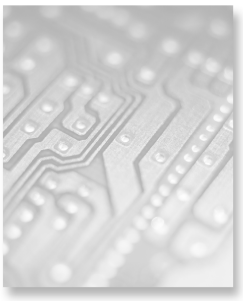

14 DIFFERENTIAL TRACES AND IMPEDANCE

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.

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.

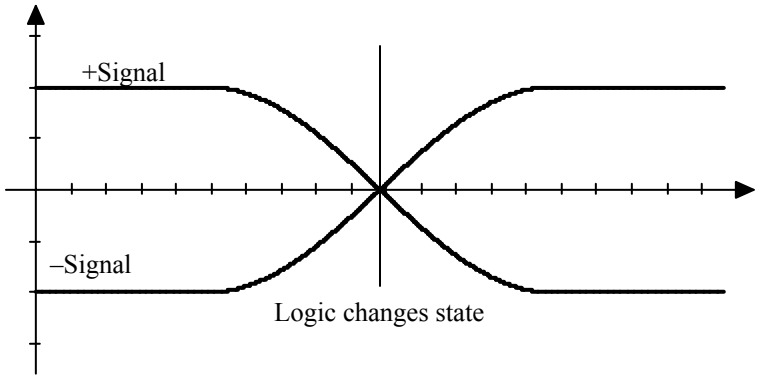

Figure 14-1 Timing with differential signals can be very precise.

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-

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.

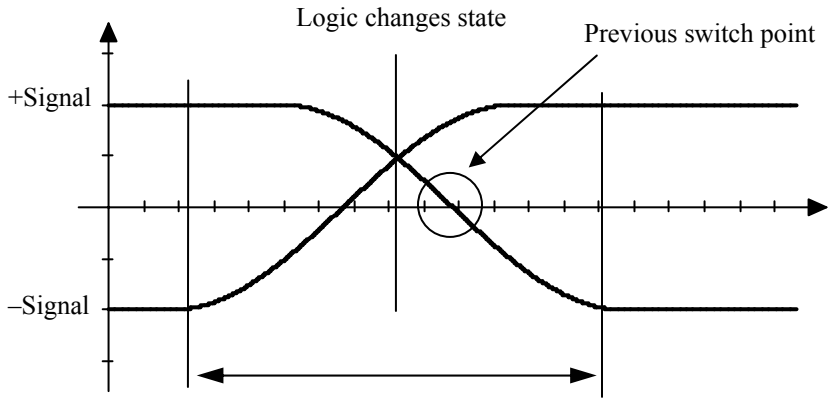

Figure 14-2 The (–) trace is shorter than in Figure 14-1, and it is no longer true that the differential signals are equal and opposite over the range indicated by the arrow. Thus, there will be current flowing through the power system during this time frame.

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

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.

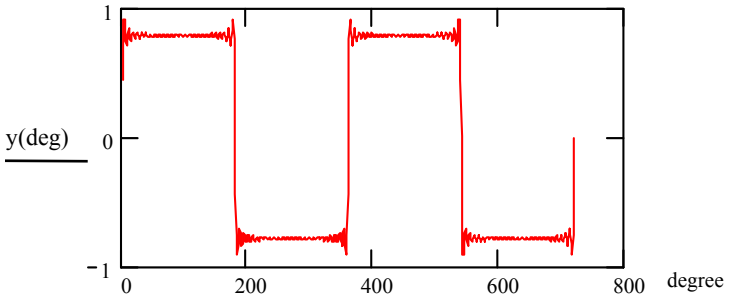

Figure 14-3 Square wave repeated from Figure 1-13.

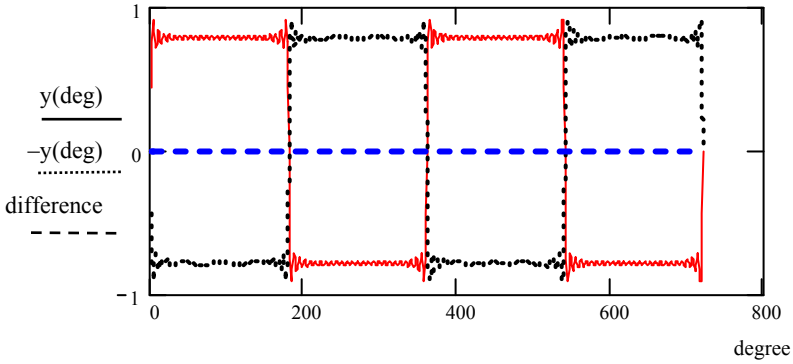

Figure 14-4 The square wave is on one trace and its exact inverse is on the return trace. They combine to zero.

