creases, up to a maximum amount (reflected in the coupling coefficient formula). At that point the amplitude levels out and stays constant.

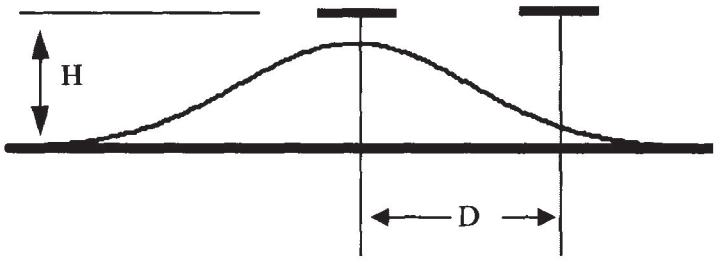


Figure 12-3 Current density under the aggressor trace can couple into an adjacent trace with a coupling coefficient proportional to H^2 and inversely proportional to D^2.

The length of the coupled region that just results in the full amplitude of the crosstalk signal is called the critical length. It is the same critical length we talked about with transmission lines, the length where the round-trip propagation time equals the rise time of the signal. Another way to look at this same question is to look at the round-trip propagation time of the coupled region and test whether that time is greater than the rise time. That is, calculate the round-trip propagation time, T_{RT} (in ns), using the formula discussed in Chapter 2:

For stripline

$$T_{RT} = .085\sqrt{\varepsilon_r} \text{ x coupled length x } 2 \qquad [12\text{-}2]$$

where ε_r = relative dielectric coefficient
and the coupled length is measured in inches.

Then, test if T_{RT} is greater than the rise time, T_r, of the aggressor pulse in nanoseconds (Equation 12-3).

$$T_{RT} > T_r \qquad [12\text{-}3]$$

If $T_{RT} < T_r$, we decrease the coupling ratio in Equation 12-1 by the ratio given in Equation 12-4.

INTEL-1010
U.S. Patent 7,151,430

$$T_{RT}/T_r \qquad [12\text{-}4]$$

Otherwise we leave the coupling ratio as it is.

Terminations Finally, we look and see if the victim trace is terminated, and if so how effectively. If two traces are coupled, backward crosstalk will develop. But if the victim trace is a transmission line that is terminated in its characteristic impedance at the near end (Figure 12-4), then the backward crosstalk pulse will be completely absorbed by the terminating resistor and will not reflect forward.

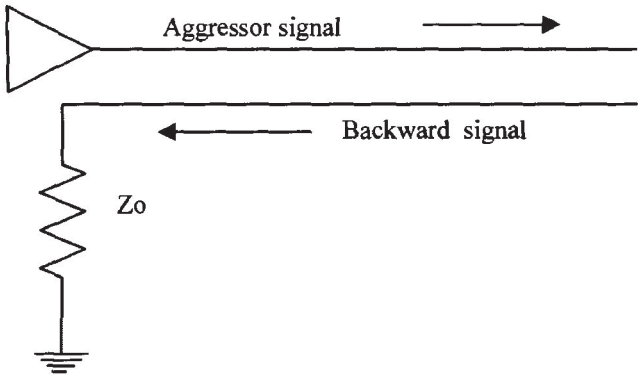


Figure 12-4 Near end termination of the victim trace can absorb backward crosstalk.

In this illustration, there will be no crosstalk at the far end of the victim line because:

- There is minimal or no forward crosstalk anyway.
- The backward crosstalk is absorbed by the terminating resistor (Zo) and will not reflect forward.

This solution may or may not be an effective alternative for your design. First, it may be difficult or impossible to calculate the correct value for a terminating resistor at the near end of the victim line. It may also be impossible to place a terminating resistor there without adversely affecting the operation of the circuit. And, it is frequently the case that individual traces can be both aggressor and victim, depending on when switching occurs. So, although this knowledge may have

limited use in eliminating crosstalk, it may be very helpful in making judgments about how destructive crosstalk signals may be to the normal circuit operation.

UltraCAD Calculator UltraCAD has created a freeware calculator for estimating crosstalk based on the preceding discussion (see Appendix H). It is illustrated in Figure 12-5. A detailed Help file explains the operation and the formulas used. UltraCAD does not guarantee any level of accuracy with this calculator, and suggests that it be "calibrated" first against prior designs known to be good and those known to have crosstalk problems. Nevertheless, even though the company does not guarantee any level of accuracy, it is still fun to sometimes compare it against other tools.

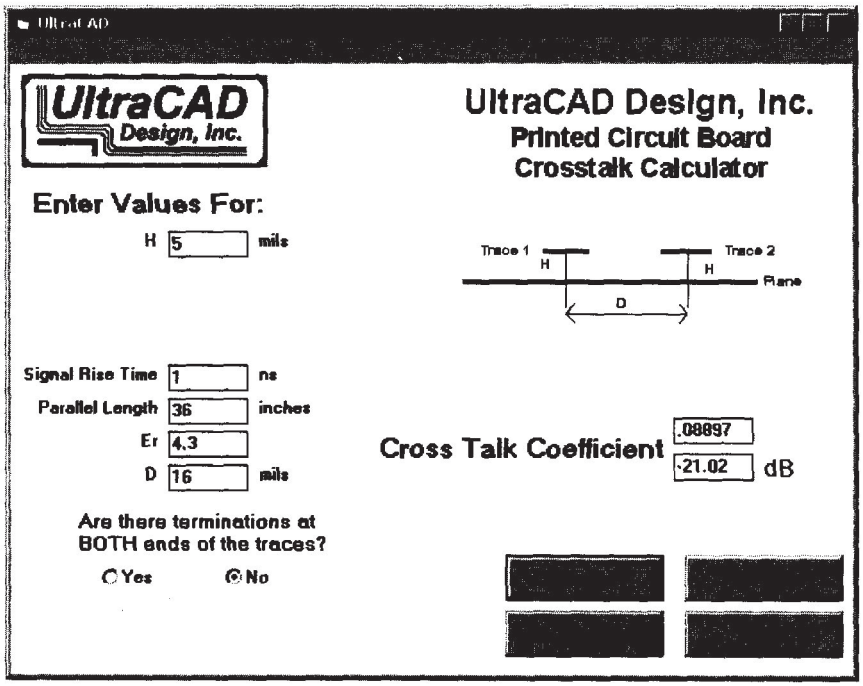


Figure 12-5 UltraCAD's freeware crosstalk calculator.

Figure 12-5 shows the calculated crosstalk coefficient for two 36-inch traces spaced 16 mils apart, each 5 mils above the reference plane, with $\varepsilon_r = 4.3$. The calculated coefficient is .08897. This means that you would multiply the aggressor voltage by this coefficient to calculate the crosstalk voltage on the victim line.

HyperLynx Simulation Tool Figure 12-6 illustrates a HyperLynx model with the same parameters as used with the UltraCAD calculator. (More about how to work with HyperLynx crosstalk models is presented in the next chapter.) The driver voltage in the HyperLynx model is measured as 4.4 volts. Resistor RS(A1) is placed simply to make sure there is no attenuation in the backward crosstalk signal caused by U(A1). Similarly, the load resistor RD(B1) is set to a high value so as not to load down the far-end crosstalk signal. The reverse crosstalk noise is measured as 406 mV at the resistor at the far end (RD(B1)). If we take the 4.4 volts and multiply it by the crosstalk coefficient calculated from the UltraCAD crosstalk calculator (.08897), we get 4.4 x .08897 = 391 mV as the predicted magnitude of the coupled signal. The HyperLynx model predicts 406 mV. Not too bad for a coarse estimating tool.

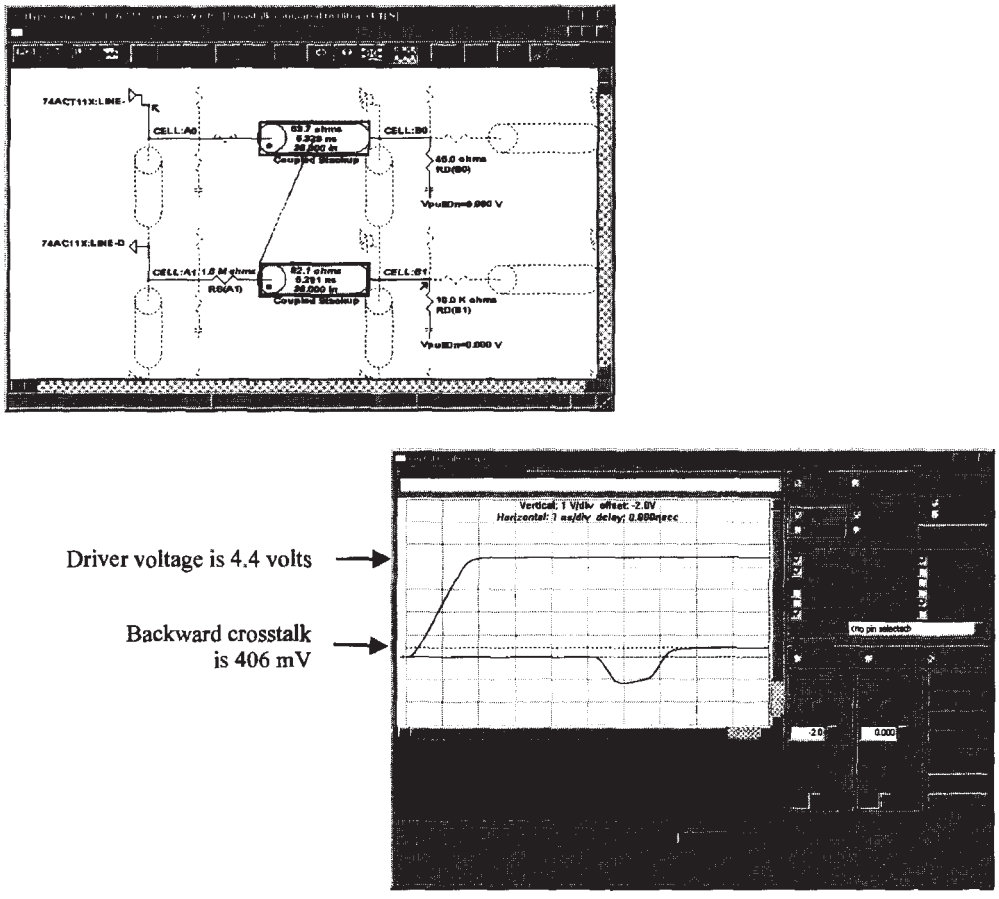


Figure 12-6 HyperLynx model for a pair of 36-inch traces placed 16 mils apart. The model predicts 406 mV backward crosstalk.

DESIGN CONSIDERATIONS

It should be clear by now that the designer has only a few tools to work with in controlling crosstalk. Crosstalk can be reduced by:

1. Using stripline configurations (eliminating forward crosstalk).
2. Using stackups that place traces as close as possible to their reference planes.
3. Separating traces as far apart as possible.
4. Making judicious use of terminations for further reducing crosstalk (but determining the extent is a very complicated process).

Slots in Planes, 3 There is one other crosstalk design consideration that designers need to remember when laying out boards. Figure 12-7 illustrates yet another case where routing traces across a slot in a plane can have serious effects (refer back to Figures 9-7 and 10-18). The signal traces may meet all the criteria specified for good crosstalk control, especially separation. Remember, however, that the signal *returns* want to travel directly underneath the traces. When the returns meet the slot, they must travel around the slot and then come back to their position under the traces. During this path, the returns are very close together. So, while it is true that the traces do not crosstalk with each other, the returns might!

