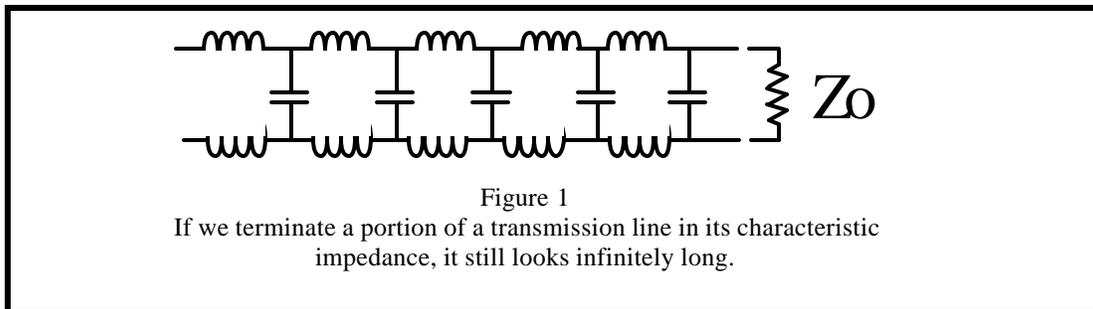


Transmission Line Terminations

It's The End That Counts!

In previous articles¹ I have pointed out that signals propagating down a trace reflect off the far end and travel back toward the source. These reflections can cause noise, and therefore signal integrity problems. These reflections can be controlled if we design our traces to look like transmission lines. Then, for “long” traces (those longer than the critical length) we can control reflections using transmission line termination techniques. This paper describes the five most common types of termination techniques used on printed circuit boards.

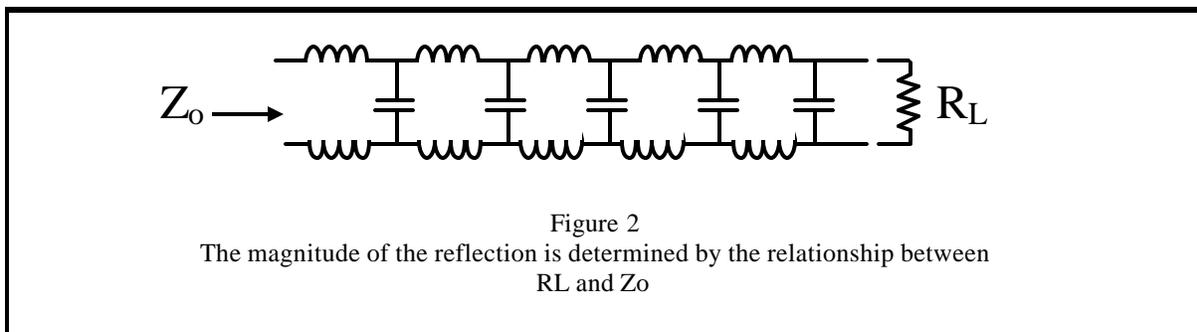
As suggested in the first of the referenced papers, the ideal model of a transmission line is one that is infinitely long. If we take just a portion of that line, and terminate it in its characteristic impedance, it still looks infinitely long when viewed from the front (Figure 1). And, there is no reflection from the far end. If we terminate a transmission line in anything *other than* its characteristic impedance, a reflection will occur. The magnitude of that reflection can be determined from a parameter known as a reflection coefficient, ρ .



The reflection coefficient, ρ , is calculated as:

$$\rho = \frac{R_L - Z_0}{R_L + Z_0}$$

where R_L , the load resistor, and Z_0 , the characteristic impedance of the transmission line, are as shown in Figure 2. It can be observed that the reflection coefficient has a value between -1 and $+1$. If the trace is left open circuited, the reflection coefficient is $+1$ and there will be a 100% reflection back toward the





Transmission Line Terminations

It's The End That Counts!

driver. If the trace is shorted, the reflection coefficient is -1 and there will be a 100% reflection of the *opposite sign* back toward the driver. If the line is terminated with a resistor whose value is the *same* as the characteristic impedance of the trace (Z_0), the reflection coefficient is zero and there will be *no reflection* at all.

It should be noted that if a reflection does occur, and propagates back to the driver (source), it can reflect again off the source. The driver has an output impedance. If that output impedance is exactly equal to Z_0 , then there will be no further reflection from the source. But if the output impedance of the driver is different than Z_0 , an additional reflection will occur. The magnitude of that reflection is again determined by a reflection coefficient. Simply substitute R_s , the output impedance of the driver, for R_L in the reflection coefficient formula.

UltraCAD has created a simple Transmission Line Simulator that can illustrate these reflections under different conditions (Figure 3). It is described in the Appendix and is freely available for download from UltraCAD's web site.

Although when we talk about terminations and reflections we usually talk about a single resistor at the end of the trace, there are actually five common termination techniques seen in typical circuits.² These are summarized in Figure 4.

