THE SWITCHNG BEHAVIOUR OF A SERIES TERMINATED TRANSMISSION LINE

Series terminated transmission lines are the primary method of connecting CMOS logic devices. Understanding how these transmission lines work is vital to making sure signals are properly delivered to each receiver. How all of this works is not intuitive and baffles some of us until it is explained. This short write up is intended to clear up some of the confusion..

Figure 1 is a typical 5V CMOS driver with a 50 ohm transmission line connected to a CMOS receiver that is passive, meaning that it simply responds to the voltage waveform presented at its input. For purposes of this explanation, CMOS receivers look like very small capacitors that can be considered to be open circuits. In this example, the line is 12" or about 30 cm long. In a PCB, energy travels at approximately six inches per nanosecond, so this line is about 2 nanoseconds long.

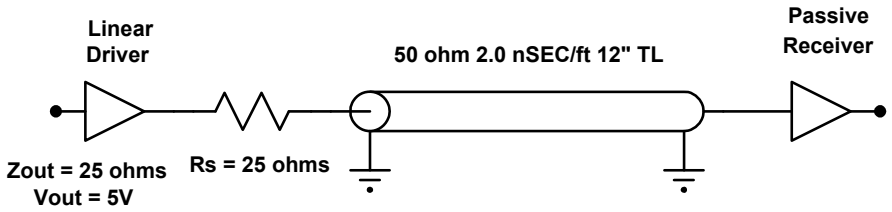

SERIES TERMINATED TRANSMISSION LINE

Figure 1. A Typical Series Terminated 5V CMOS Circuit

Figure 2 is an equivalent circuit for a transmission line like that shown in Figure 1.

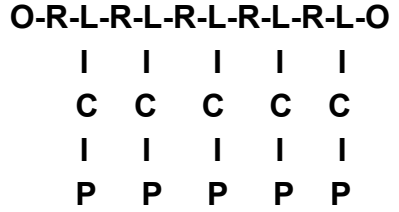

O is input or output, C is capacitance per unit length
R is resistance per unit length L is inductance per unit length
P is plane layer

SCHEMATIC REPRESENTATION OF A TRANSMISSION LINE

Figure 2. An Equivalent Circuit for the Transmission Line in Figure 1

Notice that there is capacitance, resistance and inductance distributed along the length of the transmission line. These elements are called parasitics and determine the behavior of a transmission line with the ratio of inductance per unit length to capacitance per unit length determining the impedance of the line as shown in equation 1. Lo is inductance per unit length and Co is capacitance per unit length. These two variables are determined for a particular type of transmission line using a tool such as a 2D field solver. There are many field solvers available as parts of signal integrity tools.

In almost all cases, the value of R is so small compared to the L and C that it can be ignored. Until frequencies involved go above a GHz, this is a reasonable assumption.

$$\mathbf{Zo} = \sqrt{\boldsymbol{Lo/Co}}$$

Equation 1, Impedance as a Function of Distributed Capacitance and Inductance

When the driver in Figure 1 wishes to move the logic level on the transmission line from a logic 0 to a logic 1 it must charge up the distributed parasitic capacitance of the transmission line. This is the primary power consumed by CMOS logic circuits. When the same driver wishes to move the logic level from a logic 1 to a logic 0 it must remove that charge.

> When a signal is sent along a wire or transmission line it is energy in the form of an electromagnetic field. This energy will travel along the path and be reflected at the ends of the path forever unless it is absorbed by a terminating resistor or is slowly lost in the resistance of the conductor. If the ends of the path are open circuits, the reflected energy will be the same polarity as the incident energy. If the ends of the path are short circuits, the reflected energy will be inverted.

How Charge Is Put On a Logic Line to Move It from a Zero to a One

Figure 3 is the equivalent circuit at the moment that the driver starts to move the logic line from a zero to a one. Notice that a voltage divider has been formed by the combination of the driver output impedance and the series termination in the upper part and the impedance of the transmission line in the lower part. When the series termination has been properly chosen the combination of Zout and Zst will be the same as Zo. In this example, both will be 50 ohms and so the voltage at the input to the transmission line will be V/2.

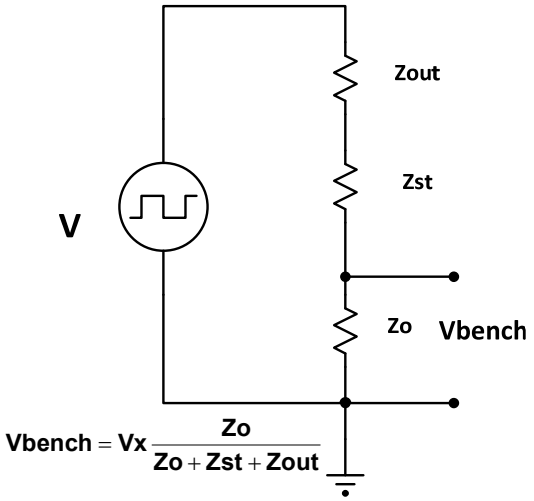

Figure 3. Equivalent Circuit of Figure 1 When A Transition From a Zero to a One Begins.