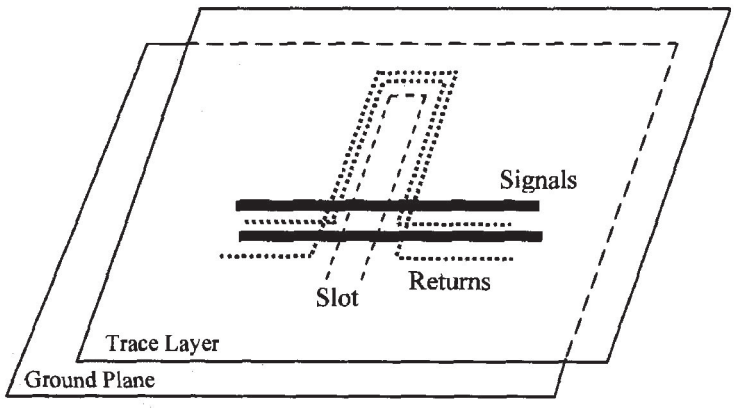


Figure 12-7 When return signals have to divert around a slot or discontinuity in a plane, crosstalk problems can result.

INTEL-1010
U.S. Patent 7,151,430

When this board goes to test, there might be a crosstalk problem. The engineers may spend hours going over the board, trying to find that problem. If they are not alert to the existence of the slot in the plane, they may never find the fundamental problem. After several other fixes and a redesign that no longer has the slot, the board may finally meet the desired specifications. All the other fixes may be credited with solving the problem, and a new set of erroneous design rules will result. Don't cause these kinds of problems for your engineers and companies during the design stages of your boards.

Guard Bands Guard bands are sometimes used to isolate traces from each other. The idea is to place a trace, called a *guard band*, between the aggressor and the victim traces. The guard band trace is grounded. Thus, it acts as an electrical shield between the other two traces, much like the shield in a shielded cable might.

The most common application for guard bands is when crosstalk is an issue. The idea is that any coupling that takes place will be to ground (the guard band) and not into the next, more sensitive signal trace.

Although there is not uniform agreement on this issue, most people feel that guard banding is ineffective at best and can actually be harmful at worst. There are several issues that can be of concern here. The first is how and where to ground the guard trace. It certainly must be grounded at one end, but many people say it should be grounded at both ends. The problem with this solution is that now the possibility for ground loops has developed. That is, there is the possibility that stray currents may find their way onto the guard trace at one end and propagate down the guard trace and back to ground at the other end. Ground loops can have very serious effects on performance, not the least of which is the creation of EMI problems.

Still others say that guard traces should be "stitched" to ground every so often along their length. The stitch spacing is an additional point of disagreement among various engineers.

Finally, there are people who say that guard banding actually can make crosstalk worse. The signal from one trace will couple into the guard trace and then that signal will couple from the guard trace to the next signal, making the crosstalk between the two signal traces worse than if there had been no guard band at all.

We know of no definitive studies that clarify this issue one way or the other. The anecdotal evidence that exists supporting guard banding is often explained away by saying that the benefit does not come from the guard trace itself,

INTEL-1010
U.S. Patent 7,151,430

but from the simple act of spreading apart the two signal traces to make room for the guard band. The benefit comes from separating the traces, not by putting another trace in between them.

We do not have a clear answer here, but we can observe that the two times our customers have insisted on placing significant amounts of copper between signal traces, the boards didn't work. They had serious crosstalk problems. The redesigned boards, without the copper floods, worked significantly better. We don't recommend guard banding to our customers.

INTEL-1010
U.S. Patent 7,151,430



13 CROSSTALK SIMULATIONS

*I*n this chapter we look at a series of crosstalk simulations based on the HyperLynx LineSim tool. This tool has an excellent capability for simulating a hypothetical crosstalk situation without the necessity of building up a board design (model) or requiring actual IC models. Therefore, it is a useful tool for looking at and explaining concepts.

BASIC MODEL

Figure 13-1 illustrates a basic coupled model. The HyperLynx model work surface has the various icons already placed on it. You activate them by either selecting or deselecting them. As constructed in Figure 13-1, there are two nets in this basic model. The first consists of a 74ACT11X line driver (U(A0)) connected through a transmission line to a 53-Ω resistor (RD(C0)). The second net consists of a 74ACT11X device set up in receive mode, connected through a transmission line to a 53-Ω resistor (RD(C1)).

There is provision for another transmission line between each IC and the actual transmission lines shown in the model. This provision is shorted out in this model. The reason for that will become evident later. Each coupled transmission line has a 53.3-Ω characteristic impedance and is 67.5 inches long. HyperLynx calculates the delay through each line as approximately 10 ns. The reason for the extraordinary length will also become apparent later.

In this model the transmission lines are based on a hypothetical stackup illustrated in Figure 13-2. The important parameters to note are that this is a microstrip configuration, with 8 mil trace widths separated by 5 mil spacing, with the

INTEL-1010
U.S. Patent 7,151,430

traces 5 mils above the underlying plane. The relative dielectric coefficient is set at 4.3.

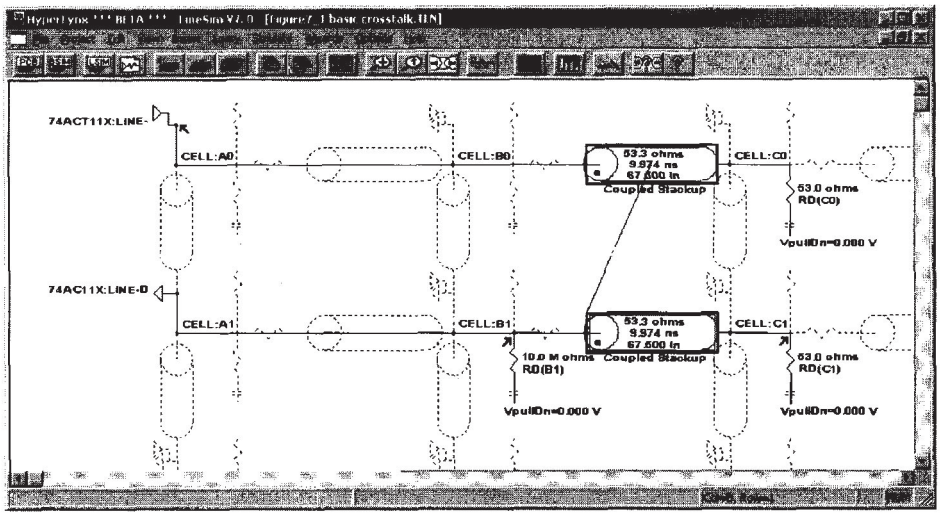


Figure 13-1 A basic HyperLynx model for crosstalk.

You will note that there are three "arrows" in Figure 13-1, pointing to U(A0), RD(B1), and RD(C1). These correspond to oscilloscope probes, or points that will be displayed on the oscilloscope screen during the modeling. The 10.0-MΩ resistor (RD(B1)) is simply placed on the model as a place to put a probe so we can display the signal at that point

Figure 13-3 shows the result of the model. The model is set to use an edge-driven driver waveform with a rising edge and a "Fast-Strong" (fast rise time) step-function signal. The vertical scale is set to 1 volt/division and the horizontal scale to 5 ns/division. Three waveforms, corresponding to the three probe points (arrows), are shown on the screen.

Recall in Chapter 12 we said that forward crosstalk was small in microstrip configurations and virtually nonexistent in stripline configurations, and also that forward crosstalk increases with increasing length. The reason for the extraordinary lengths of these traces is to generate forward crosstalk signals that are clearly visible. Later in this chapter we use more realistic lengths to see what might really happen on normal circuit boards.

INTEL-1010
U.S. Patent 7,151,430

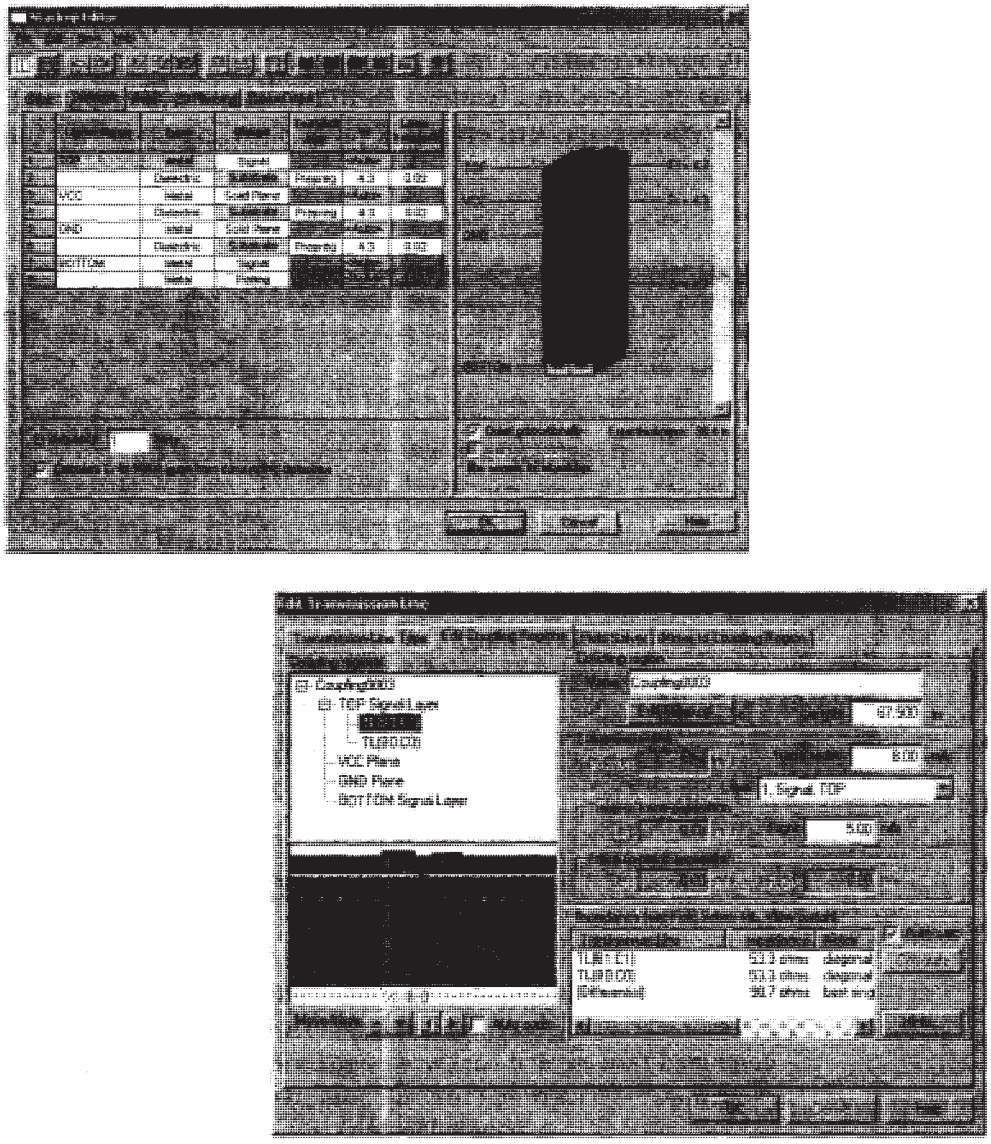


Figure 13-2 Screens for defining and editing the stackup and coupling region between the coupled transmission lines in Figure 13-1.

The driver signal (a) is shown rising immediately on the left axis. There is a coupled backward crosstalk signal (b) that immediately reflects off the receiver and starts forward down the victim trace. Note that this signal is approximately 20 ns wide, just twice as wide as the coupled region (exactly as predicted in Chapter 12).

Ten ns after the aggressor pulse rises (the delay of the coupled region) the forward crosstalk pulse (c) shows up as a negative-going signal at RD(C1). This is a sharp pulse whose width is approximately 2 ns (approximately the rise time of the aggressor pulse, again just as predicted in Chapter 12). Immediately following the forward crosstalk signal is the 20-ns backward crosstalk pulse showing up at RD(C1).