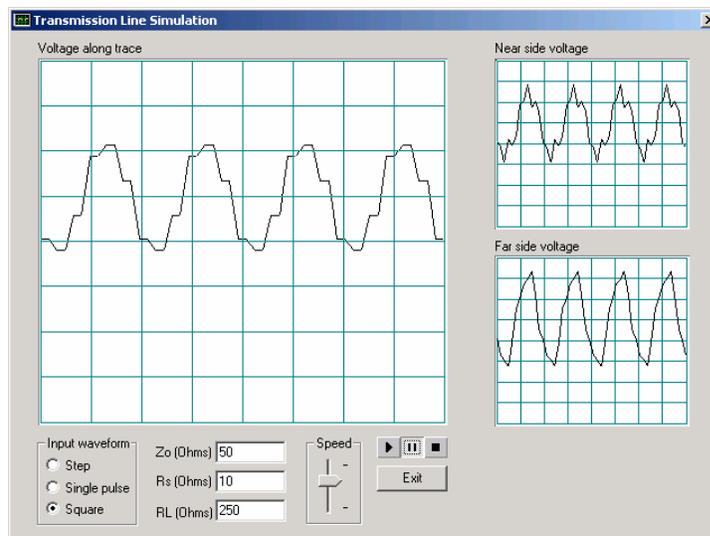


Figure 3
UltraCAD's Transmission Line Simulator

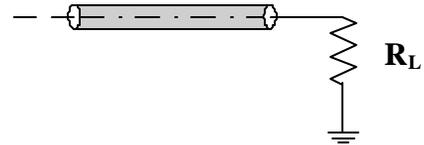
Transmission Line Terminations

It's The End That Counts!

Parallel

Pros: Terminate to either Gnd or Vcc
 R easy to determine
 Only one additional component
 Performs well with distributed loads

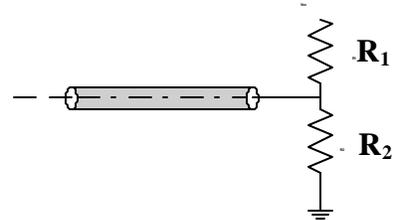
Cons: Power dissipated in R_L at all times
 Power requirement high



Thevenin

Pros: Properly chosen, pull up/down resistors can improve noise margins
 Performs well with distributed loads

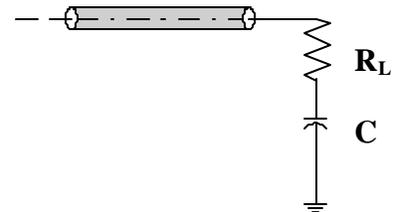
Cons: Results in steady flow of current through R's
Optimum selection of R_1 and R_2 can be complicated
 Complicated if used with tri-state devices



AC

Pros: Performs as well as parallel without the dc power drain

Cons: C is difficult to optimize
 Requires 2 components
 Can lead to timing problems



Series

Pros: One component
 No dc load

Cons: Can be difficult to optimize R_S
 There IS a reverse reflection.



Diode

Pros: Does not depend on Z_0
 Little increase in power
 Can be placed anywhere on line

Cons: Reflections still exist
 Requires 2 devices
 Diodes must be FAST with low forward voltage

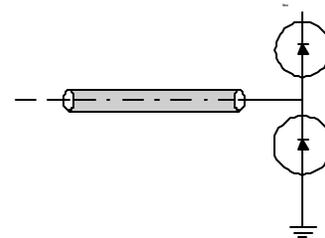


Figure 4
 The five most common termination strategies.



Transmission Line Terminations

It's The End That Counts!

Parallel Termination: The first of these is the most intuitive and possibly the most common. It is called “parallel termination” and simply consists of a single resistor from the trace to ground (or to Vcc). It is the technique we have been referring to throughout this paper so far.

This technique has several advantages: (1) the value of the resistor is relatively easy to determine, (2) there is only a single component, (3) it is easily connected, and (4) it performs well with distributed loads (i.e. loads that are distributed along the trace.) There is only one drawback to this type of termination: it provides a continuous DC path to ground. Therefore, continuous DC current can flow through it at all times. This may or may not be an issue for a single trace. But if your design has a thousand or so impedance controlled nets, the total power dissipation in a thousand terminating resistors can become quite significant!

Thevenin Termination: A closely related variation of the parallel termination strategy is the Thevenin termination. This consists of a *pair* of resistors, one going to ground and one going to Vcc. The pair of resistors provide a pull-up/pull-down function as well as a termination function. Therefore, it can improve noise margins in certain situations, and performs as well as the parallel termination with distributed loads.

On the one hand, it would appear that the selection of resistor values would be relatively straightforward. The parallel combination of the two resistors must simply equal Z_o , or:

$$Z_o = \frac{1}{\frac{1}{R1} + \frac{1}{R2}}$$

But Ethirajan and Nemeč² show that there is an optimum value for R1 and R2 based on the specific characteristics of the driver. The optimum values for R1 and R2 primarily optimize (minimize) the power dissipated in the circuit. This optimum value can be difficult to determine. Other drawbacks to this strategy include the addition of an additional component, and the fact that DC current still flows through the resistor pair at all times. And this strategy is only well suited for bipolar (two-state) devices, not for tri-state logic families.

Some people have reported that there can be an EMI problem with this termination strategy. Note that there are different currents flowing between Vcc and Gnd, through R1 and R2, depending on the logic state. And this current changes at the same rate (di/dt) as the logic state changes (the rise/fall time of the signal). In this respect, the two resistors look *exactly* like a switching logic gate. Therefore, they might need to be decoupled (with bypass capacitors) just as a logic gate might need to be decoupled (with bypass capacitors). The issue really is related to loop areas and currents, which are one of the primary causes of EMI. Thus, you may occasionally hear people say that Thevenin terminations may also require decoupling capacitors in order to quiet down EMI emissions.