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.

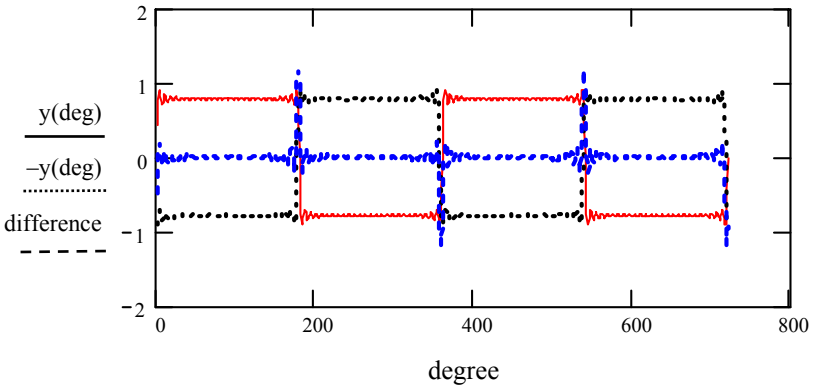

Figure 14-5 If one trace in the differential pair is a slightly different length than the other, a noise signal will be present when they change states.

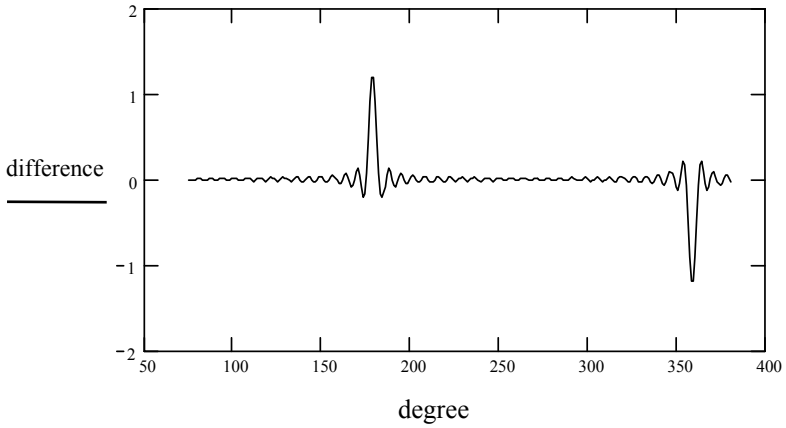

Figure 14-6 A closer look at the noise signal from Figure 14-5.

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

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

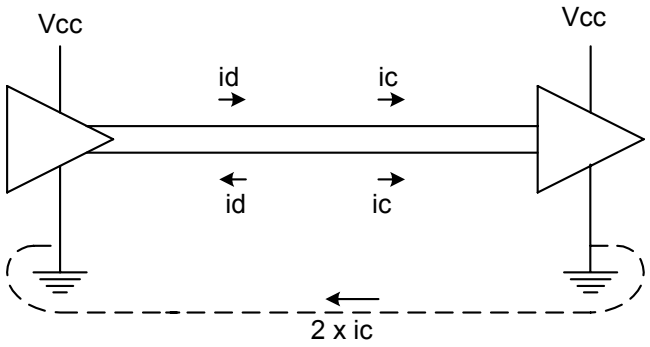

Figure 14-7 Common mode current flows with differential signals.

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.

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.

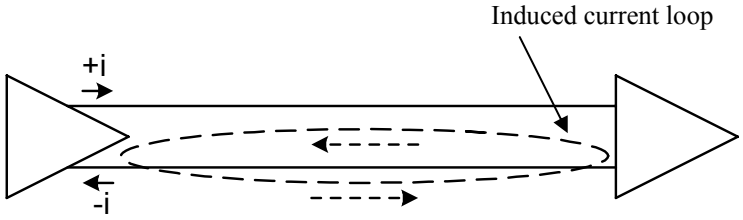

Figure 14-8 Differential traces will couple to a power system plane, even if there are no return currents flowing there.