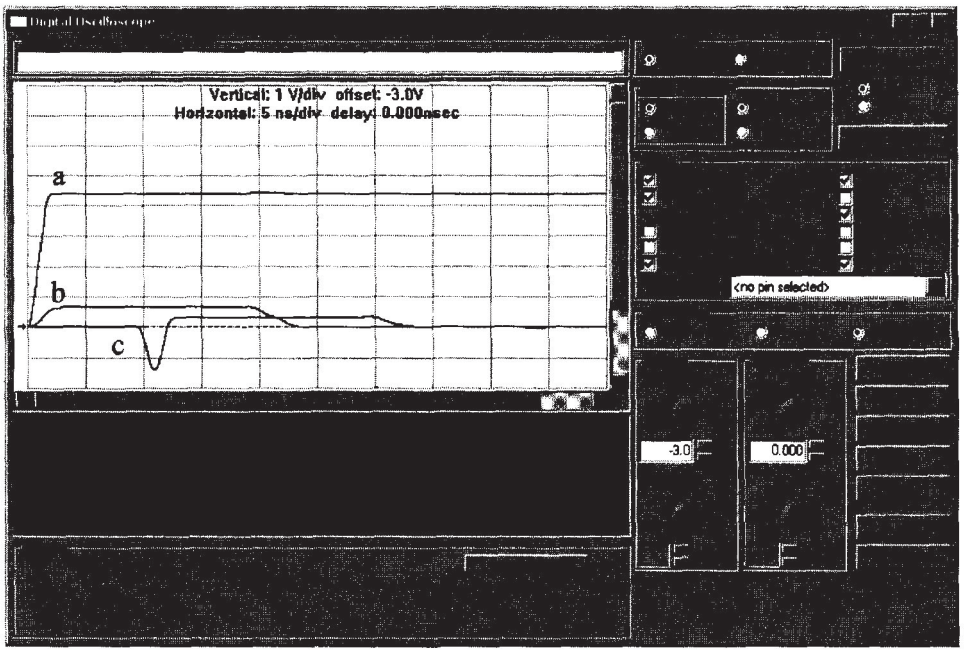


Figure 13-3 Actual results of the model. Traces a, b, and c represent the aggressor signal, the backward crosstalk signal at at RD(B1), and the combined forward and backward signals at RD(C1), respectively. Note how the backward crosstalk pulse width (20 ns) is twice the coupled length (10 ns).

The backward crosstalk signal at RD(C1), the far end of the victim trace, is exactly one-half the amplitude of the backward crosstalk signal at RD(B1), the near end of the victim trace. They are the same signal. But the signal at RD(C1) (far end) is delayed by 10 ns (the length of the line) and is reduced by 50% at RD(C1) because of the voltage divider action formed by the 53-Ω transmission line and the 53-Ω terminating resistor. That is, one-half the voltage at RD(B1), at the near end, ap-

pears across RD(C1) and the other half appears across the trace (transmission line) itself.

ADD AN UNCOUPLED REGION

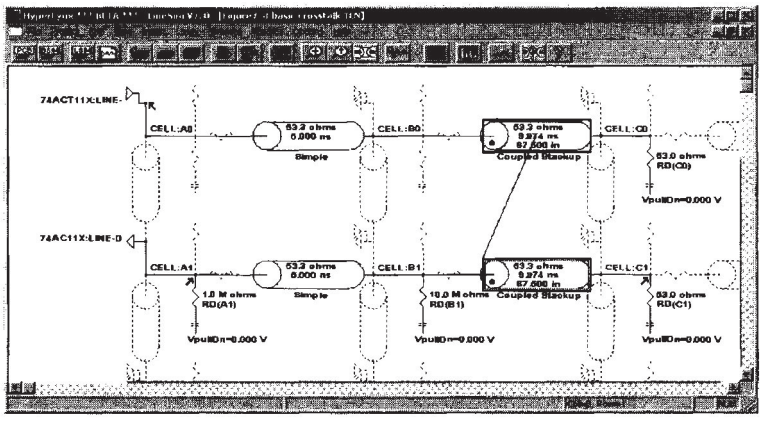
Now let's add a little complexity. Figure 13-4 is the same model with the addition of another 5-ns long uncoupled transmission line section placed in front of the coupled region. Think of this as a pair of traces that are widely separated over part of their length (5 ns), then coupled together over another part of their length (10 ns).

Note how the backward crosstalk signal (b) shows up at RD(A1) 10 ns after the driver signal (a) rises. That 10 ns consists of 5 ns for the driver signal to get to the coupled region, and then another 5 ns for the backward crosstalk signal to travel *back* from the coupled region to RD(A1). This backward crosstalk pulse width is still 20 ns wide, twice the length of the coupled region. It is also the same magnitude as before.

The forward crosstalk signal arrives (c1) at RD(C1) 15 ns after the driver signal rises. That 15 ns is made up of the 5 ns it takes the driver signal to get to the coupled region and then 10 ns for the forward crosstalk signal to travel through the coupled region. The forward crosstalk signal is the same magnitude as before.

It takes about 25 ns for the backward crosstalk signal (c2) to arrive at RD(C1). That 25 ns consists of the 5 ns for the driver signal to get to the coupled region, plus 5 ns for the backward crosstalk pulse to travel backward to U(A1) where it reflects forward again, plus 5 ns back through the uncoupled length of transmission line, plus the final 10 ns to travel through the coupled region. By inserting this 5-ns length of uncoupled transmission line into the model, we have completely separated the forward and backward pulses so that they arrive at the far end on the victim trace at entirely different times. Think, now, of the engineer who might be troubleshooting this simple circuit. The crosstalk signals are showing up at the far end of the victim trace at different times, and the backward crosstalk signal is a different amplitude depending on where you measure it.

INTEL-1010
U.S. Patent 7,151,430



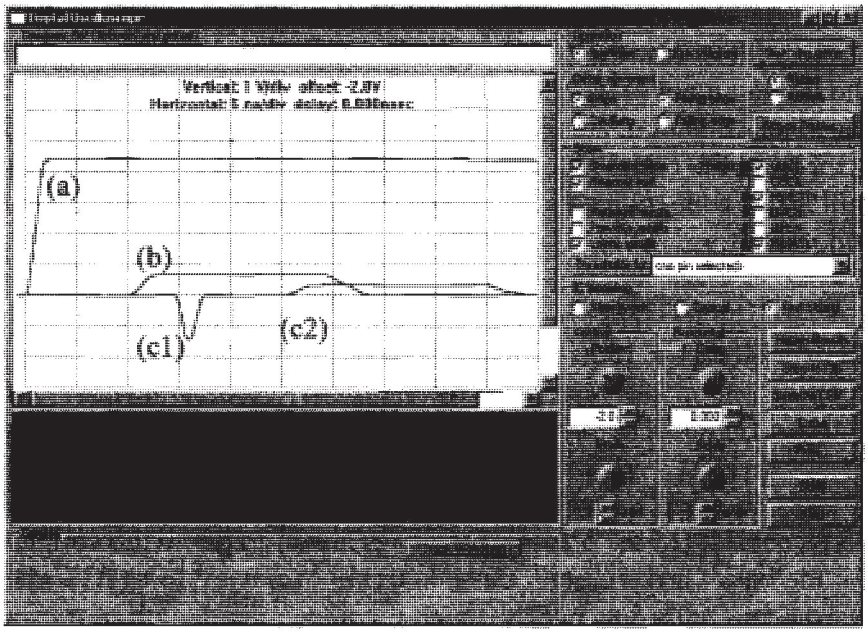


Figure 13-4 We are adding a 5-ns uncoupled region in front of the coupled region. This effectively separates the backward and forward crosstalk signals at the far end of the victim line.

EFFECT OF LENGTH

In Chapter 12 we pointed out that as the coupled region increases:

- The forward crosstalk pulse amplitude increases.
- The forward crosstalk pulse width remains constant.
- The backward crosstalk amplitude remains constant.
- The backward crosstalk pulse width increases with the coupled region.

These effects are easily simulated with the HyperLynx simulator. Refer back to the model shown in Figure 13-1. This time we are going to run the exact same model three times, with coupled region lengths equal to (a) 33.75 in., (b) 67.5 in., and (c) 135 in., respectively. This will correspond to coupled regions of approximately 5 ns, 10 ns, and 20 ns, respectively. Figure 13-5 illustrates the results.

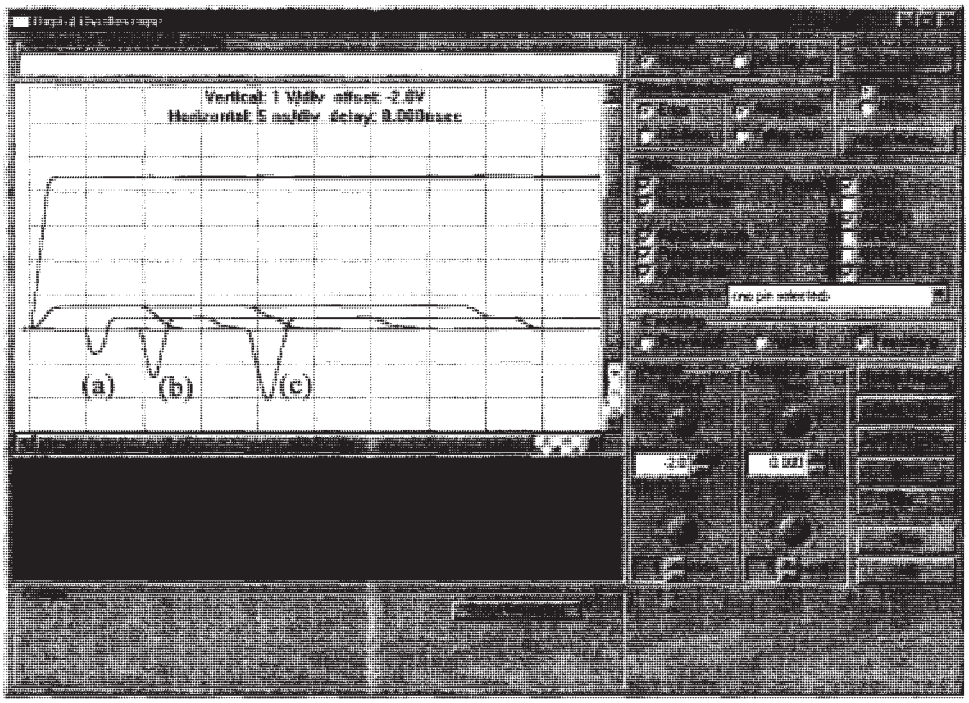


Figure 13-5 Changing the coupled length causes the backward crosstalk pulse to increase in width and the forward crosstalk pulse to increase in amplitude. The modeled lengths are (a) 33.75 in., (b) 67.5 in., and (c) 135 in.

Note how the amplitude of the forward crosstalk pulse increases as the length of the coupled region increases. Note also, however, that the pulse width of the forward crosstalk pulse remains constant, regardless of the coupled length.

Notice how the amplitude of the backward crosstalk pulse does not change as the coupled region changes, but its pulse width does. The backward crosstalk pulse widths are 10 ns, 20 ns, and 40 ns, respectively, twice the length of the coupled region for each case.

STRIPLINE

All the previous models assumed a microstrip configuration. Let's now do similar simulations, this time using a stripline configuration. Figure 13-6 illustrates the setup. We are using effectively the same trace dimensions as before (8 mil wide, 5 mil spacing, 67.5-inch length), but this time the traces are on an inner stripline layer (part of a dual stripline stackup). Because the traces are now in a stripline environment, some of the trace characteristics change. The characteristic impedance is now shown to be 57.3 Ω, and the delay down the line is just under 12 ns.