AC Termination: Yet another variation is the addition of a capacitor in series with the parallel terminating resistor. The primary advantage of this is that the capacitor blocks DC current, so there is no steady-state current flowing through the termination. At first glance it would appear that this strategy otherwise has all the advantages of the parallel termination strategy. However, the “cost” of this, of course, is the added component.



Transmission Line Terminations

It's The End That Counts!

But there are some subtle and some not-so-subtle problems with AC termination. If a large capacitor value is used, there can be considerable power dissipation and the strategy is little different from normal parallel termination. If a very small capacitor is used, it will cause overshoot and may interfere with the rise and fall times of the signal.

There are some subtle interactions that take place with AC termination using a small capacitor. As the voltage across the terminating resistor changes, the current through the capacitor changes. This causes the capacitor to charge or discharge with an RC time constant related to the component values. If the time constant is too short, the capacitor charges during the half-cycle and the voltage at the receiver changes accordingly.

Figure 5 illustrates this effect. Figure 5(a) shows the waveforms associated with a relatively large capacitor. The scope is a HyperLynx simulation of a 50 Ohm transmission line, 5 ns long, terminated with a 50 Ohm resistor in series with a .002 uF capacitor. The green line is the driver voltage. The red line, trailing 5 ns later, is the voltage at the receiver (and at the resistor). The blue line is the voltage across the .002 uF capacitor. As can be seen, the waveforms look very clean.

Figure 5(b) illustrates the same simulation with a 200 pF capacitor (.0002 uF). Note that there is overshoot and undershoot in the receiver waveform. Each transition of the driver signal is followed 5 ns later with the same (magnitude) transition at the receiver. But then the receiver voltage continues to rise (or fall, as the case may be) as the capacitor charges and discharges. If this over/undershoot is severe enough, logic errors may result.

Note that this is *not* a reflection. This is not an impedance matching issue. It is a changing reference voltage that shifts the value of the waveform at the receiver. As is seen, the driver signal is still relatively clean. In extreme cases the voltage at the driver may also start to change. But again this is better classified as a shifting reference rather than an actual reflection.

Line length: An interesting thing happens as we increase the line length. The line has an intrinsic capacitance (C_0) and an intrinsic inductance (L_0) that is expressed in units per unit length. In absolute terms, the values increase with increasing line length. At some point the values become large enough that *ringing* occurs as the average value across the capacitor changes.

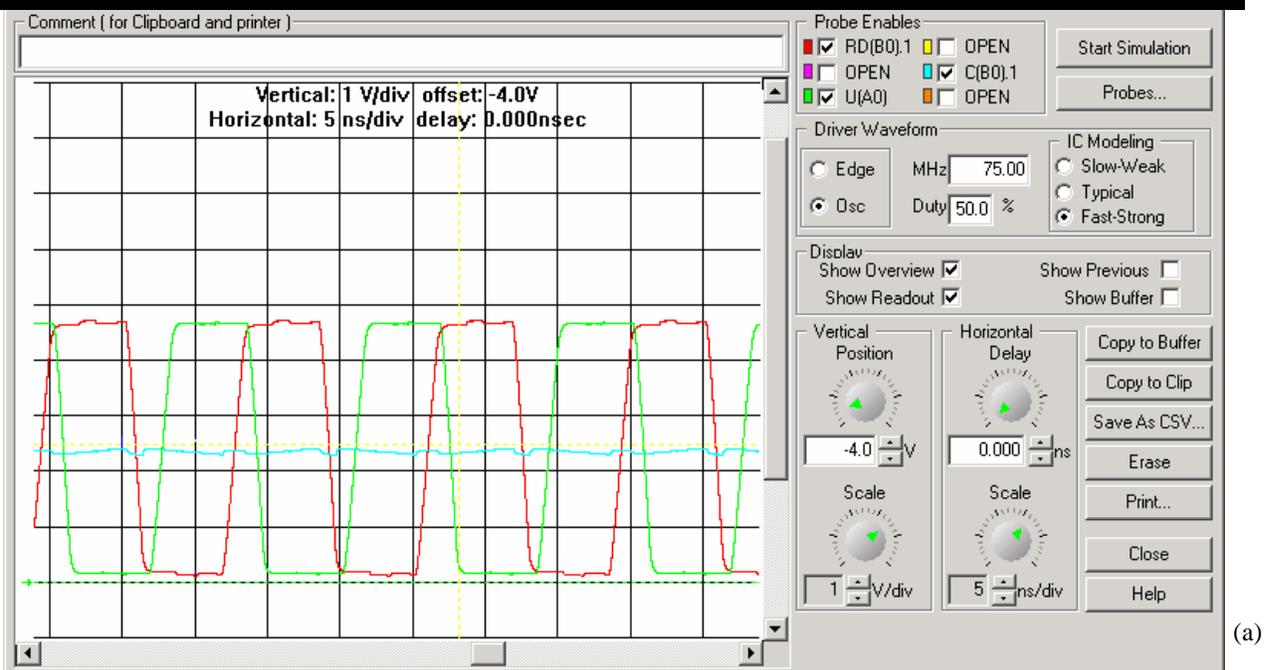
Figure 6 illustrates this. Figure 6(a) is a 5 ns long line with a .002 uF capacitor. The voltage across the capacitor (black line in figure 6(a)) rises in a controlled manner to 50% of the voltage at the receiver (red line). The sample time of the image on the scope is 1000 ns. Figure 5(a) is taken from a central part of this simulation waveform.

