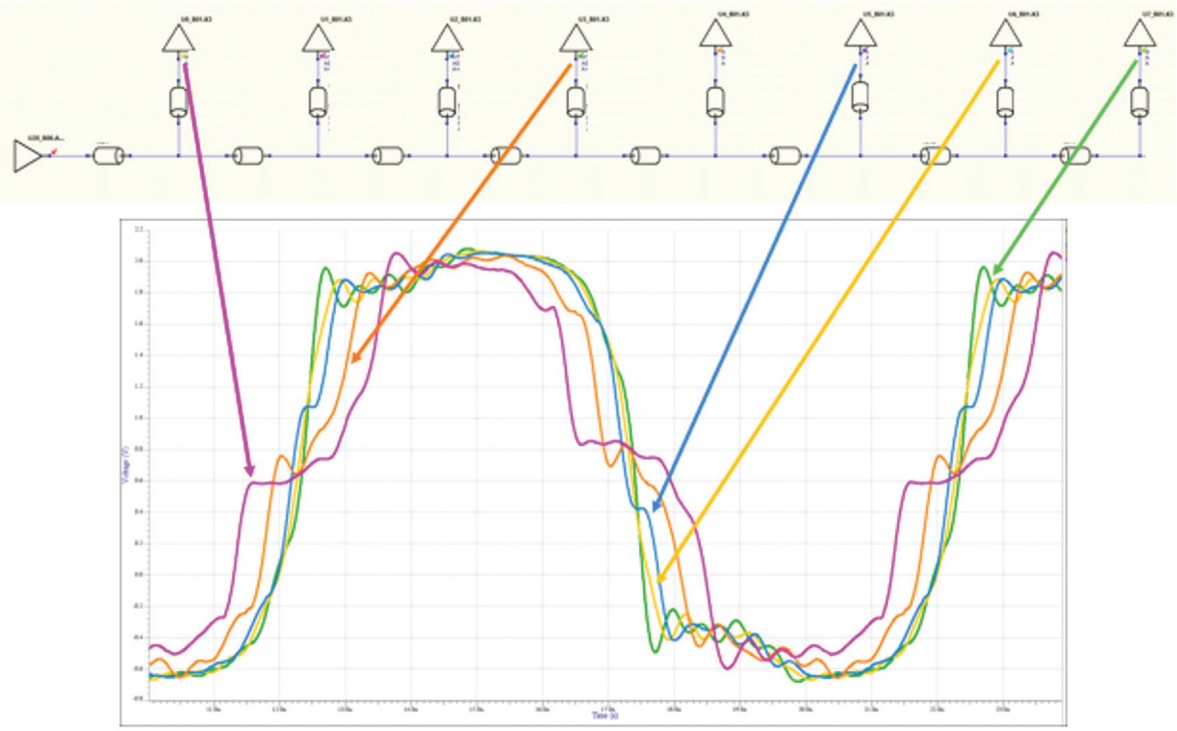

Figure 33. Multi-load example.

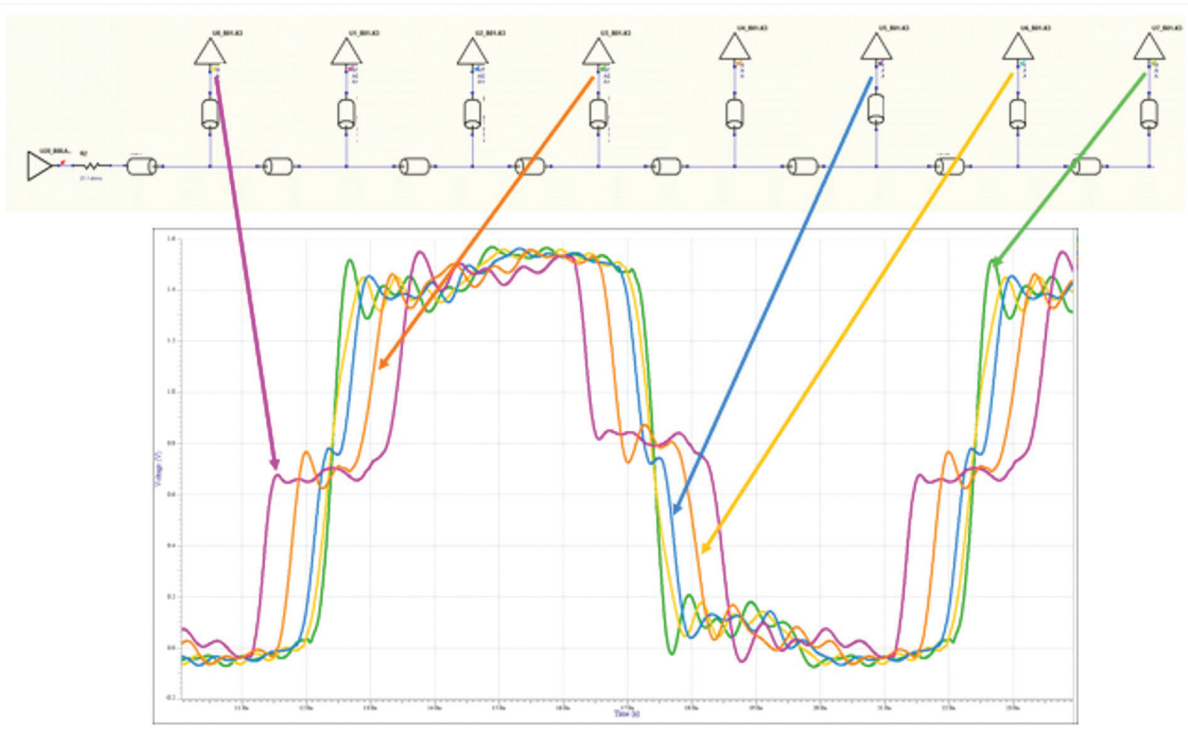

Figure 34. Multi-load example – series termination.

one (green). Continue moving closer to the transmitter and notice the waveforms look worse the closer we get to the transmitter. We start to see some non-monotonicity in these waveforms depending on when the reflections from the farther loads arrive at the loads closer to the transmitter. Even though the delay from the transmitter to the receiver closest to the transmitter (pink) is the smallest delay, the waveform there is the last to reach valid logic levels.

The waveform at the first load in the previous example looks a bit like the waveform at the transmitter we saw in our first reflection example. In figure 34, we try a series termination to see if it will work in this case. The result is that we see similar waveforms to those we saw with no termination. A series termination does not really help in this case. We still get the reflections from the loads and, even though they are terminated at the transmitter, they still affect the other loads.

Now, let's try a parallel termination as shown in figure 35. The incident signal propagates to each load and is terminated at the end after the last load so there is no reflection there. Note that the branches, or stubs, to each load are very short in this case. If they were not, we could get reflections coming from each of those branches. The waveforms at each load look good. The static voltages are a bit less due to the parallel termination to a termination voltage but the waveshapes look good. We also notice that the first load is the first to reach valid logic levels compared to when it was the last to reach valid logic levels in the unterminated and series terminated cases.

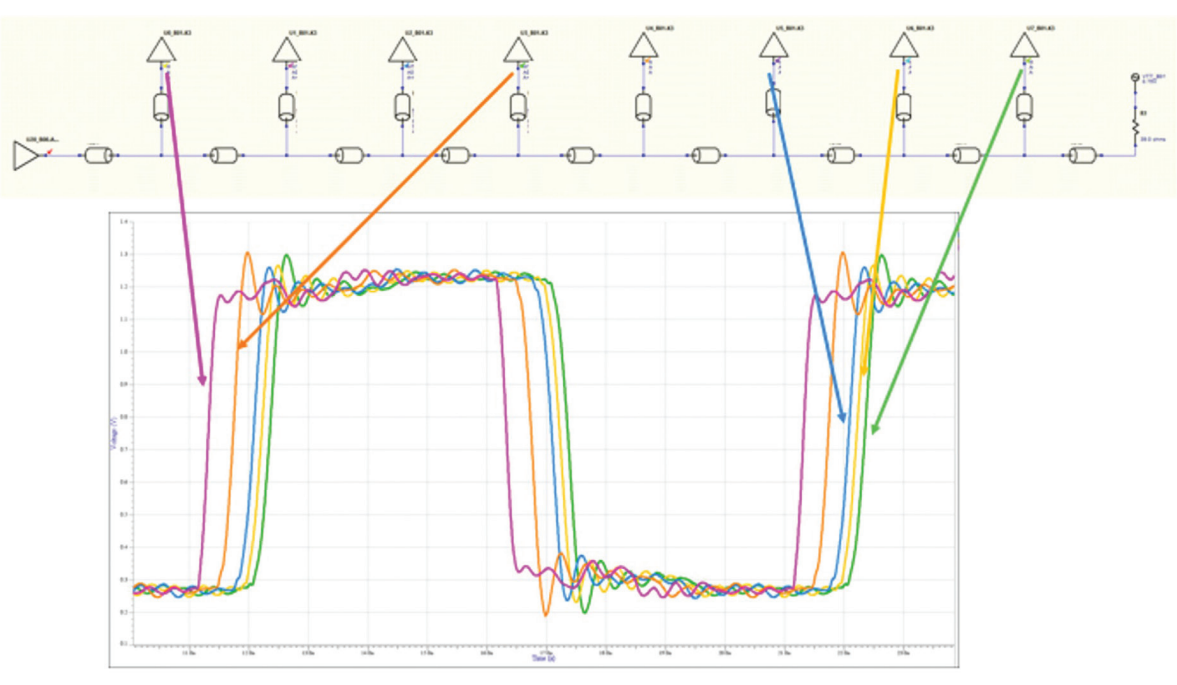

Figure 35. Multi-load example – parallel termination.

The previous example is not the only possibility of how to connect multiple loads. Though we won't cover them here, there could be others, such as star and tee topologies with longer branches and stubs.

A simulator can be used to evaluate and select the best option for interconnect topology and termination technique and value for your design needs.

Crosstalk

Crosstalk happens when we have unwanted energy coupled from one signal (we call the aggressor) onto another signal (we call the victim). Similar to the voice example in the beginning, if others are talking, then someone listening (the receiver) may not be able to hear completely, or correctly, the speaker (transmitter) they intend to hear. In signal integrity, we also consider signals interfering with other signals. In figure 36, we see some signal traces routed next to each other on a PCB.

Previously we saw there is capacitance and inductance associated with a transmission line's signal and return path. There is also mutual capacitance and inductance to other nearby metal – like other signals. This mutual reactance allows for one signal to switch and inject noise onto the other. A voltage change in the aggressor will induce a current to flow in the victim proportional to $C_m(dV/dt)$. A current change in the aggressor will induce a voltage in the victim (which will establish a current in the victim)

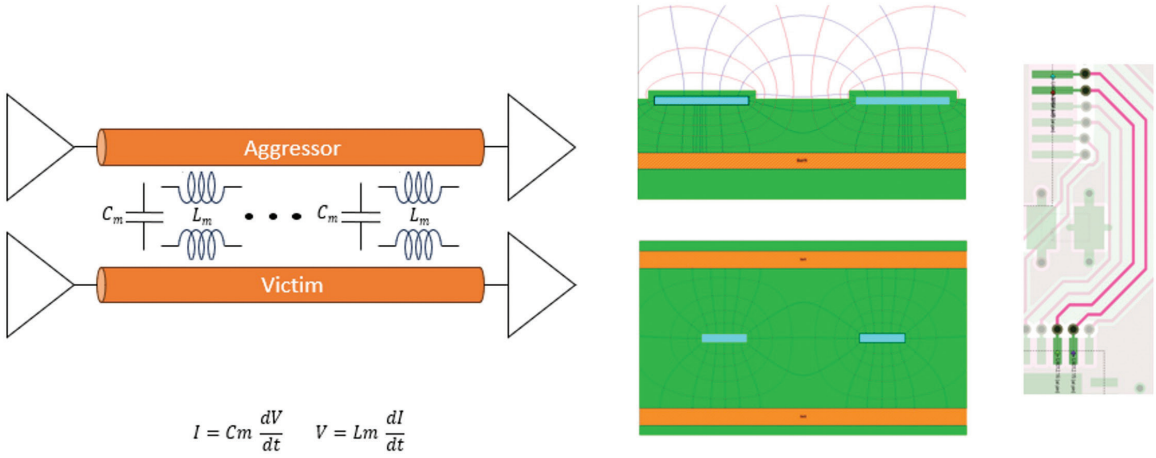

Figure 36. Coupled signals.

proportional to $L_m(dI/dt)$. The figure also shows microstrip and stripline transmission line cross sections with electric and magnetic field lines coupling between the signals.

In figure 37, we have an example with microstrip traces routed close to each other. The aggressor is being actively driven and its waveform at the receiver is the green waveform on the top. This is a series terminated example, so we have a generally clean signal with little overshoot and ringing. The victim is present but is not switching so the only voltage changes we observe on the victim are from crosstalk. The voltage at the victim's receiver is the red waveform and the voltage at the victim's driver is the blue waveform. Even though this is a series terminated example, there are reflections at the receiver ends that affect the crosstalk waveforms. The reflected aggressor signal can cause crosstalk in the victim and the crosstalk voltage on the victim can reflect off the receive end. If there were no series terminations, the crosstalk would be even more pronounced.

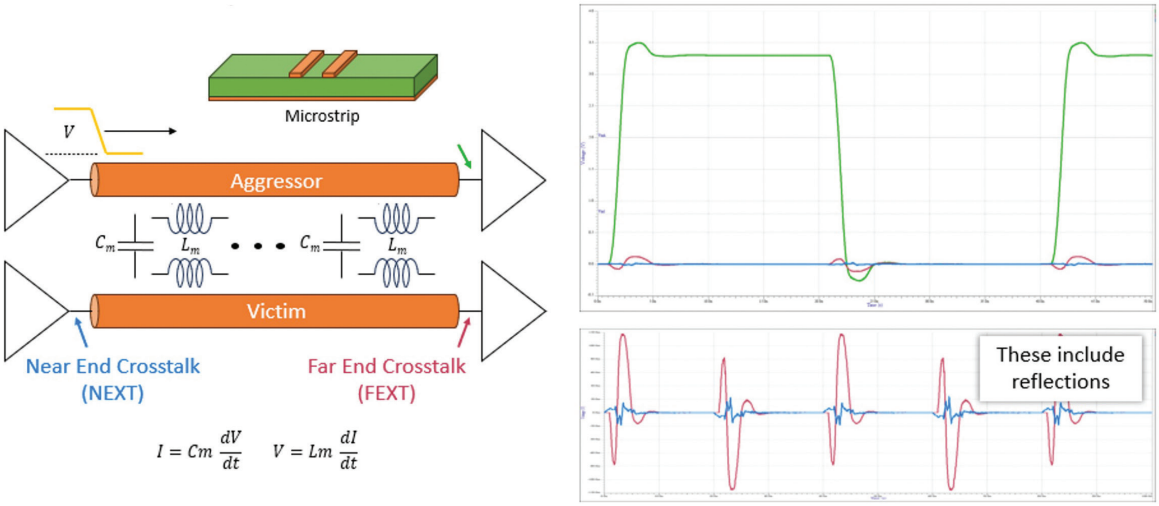

Figure 37. Crosstalk example.