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.

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:

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

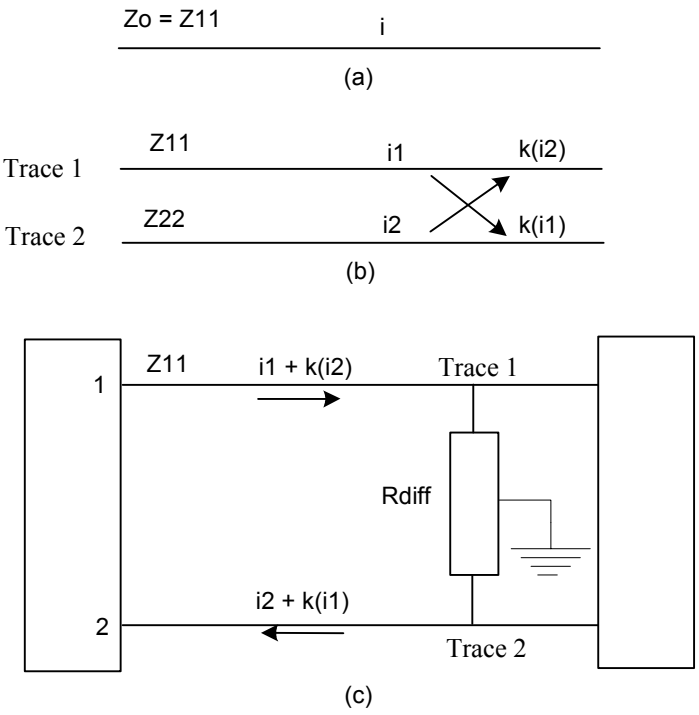

Figure 14-9 Signals on differential traces couple into each other, exactly as crosstalk couples between adjacent traces.

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.

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

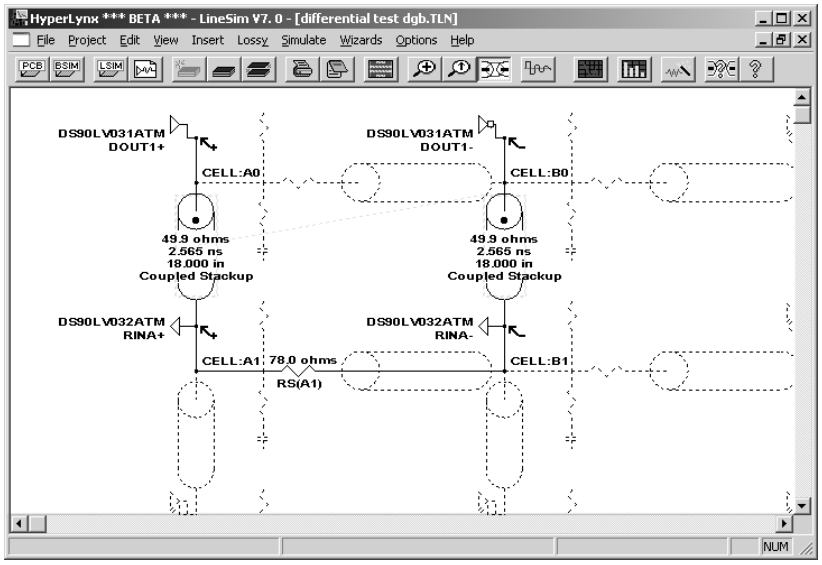

Figure 14-10 HyperLynx LineSim simulation of a stripline differential net.

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.