Figure 4 shows the voltage waveforms at the input to the transmission line and at the input to the receiver as time goes by. The red waveform is the input to the transmission line and the orange waveform is the input to the receiver at the end of the transmission line. Notice that the voltage level immediately after the transition from zero to one is only half size. This is because of the voltage divider shown in Figure 3. This voltage level is often referred to as the “bench” voltage.

What has been launched into the transmission line is energy in the form of an electromagnetic field (EM) the voltage component of which is V/2. This energy is charging the parasitic capacitance of the transmission line to a voltage level of V/2 as the field travels out the transmission line.

After 2 nanoseconds (the electrical length of the transmission line) the line has been fully charged to V/2 and the electromagnetic field encounters an open circuit at the receiver. When such a field encounters an open circuit none of the energy in the field is absorbed and it is reflected back at the same magnitude it had when it was outbound.

At the moment of total reflection, the voltage level on the end of the line is V/2. Since the voltage magnitude of the electromagnetic field is V/2, after total reflection the amplitude will be V. Notice that the orange waveform has an amplitude of V as soon as the EM field arrives at the end of the line. On the return trip, the parasitic capacitance of the transmission line is charged all the way up to V. Once the electromagnetic field returns to the driver it encounters the equivalent circuit shown in Figure 5.

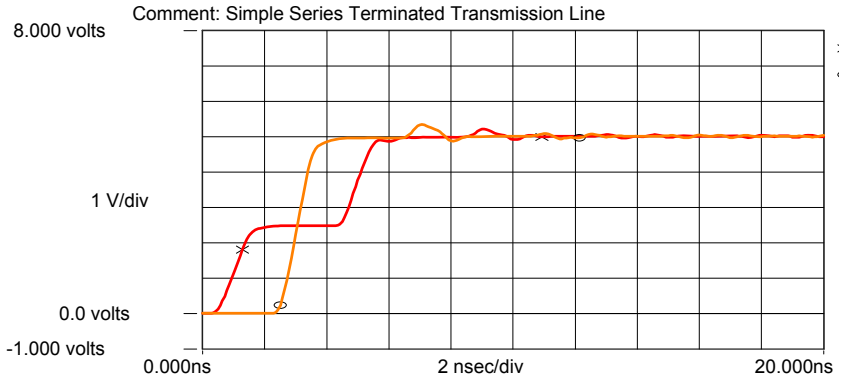

Figure 4. Voltage Waveforms At The Two Ends of The Transmission Line In Figure 1.

When a resistor of the same value as the impedance of a transmission line is placed across the ends of the transmission line, all of the energy in an electromagnetic field will be absorbed by the resistor and there will be no further reflections. Such a resistor is called a parallel termination.

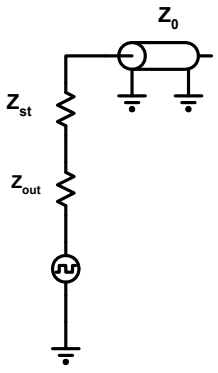

Figure 5. Equivalent Circuit of the Driver in Figure 1 Seen By the Reflected Electromagnetic Field

A voltage source, like that shown in Figure 5 has zero impedance.

Since the sum of Z_{out} and Z_{st} is 50 ohms, and the voltage source is a short circuit, together they constitute a parallel termination that has the same value as the line impedance. As a result, all of the energy in the electromagnetic field is absorbed and the voltage level on the transmission line stabilizes at 5 volts which is an ideal logic 1 for this circuit.

Switching From a Logic 1 to a Logic 0

When the circuit in Figure 1 switches from a logic 1 to a logic 0 the driver has the task of removing the charge on the line capacitance that was put there in order to move it from a logic 0 to a logic 1. To do this the driver level moves internally from 5V to 0V. As with the transition from a logic 0 to a logic 1 the equivalent circuit is like that shown in Figure 3, but, now, the line is at 5V and the output impedance and series terminating resistor are connect to 0V. The voltage divider is at work as it was before.