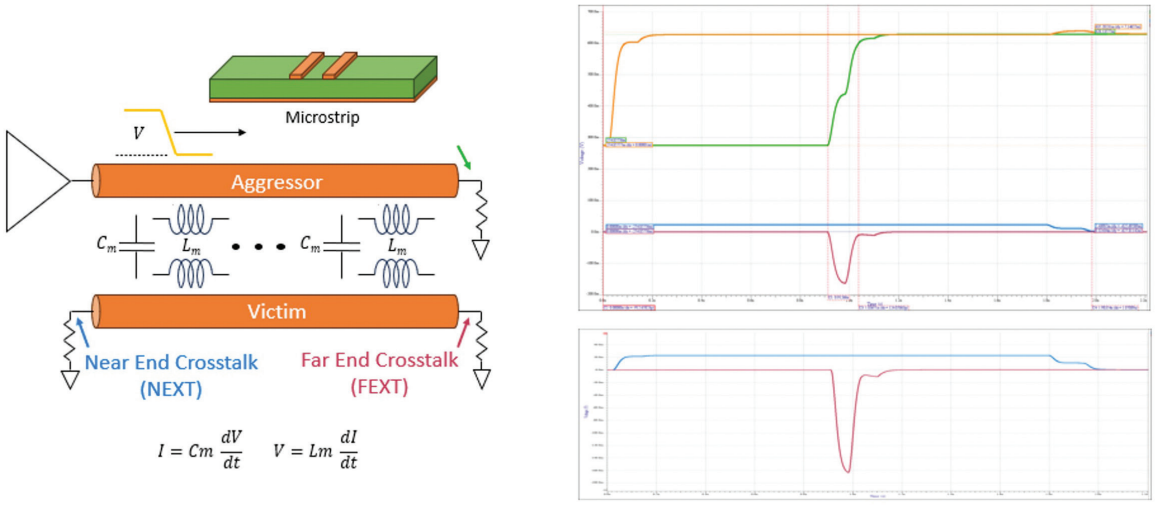

Figure 38. Crosstalk example – NEXT/FEXT.

Figure 38 shows a similar microstrip case. This driver has faster rise/fall times (with no series termination) and the far end of the transmission line was terminated to eliminate reflections. Both ends of the victim transmission line were also terminated to eliminate reflections. The top waveforms are the aggressor signal at the transmit (yellow) and receive (green) ends. Energy is coupled onto the victim through the mutual capacitance and inductance. That coupled energy then propagates backward and forward on the victim transmission line and arrives at the near and far ends as near end crosstalk (NEXT) and far end crosstalk (FEXT). These are the blue waveform and red waveforms, respectively. The bottom waveform view is a zoomed in view of the NEXT and FEXT. The FEXT and NEXT waveforms have different profiles. NEXT increases to a fixed level in a time relative to the aggressor's rise/fall time and then remains at that level for a time twice that of the transmission line delay. FEXT increases in voltage with longer coupled length and arrives at the far end at the same time the aggressor signal arrives at the far end. Its time duration is the same as the aggressor's rise/fall time.

Now consider what happens if we space the microstrip traces farther apart as shown in figure 39. Both the NEXT and FEXT have the same profile, but their amplitudes have been reduced by moving them away from each other.

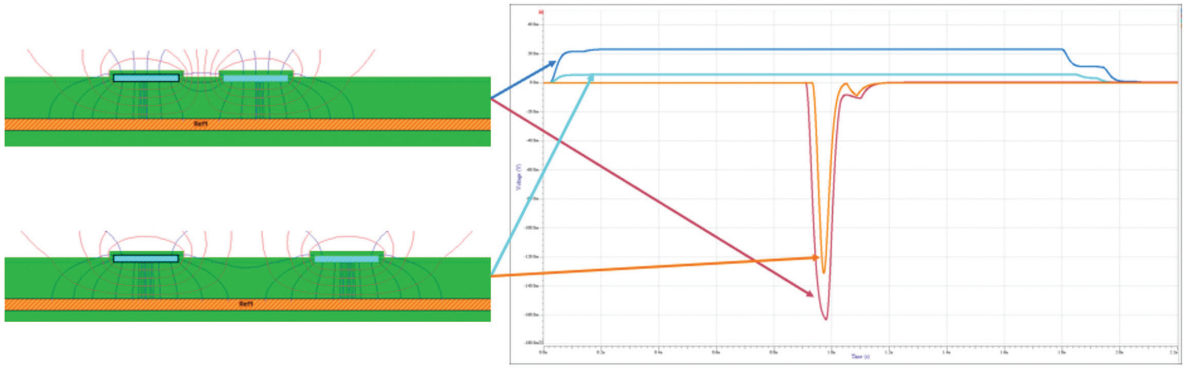

Figure 39. Crosstalk example – trace spacing effect on NEXT/FEXT.

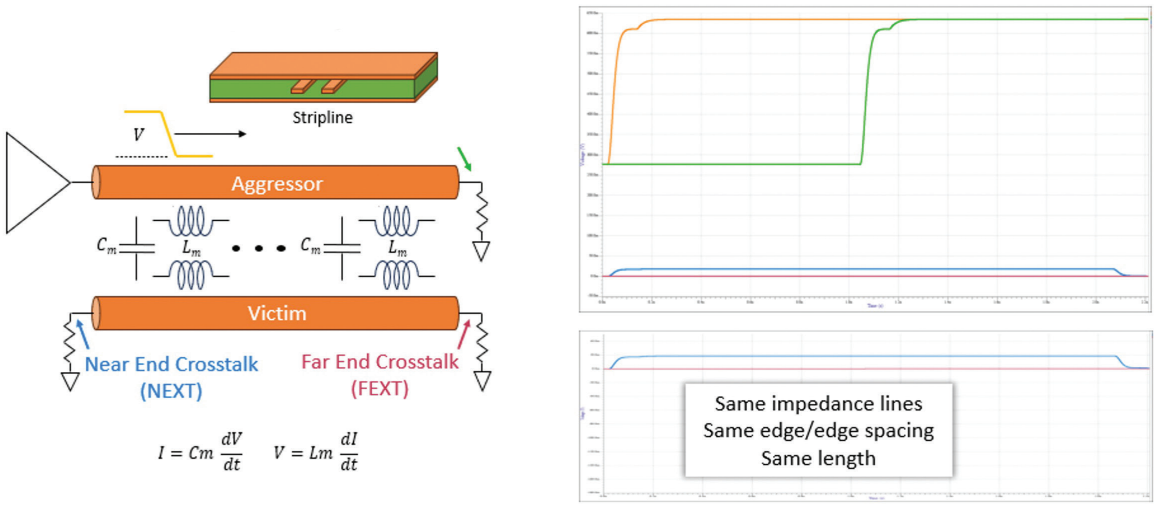

Figure 40. Crosstalk example — NEXT/FEXT in stripline.

Notice in figure 40 that using stripline transmission lines can significantly reduce FEXT. We're using the same impedance lines with the same spacing and same length but on stripline this time. Notice now there is no FEXT. A homogeneous stripline environment eliminates FEXT (forward crosstalk).

Other changes could also be considered to reduce the crosstalk. For example, the coupling between the signals could be reduced by spacing them farther apart, minimizing their parallel coupled length, using guard traces between them, or moving the signals closer to the reference plane. Slower rise/fall times could also reduce crosstalk. Implementing crosstalk mitigation/management strategies is critical for signal integrity and can help ensure reliable signal transmission. While considering these techniques, keep in mind that they have other impacts. For example, moving signals farther apart or using guard traces takes more space on the PCB and decreases routing density. Using thinner dielectrics so the signals are closer to the reference affects the characteristic impedance, and this needs to be considered. Figures 41 through 44 show crosstalk examples while varying trace spacing and coupling length. Figure 41 is a microstrip example where the trace separation is kept constant, and the transmission line length is varied. Figure 42 is a microstrip example where the transmission line length is kept constant, and the trace separation is varied. Figure 43 is a stripline example where the trace separation is kept constant, and the transmission line length is varied. Figure 44 is a stripline example where the transmission line length is kept constant, and the trace separation is varied. In all of these cases notice that decreasing the coupling length and increasing the trace separation decreases the crosstalk.

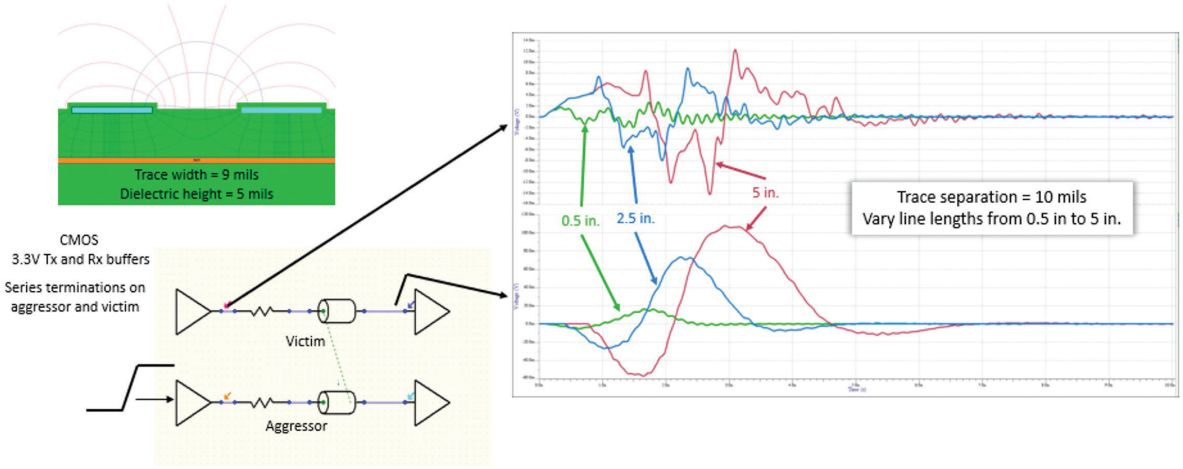

Figure 41. Crosstalk example – microstrip with constant separation and varying length.

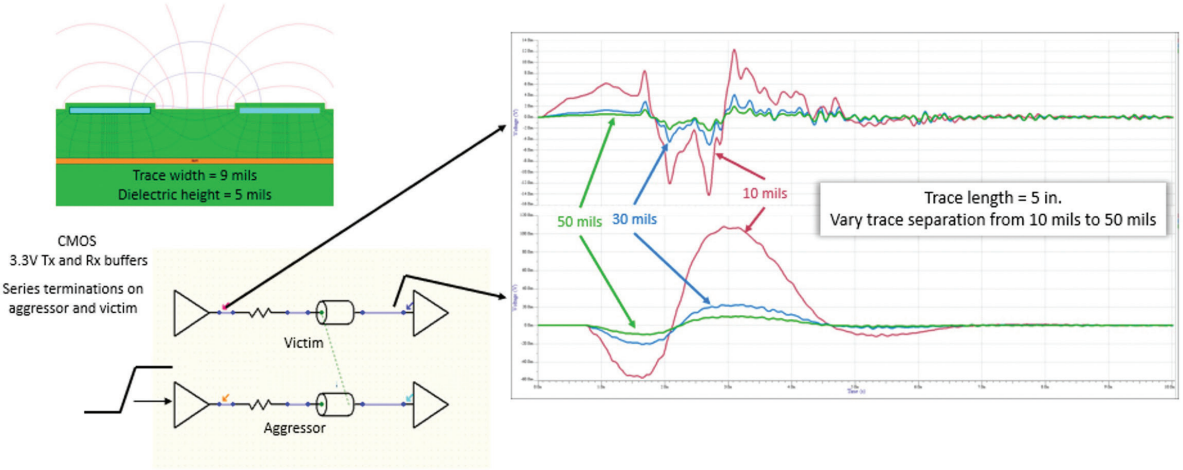

Figure 42. Crosstalk example — microstrip with constant length and varying separation.

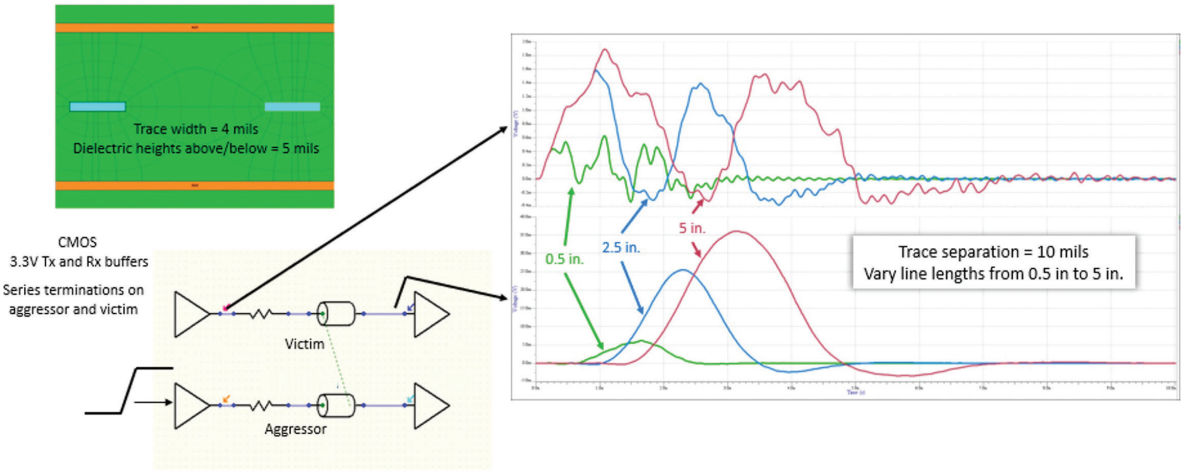

Figure 43. Crosstalk example – stripline with constant separation and varying length.

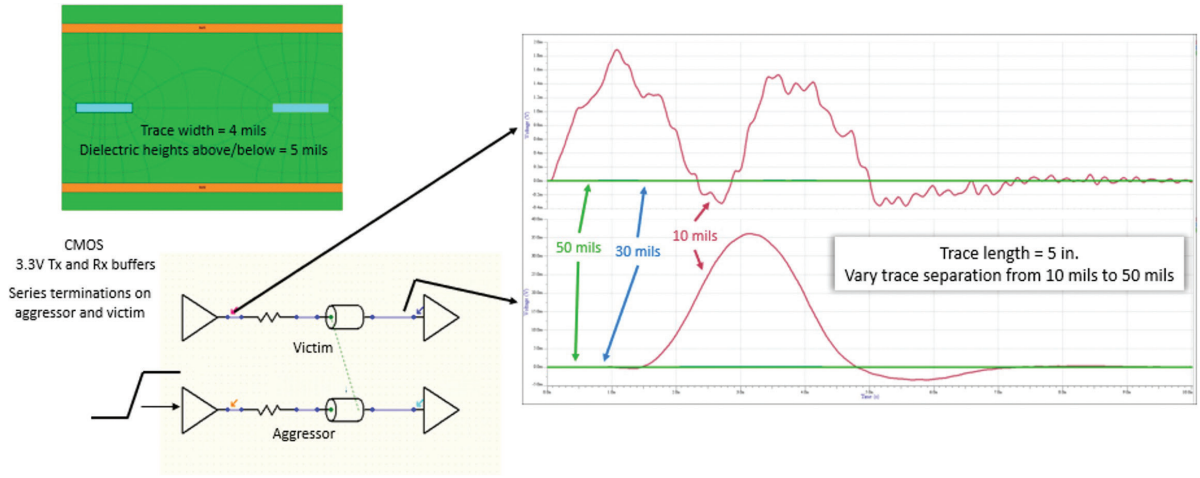

Figure 44. Crosstalk example — stripline with constant length and varying separation.