Figure 6(b) is the same 5 ns line with a smaller, 200 pF capacitor. Note that there is a short stabilization time, but then the system settles down. Figure 5(b) is taken from this simulation.

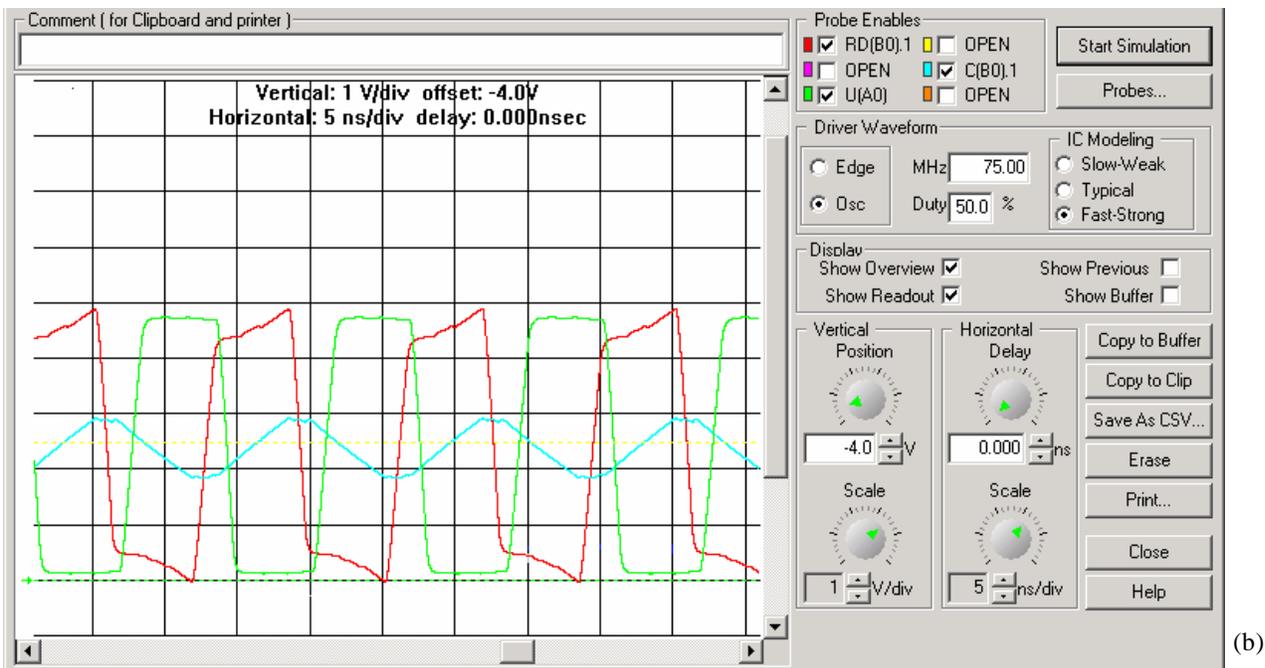
Figure 6(c) is the 200 pF capacitor but with a much longer (30 ns) line. The value of the capacitor interacts with the parameters of the transmission line to cause oscillations that take a finite time to die out.

Transmission Line Terminations

It's The End That Counts!



(a)



(b)

Figure 5

HyperLynx simulation of a 50 Ohm transmission line, 5 ns long, with AC termination with 50 Ohm resistor. Figure (a) is with 2000 pF (.002 uF) cap, (b) is with 200 pF cap. The driver signal (green) is clean in both cases. Each transition of driver signal is accompanied by the same transition at the receiver, but then an additional overshoot or undershoot as the smaller capacitor (blue) charges or discharges.

Transmission Line Terminations

It's The End That Counts!