As a result, the line voltage is moved to V/2 and charge is removed from the line capacitance to this level as the energy moves down the line. (The voltage level of this transitions is –V/2.) When the EM field arrives at the end of the transmission line 2 nanoseconds later it encounters an open circuit and is reflected back down the line. The result after the reflection takes place is the line is now at 0V. Two nanoseconds later the EM field arrives back at the driver and encounters the circuit shown in Figure 5 and is absorbed. The resulting waveform is shown in Figure 6.

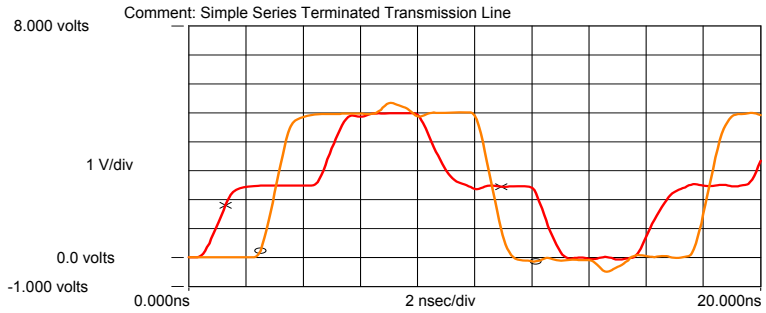

Figure 6. Voltage Waveforms at the Two Ends of the Transmission Line after Switching From 1 to 0

Notice that the voltage waveform at the receiver (orange) is a proper square wave logic signal which is the goal of this signal path. This signaling method is known as “reflected wave” switching because the correct logic level is created by the reflected wave as it makes its round trip along the transmission line. This is the lowest power consumption method of logic signaling because current is only being drawn from the power system while the line is being charged. Once the line has been fully charged to a logic 1 the current draw goes to zero.

This is the switching method that is employed with the PCI bus that is incorporated in most personal computers.

Also, notice that the voltage waveform at the driver output is at an indeterminate logic state for a time that is the round trip delay along the transmission line each time switching takes place. If loads are placed along the length of the transmission line, as is done with the PCI bus, they do not experience a “data good” condition until the reflected wave passes by them on the return trip. Therefore, clocking of data at these inputs must

be delayed until data is good at all inputs. This is how data is clocked on the PCI bus and other bus protocols that rely on reflected wave switching.

What Happens When The Driver Impedance Does Not Match The Line Impedance?

The circuit shown in Figure 7 is the same as that shown in Figure 1 except that the series termination has not been inserted in series with the output.

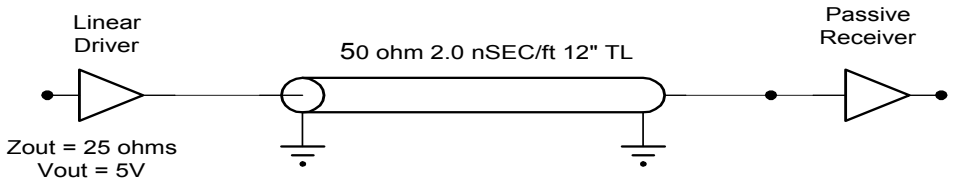

Figure 7. An Unterminated 5V CMOS Transmission Line

Figure 8 shows the switching waveform for the transition from a logic 0 to a logic 1. Notice that the bench voltage is much higher than V/2. In fact it is 2V/3 or 2/3 of the total of 5 volts or 3.33V. Why is this? If you refer to the voltage divider in Figure 3 in this example the upper resistance is 25 ohms or Z_{out} of the driver and the lower resistance or impedance is 50 ohms producing the 2/3 voltage level.

The EM field is charging the line capacitance to this value as before. When the EM field arrives at the receiver 2 nanoseconds after being generated it is reflected back, doubling the voltage to 6.66V. As before, the EM field charges the line capacitance up to 6.66V. After another 2 nanoseconds the EM field arrives back at the driver and encounters a termination like that shown in Figure 5. However, the parallel termination is not 50 ohms. Instead it is 25 ohms. Two things will happen. First, the voltage divider this time is 50 ohms on top and 25 ohms on the bottom as shown in Figure 9 with the series terminator value being zero ohms, so the voltage is divided down. Second, not all of the energy will be absorbed.

When an EM filed encounters a parallel termination that is lower in value than the TL, the reflected energy will be the opposite polarity of the incident waveform. This cannot be seen at the driver. Two nanoseconds later the energy arrives at the receiver and, as can be seen, it is inverted or is negative going.

As before, the amount of energy will double the voltage level at the receiver and travel back toward the driver. When it arrives at the driver some of it is absorbed and the rest is reflected inverted. This goes on until such time as all of the energy has been absorbed in the driver output impedance and the logic level settles out at 5V. This can be seen in Figure 10.