Differential pairs

Until now we've focused on single-ended signals but now consider differential signals and differential pairs which are another critical concept in signal integrity. First, we consider differential signals as shown in figure 45. V_{pos} has special significance in its relation to the differential signal, but it is a single ended signal because its voltage is relative to the return. Similarly, V_{neg} has special significance in its relation to the differential signal, but it is a single ended signal because its voltage is relative to the return. These signals have an intended relationship with each other and together they make up a differential pair. We call the difference voltage between them the differential voltage (V_{diff}), or the differential signal. Notice the differential signal is biased at 0 V. There is also a common signal shared

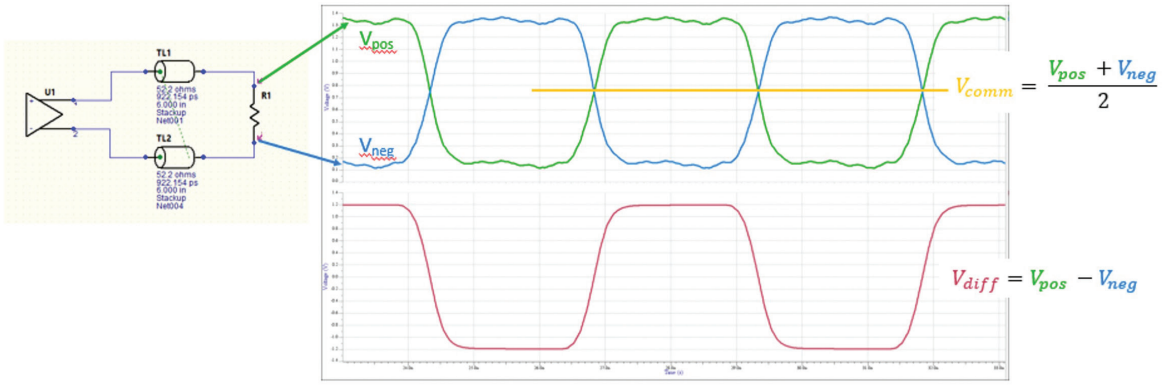

Figure 45. Differential signals.

by the two single ended signals (V_{comm}), which we usually want to be a constant. For well-balanced signals, the common signal is at the bias of the single-ended signals. If those signals are not well-balanced, the common signal can vary. As an aside, when we have elements in the channel carrying a differential signal that are imbalanced, we get what is called mode conversion as some of the differential signal is converted to a common signal. The details are beyond the scope here, but it's an important consideration and generally something to be avoided.

Some of the coupling concepts we learned in the previous chapters will help us here regarding differential pairs. Typically, the signals making up a differential pair are routed together and near each other on a PCB, connector, or cable. Similar to the discussion on crosstalk, there is coupling between those signals. This is shown in figure 46 generically as mutual impedances Z_{12} and Z_{21}. There are two important ways, or modes, in which we could drive these transmission lines. One mode is called the even mode and the other is called the odd mode. The even mode is when we drive the same signal onto both transmission lines. The odd mode is when we drive opposite signals on the transmission lines. In figure 45 we are driving in the odd mode as can be seen by the V_{pos} and V_{neg} signals.

When a signal is coupled to another, the characteristic impedance of each signal is influenced by the other. This influence depends on how close they are to each other and how they are being driven. Qualitatively, we can understand this considering the coupling between signals as described in the discussion of crosstalk and by remembering that impedance is the ratio of voltage to current. For example, if we drive a signal onto an isolated transmission line, the voltage on the transmission line will be the characteristic impedance of line times the current in the line. However, if we have coupling with another signal, the cross section changes due to the nearby metal and there is coupled current/voltage between the lines through the mutual impedance.

The odd mode impedance, Z_{odd}, is the characteristic impedance of one line when the pair is driven in the odd mode. The even mode impedance, Z_{even}, is the characteristic impedance of one line when the pair is driven in the even mode. When there is no coupling, $Z_{odd}=Z_{even}=Z_0$. The differential impedance is the impedance seen by the differential signal and is twice the odd mode impedance ($Z_{diff} = 2Z_{odd}$). For loosely coupled (minimal coupling between signals) 50 Ω transmission lines, the differential impedance would be approximately 100 Ω. The common impedance is the impedance seen by the common signal and it is half the even mode impedance ($Z_{com} = Z_{even}/2$). As the amount of coupling between the signals increases, the odd mode impedance and the differential impedance decrease, while the even mode impedance and common impedance increase.

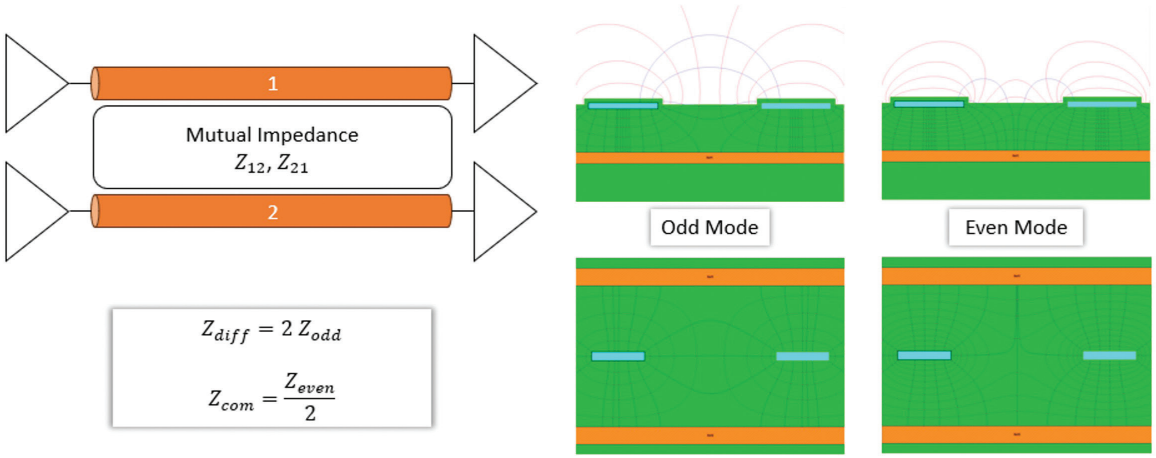

Figure 46. Differential pair impedances.

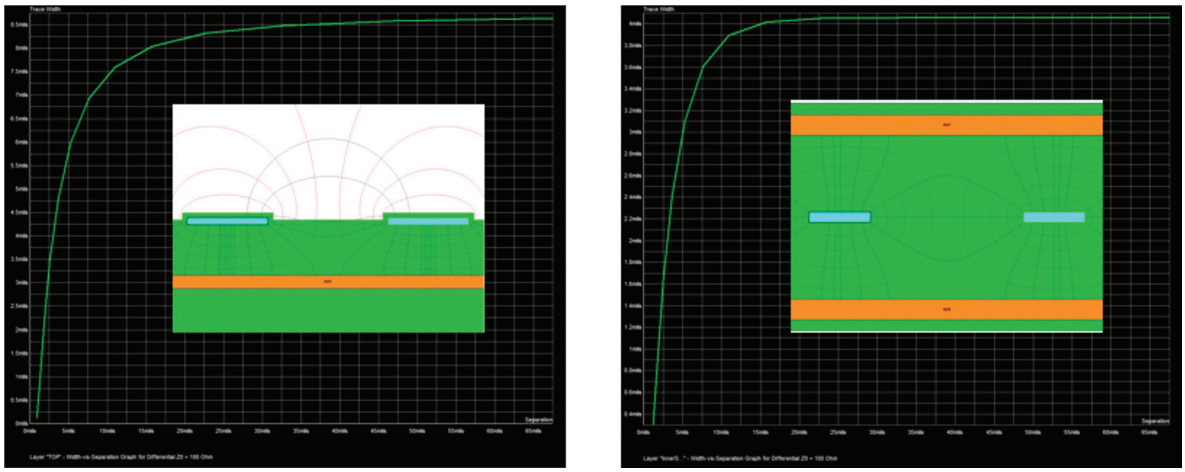

Figure 47. Differential impedance – trace width vs. separation..

Figure 47 shows cross sections of stripline and microstrip differential pairs, along with plots of trace width and trace separation to achieve 100 Ω differential impedance. Trace separation is on the horizontal axis and trace width is on the vertical axis. Anywhere along the green curves represents 100 Ω differential impedance. Starting in the lower left corners, the separation is small, and the trace widths are small. This is the area with greatest coupling. The trace width and separation have a strong dependence on each other on this part of the curve as can be seen by the steep slope there. If we move along the curve to the right, the trace separation increases so the trace width also increases to maintain 100 Ω. Moving farther to the right, the slope of the curve starts to decrease as the coupling is less (loosely coupled). Continuing to the right the slope flattens off and this is the region with very little or negligible coupling. In this part of the curve, we can continue to the right with even more trace separation with little to no need to change the trace width to maintain 100 Ω.

Figure 48 shows some examples of termination possibilities for differential pairs. These examples are not meant to cover every possible case but to get you thinking about some of the possibilities and to help solidify the concepts of common and differential signals. The left-most example is a differential termination with a resistor equal to the differential impedance. This will terminate a differential signal nicely. However, it doesn't terminate a common signal. The middle example uses separate resistors on each signal equal to the odd mode impedance. This is similar to the parallel termination strategies discussed with single-ended signals. The odd mode impedance is used for the individual resistors because we are intending to have a differential signal and are therefore driving the signals in the odd mode. The right-most example uses three resistors. The two resistors that are equal to the odd mode impedance are connected together and, combined in series, are like the single differential resistor. Remember Z_{diff}=2 Z_{odd} so a differential signal will see a termination equal to the differential impedance. The third resistor is used to terminate a common signal. The even mode signal will see the parallel combination of Z_{odd} in series with this resistor and that equals Z_{com}. So this last termination technique can terminate both a differential and a common signal. There are other termination possibilities. Using a simulator can help to evaluate the best option for your design needs.

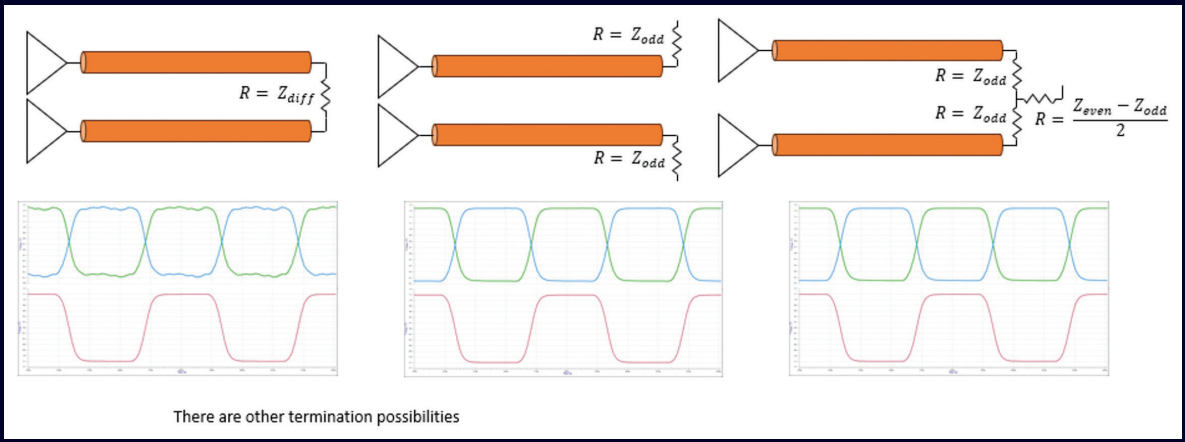

Figure 48. Differential pair termination examples.

Vias and impairments

Now we'll shift our focus to some other common impairments to signal integrity. In particular, we'll look at vias, return path discontinuities, loss, and power integrity. We begin with vias.

We can't review all via types and implications here but will raise some common ideas, so you have some foundation later to analyze and consider your applications. In figure 49 we have an example of a basic via. While not all vias are created equal, this one has some typical elements. The lower-right picture is the top view, and the center picture is a cross section. The signal via is the red via in the figure that allows the signal to change layers in the PCB. The green vias are called stitching vias and they connect metal on the reference nets together between layers. This signal via (and stitching vias)

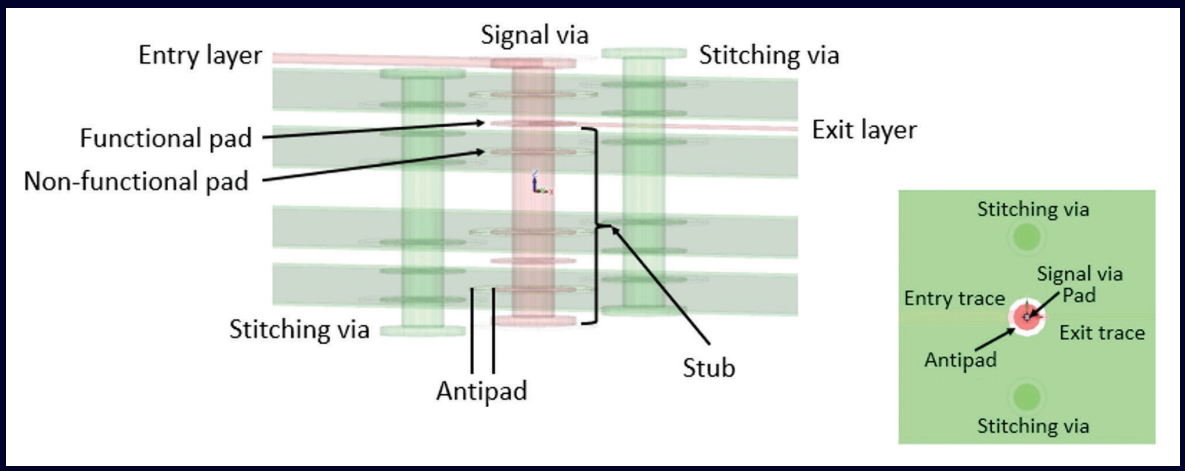

Figure 49. Basic via structure.

have pads on all layers in this case. Pads that connect to something on a layer are called functional pads and pads that do not connect to something on a layer are called non-functional pads. Antipads are holes in the reference planes surrounding the signal via separating signal and reference. A stub is the part of the via that extends beyond the layers on which the signal has connections. There are entry and exit traces on the entry and exit layers where the signal connects. Notice in the top view of the example that the entry and exit traces have different widths. This is to achieve the same characteristic impedance on both layers which have different geometries and materials.

We're going to examine the performance of this and other vias in both the time and frequency domains. In the frequency domain we will use S-parameters. In the time domain we will use TDR and eye diagrams. We discussed eye diagrams earlier. We'll very briefly describe S-parameters and TDR here before analyzing via performance.

S-parameters

S-parameters, or scattering parameters, are a set of frequency-domain measurements used to characterize the behavior of electrical networks and are frequently used in signal integrity analysis. S-parameters use the concept of a port which is a location where the signals are injected and measured. They describe the relationship between input and output signals of a device or a network by quantifying how much signal comes out of each port relative to what is injected into another port, for each frequency. That is, when a signal is injected to one port, how much of that signal comes out of all the other ports, including the input port (reflected signal). In signal integrity analysis, S-parameters are important for understanding how signals propagate through interconnects including reflection, transmission, and crosstalk. S-parameters are represented as matrices of complex numbers, typically denoted as S_{ij}, where i is the observation port where the response is measured, and j is the excitation port where the stimulus is injected. In our via examples, we only have two ports. One is connected to the entry trace and the other is connected to the exit trace. S_{11} is a measure of reflection and is the ratio of signal that comes out of port 1 when we inject a signal into port 1. It is often called Return Loss. The same principle applies to S_{22} relative to port 2. S_{21} is a measure of transmission and is the ratio of signal that comes out of port 2 when we inject a signal into port 1. It is often called Insertion Loss. The same principle applies to S_{12} with the excitation/observation ports reversed. As complex numbers, S-parameters contain magnitude and phase (or real and imaginary parts). For this analysis, we will only be looking at magnitude and we use the dB scale.