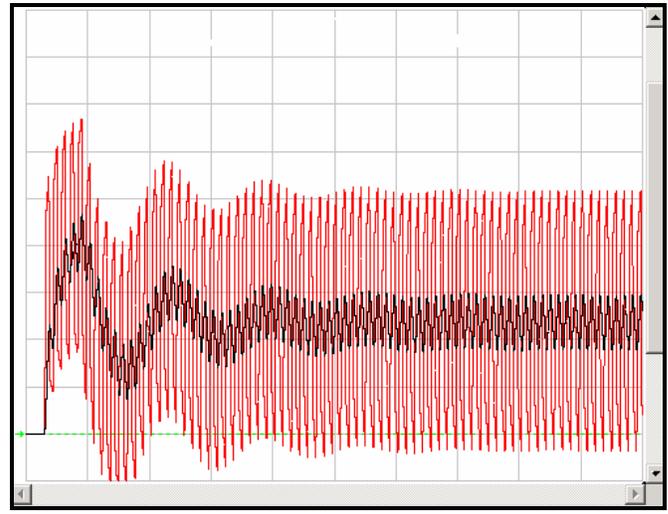
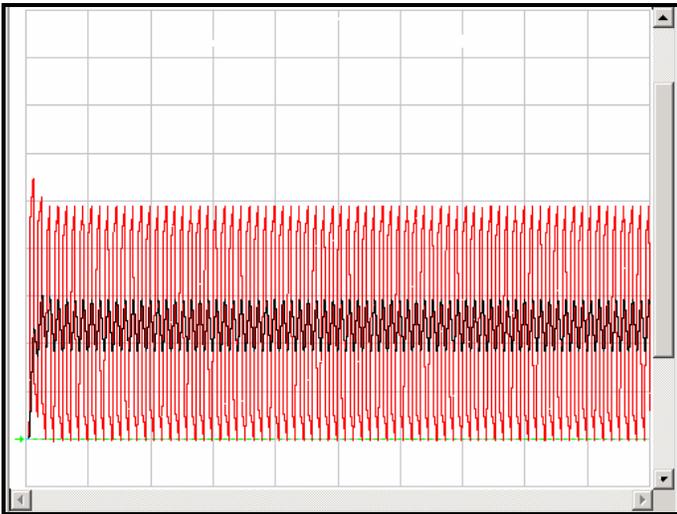
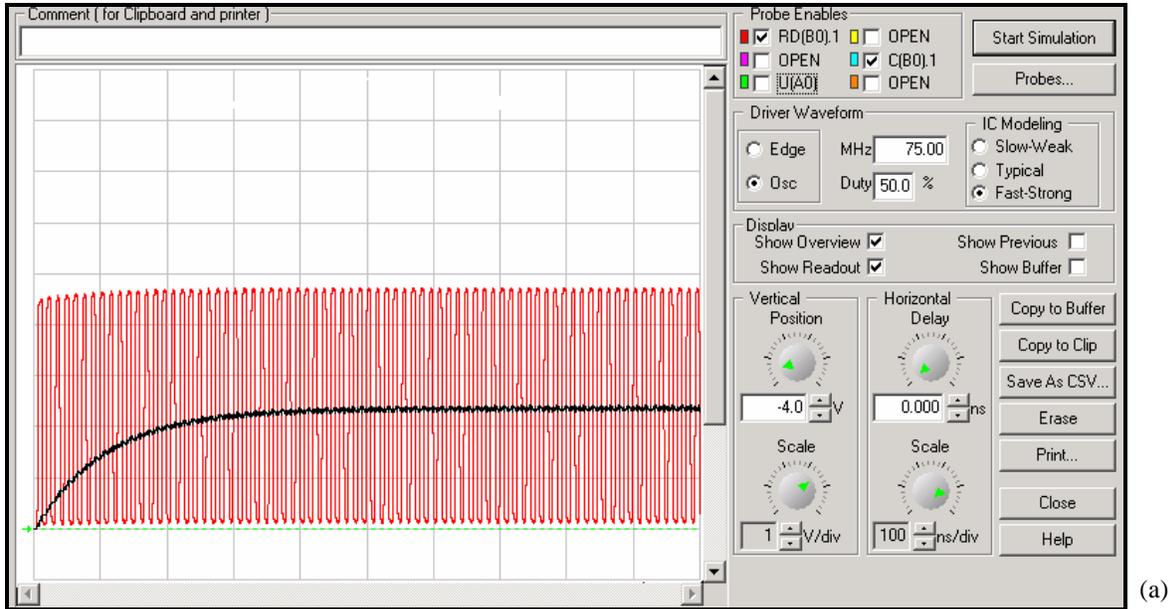


Figure 6

Simulations showing how the voltage across the capacitor stabilizes with time. Figure 6(a) is a 5 ns long line with .002 uF capacitor. Figure 6(b) is the same line with a 200 pF capacitor. Figure 6(c) is a 200 pF capacitor with a much longer (30 ns) line. Note the oscillations before the voltage at the receiver stabilizes.

Red Line: Voltage at the receiver.
 Black Line: Voltage across the capacitor.

Transmission Line Terminations

It's The End That Counts!