Figure 13-7 illustrates the result of the simulation. The result looks almost exactly like Figure 13-3, with slightly different pulse widths and voltage levels because the stripline parameters are slightly different. There is one very important exception. There is no forward crosstalk pulse, even with a 67.5-inch (5.6-*foot*) line.

This illustrates the point made earlier that in stripline environments the capacitive and inductive components of forward crosstalk are almost exactly equal and opposite—and they cancel out. Thus, if we are concerned about crosstalk, put all the traces in a stripline environment and one of the crosstalk components goes away. It is as simple as that.

STRIPLINE WITH TERMINATIONS

Now, if you liked that, you will positively love what comes next. In Figure 13-8 we modify the stripline environment a little further and add a terminating resistor, shown as RD(A1), at the near end of the victim line. This terminating resistor will absorb the backward crosstalk flowing toward U(A1), so there will be no reflection from the near end. Now, when we do the simulation (Figure 13-8) we see (a) the aggressor pulse leaving the driver, and (b) the backward crosstalk pulse as it propagates back to RD(A1). However, virtually no pulse is reflected toward the far end of the coupled region (RD(C1)).

INTEL-1010
U.S. Patent 7,151,430

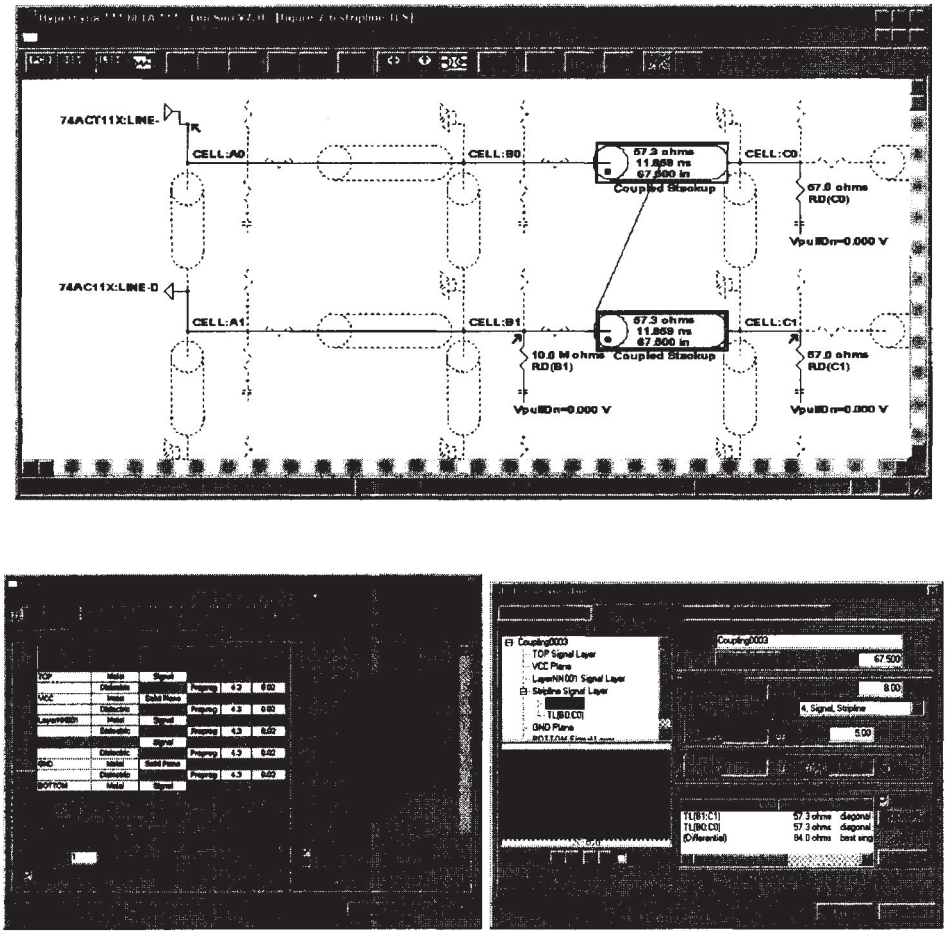


Figure 13-6 HyperLynx model set up for stripline.

Look at what we have here. Traces that are 5.6 *feet* long, spaced 5 mils apart, and there is virtually no crosstalk at the far end, (and you thought crosstalk was a difficult problem).

The practical reality is that we may not be able to place terminating resistors on our boards or in our circuits as easily as we can in this model. If not, we probably can't eliminate crosstalk at the far end completely. And we really haven't eliminated backward crosstalk at the near end, just absorbed it so it won't reflect back. So it is not correct to say that we just eliminated crosstalk, but we certainly have knocked it down some.

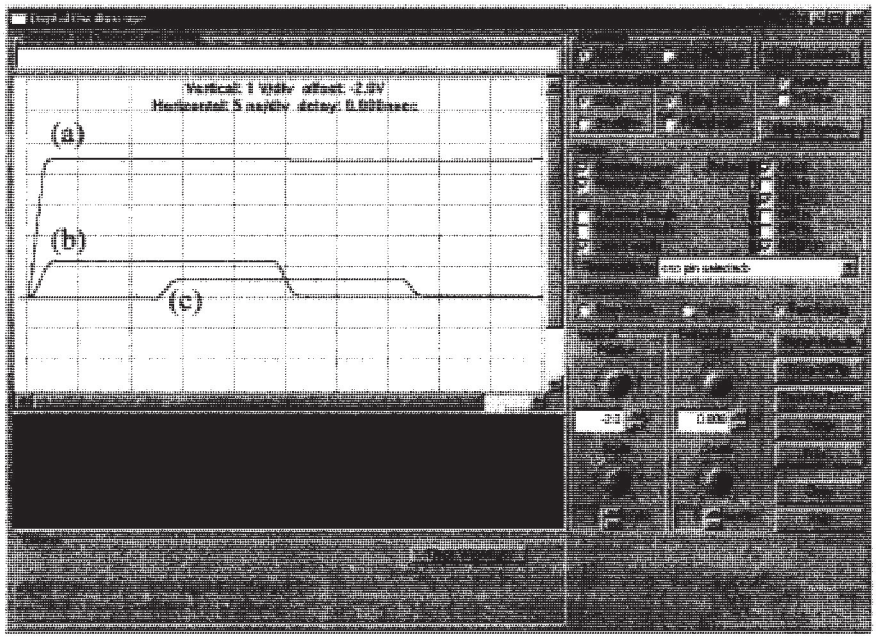


Figure 13-7 Stripline results showing drive signal (a), backward crosstalk signal at RD(B1) (b), and backward crosstalk signal at RD(C1) (c). Note the absence of a forward crosstalk pulse.

It is this final interaction of termination strategies and backward crosstalk reflections that make practical assessments of crosstalk very difficult for circuit design engineers to estimate and plan for.

MORE REALISTIC EXAMPLE

All of these illustrations used unrealistic trace lengths to create noise signals big enough to easily see and analyze, but what do real signals look like? Figure 13-9 shows a simulation of exactly the same setup as Figure 13-6 but with only a 12-inch (2.1-ns) coupled region. The backward crosstalk reflection, when it shows up at the far end (c), is pretty significant, approximately 600 mV. This design, however, is not a very good design for crosstalk control. The traces are 10 mil off the plane and there is only 5 mil spacing between the traces.

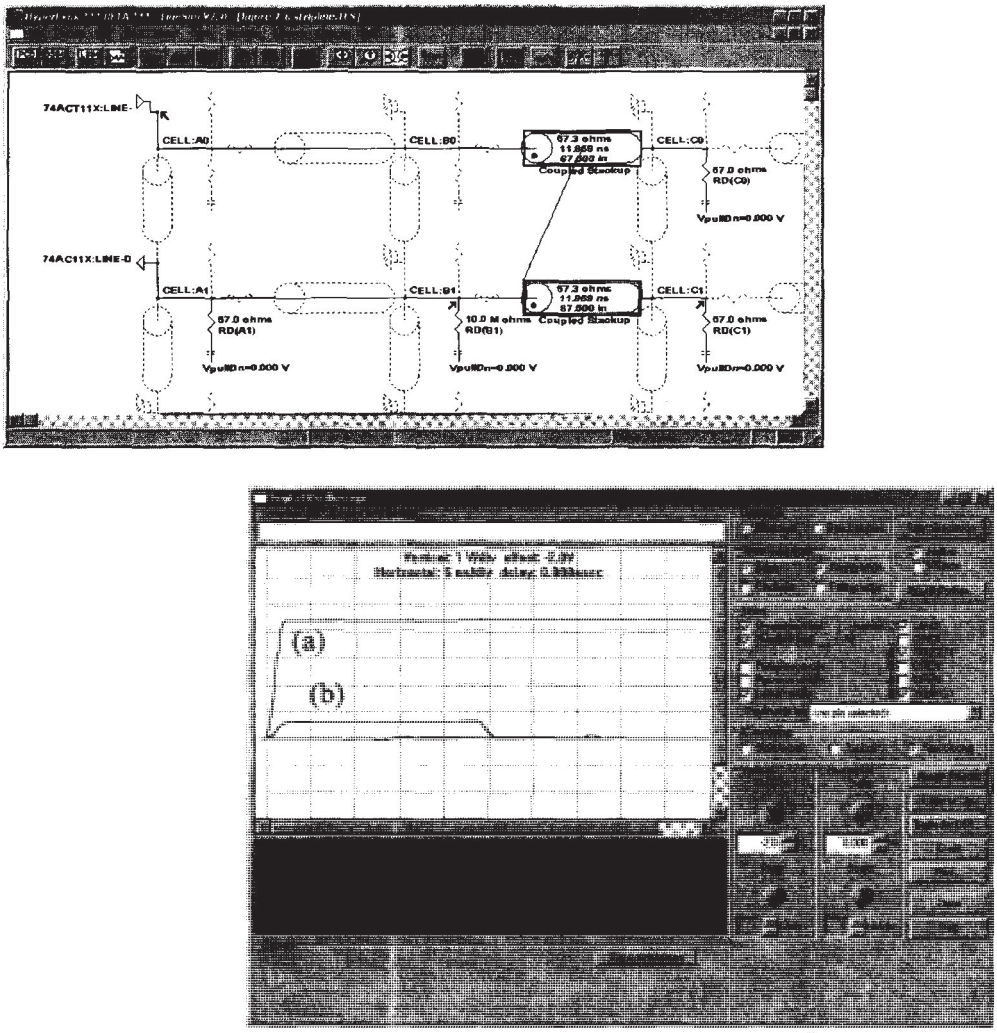


Figure 13-8 Combining a stripline with near-end termination can completely eliminate all crosstalk from the far end of the victim trace.

A better design would have the traces perhaps 5 mil above the planes and then use 16-mil spacing (with our 8-mil traces). Figure 13-10 illustrates this simulation and shows how dramatically the crosstalk is reduced. Crosstalk at the far end is now approximately 100 mV. (Remember, this simulation is using no terminating resistors at the near end.) Whether this level of crosstalk is "good enough" is up to the circuit design engineer to decide.