If we had an ideal, lossless interconnect then everything we inject into port 1 would come out of port 2. In that case the ratio of output signal to input signal would be 1 and that would mean S_{21}=0 dB. And this would be for all frequencies. Because there would be no reflections, the ratio of reflected signal to input signal would be 0 and that would mean S_{11}=-∞ dB. If we had a transmission line with loss but still with no impedance discontinuities, then S_{21} would be something less than 0 dB and, in a typical PCB, the value decreases with increasing frequency. Because there are still no impedance discontinuities and no reflections, we would still have S_{11}=-∞ dB.

TDR

Time Domain Reflectometry, or TDR, is a powerful technique used in signal integrity to analyze interconnects and identify issues such as impedance discontinuities, their values, and locations along the interconnect. As the name implies, it is a reflected measurement. It works by injecting a very fast edge into the interconnect and then measuring reflections that come back. Knowing the source reference impedance and the size of the reflections relative to the incident voltage step, impedance discontinuities can be calculated. Also, knowing the time at which the reflections arrive back at the source, the location

of the discontinuities can be determined. If we know a time delay and the propagation velocity of the interconnect, then we can calculate a length. We have to remember that since this is a reflected measurement, the TDR delay is the round-trip delay so the actual length to the discontinuity will be half of that calculated using the round-trip delay.

Now, let's return to the basic via example and examine its performance as shown in Figure 50. We see at low frequencies, S_{21} is close to 0 dB and S_{11} is close to $-\infty$ dB, so most of the energy is transmitted through and very little gets reflected back. At higher frequencies we see some increasing loss (decreasing S_{21}) and we also start to see S_{11} increasing because some energy is being reflected back to the source. We also have a dip in S_{21} at approximately 15 GHz. Without that, the loss would be constantly decreasing at the same rate over frequency. The dip at approximately 15 GHz is caused by the via's stub (the length of the signal via extending below the third layer) which acts like a quarter-wave resonator. The via stub length is a quarter of the wavelength at this frequency. As the signal arrives at the stub, energy propagates down the stub one quarter wavelength which is 90 degrees. Then it encounters an open circuit at the bottom of the via so the signal arriving there reflects with a 100% reflection coefficient. The reflected signal then propagates one quarter wavelength back up the via stub which is another 90 degrees. The reflected signal has now traveled 180 degrees (90 degrees down + 90 degrees up) and arrives at the incident signal point 180 degrees out of phase with the incident signal. It cancels the incident signal, and we get no (or very little) transmission at that frequency.

Note that the via structure is in between short 50 Ω transmission lines. These transmission lines are known length and impedance and are just added for convenience so we can clearly see where the via begins and ends in the TDR plot. In the TDR plot, the flat lines at 50 Ω on the right and left represent those transmission line segments. In the middle we see the via's impedance. As a reminder, the flat sections are approximately 400 ps long each because we have the round-trip delay for transmission line segments with 200 ps delay one way. Approximately 400 ps into the plot we start to see a change from 50 Ω where the via is. The impedance initially decreases. A decreasing impedance indicates a capacitive discontinuity because an increase in capacitance causes a decrease in impedance. After the dip, the impedance increases which indicates an inductive discontinuity because an increase in inductance causes an increase in impedance.

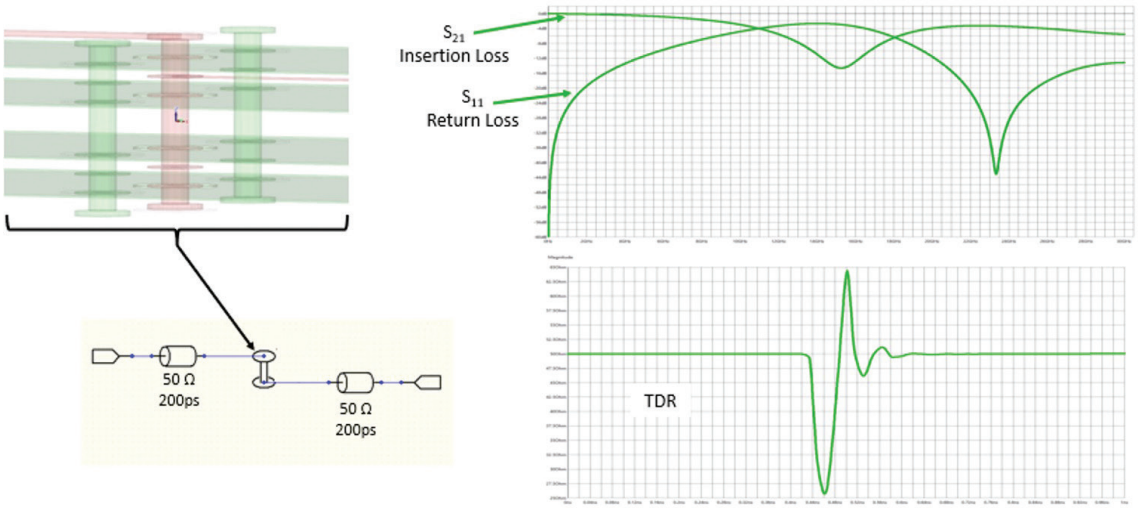

Figure 50. Basic via performance.

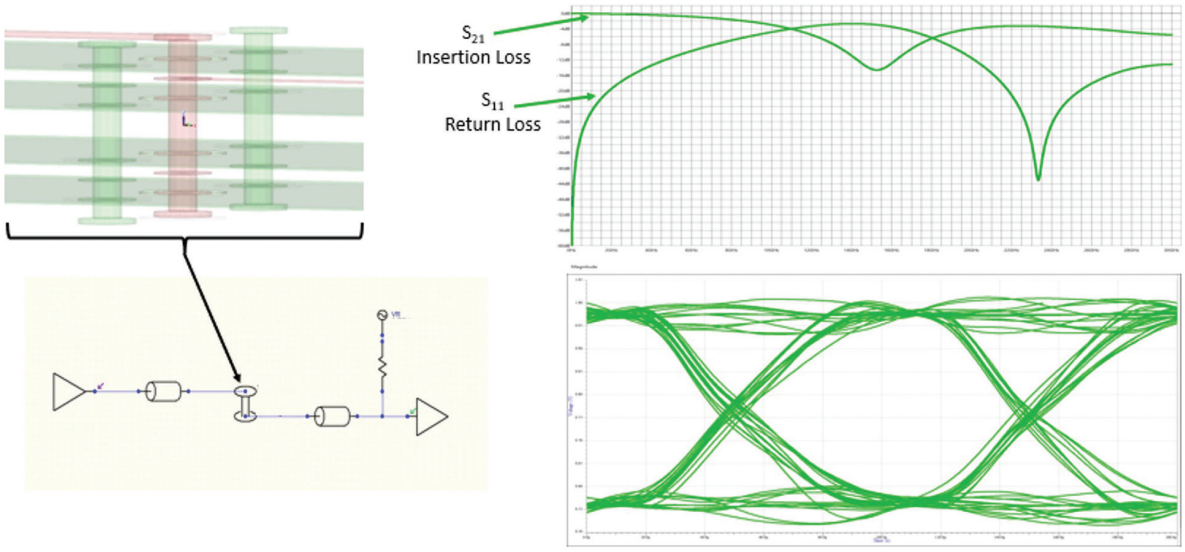

Figure 51. Basic via performance with eye diagram.

Observe the eye diagram for this via in figure 51. We won't comment on it here but will use it as a reference for comparison in future examples. We will see if we can improve the via's performance and the corresponding S-parameters, TDR, and eye diagram by making changes to some of the basic via elements. In all of the following comparisons, the performance of the original via will be plotted in the green color shown here so we can easily compare the changes.

Figure 52 shows the original basic via just analyzed along with four other vias that we will analyze. Each via includes one or more modifications to the original via. The modifications include removing non-functional pads, using larger antipads, using a different exit layer farther down the via, and removing the via stub. We'll compare the effects of these changes one-by-one.

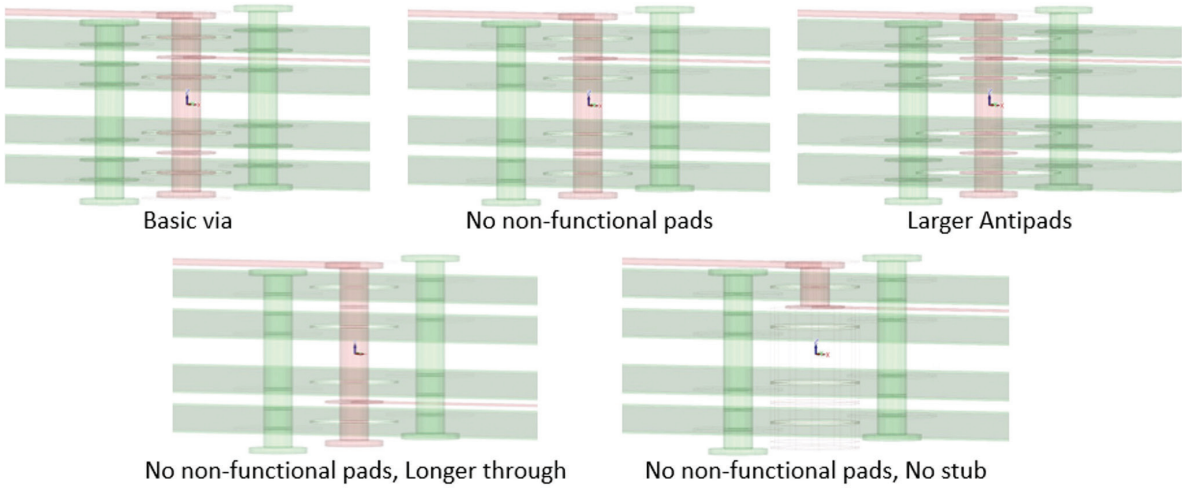

Figure 52. Original basic along with four via modifications.

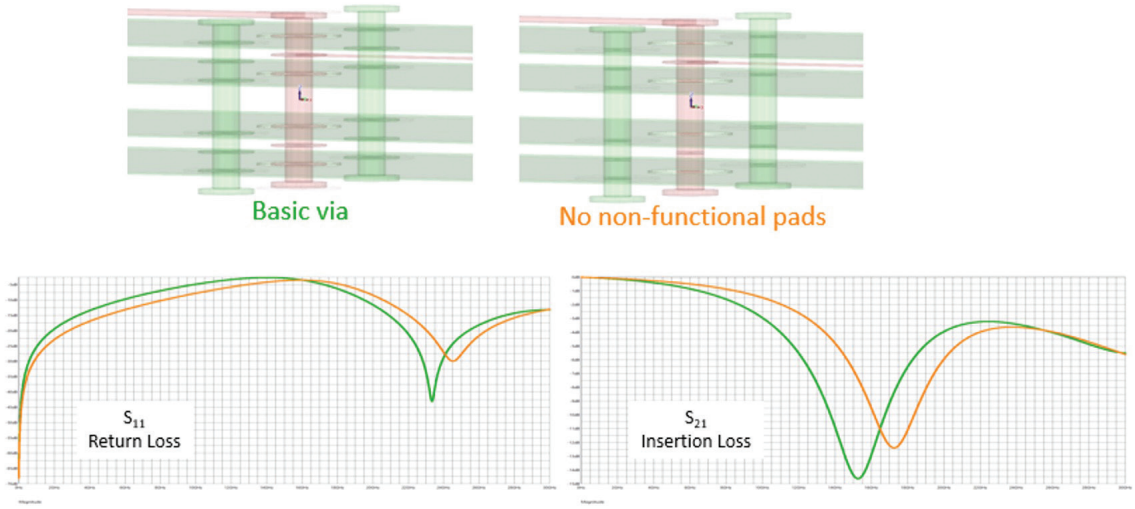

Figure 53. Via performance comparison 1.

The stitching vias will be consistent for all of these vias. In general, stitching via presence and location does make a difference in via performance and signal integrity. We'll see that briefly in an example later but for now, the stitching vias will be constant.

The first experiment is to remove the non-functional pads as shown in figure 53 and compare the S-parameters with those of the original via. Again, the original via performance is plotted in green and the new via performance is plotted in orange. By removing the non-functional pads, we have a noticeable improvement in return loss and insertion loss (at least for a large part of the frequency range). The resonant dip seen in S_{21} has been pushed out to a higher frequency. You could imagine a case where that simple change could be enough of an improvement to pass your signals. There could also be reasons – perhaps mechanical reliability – to keep the non-functional pads. But at least we can see some effect they have and use that information to help make engineering decisions and tradeoffs.

We're going to continue comparing the various via modifications mentioned earlier. We'll stick with the frequency domain first and then come back to compare the time domain performances. The original via performance will still be plotted in green and the modified via's plots will each have a new color.

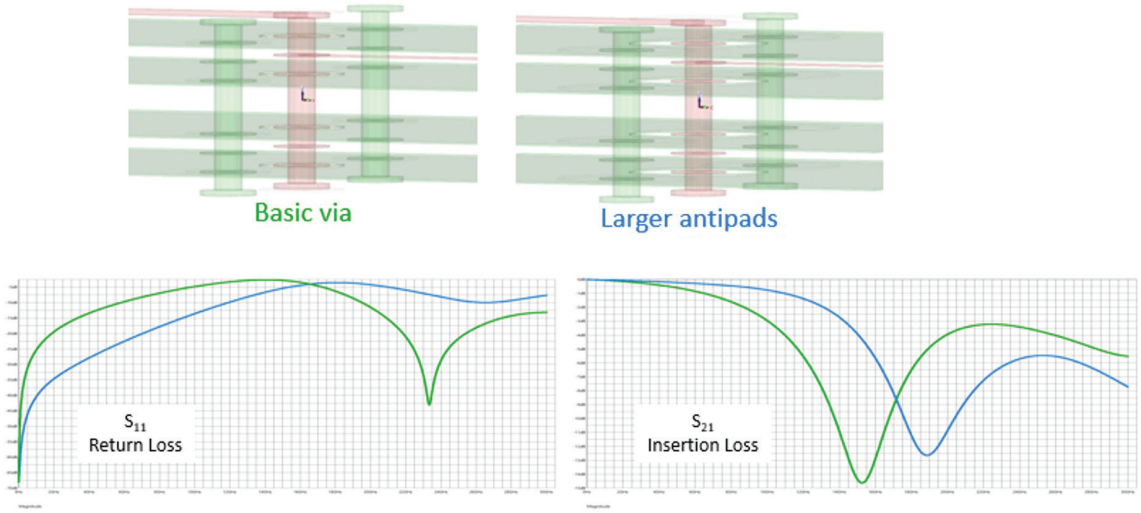

Figure 54. Via performance comparison 2.

The modified via in figure 54 has the non-functional pads but it now has larger antipads. It performs better than the via without the non-functional pads. By making the antipads larger, we have a noticeable improvement in return loss and insertion loss. The resonant dip in the insertion loss has been pushed out to an even higher frequency.