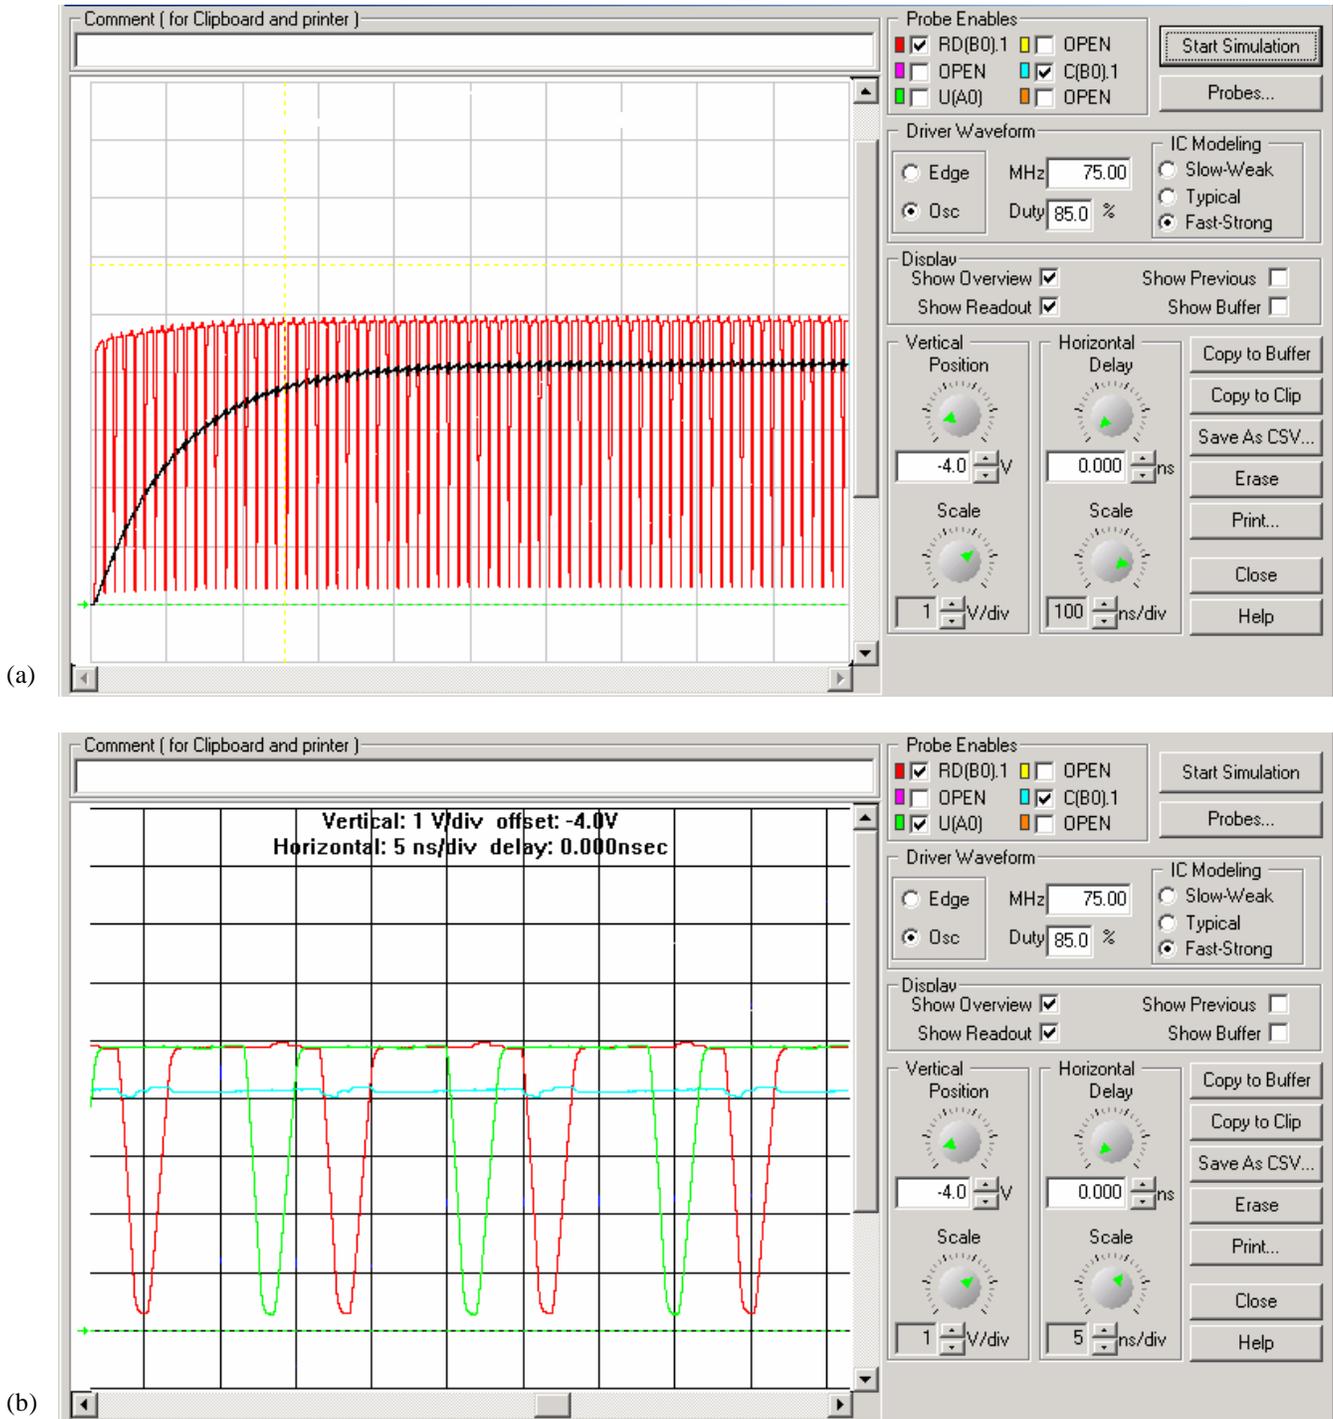


Figure 7
Same simulation as Figure 5(a) and 6(a) but with an 85% duty cycle.



Transmission Line Terminations

It's The End That Counts!

Admittedly this 30 ns line length might be an extreme example from a PCB standpoint. But it may not be extreme if we are talking about transmission lines that connect two pieces of equipment separated by some distance.