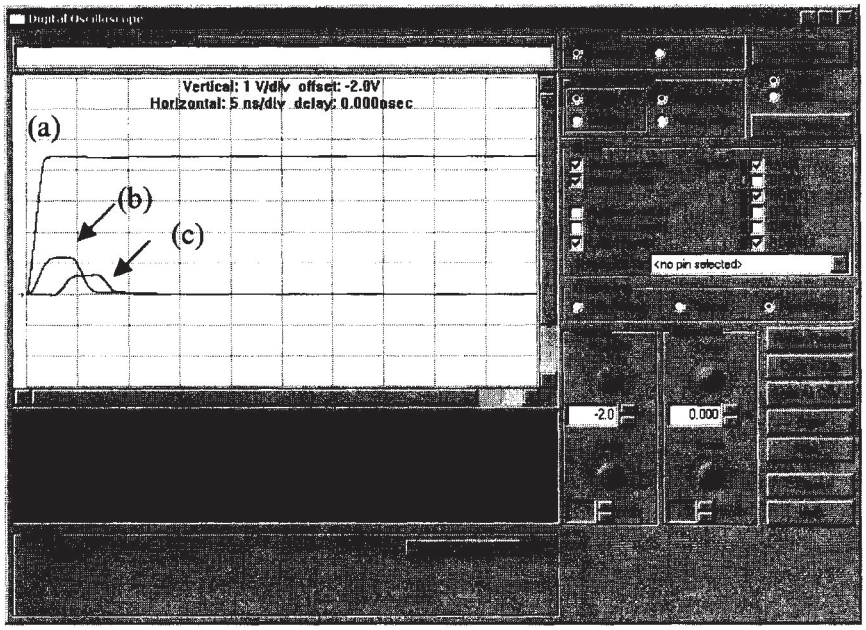


Figure 13-9 Simulation of the stripline example with only a 12 inch coupled region. Backward crosstalk (c) is still somewhat large.

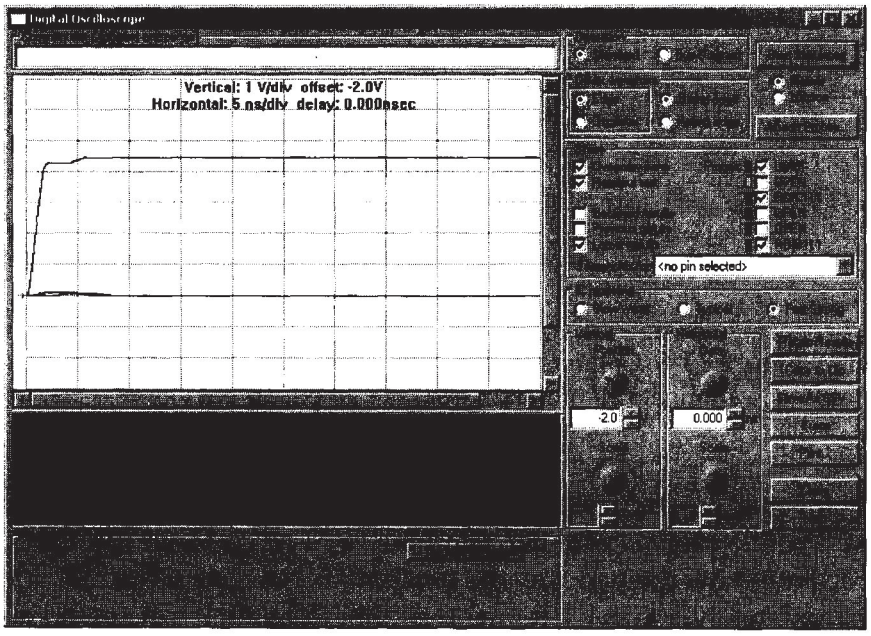


Figure 13-10 Crosstalk model results from a more normal design made with crosstalk control in mind.

SUMMARY

These simulations confirm the steps we need to take to control crosstalk:

- Place traces in stripline environments.
- Minimize the distance between the trace and the plane.
- Maximize the separation between the traces.
- Consider the beneficial effects of near-end terminations.

INTEL-1010
U.S. Patent 7,151,430



14 DIFFERENTIAL TRACES AND IMPEDANCE

BACKGROUND

We generally think of signals propagating through our circuits in one of three commonly understood modes: single-ended, differential mode, or common mode.

Single-ended mode is the mode we are most familiar with. It involves a single wire or trace between a driver and a receiver. The signal propagates down the trace and returns through the ground system.[1]

Differential mode (more properly called *odd mode*) involves a pair of traces (wires) between the driver and receiver. We typically say that one trace carries the positive signal and the other carries a negative signal that is both equal to, and the opposite polarity from, the first. Since the signals are equal and opposite, there is no return signal through ground; what travels down one trace comes back on the other.

Common mode (more properly called *even mode*) signals are those that travel in the same direction on both traces. They are generally created by some sort of unwanted noise or variation from unexpected conditions.

Advantages Differential signals have one obvious disadvantage over single-ended signals. They require two traces instead of one, or twice as much board area. However, they also possess several advantages:

[1]In truth the signal can return through either or both the ground or power system. I use the singular term *ground* throughout this chapter simply for convenience.

INTEL-1010
U.S. Patent 7,151,430

1. Differential circuits can be very helpful in low-signal-level applications. If the signals are very low level, or if the signal/noise ratio is a problem, then differential signals effectively double the signal level:

   $$(+v - (-v) = 2v)$$

   Differential signals and differential amplifiers are commonly used at the input stages of very low-signal-level systems.

2. Since differential signals are (by definition) equal and opposite, there is no return signal through any other path. If there is no return signal through ground, then the continuity of the ground path becomes relatively unimportant. If we have, for example, an analog signal going to a digital device through a differential pair, we don't have to worry about crossing power boundaries, plane discontinuities, and so on. Separation of power systems can be made easier with differential devices.

3. Differential receivers tend to be sensitive to the *difference* in the signal levels at their inputs, but they are usually designed to be *insensitive to common-mode* shifts at the inputs. Therefore, differential circuits tend to perform better than single-ended ones in high-noise environments.

4. Switching timing can be more precisely set with differential signals (referenced to each other) than with single-ended signals (referenced to a less precise reference signal subject to noise at some other point on the board). The crossover point for a differential pair is very precisely defined (Figure 14-1). The crossover point of a single-ended signal between a logical one and a logical zero, for example is subject to noise, noise threshold, and threshold detection problems.

INTEL-1010
U.S. Patent 7,151,430

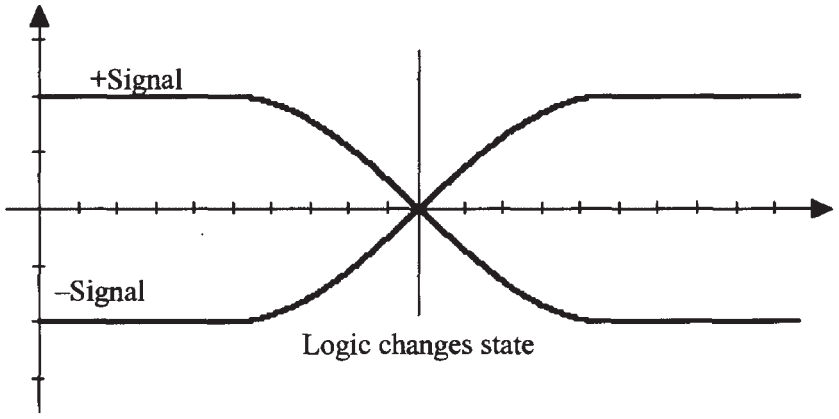


Figure 14-1 Timing with differential signals can be very precise.

Key Assumption There is one very important aspect to differential signals that is frequently overlooked, and sometimes misunderstood, by engineers and designers. Let's start with the two well-known laws that (a) current flows in a closed loop and (b) current is constant everywhere within that loop. Consider the positive trace of a differential pair. Current flows down the trace and must flow in a loop, normally returning through ground. The negative signal on the other trace must also flow in a loop and would also normally return through ground. This is easy to see if we temporarily imagine a differential pair with the signal on one trace held constant. The signal on the other trace would have to return somewhere, and it seems intuitively clear that the return path would be where the single-ended trace return would be (ground). We say that, with a differential pair, there is no return through ground not because it can't happen, but because the returns that do exist are equal and opposite and therefore combine to zero and cancel each other out.

This is a *very* important point. If the return from one signal (+i) is exactly equal to, and the opposite sign from, the other signal (–i), then they combine to zero and there is no current flowing anywhere else (and in particular, through ground). Now assume the signals are not *exactly* equal and opposite. Let one signal be +i1 and the other be –i2 where i1 and i2 are similar, but not exactly equal, in magnitude. The combination of their return currents is (i1 – i2). Since this is not zero, then this incremental current must be returning somewhere else, presumably ground.

So what, you say? Well, let's assume the sending circuit sends a differential pair of signals that are, in fact, exactly equal and opposite. Then we assume they will still be so at the receiving end of the path. But what if the path lengths are dif-

INTEL-1010
U.S. Patent 7,151,430

ferent? If one path (of the differential pair) is longer than the other path, then the signals are no longer equal and opposite during their transition phase at the receiver (Figure 14-2). If the signals are no longer equal and opposite during their transition from one state to another, then it is no longer true that there is no return signal through ground. If there is a return signal through ground, then power system integrity *does* become an issue, and EMI may become a problem.

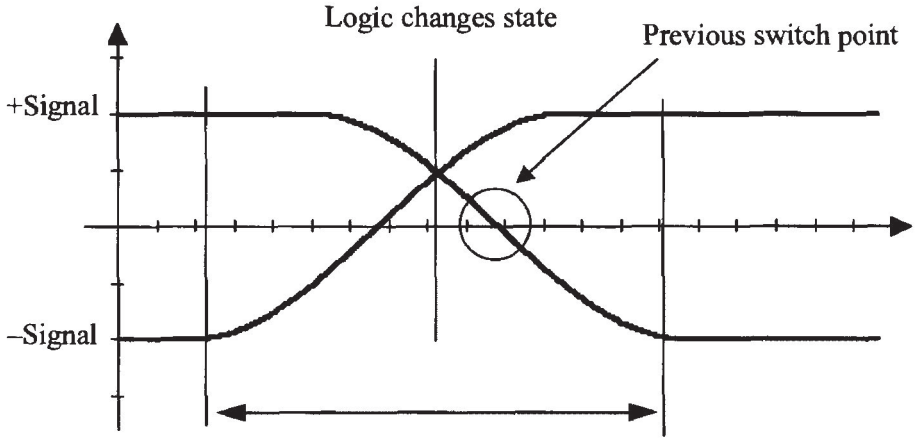


Figure 14-2 The (–) trace is shorter than in Figure 14-1, and it is no longer true that the differential signals are equal and opposite over the range indicated by the arrow. Thus, there will be current flowing through the power system during this time frame.

Let's look at this another way. The square wave in Figure 14-3 is similar to that in Figure 1-13. In a differential signal pair, we might have this signal on one trace, and the opposite signal on another trace. These two signals would then sum to zero (Figure 14-4).

Now consider what happens when we let one trace be slightly longer than the other trace. This is the same thing as having the two signals (the positive and negative signals) being slightly out of phase at the receiving end. Figure 14-5 illustrates the resulting difference signal when this happens. Figure 14-6 illustrates just

INTEL-1010
U.S. Patent 7,151,430

this difference signal, showing more clearly that it can be very pronounced and also of considerable magnitude for just a very minor difference in phase. The "noise" pulse width in this illustration is equal to the phase shift between the two signals.

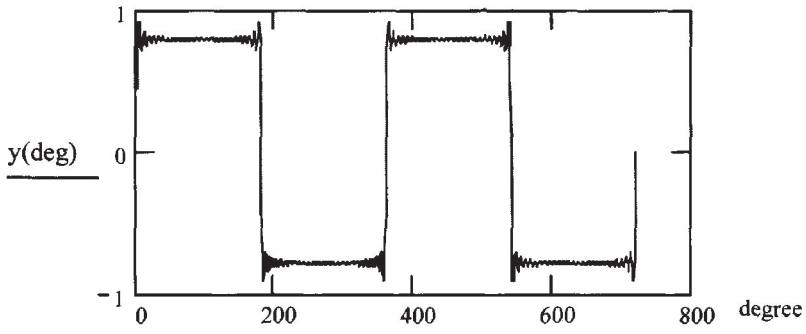


Figure 14-3 Square wave repeated from Figure 1-13.