The new modification in figure 55 moves the exit trace to a layer farther along the length of the via so we have a longer via transition and a shorter stub. This via has the original sized antipads, non-functional pads are removed, and the exit layer is now deeper in the board. It performs even better than the previous two vias. Again, there is a noticeable improvement in return loss and insertion loss and the resonant dip in the insertion loss has been pushed out to an even higher frequency. Notice that the via stub is much shorter in this case which causes the quarter wave resonance to be at a higher

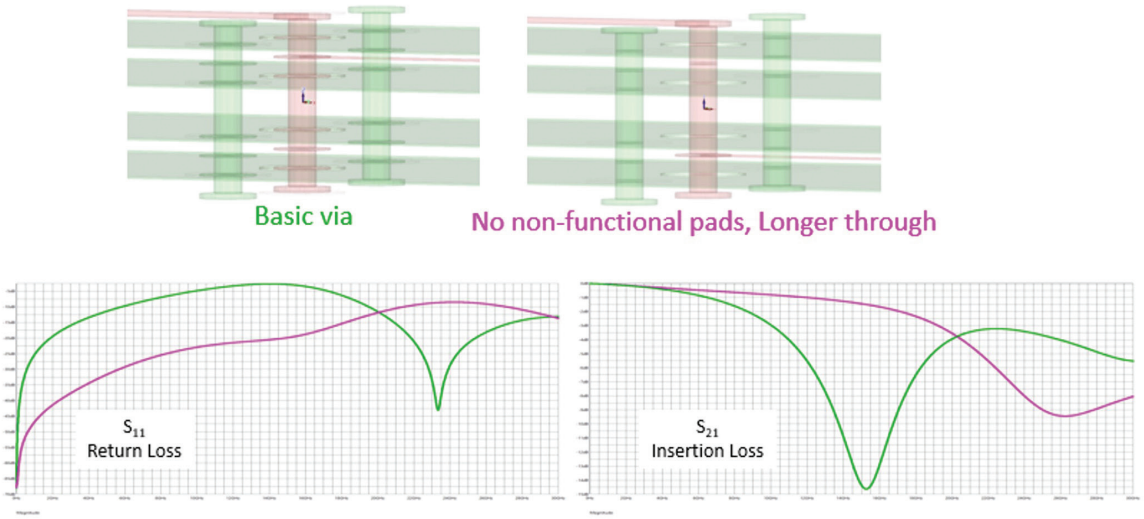

Figure 55. Via performance comparison 3.

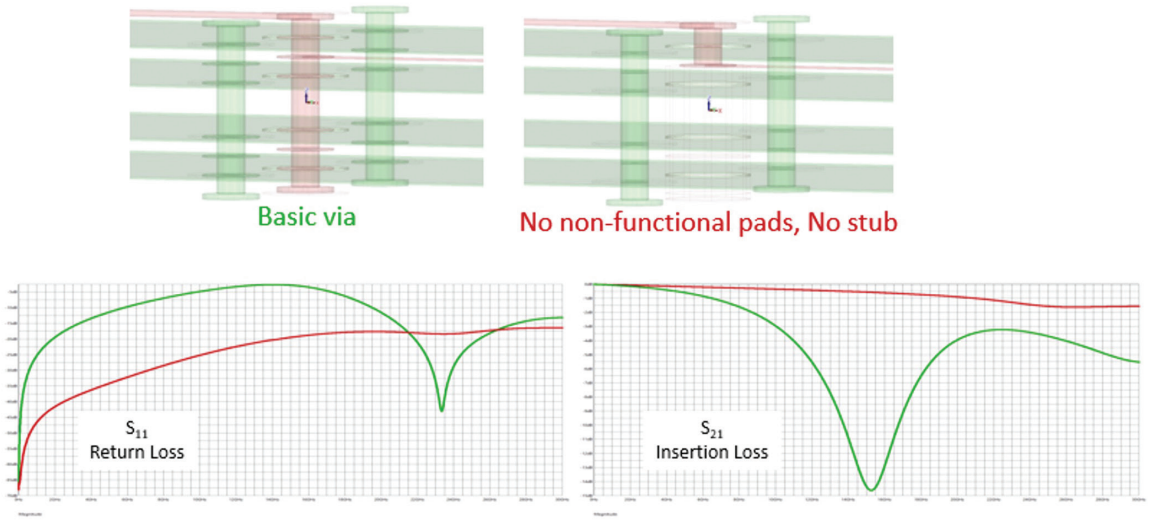

Figure 56. Via performance comparison 4.

frequency. If your signal had required frequency content in the frequency range of the S_{21} dip from the original via, then those frequencies could pass through this new via with some loss but not be largely blocked.

The last via modification we'll examine is shown in figure 56. This new via has the original sized anti-pads, non-functional pads are removed, the original exit layer, but we have backdrilled and removed the via stub. It performs better than all of the previous vias with better return loss and insertion loss. The resonant dip is essentially eliminated which makes sense because we no longer have a stub on the via.

Now we'll review the performance of the same vias in the time domain with eye diagrams and TDR plots. In figure 57 we're comparing the original via with the via that has the non-functional pads removed. The original via performance is plotted in green and the new via performance is plotted in orange. Notice the improvements in the TDR and eye diagram plots. The TDR shows less impedance deviation from 50 Ω. The eye diagram improvement

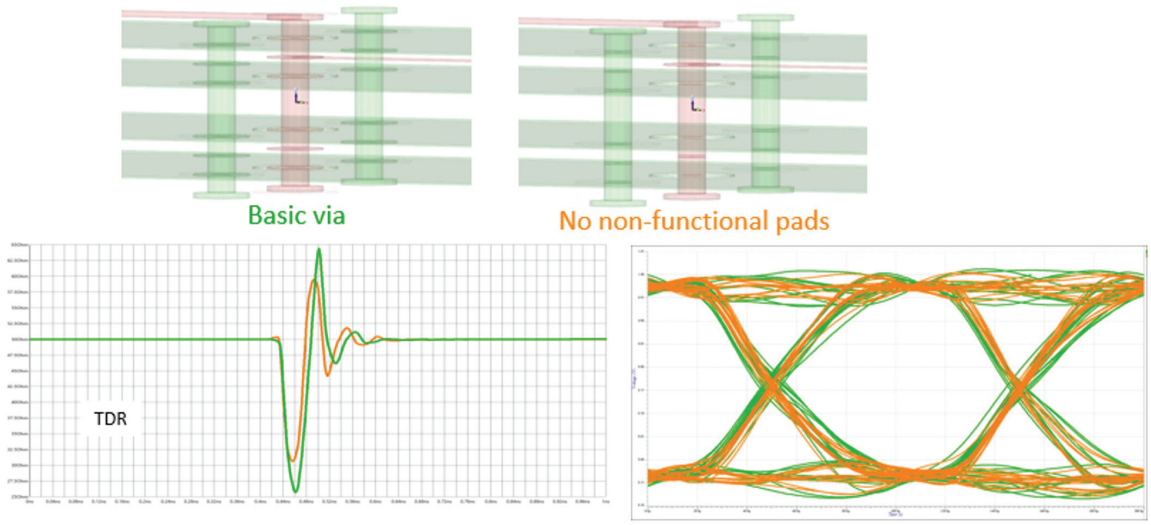

Figure 57. Via performance comparison 5.

is most notably seen in having less time variation at the crossing points.

We're going to continue comparing the performance of the various via changes in the time domain, as we did in the frequency domain. The original will still be plotted in green and the modified via's plots will each have a new color.

Figure 58 shows the time domain performance of the via with larger antipads. It performs a little bit better than the via without the non-functional pads. The TDR plot at the via now starts out inductive, rather than capacitive. The eye diagram is improved again as we can see the time variation improvement at the crossing points. There is also improvement in the amplitude as we can see thinner bundles (less ringing) at the high and low levels.

Figure 59 shows the time domain performance of the via with the original sized antipads, non-functional pads are removed, and the exit layer is now deeper in the board. It performs even better than the previous vias. There is a much lower impedance variation in the TDR and the eye diagram is much cleaner. There is very little time variation at the crossing points and very little amplitude variation (ringing) at the high and low levels.

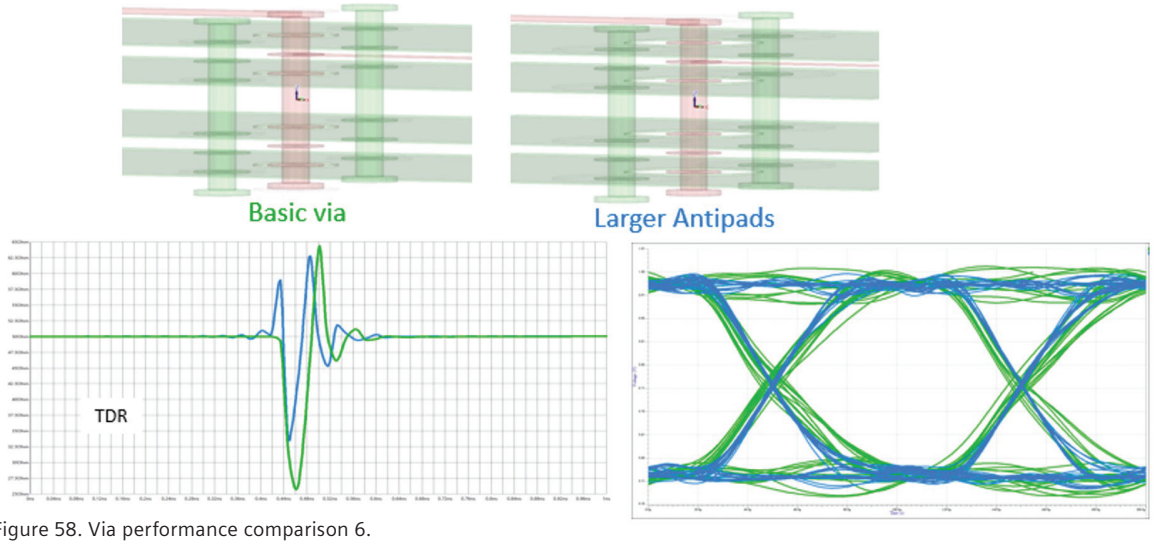

Figure 58. Via performance comparison 6.

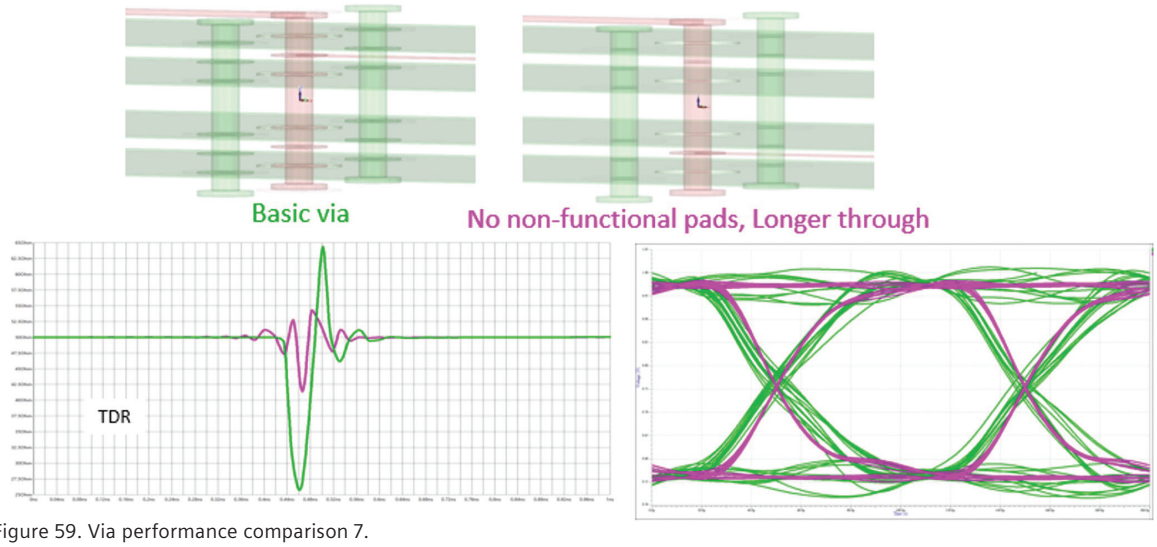

Figure 59. Via performance comparison 7.

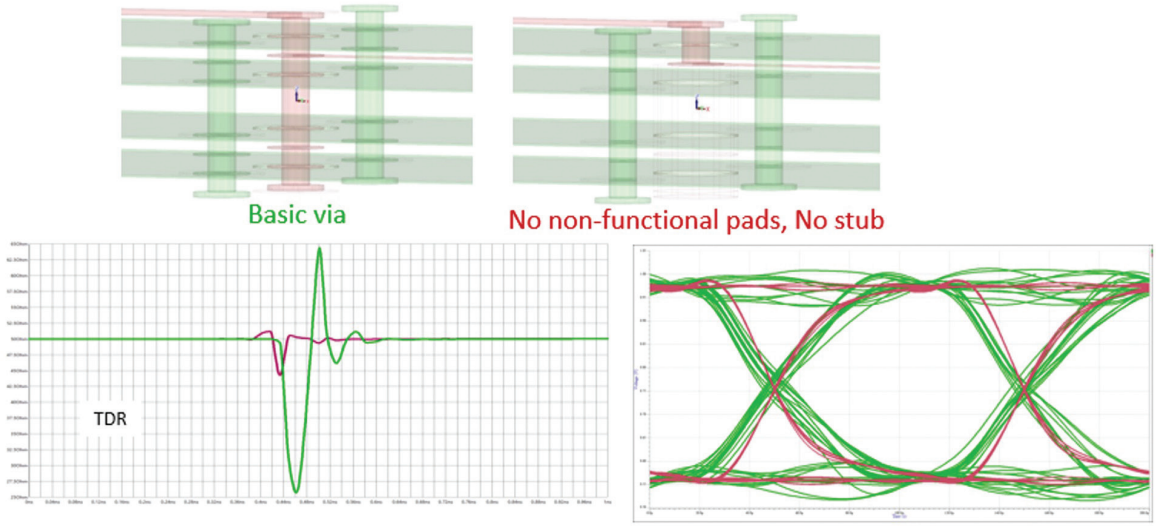

Figure 60. Via performance comparison 8.

Finally, we can examine the performance of the via that has the original sized antipads, non-functional pads are removed, the original exit layer, but we have backdrilled and removed the via stub, as shown in figure 60. It performs better than all of the other vias. There is very little impedance variation seen in the TDR plot. In the eye diagram there is very little time variation at the crossings and very little amplitude variation on the high and low levels. With smaller impedance discontinuity, we expect less reflections and ringing and therefore a cleaner eye diagram, which is what we see here.

There are other things we could consider like the via drill and pad sizes (even micro, blind, and buried vias) and various combinations of these. But the intent here was not to be comprehensive with designing good via transitions. The intent was to introduce the concepts of how via performance and various via elements can affect signal integrity. Also, there could be tradeoffs to be made in designing vias for a given PCB. There may be layer and routing density restrictions, mechanical reliability, cost, and other considerations that have to be taken into account when designing and implementing vias.

Using a simulator is a powerful way to help design good performing vias and make design tradeoffs. It should be noted that for lower frequency applications, an algorithmic approach to via modeling may be sufficient to make design decisions. But for high frequency applications, a full-wave 3D electromagnetic field solver may be required to make accurate trade-offs on a via design.

A few additional considerations for differential via structures are shown in figure 61. We aren't going to be exhaustive, review experiments, or compare performances. The intent here is simply to expand the concepts that were discussed with single vias to differential via structures. Regarding antipads, they could be implemented as individual antipads around each signal via or as a common antipad surrounding both signal vias, and with different size gaps. The spacing between signal vias is a consideration, as is the presence and location of stitching vias. The other elements discussed with single vias are also relevant for differential vias, such as non-functional pads, layer transitions – entry/exit layers, stubs, backdrilling, blind/buried micro vias, pad size, and drill size.