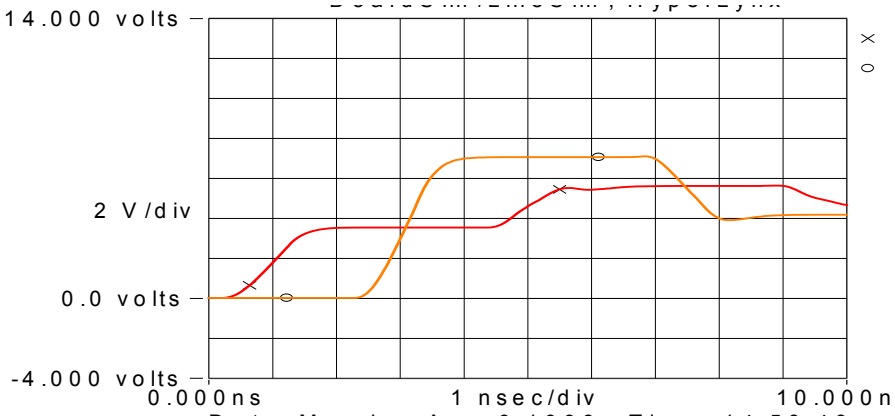

Figure 8. Switching Waveform for an Unterminated CMOS Transmission Line

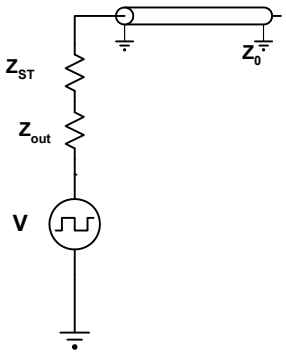

Figure 9. Switching Waveform for an Unterminated CMOS Transmission Line

When a parallel termination does not match the impedance of the transmission line across which it is placed, it will not absorb all of the energy reflected back down the TL.

If the value of the termination is larger than the TL impedance the energy will be reflected back with the same polarity as the incident waveform and is often called “overshoot”.

If the value of the termination is smaller than the TL impedance the energy reflected back will be inverted or the opposite polarity of the incident waveform and is often called “undershoot”.

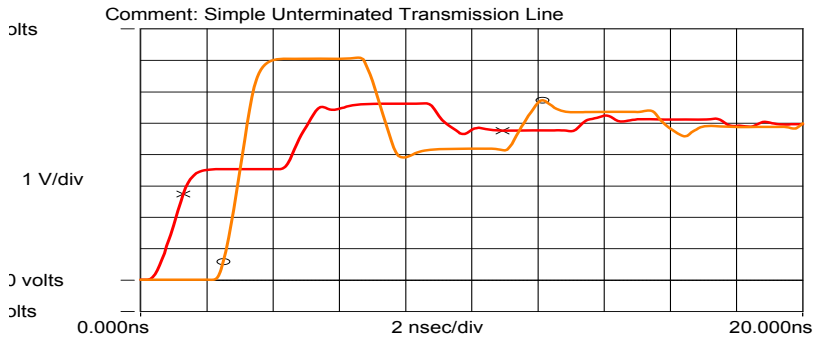

Figure 10. Switching Waveform for an Unterminated CMOS Transmission Line

There are two problems with the waveform in Figure 10. First, the voltage goes 1.66 volts above Vdd. This excess voltage can cause logic failures or damage the receiver. Second, after the signal arrives back at the drive and is inverted, it cause the logic 1 at the receiver to drop to below 4 volts. This diminishes the logic one to a level that could result in a logic failure. Neither of these is good. That is why a series termination is added to a circuit like this.

Figure 11 shows the waveform when the signal switches to a logic zero. As you can see, the same level violations happen in this logic state.

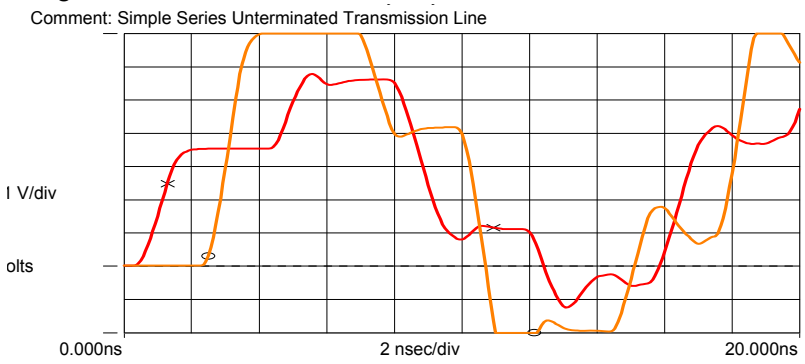

Figure 11. Another Switching Waveform for an Unterminated CMOS Transmission Line