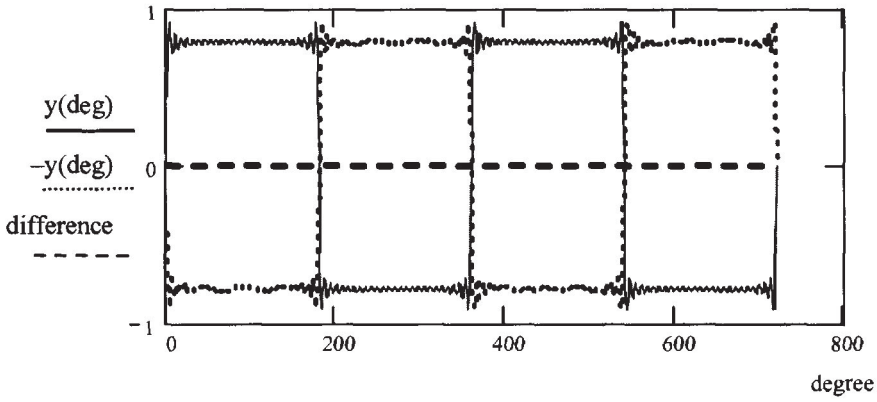


Figure 14-4 The square wave is on one trace and its exact inverse is on the return trace. They combine to zero.

This signal might now be showing up on the ground plane. Not only is it not consistent with our assumption that there are no currents on the ground plane, but the current that now shows up on the plane has sharp rise times, is of considerable magnitude, and can be a serious EMI problem.

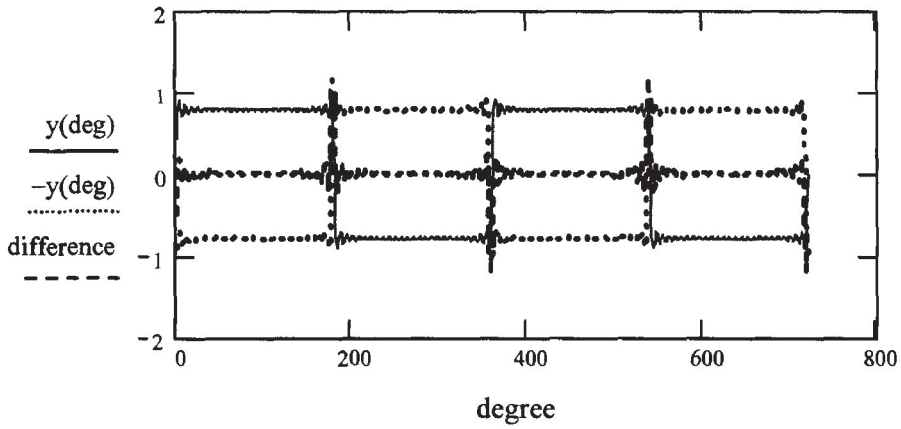


Figure 14-5 If one trace in the differential pair is a slightly different length than the other, a noise signal will be present when they change states.

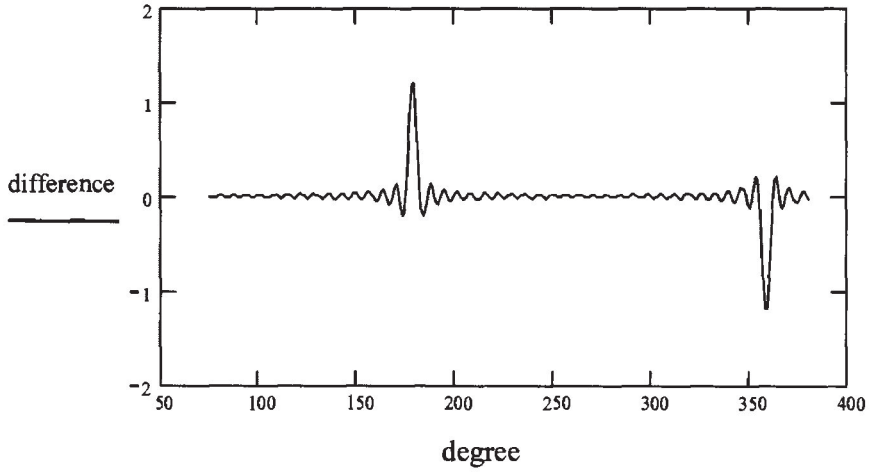


Figure 14-6 A closer look at the noise signal from Figure 14-5.

An interesting question is this: What kind of dimensions are we looking at here before this becomes a problem? A significant part of the answer depends on the rise time of the signal. Even for a poorly defined square wave, a 1- or 2-degree phase shift could be significant.

Assume we have a 50-MHz square wave. That means there are 50 x 10^6 cycles in one second, or there is a single cycle every 20 ns. If propagation time is 6 inches per nanosecond in FR4, then a one-degree phase shift equates to a 333-mil distance. If we set one degree as our threshold, which might be too much, then the corresponding offsets for selected frequencies would be as follows:

Frequency (MHz)	Offset (mils, or thousandth in.)
50	333
500	33
5 GHz	3

DESIGN RULES

Design Rule 1 This brings us to our first design guideline when dealing with differential signals: *The traces should be of equal length.*

There are some people who argue passionately against this rule. Generally, the basis for their argument involves signal timing. They point out in great detail that many differential circuits can tolerate significant differences in the timing between the two halves of a differential signal pair and still switch reliably. Depending on the logic family used, trace length difference of 500 mil can be tolerated. These people can illustrate their points very convincingly with parts specs and signal timing diagrams. But these people miss the point. The reason differential traces must be of equal length has almost nothing to do with signal timing. It has everything to do with the assumption that differential signals are equal and opposite and what happens when that assumption is violated. What happens is this: Uncontrolled ground currents start flowing that at the very best are benign but at worst can generate serious common-mode EMI problems.

So, if you are depending on the assumption that your differential signals are equal and opposite, and that therefore there is no signal flowing through ground, a necessary consequence of that assumption is that the differential pair signal lengths must be equal.

Common Mode Implications Refer back to the common mode discussion in Chapter 9 and in particular to Figure 9-14. Figure 14-7 is the differential signaling equivalent to the single-ended case shown in Figure 9-14. The currents id and ic represent the signals commonly referred to as differential mode and common mode

INTEL-1010
U.S. Patent 7,151,430

in differential signaling. In reality, they are correctly called odd mode and even mode, respectively.

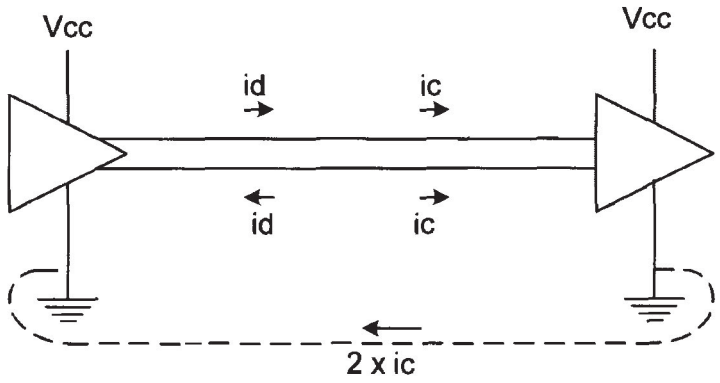


Figure 14-7 Common mode current flows with differential signals.

If the signals are not exactly equal and opposite on the differential traces, we can separate them into their corresponding differential and common (odd and even) mode components as before. Assume again (as we did in Chapter 9) that the plus signal is 10,000 μA (10 mA) and the return signal is 9,950 μA. This is equivalent to the differential (odd) mode component being 9,975 μA and the common (even) mode component being 25 μA. It is easy to see, looking at the problem this way, that if the signal and the return are not exactly equal, there must be a common mode current component that returns through the power system *somewhere*. That current component is uncontrolled, and even though it might be small, it can generate some significant EMI radiations.

Differential Signals and Loop Areas If our differential circuits are dealing with signals that have slow rise times, high-speed design rules are not an issue. Let's assume, however, we are dealing with fast rise time signals. What additional issues then come into play with differential traces?

Consider a design where a differential signal pair is routed across a plane from driver to receiver. Let's also assume that the trace lengths are perfectly equal and the signals are exactly equal and opposite. Therefore, there is no return current path through ground. But there *is* an induced current on the plane, nevertheless.

INTEL-1010
U.S. Patent 7,151,430

Any high-speed signal can (and will) induce a coupled signal into an adjacent trace (or plane). The mechanism is exactly the same mechanism as crosstalk. It is caused by electromagnetic coupling, the combined effects of mutual inductive coupling and capacitive coupling. So, just as the return current for a single-ended signal trace tends to travel on the plane directly under the trace, a differential trace will also have an induced current on the plane underneath it.

This is *not* a return current, however. All the return currents have cancelled. So this is purely a coupled noise current on the plane. The question is this: If current must flow in a loop, where is the rest of the current flow? Remember, we have *two* traces, with equal and opposite signals. One trace couples a signal on the plane in one direction, the other trace couples a signal on the plane in the other direction. These two coupled currents on the plane are equal in magnitude (assuming otherwise good design practices). So the currents simply flow in a closed loop underneath the differential traces (Figure 14-8). They look like eddy currents. The loop these coupled currents flow in is defined by (a) the differential traces themselves, and (b) the separation between the traces at each end. The loop "area" is defined by these four boundaries.

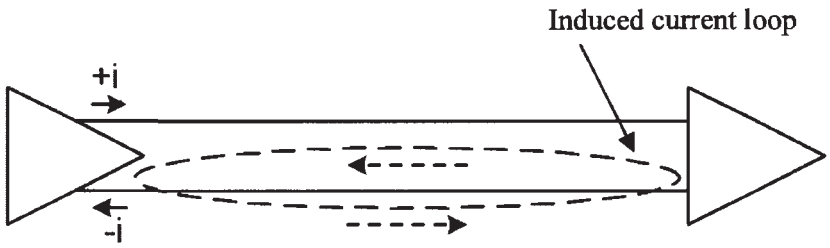


Figure 14-8 Differential traces will couple to a power system plane, even if there are no return currents flowing there.

Design Rule 2 We showed in Chapter 9 that EMI is related to loop area. Therefore, if we want to keep EMI under control, we need to minimize this loop area. And the way we do that brings us to Design Rule 2: *Route differential traces close together*. There are people who argue against this rule, and indeed the rule is not necessary if rise times are slow and EMI is not an issue. However, in high-speed environments, the closer we route the differential traces to each other, the smaller will be their own loop area and the loop area of the induced currents under the traces, and the better control over EMI we will have.

INTEL-1010
U.S. Patent 7,151,430

It is worthwhile to note that some engineers ask designers to remove the plane under differential traces. Reducing or eliminating the induced current loops under the traces is one reason for this. Another reason is to prevent any noise that might already be on the plane from coupling into the (presumably) low signal levels on the traces themselves. (I know of no definitive studies that either support or refute this practice.)

There is another reason to route differential traces close together. Differential receivers are designed to be sensitive to the difference between a pair of inputs, but also to be insensitive to a common-mode shift of those inputs. That means if the (+) input shifts even slightly in relation to the (–) input, the receiver will detect it. If the (+) and (–) inputs shift together (in the same direction), however, the receiver is relatively insensitive to this shift. Therefore, if any external noise (e.g., EMI or crosstalk) is coupled equally into the differential traces, the receiver will be insensitive to this (common-mode coupled) noise. The more closely differential traces are routed together, the more equal any coupled noise will be on each trace. Therefore, the better the rejection of the noise in the circuit will be.