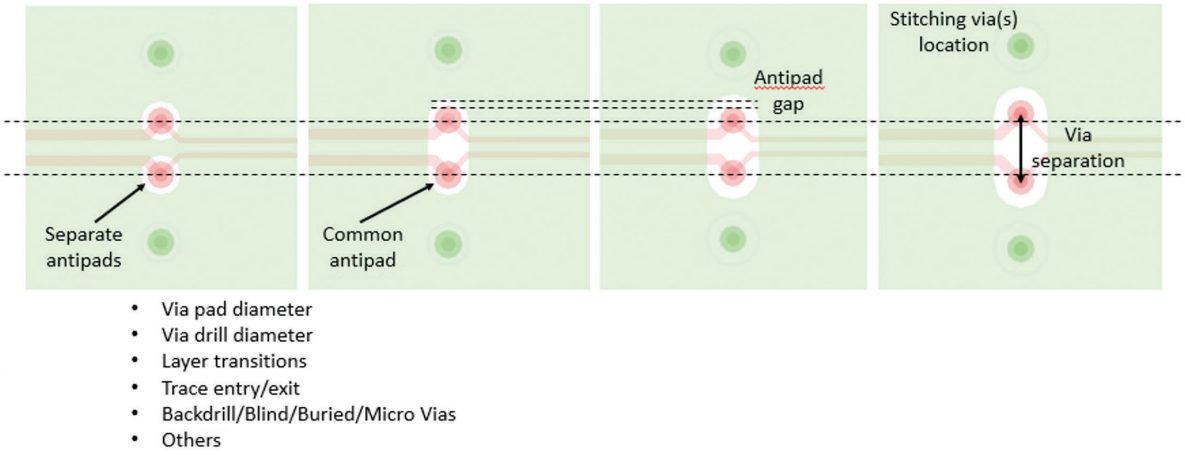

Figure 61. Additional differential via considerations.

To conclude the topic of vias, observe the example in figure 62 which shows the effect of stitching via location. As mentioned previously, the presence and location of stitching vias can signal affect signal integrity. They provide a path for return current to change layers when a signal changes layers. This is mostly a qualitative example that will also help us begin to transition to the next topic signal integrity impairments related to return path discontinuities. This example has 4 layers where the top and bottom layers are signal layers, and the two middle layers are reference layers. There is one signal that transitions through a via between the top to bottom layers.

The figure plots signal and return currents at 100 MHz. Dark blue represents no current and the other colors show us where current is flowing. It is zoomed in the z-direction so we can see inside between the layers better. The top-left picture has a stitching via close to the signal via. Notice the return current flows on the return layers, closely following the signal path and transitions through the stitching via. The bottom-right picture has the stitching via placed far from the signal via. Notice the return current shown by the green/yellow color on the return layers deviates far from the signal path to reach and transition on the stitching via. With a structure like this, the image helps provide a visual example of how return current could flow to unexpected or unwanted places resulting in degraded signal quality as well as other electromagnetic problems.

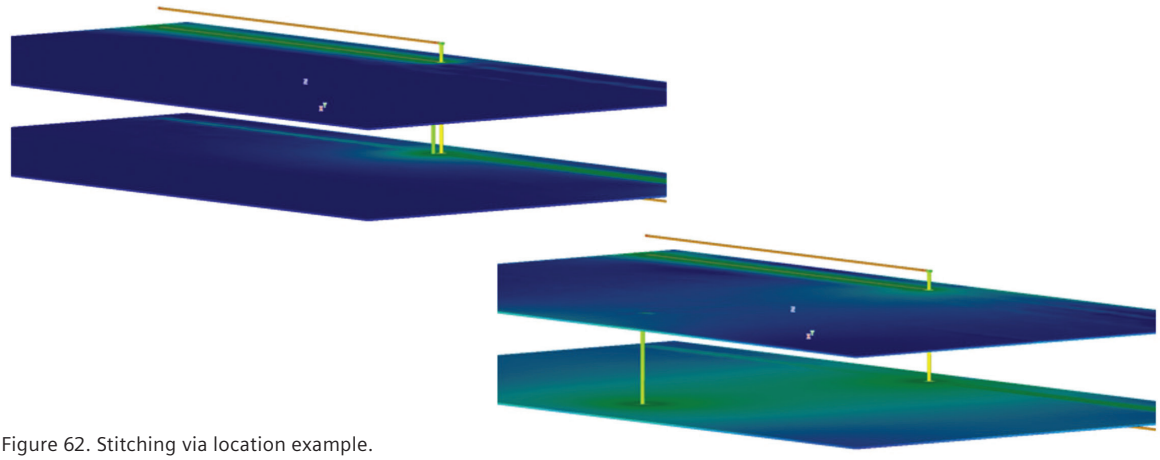

Figure 62. Stitching via location example.

We will now consider two additional examples of return path discontinuities as signal integrity impairments. First, observe another mostly qualitative example where signal traces cross a gap in the return path as shown in figure 63. On the left we see several signals crossing a gap – three of which are highlighted in red. The solid green areas are two area fills on an internal layer in the board. The signals are on an adjacent internal layer. One area fill is a power supply net and the other area fill is a filtered power supply net. They are connected together by a ferrite bead shown in the lower right corner of the slot. The ferrite bead is placed on the bottom of the board and is connected to the area fills with vias.

The signal on the far left is excited with a current. The vertical colorful bars are the port currents and can be ignored for this discussion. Dark blue represents no current and the other colors show us where current is flowing. The figure plots 100 MHz current on a few adjacent layers and is zoomed in the z-direction so we can see inside better.

Notice in the top right picture, there is current on the vias connecting the ferrite bead and through the ferrite on the bottom layer of the board. Because the signals are crossing the gap, the return current deviates from closely following the signal path. It flows on one of the reference area fills to the vias and ferrite bead to transition to the other reference area fill and get back to closely following the signal path on the other side of the gap. We can see from this example that when there is an interruption in the return path, return current could flow to unexpected or unwanted places.

In the picture on the bottom right, the structure was modified to close the gap with a piece of metal, so the signals have a consistent reference. Notice in this case, the vias and ferrite have essentially no return current because now there is no longer any break in the return path and the return current can closely follow the signal path.

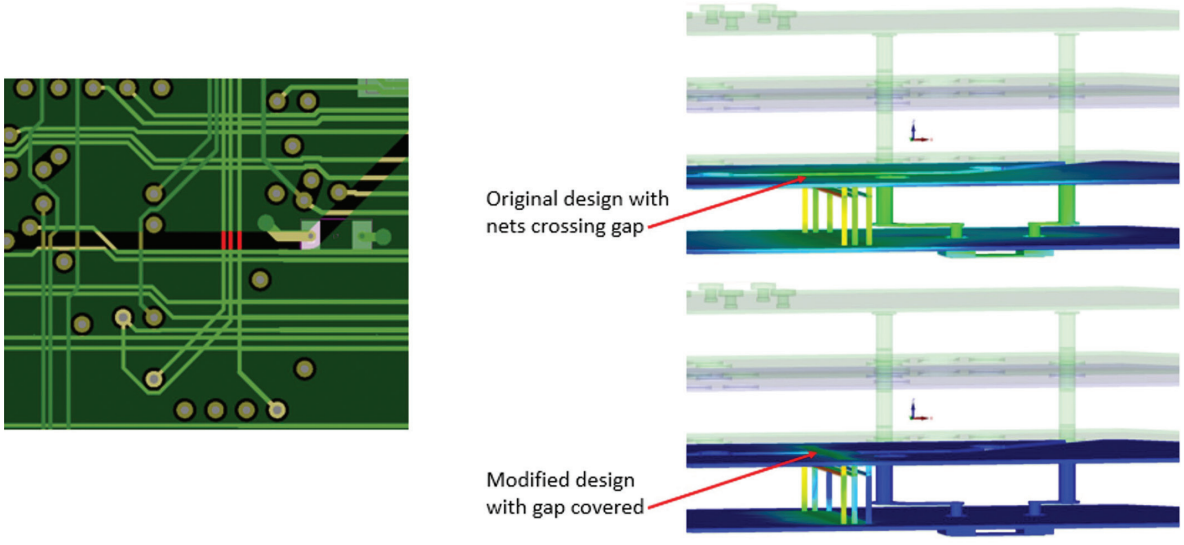

Figure 63. Return path gap example.

For the last example of return path discontinuity, see figure 64. This is an example with two signal traces crossing breaks in the return path. The signal traces themselves are short – only approximately 0.2 inches long. The picture on the left shows a top view and the two pictures on the right show two different variations of the structure. Only one of the signals is excited with a current stimulus. It is the top signal in the top-view picture. The signal is driven from the pin on the left to the pin on the right. The right pin of the driven signal is terminated. The undriven signal is terminated on both ends (this can be seen schematically in figure 65).

In the original case shown in the top right picture of figure 64, the signals cross a gap in the reference plane on the layer directly beneath them. There is a solid reference layer below the split layer. There are return vias very close to the signal end points. On the right side, the return via connects to the area fill just below the signal and also to the solid layer below that. The return via near the signals on the left side only connects to the closest reference area fill beneath the signals and not to the solid layer. There is a third reference via on the far left that stitches the area fill on the left to the solid plane beneath. Dark blue represents no current and the other colors show us where current is flowing. We can see a lot of current on the solid return plane and on the stitching vias connecting the split fills to the solid return plane. This is zoomed in the z-direction so we can see inside better. The signal is being driven only from the pin on the left of the slot a short distance to the pin on the right of the slot. But return current shows up on the left side of the left area fill, through the far left return via, across the solid plane, and through the far right return via.

In the bottom right picture of figure 64, the gap in the top reference layer has been closed so it is now solid and there is no longer a break in the reference under the signals. Under the same excitation conditions, notice there is no (or very little) current on the bottom reference layer or in the return vias connecting the returns. In this case, the return current stays directly below the signals on the top reference layer.

These were again plotted using 100 MHz current, but the return current path can, in general, be frequency dependent, as the current follows the path of least impedance. However, this example, illustrates again the concept that when there is a break in the return current path, the return current could flow to unexpected or unwanted places.

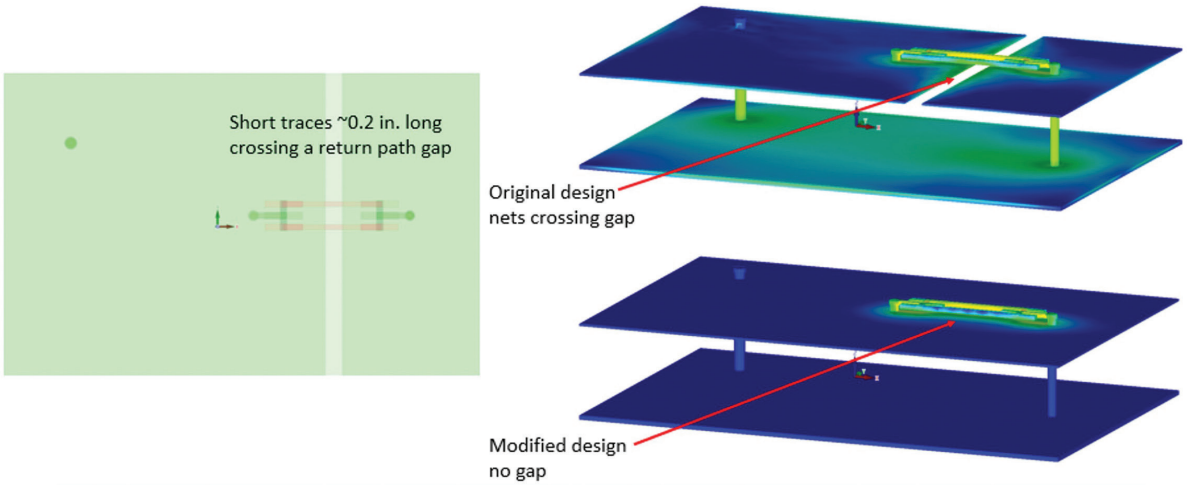

Figure 64. Return path gap example 2.

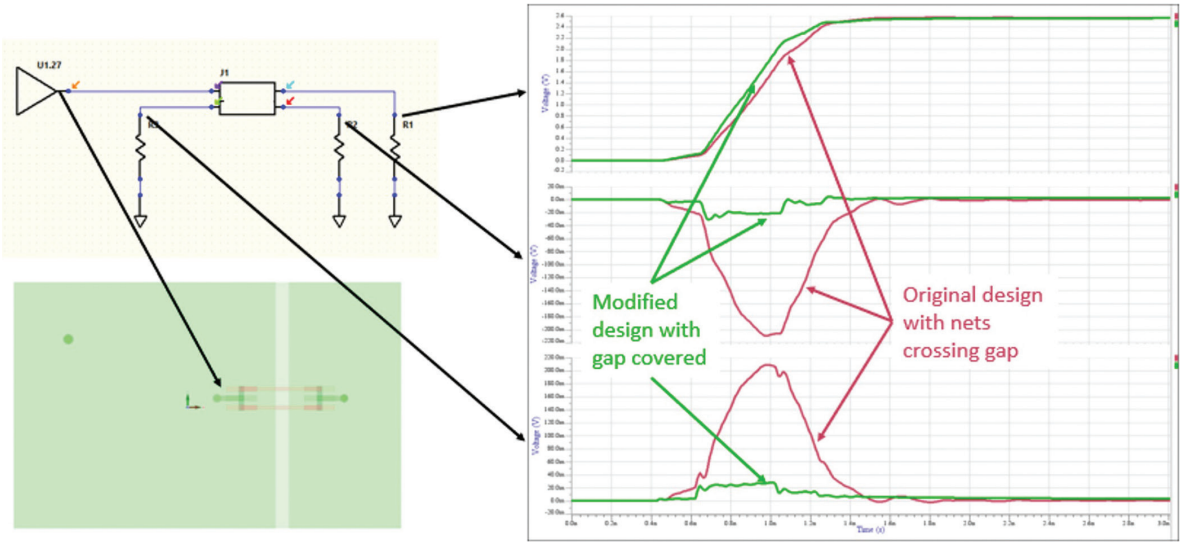

Figure 65. Return path gap example 2 – time domain results.

Now let's look at the previous examples with the short signals crossing the gap and when there is no gap in the time domain. See figure 65. Instead of driving the signal with 100 MHz current, it is being driven now with a rising edge (from the left side). The right end of the driven net is terminated and both ends of the undriven net are also terminated. Observe the signals at the various locations. The red waveforms are for the case with the gap in the reference. The green waveforms are for the solid reference case. The top waveforms are the waveforms at the right side of the driven net. The middle waveforms are at the right side of the undriven net. The bottom waveforms are at the left side of the undriven net. We may not normally expect crosstalk for signals with a very short coupling distance but notice the voltage showing up on the undriven net when there is a gap in the reference plane. When there is no gap, we have much less voltage on the undriven net. In this example there is approximately 10 times more voltage on the undriven net when there is a gap in the reference compared to having a solid reference. Notice also the rise time of the driven signal at the receiving end is affected (slowed down) by the broken reference. In this simple example we see effects (impairments) on a driven signal and also on other signals by having a return path discontinuity.

Another consideration in signal integrity is loss. Loss can often be an impairment, especially for long interconnect and high frequency signals. For short interconnect, or low frequency signals, or where some damping is needed, loss may not be a concern.