Duty cycle: If we have a uniform bit-stream (50% duty cycle) traveling down a trace, the capacitor will charge to the average value between a logical zero and a logical one (Figure 6(a)). If the duty cycle changes, that voltage level will adjust to a new average value. One question is how fast that adjustment takes place. Another question is whether the change in duty cycle has an effect on the waveforms.

Figure 6(a) shows a 50% duty cycle. Figure 7(a) shows the same simulation with an 85% duty cycle. The voltage across the capacitor rises to a higher level, but the waveforms themselves (Figure 7(b)) are still clean and undistorted.

Conclusion: AC termination can be effective if the capacitor is large enough. But if the capacitor is too large, the benefits (of blocking the DC component of the signal) may be reduced. Long lines may introduce unexpected oscillations when the *average* DC voltage component of the signal changes.

Series Termination: Series termination is becoming more and more common in today's high-speed designs. It has the two desirable attributes of a single component and no DC current draw at all. But the series termination resistor is placed at the front of the trace, not the end; the end is left open-circuited. Therefore, there is a 100% positive reflection from the far end of the trace that reflects back towards the front.

The value for the series termination resistor is set so that the *SUM* of it and the output impedance of the driver totals to the impedance of the trace. That way, there is no secondary reflection from the driver. The voltage pattern we expect to see is a single, positive reflection from the far end of the trace which travels to the front of the trace and is completely absorbed in the source resistor and driver.

That means that we have to know the output impedance of the driver. That isn't too serious a problem, since the output impedance should be specified by the manufacturer. But an additional complication is that the output impedance is often different depending on the output state.

Figure 8(c) illustrates a Signal Vision series termination simulation. The *tech-als* IBIS model supplied with Signal Vision has a *logical high* output impedance of approximately 19 Ohms and a *logical low* output impedance of about 5 Ohms. Simulation results (a) and (b) are for termination resistor values of 31 Ohms and 45 Ohms, respectively. Note how the rising waveform in (a) and the falling waveform in (b) appear to be terminated correctly. But both the rising and falling edges cannot be terminated correctly at the same time!

Diode Termination: This termination strategy is really just a reflection *limiting* strategy. No attempt is made to absorb reflections or to prevent them. The intent is to let them occur and simply limit them to (a) one diode drop above V_{cc} and (b) one diode drop below Gnd . This approach is not very common. UltraCAD has never designed a board with diode terminations!

Transmission Line Terminations

It's The End That Counts!