Rule 2 Consequence Again assuming a high-speed environment, if differential traces are routed close to each other (to minimize their own loop area and the loop area of the induced currents underneath them), then the traces will couple into each other. If the traces are long enough that termination becomes an issue, this coupling impacts the calculation of the correct termination impedance. Here's why:

Figure 14-9a illustrates a typical, individual trace. It has a characteristic impedance, Zo, and carries a current, i. The voltage along it, at any point, is (from Ohm's Law) V = iZo.

Figure 14-9b illustrates a pair of traces. Trace 1 has a characteristic impedance Z11, which corresponds to Zo, and current i1. Trace 2 is similarly defined. As we bring Trace 2 closer to Trace 1, current from Trace 2 begins to couple into Trace 1 with a proportionality constant, k. Similarly, Trace 1's current, i1, begins to couple into Trace 2 with the same proportionality constant. The voltage on each trace, at any point, again from Ohm's Law, now is as shown in Equation 14-1:

$$V1 = Z11 \times i1 + Z11 \times k \times i2 \qquad [14\text{-}1]$$
$$V2 = Z22 \times i2 + Z22 \times k \times i1$$

INTEL-1010
U.S. Patent 7,151,430

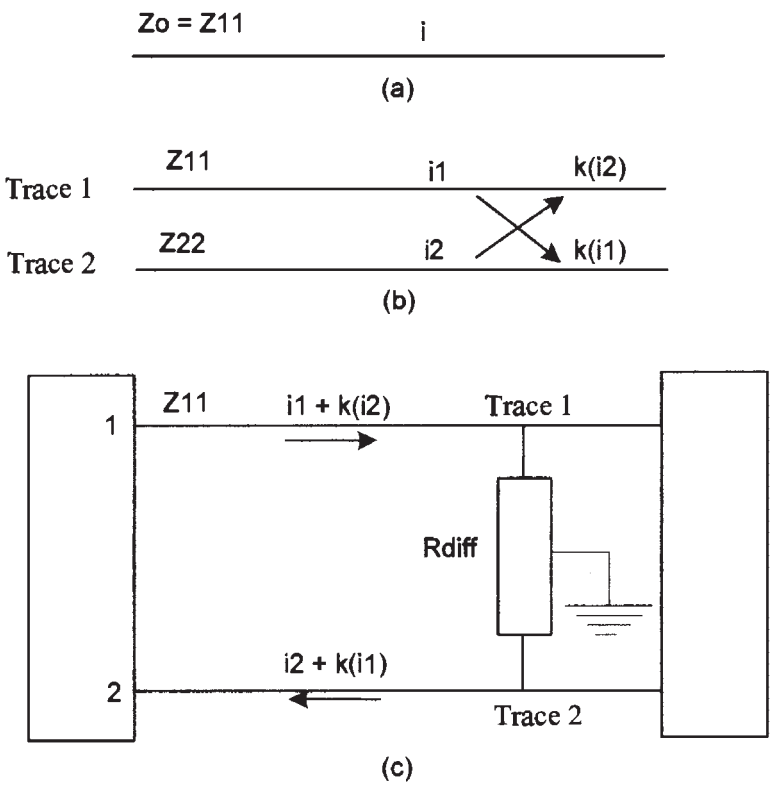


Figure 14-9 Signals on differential traces couple into each other, exactly as crosstalk couples between adjacent traces.

Now let's define Z12 = k x Z11 and Z21 = k x Z22. Then, Equation 14-1 can be written as shown in Equation 14-2:

$$V1 = Z11 \times i1 + Z12 \times i2 \qquad [14\text{-}2]$$
$$V2 = Z21 \times i1 + Z22 \times i2$$

This is the familiar pair of simultaneous equations we often see in texts. The equations can be generalized into an arbitrary number of traces, and they can be expressed in a matrix form that is familiar to many of you.

Figure 14-9c illustrates a differential pair of traces. Recall Equation 14-1:

$$V1 = Z11 \times i1 + Z11 \times k \times i2$$
$$V2 = Z22 \times i2 + Z22 \times k \times i1$$

Now, note that in a carefully designed and balanced situation,

$$Z11 = Z22 = Zo, \text{ and}$$
$$i2 = -i1$$

This leads (with a little manipulation) to Equation 14-3:

$$V1 = Zo \times i1 \times (1 - k) \quad [14\text{-}3]$$
$$V2 = -Zo \times i1 \times (1 - k)$$

Note that V1 = –V2, which we already knew, of course, since this is a differential pair.

The voltage, V1, is referenced with respect to ground. The effective impedance of Trace 1 (when taken alone this is called the odd-mode impedance of a single trace of a differential pair, or single-ended impedance in general) is voltage divided by current, or, as given in Equation 14-4:

$$Zodd = V1/i1 = Zo \times (1 - k) \quad [14\text{-}4]$$

Since (from above) Zo = Z11 and k = Z12/Z11, this can be rewritten as Equation 14-5:

$$Zodd = Z11 - Z12 \quad [14\text{-}5]$$

which is a form also seen in many textbooks.

The proper termination of this trace, to prevent reflections, is with a resistor that has a value of Zodd connected between the trace and ground. Similarly, the odd mode impedance of Trace 2 turns out to be the same (in this special case of a balanced differential pair).

Design Rule 3 The consequence of Design Rule 2 is that the differential pair of traces couple into each other. This coupling impacts the proper termination required if we want to prevent reflections from occurring at the far end of the line. The effect of this is slightly different depending on whether we are looking at the proper differential (odd) mode termination or common (even) mode termination.

Differential Mode Impedance Assume for a moment that we have terminated each trace of a differential pair with a resistor to ground. Since i1 = –i2, there

would be no current at all through ground. Therefore, there is no real reason to connect the resistors to ground. In fact, some people would argue that you must *not* connect them to ground in order to isolate the differential signal pair from ground noise. The normal connection would be as shown in Figure 14-9c, a single resistor from Trace 1 to Trace 2. The value of this resistor would be the sum of the odd-mode impedance for Trace 1 and Trace 2, or as given by Equation 14-6:

$$\begin{aligned} Zdiff &= 2Zo(1-k) \text{ or} \\ Zdiff &= 2(Z11-Z12) \end{aligned} \qquad [14\text{-}6]$$

This is why you often see references to the fact that a differential pair of traces might have a differential impedance of around 80 Ω when each trace, individually, is a 50-Ω trace.

Note this very important consequence: As the traces become closer together, the coupling between them increases. As the coupling increases, Z12 becomes larger, since it is directly related to the coupling coefficient, k. As Z12 gets larger, the differential impedance gets smaller, even though the single-ended impedance (Zo) is not changing. Thus, the more closely two traces are coupled, the lower is the differential impedance. This leads to Design Rule 3: D*ifferential impedance calculations are necessary with differential signals and traces.*

Common Mode Impedance Just to round out the discussion, common-mode impedance differs only slightly from differential mode. The first difference is that i1 = i2 (without the minus sign). Thus Equation 14-3 becomes Equation 14-7:

$$\begin{aligned} V1 &= Zo \times i1 \times (1+k) \\ V2 &= Zo \times i1 \times (1+k) \end{aligned} \qquad [14\text{-}7]$$

and V1 = V2, as expected. The individual trace impedance, therefore, is Zo(1 + k). In a common-mode case, both trace terminating resistors *are* connected to ground, so the current through ground is i1 + i2 and the two resistors appear (to the device) to be in parallel. Therefore, the common-mode impedance is the parallel combination of these resistors, or as given by Equation 14-8:

$$\begin{aligned} Zcommon &= (1/2) \times Zo \times (1+k), \text{ or} \\ Zcommon &= (1/2) \times (Z11+Z12) \end{aligned} \qquad [14\text{-}8]$$

Note, therefore, that the common-mode impedance is approximately one-quarter the differential mode impedance for trace pairs.[2]

Design Rule 4 Differential impedance changes with coupling, which changes with trace separation. Since it is always important that the trace impedance remain constant over the entire length, this means that the coupling must remain constant over the entire length. This leads to our fourth rule: T*he separation between the two traces (of the differential pair) must remain constant over the entire length.*

Note that these differential impedance impacts are merely consequences of Design Rule 2. There is nothing really inherent about them at all. The reason we want to route differential traces close together is because of EMI and noise immunity. The fact that this has an impact on the correct termination of "long" traces, and this in turn has an impact on the uniformity of trace separation, is simply a consequence of routing the traces close together for EMI control. (Note: The reason this doesn't happen with other closely routed traces (those subject to crosstalk for example) is that other traces don't have a coupling between them that is perfectly correlated—that is, equal and opposite. If the coupled signals are simply randomly related to each other, the average coupling is zero and there is no impact on the impedance termination.)

DIFFERENTIAL SIMULATIONS

The HyperLynx LineSim simulator can be used to simulate differential waveforms under different conditions. Figure 14-10 shows a typical LineSim setup for a differential signal. The differential driver (U(A0) and U(B0)) at the top of the model drive a pair of traces, each of which has a single-ended impedance (Zo) equal to 50 Ω. The traces are 8 mil wide and are separated by 4 mil, so they are tightly coupled. The traces are each 18 inches (2.5 ns) long.

A differential receiver (U(A1) and U(B1)) is connected at the other end of the traces. The termination of the trace is provided by the resistor RS(A1). The

[2]For an interesting discussion about how to terminate both the differential mode and common mode components of a pair of traces, see "Terminating Differential Signals on PCBs," Steve Kaufer and Kellee Crisafalu, *Printed Circuit Design* magazine, March 1999, p. 25.

INTEL-1010
U.S. Patent 7,151,430

question is: What value should this resistor be and what are the consequences of picking the wrong value?

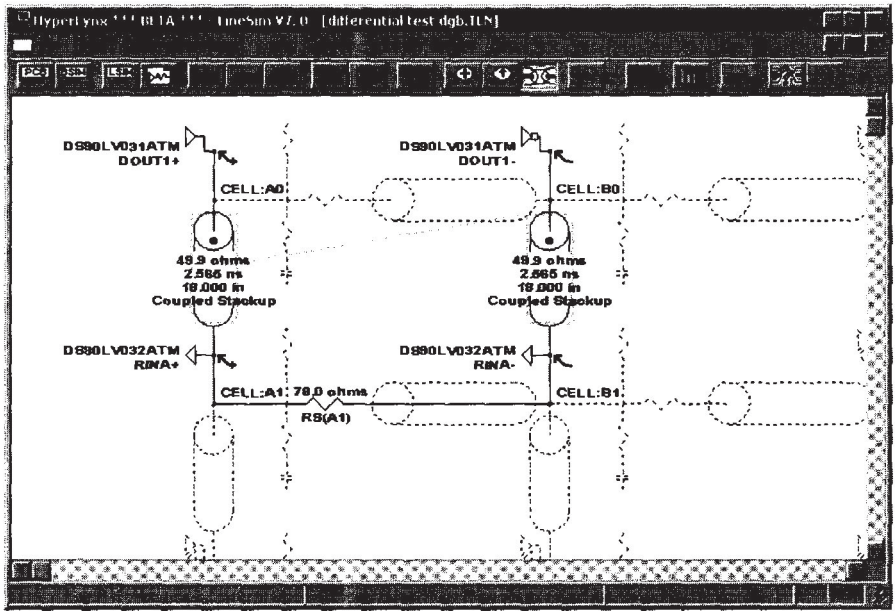


Figure 14-10 HyperLynx LineSim simulation of a stripline differential net.

The HyperLynx termination wizard (part of the LineSim tool) predicts that a termination resistor equal to 78 Ω should be used. Note that this is something less than 2Zo, or 100 Ω. The result of the simulation using 78 Ω is shown in Figure 14-11. The 25-MHz driver signal has a 40 ns period, or is "high" for 20 ns. The signal at the receiver follows that of the driver 2.5 ns later. The signals shown in Figure 14-11 are very clean signals.