Looking at the plot on the left of figure 66, we see loss for a transmission line plotted as attenuation in dB per unit length. That loss has conductive and dielectric contributors. Conductive loss comes from the resistance of the metal used in the interconnect. It increases with frequency as the skin depth decreases and resistance increases with increasing frequency. There is loss associated with the dielectrics as well and this loss also increases with frequency. In signal integrity we're concerned with signals with wide (broadband) and high frequency content, and we see here that higher frequencies experience higher loss while lower frequencies experience lower loss. This is the low pass filter effect as described in the section on Intersymbol Interference. In addition, dielectric material properties vary over frequency such that the dissipation

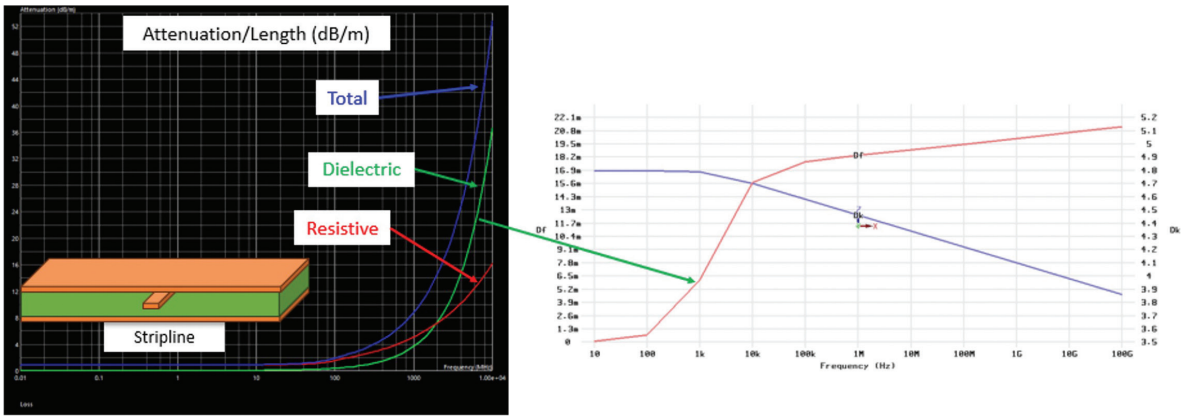

Figure 66. Loss – conductive and dielectric loss vs. frequency.

factor Df (or loss tangent) increases and the dielectric constant (Dk) decreases with frequency as shown in the picture on the right of the figure.

In figure 67 we have an example showing how loss is affected by transmission line length. We're looking at insertion loss (S_{21}) to compare the impact length has on loss. We have three different transmission line lengths, which are 1 inch, 6 inches, and 12 inches. Everything else is the same between the cases and we can see big differences in loss depending on the length. At 10 GHz, for example, the 1 inch trace experiences about 1dB loss (~89% of the transmitted signal arrives at the far end). The 6 inch trace experiences about 7.5dB loss (~42% of the transmitted signal arrives at the far end). The 12 inch trace experiences about 17.5dB loss (~13% of the transmitted signal arrives at the far end). Note that this loss behavior is specific to the materials and trace dimensions used in this example to demonstrate how the effects of loss are cumulative with increased length.

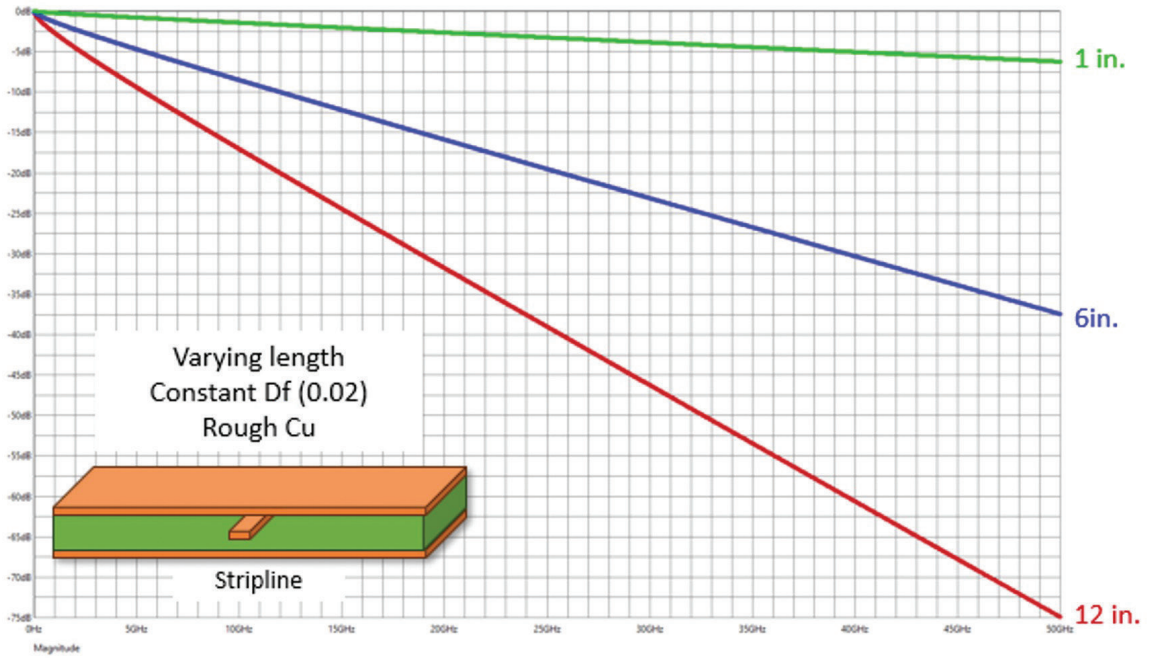

Figure 67. Loss – comparing impact on loss when varying length.

In figure 68, we have an example showing how loss is affected by the dielectric loss tangent. This is a 12 inch transmission line like we saw in the last figure but here we have two different loss tangents. They are 0.02 and 0.002. Everything else is the same and we can see big differences in loss depending on the loss tangent. At 10GHz, for example, the trace in the dielectric with loss tangent of 0.02 experiences about 17.5dB loss (~13% of the transmitted signal arrives at the far end). This is the same as before. However, the trace in the dielectric with 0.002 loss tangent experiences about 7dB loss (~44% of the transmitted signal arrives at the far end). This is similar to the performance we saw earlier for the 6 inch trace on the higher loss dielectric. Here though, we still have a 12 inch trace but used a lower loss material to achieve a lower loss transmission line. Note that this loss behavior is specific to the materials and trace dimensions used in this example to demonstrate how loss tangent can impact loss.

For the last example on loss, refer to figure 69, which shows that loss can be affected by the conductor surface roughness. This is the 12 inch transmission line we saw previously using the lower loss tangent dielectric. In this case the only difference is conductor surface roughness which produces different loss. Note that there are multiple possible surface roughness profiles and this loss behavior is specific to the materials and trace dimensions used in this example to demonstrate how surface roughness can impact loss.

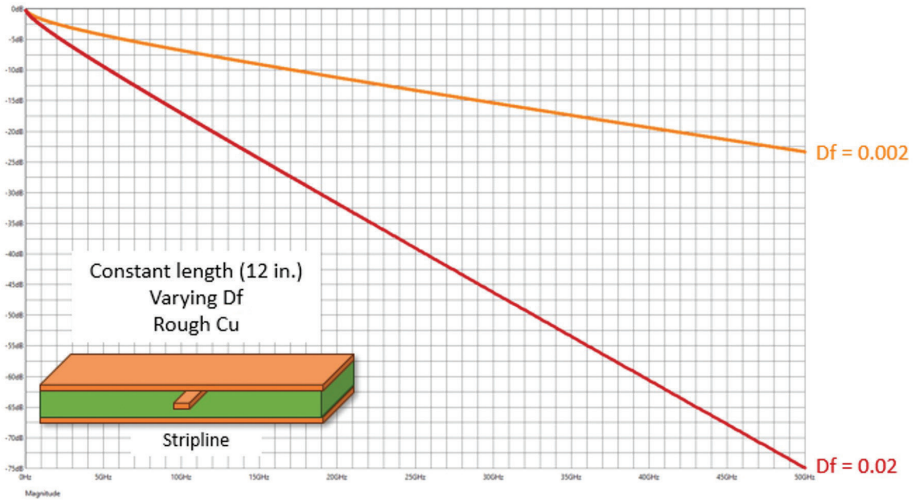

Figure 68. Loss – comparing impact on loss when varying loss tangent.

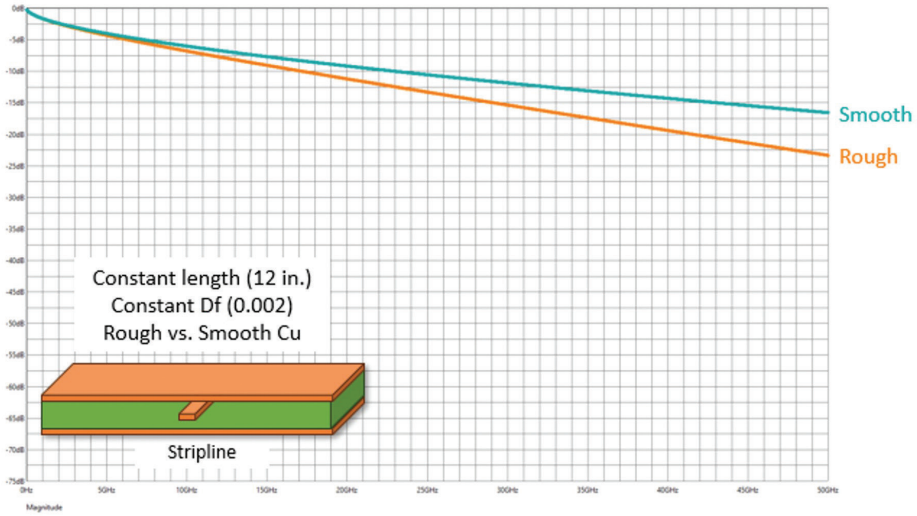

Figure 69. Loss – comparing impact of surface roughness on loss.

You can see from these examples that there are multiple contributors to loss. There could be tradeoffs in designing to meet a loss budget. There may be cost, reliability, and other considerations. Using a simulator is a powerful way to equip you with the information needed to design for loss requirements and make design tradeoffs.

While this is not a book about power integrity, power integrity is simply mentioned here as another possible signal integrity impairment. If the devices transmitting and receiving signals are not sufficiently provided with power, signal quality can be impacted. For example, if power supplies are not at the correct DC voltage or are too noisy, this can affect the signals being driven by those devices or the device's ability to correctly receive those signals. Like signals, the power delivery network has an associated impedance. If the power delivery network impedance is too high, it may not be able to sufficiently provide the required switching current to the devices. Also, as we saw in the discussion on return path discontinuities, the signals interacting with the power delivery network can be impacted by their implementation, as in the examples of signal traces crossing gaps in the return path. While the focus of this book is oriented to signals, the importance of power delivery is vital to a successful design and should not be overlooked.

Timing

And now onto our last signal integrity topic which is timing. This is not intended to be a deep or detailed description of timing. Rather, this is intended to simply point out a few common timing considerations related to signal integrity. As shown in figure 70, timing is important for synchronous buses where things like setup time, hold time, and pulse widths may be relevant. Skew is a difference in time between signals and it can be an important timing consideration when signals need to maintain a particular timing relationship to each other. There can also be timing considerations for individual bits if they need to reach voltage levels in certain time durations – like slew rate requirements. There are also timing considerations for serial channels – like eye diagrams – where we need a certain amount of eye opening.

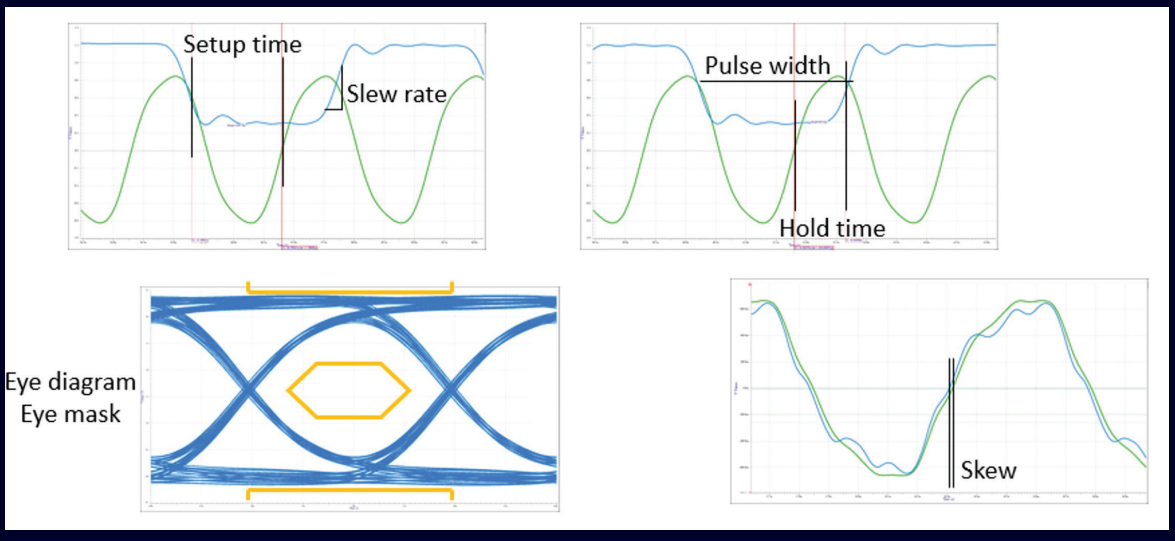

Figure 70. Examples of basic timing considerations.

In figure 71 we have an example showing propagation delay from a driver to multiple loads. This is the same multi-load example discussed in the topic of termination. Delay, and relative delay, between signals and between devices can affect things like setup and hold times. Delay, and delay differences, can come from interconnect on a PCB as shown here and it can also come from IC packages. As we can see in this example, the topology of the interconnect can affect delay and timing. Different devices receive the same signal at different times and depending on the design architecture, this difference in delay may need to be managed. The propagation velocity is a factor in delay, as well as other structures like vias. Impairments also affect delay which will be shown in the next few examples.

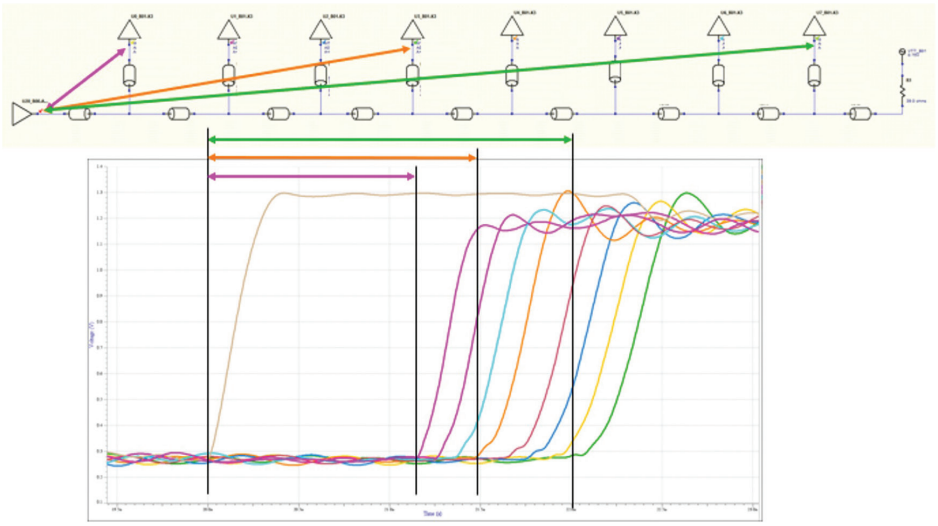

Figure 71. Propagation delay example.

The edge rate can impact timing as shown in figure 72. This example uses the same interconnect but two different rise times. Notice the slower edge rate (blue) takes a lot longer to reach a logic high threshold so the effective timing is quite different than the faster edge rate driver (green), even though the transmission line is the same for both cases. Note the faster edge also has some overshoot in this case while the slower edge rate does not. The technology being used can affect the timing as well as other signal characteristics, like overshoot as seen in this case. There may be some ringing or other issues that need to be traded off with timing to meet all the various signaling requirements.