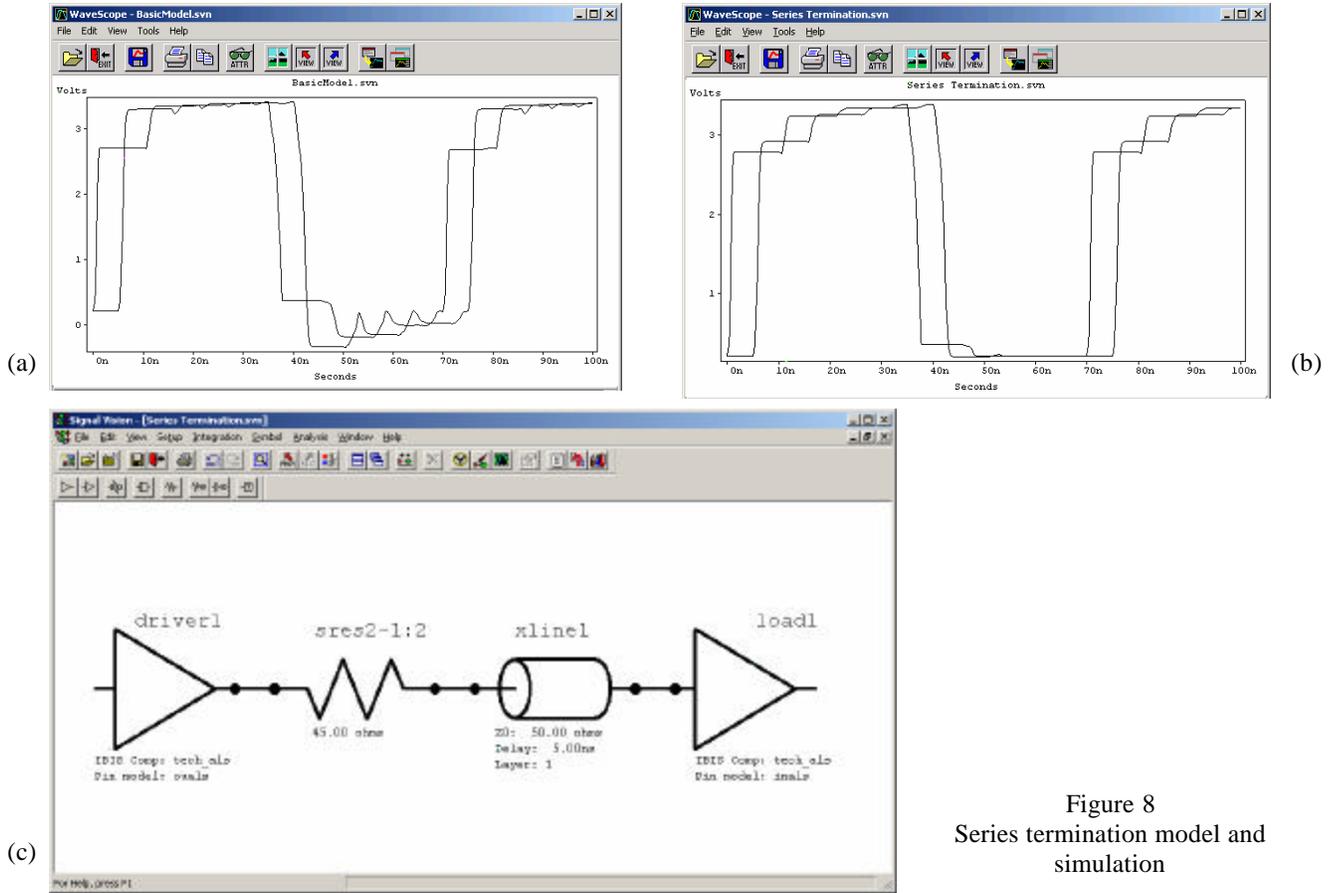


Figure 8
Series termination model and simulation

Summary: Termination strategies can be effective in eliminating, or at least minimizing, transmission line reflections. But no individual strategy is perfect. Each one has a tradeoff of some type. Fortunately, perhaps, for PCB designers, the selection of a termination strategy is not the responsibility of the board designer. It is the responsibility of the circuit design engineer. But board designers need to know what various termination strategies look like and have some feeling for how they are supposed to work. And they *do* have the responsibility of *raising* the question to their engineers if it appears that a termination strategy might have been overlooked or improperly implemented during the circuit design process.

Footnotes:

- 1 "Basic Transmission Lines, Why Use 'Em At All?" and "Propagation Times and Critical Length; How They Interrelate," both available from Mentor's web site at www.mentor.com/pcb/tech_papers.cfm.
- 2 For a good discussion of trace terminating techniques see Ethirajan and Nemeç, "Termination Techniques for High Speed Buses," EDN, 2/16/98, p. 135.



Transmission Line Terminations *It's The End That Counts!*



About the author:

Douglas Brooks has a B.S and an M.S in Electrical Engineering from Stanford University and a PhD from the University of Washington. During his career he has held positions in engineering, marketing, and general management with such companies as Hughes Aircraft, Texas Instruments and ELDEC.

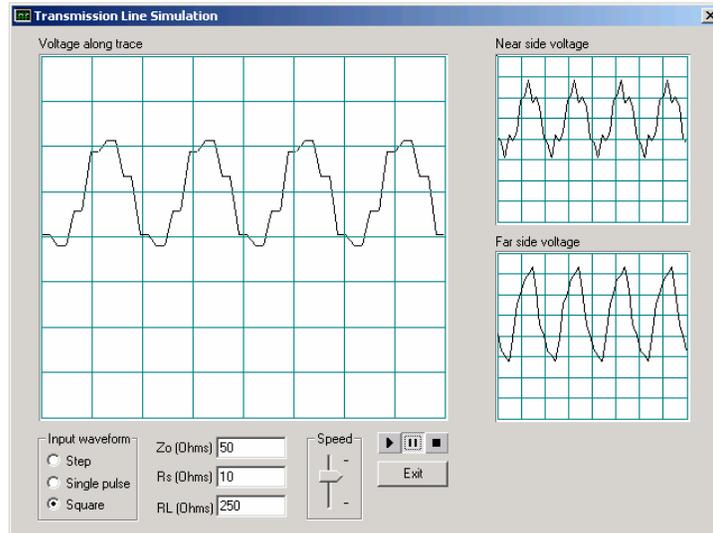
Brooks has owned his own manufacturing company, and he formed UltraCAD Design Inc. in 1992. UltraCAD is a service bureau in Bellevue, WA, that specializes in large, complex, high density, high-speed designs, primarily in the video and data processing industries. Brooks has written numerous articles through the years, including articles and a column for *Printed Circuit Design* magazine, and has been a frequent seminar leader at PCB Design Conferences. His primary objective in his speaking and writing has been to make complex issues easily understandable to those individuals without a technical background. You can visit his web page at <http://www.ultracad.com> and e-mail him at doug@ultracad.com.



Transmission Line Terminations

It's The End That Counts!

APPENDIX



UltraCAD's Transmission Line Simulator

This simulator can be downloaded from <http://www.ultracadm.com/mentor/tlinesim.zip> (108K). Unzip the file to obtain `tlinesim.exe` and execute the file. The program shown above will start.

The main screen represents a transmission line. The pattern represents the signal at every point along the line in real time. The screen at the upper right represents the instantaneous signal level at the very beginning of the trace, and the screen at the lower right represents the same thing at the far end of the trace. In this way, you can visualize what is happening at every point along the line.

The "Input Waveform" selection enables you to select a single step-function change in voltage, a single pulse, or a repetitive square wave (much like a clock signal) as the driving signal. The three text boxes allow to specify the characteristic impedance of the transmission line, and the resistive loading at each end. Having set these, the right arrow then starts the selected waveform down the trace. The waveform will reflect back and forth for at least 5 cycles, illustrating what happens to the waveform under various termination conditions. The speed at which the waveform travels is controlled by the speed selector slider.

The other two selection buttons allow you to pause or reset the simulation.