Figures 14-12 and 14-13, on the other hand, illustrate what happens when the termination resistor is selected incorrectly. For example, since each line has a single-ended impedance of 50 Ω, a designer may inadvertently select 50 Ω as the terminating differential resistor. Figure 14-12 shows the result. Or someone might simply select a differential termination resistor of 2Zo = 100 Ω. Figure 14-13 illustrates the result of that. In both cases the signal degradations caused by the poorer selection are apparent.

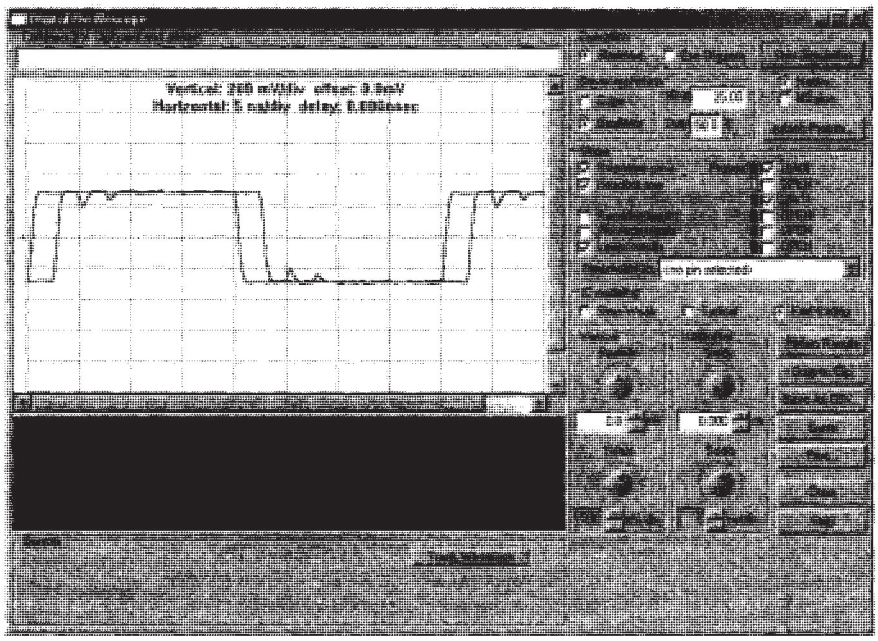


Figure 14-11 Result of differential simulation using 78 Ω.

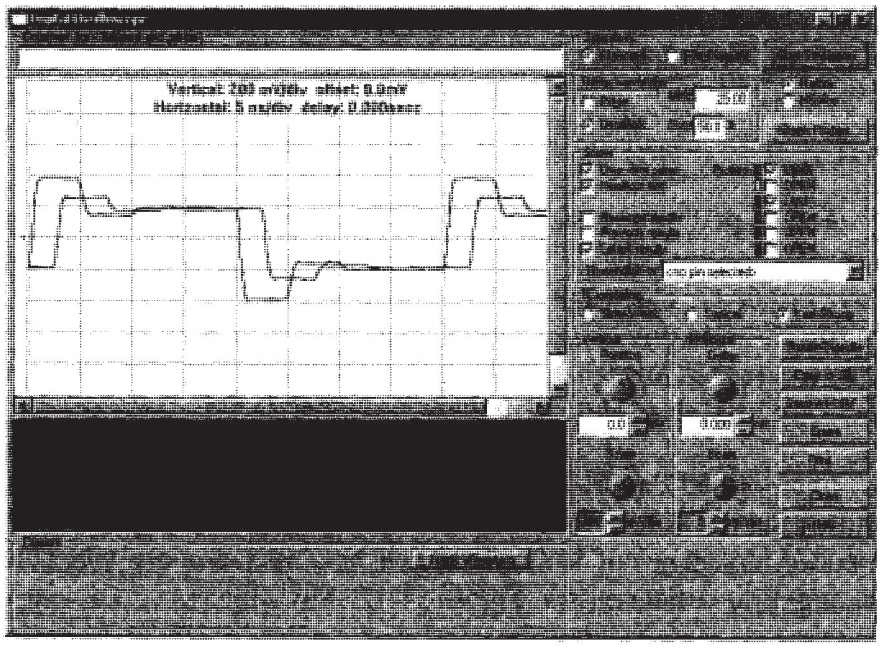


Figure 14-12 Selecting a 50-Ω terminating resistor.

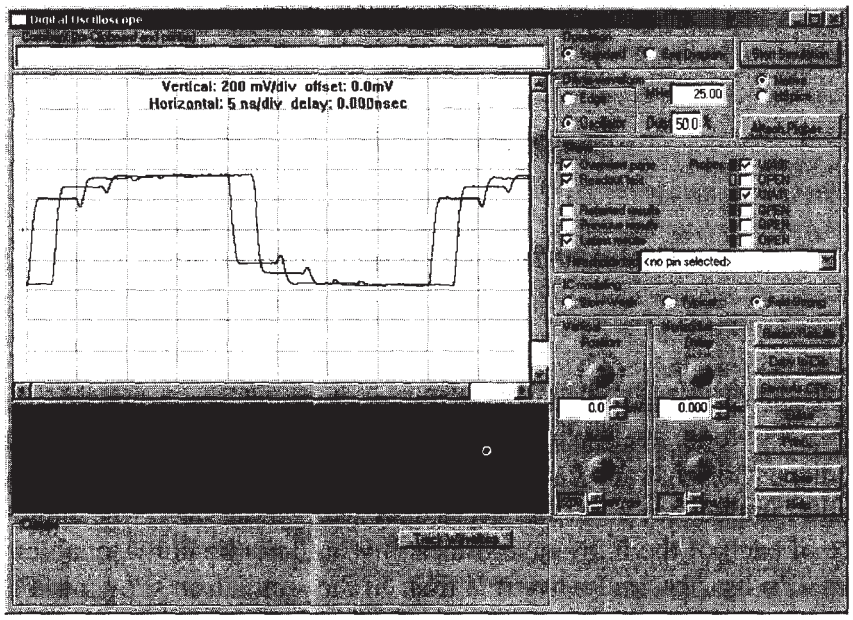
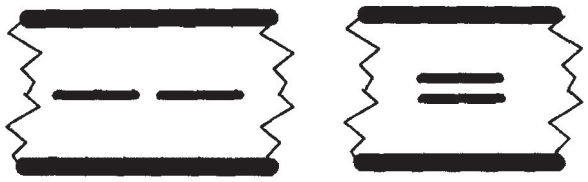


Figure 14-13 Selecting a 100-Ω terminating resistor.

CALCULATING DIFFERENTIAL IMPEDANCE

There are two fundamental types of differential trace configurations: edge coupled and broadside coupled. These are illustrated in Figure 14-14. There does not appear to be a significant advantage or disadvantage to either configuration. The broadside configuration makes it easier to route a pair of traces through a pin field keeping both the length and spacing constant. Broadside coupled configurations, however, typically restrict routing opportunities for other traces. Edge coupled configurations preserve the common X–Y trace routing strategy designers often use with adjacent layers, but then keeping trace lengths equal and traces equally spaced becomes a bigger challenge. Calculations for differential impedance are, of course, different for each configuration.



(a) (b)

Figure 14-14 Edge coupled (a) and broadside coupled (b) differential trace pairs.

Edge Coupled Calculating differential impedance is not easy. In fact, there are only a few tools available that even purport to do this. It is clear from the preceding discussion that any carefully constructed differential signal pair will have a differential impedance something less than 2Zo, where Zo is the single-ended impedance of each of the individual traces. It would not be unreasonable to simply discount 2Zo by 20% and leave it at that. In the simulation shown in Figure 14-10, that would lead to 80 Ω where the HyperLynx wizard calculated an actual value of 78 Ω (pretty close). I am aware of one complex equation for differential impedance that has a stated accuracy level of 20%. (We can *guess* it closer than that.)

National Semiconductor has published some approximate equations for edge coupled differential traces.[3] Higher-end design systems may have a differential impedance calculation ability built into them, so users can make the calculations as they design the board.

Microstrip:

$$Zdiff \cong 2Zo\left(1 - 0.48e^{-.096\frac{s}{h}}\right)$$

Centered stripline:

$$Zdiff \cong 2Zo\left(1 - 0.78e^{-2.9\frac{s}{h}}\right)$$

[3]These equations are published in several National Semiconductor publications. See in particular "Transmission Line RAPIDESIGNER Operation and Applications Guide," AN-905, National Semiconductor Corporation.

INTEL-1010
U.S. Patent 7,151,430

where:

Zo = Single-ended trace impedance
s = edge-to-edge trace separation
h = height of the trace above the reference plane (microstrip) or distance between the planes (stripline)

Polar Instruments offers an excellent stand-alone calculator for making both single-ended and differential impedance calculations. It comes in two versions: the SI6000 Quick Solver and an Excel-based spreadsheet plug-in. Figure 14-15 illustrates a Quick Solver solution for a stripline configuration.

Comparison Table 14-1 illustrates the degree of consistency between various calculational tools. For edge coupled differential impedance calculations the degree of agreement between the HyperLynx tool and the Polar calculator is quite remarkable, even recognizing the fact that they both use very advanced field-solving techniques to make these analyses. The National Semiconductor equation agrees within about 9%, which is not bad for simple, approximating equations of this type. Remember, your board house will be lucky to get within 6% or 7% of your differential targets, anyway, for reasons discussed in Chapter 10.

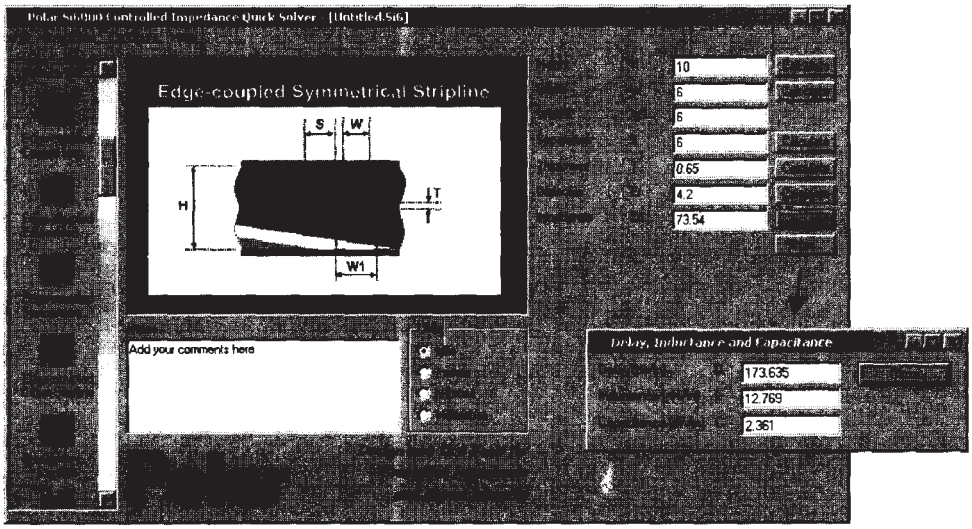


Figure 14-15 Polar Instruments' "QuickSolver" impedance calculator (see at www.polarinstruments.com).

Table 14-1. Comparison of Results

Tool	Zo (Ω)	Zdiff (Ω)
Polar	38.98	73.54
HyperLynx	39.0	73.6
National	39.0*	67.76

* taken as a given

Assumptions:

h = 10
w = 6
s = 6
t = .65
$\varepsilon_r = 4.2$

Broadside Coupled Unfortunately, I am not aware of any simple formulas for estimating broadside coupled differential impedance. Higher end tools will do this, including the Polar calculator, but there are no readily available simple tools for doing so.

INTEL-1010
U.S. Patent 7,151,430