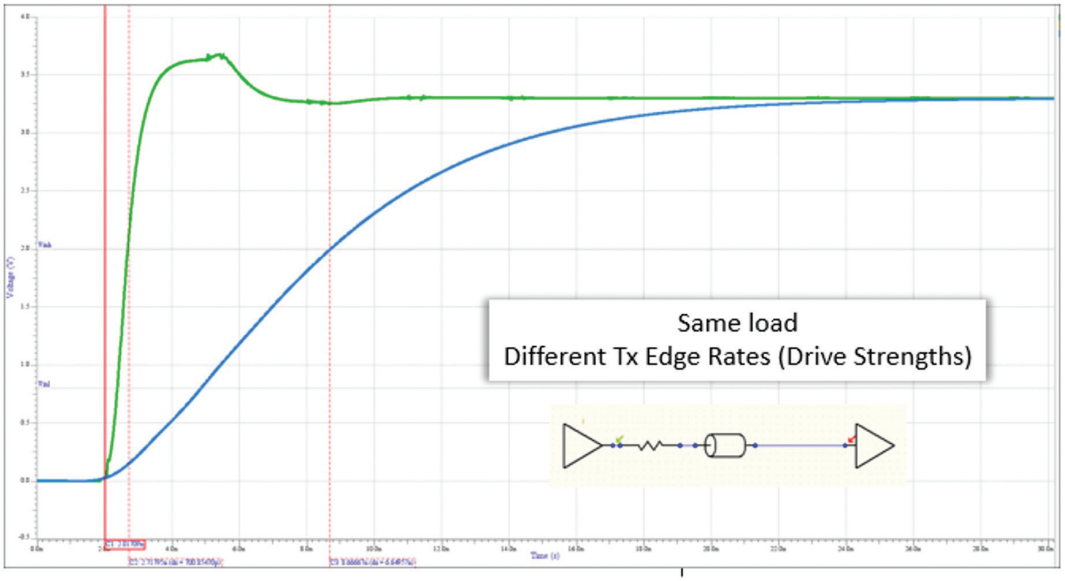

Figure 72. Impact of edge rate on timing.

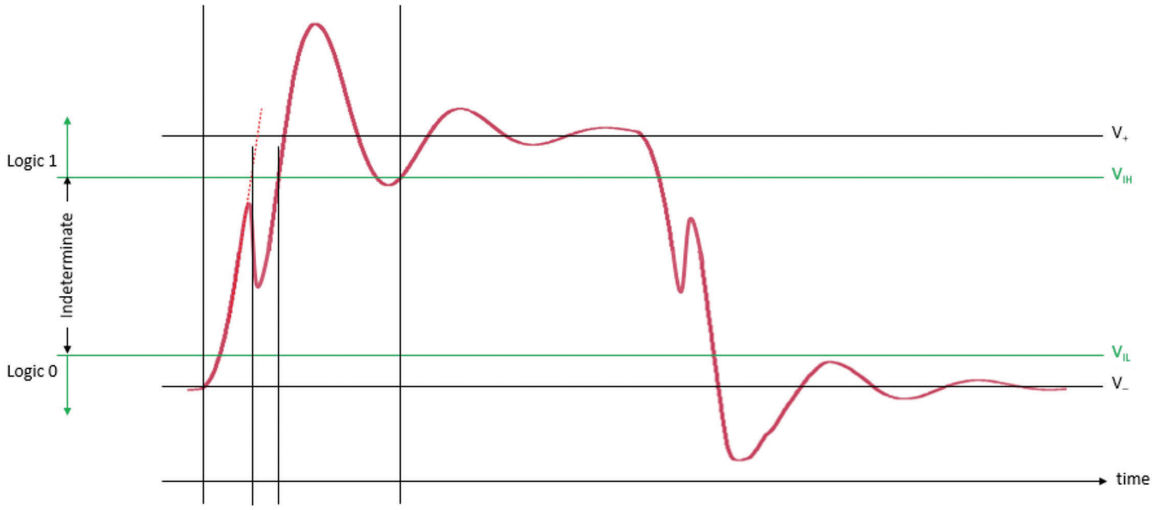

Figure 73. Waveshape example.

We saw the waveform in figure 73 previously when describing basic waveform issues. Here we want to consider those effects in relation to timing. The left-most vertical black line is our reference line when the signal begins to transition from low to high. First, notice the original rising edge before the non-monotonicity. If the signal had kept on that trajectory, it would have crossed the logic high threshold at the second from the left vertical black line. However, the signal has the non-monotonicity and does not continue on that trajectory. After the signal changes direction and returns to going high, it crosses the logic high threshold at the third from left vertical black line. But, because there is ringing, the signal leaves the logic high region and returns to the indeterminate region. Finally, after ringing back, the signal crosses the logic high threshold for the last time at the right most vertical black line. Notice the time difference between the second from the left vertical black line and the right most vertical black line. This is the difference of when the signal could have (or perhaps was expected to have) been a valid logic 1 but instead was not reliably a logic 1 until this time later. This difference in timing here is caused by the waveshape issues we discussed, not due to propagation delay.

Another example we saw earlier is shown in figure 74 that we will revisit now considering the timing implications. These signals are quite short (approximately 0.2 inch). The green waveform is for the

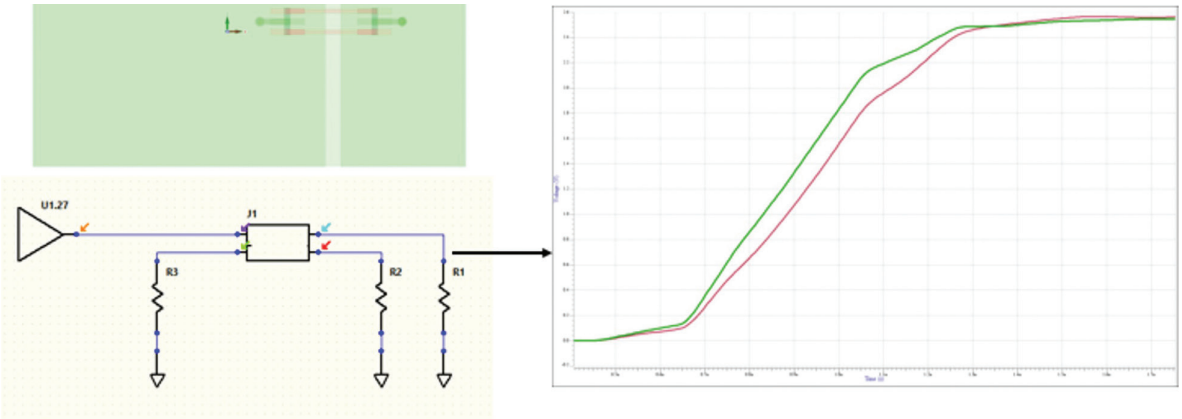

Figure 74. Return path effects.

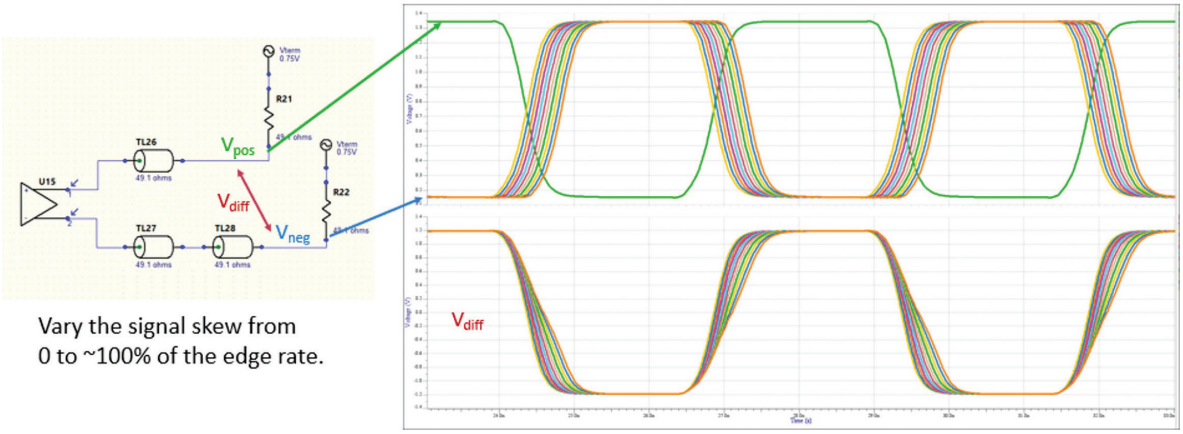

Figure 75. Differential intra-pair skew example.

signal that had a solid reference plane. The red waveform is for the signal that had a gap in the return path. This waveform has a slower rising edge. The driver's edge rate is the same in both cases but the edge rate at the receiving termination is different due to the return path difference.

In figure 75, we're considering skew (time difference) within a differential pair. We mentioned skew earlier in the general context of a time difference between signals with a required timing relationship. In this case, the skew is between the composite signals making up the differential signal. The mechanisms causing skew may be the same as those already mentioned earlier for single ended signals, but here, we're considering the differential pair and resulting differential signal. In the figure, we delay one signal making up the differential pair relative to the other and see the effect it can have on the wave shape of the differential signal. One single ended signal has additional delay that is swept from 0 to 100% of the signal's rise/fall time. As can be seen in the figure, skew between the single ended signals distorts the differential signal. We'll see the eye diagram next. The skew also creates a time-varying common signal.

Figure 76 shows the eye diagram (red) of the differential signal with intra-pair skew approximately the same as the rise/fall time. It also shows the eye diagram (green) of the same differential pair with no intra-pair skew, for comparison. Intra-pair skew

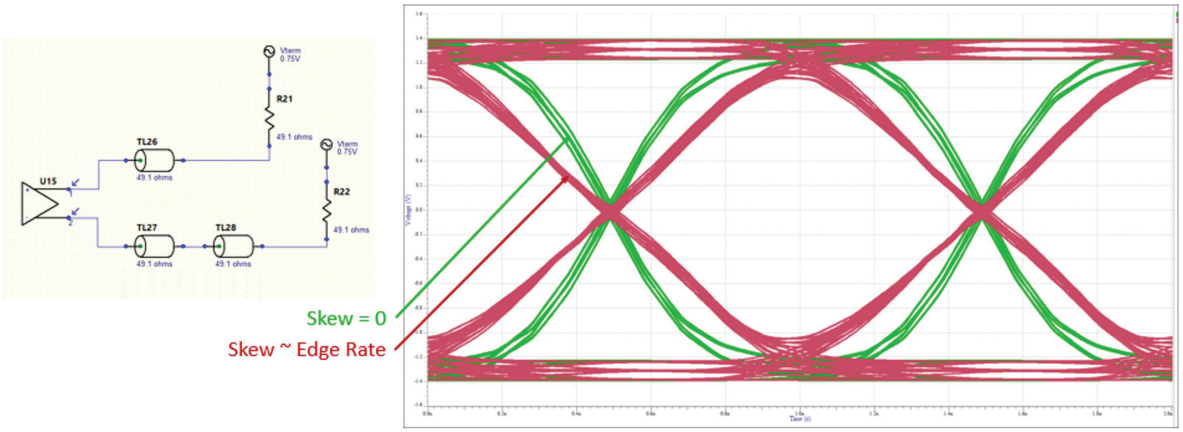

Figure 76. Differential intra-pair skew example eye diagram.

between the signals making up the differential pair can come from multiple sources, such as delay mismatch in the IC package, routing length mismatch on the PCB, etc. Even if the lengths are perfectly matched on the PCB, skew can also come from weave effect skew due to the placement of the signals relative to the glass and resin (having different dielectric constants) in the PCB laminate.

Jitter is also an important timing and signal integrity concept. The details of jitter are out of scope for this book but it is mentioned to introduce the topic. The example in figure 77 simply compares a differential signal with no jitter (the green waveform) and a differential signal with jitter (the blue waveform). Notice the top and bottom levels of both cases (with and without jitter) are the same. The only change is adding timing jitter which is a deviation of the signal in time from its ideal location. Note this is not caused by skew. Excessive jitter can cause bit errors in the received bit stream and problems in clock recovery from the received bit stream. There are multiple kinds and sources of jitter. One source of jitter that was discussed earlier is ISI. The amount of jitter caused by ISI depends on the data being transmitted.

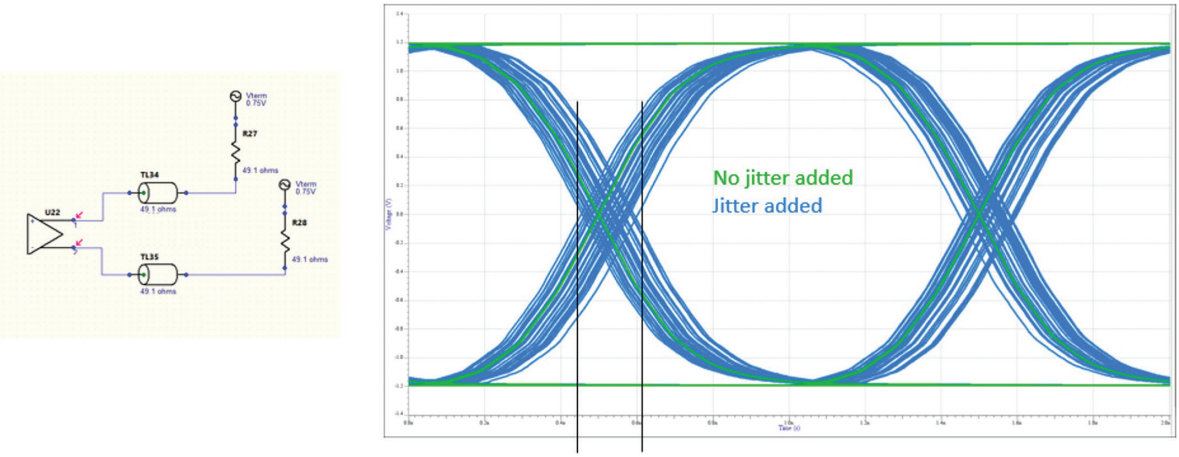

Figure 77. Jitter.

Conclusion

Understanding the fundamental concepts of signal integrity is critical for engineers striving to create high-performance electronic systems. Throughout this book, we examined principles and techniques for ensuring robust signal transmission.

From understanding transmission lines to mitigating crosstalk, each chapter has provided insights and guidance to help navigate signal integrity design. By emphasizing the importance of considerations like characteristic impedance, crosstalk, vias, return path continuity, and termination, the aim was to equip engineers with tools and knowledge to help tackle challenges inherent in high-speed digital circuits.

As technology advances and data rates continue to increase, the demand for rigorous signal integrity practices is also expected to intensify. Therefore, it is imperative for engineers to understand the fundamental concepts so they can keep up with emerging technologies and methodologies in signal integrity analysis and design.

In closing, this book serves as a resource and reference for engineers dealing with signal integrity. With a foundation in fundamental signal integrity principles, they can move forward in practice and further learning.

Next steps

HyperLynx Virtual Workshops provide a quick way to get hands-experience with HyperLynx by running analysis in a cloud-based Windows desktop environment. All workshops include simulation databases, models and step by step instructions to guide you through a complete analysis workflow. Once you've run through a HyperLynx workshop, you'll have the skills you need to run analysis on your own design!

You can register for HyperLynx Virtual Workshops here.

Virtual Workshops require a Siemens account – if you don't already have one, you can create a free one.

References

Eric Bogatin. "Signal and power integrity – Simplified, 2nd edition, Upper Saddle River, NJ," Prentice Hall, 2010.

Fadi Deek. "The printed circuit designer's guide to: Signal integrity by example., Rohnert Park, CA," BR Publishing, Inc., 2017.

Siemens Digital Industries Software helps organizations of all sizes digitally transform using software, hardware and services from the Siemens Xcelerator business platform. Siemens' software and the comprehensive digital twin enable companies to optimize their design, engineering and manufacturing processes to turn today's ideas into the sustainable products of the future. From chips to entire systems, from product to process, across all industries, Siemens Digital Industries Software – Accelerating transformation.

Americas: 1 800 498 5351

EMEA: 00 800 70002222

Asia-Pacific: 001 800 03061910

For additional numbers, click here.

86072-D3 6/24 H