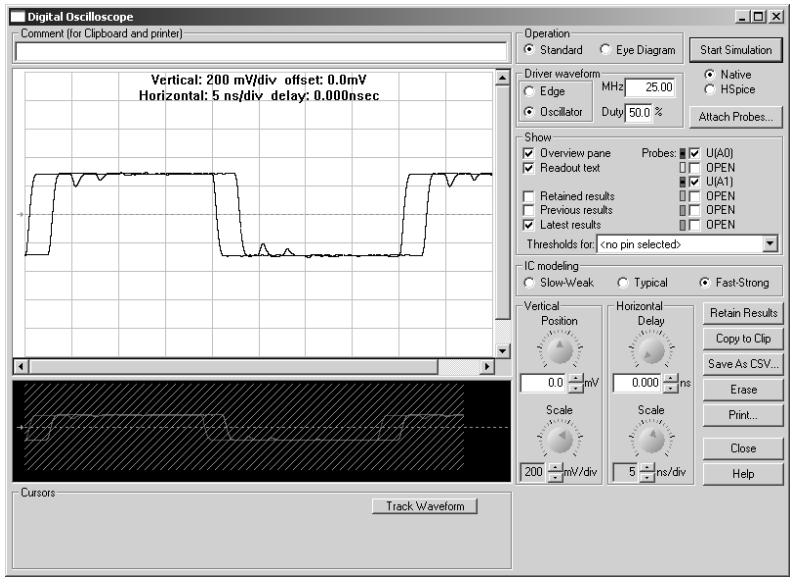

Figure 14-11 Result of differential simulation using 78 Ω.

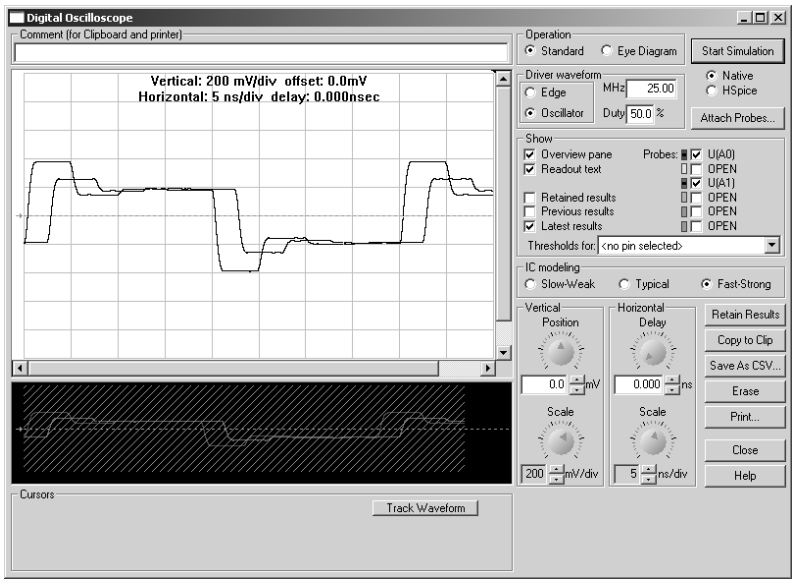

Figure 14-12 Selecting a 50-Ω terminating resistor.

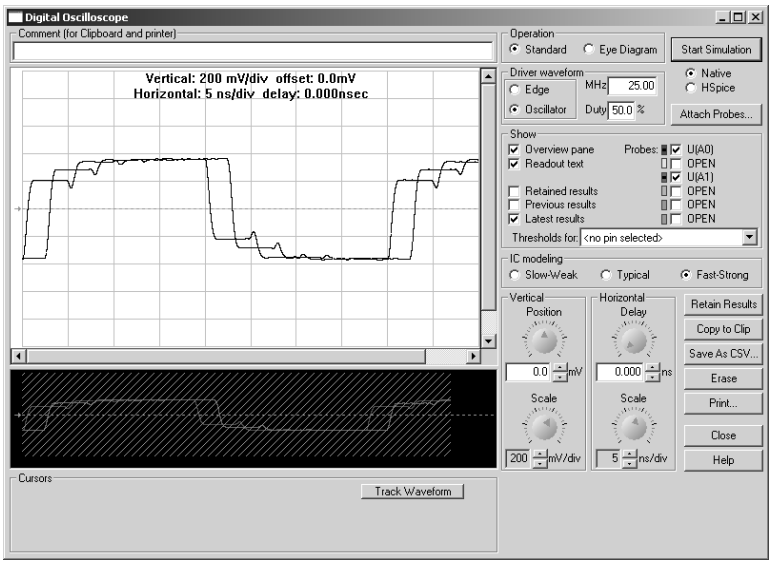

Figure 14-13 Selecting a 100-Ω terminating resistor.

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.

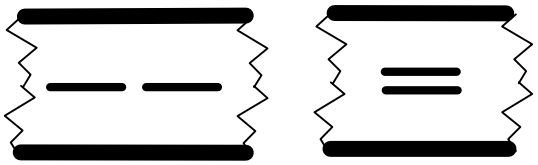

(a) (b)

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.

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.

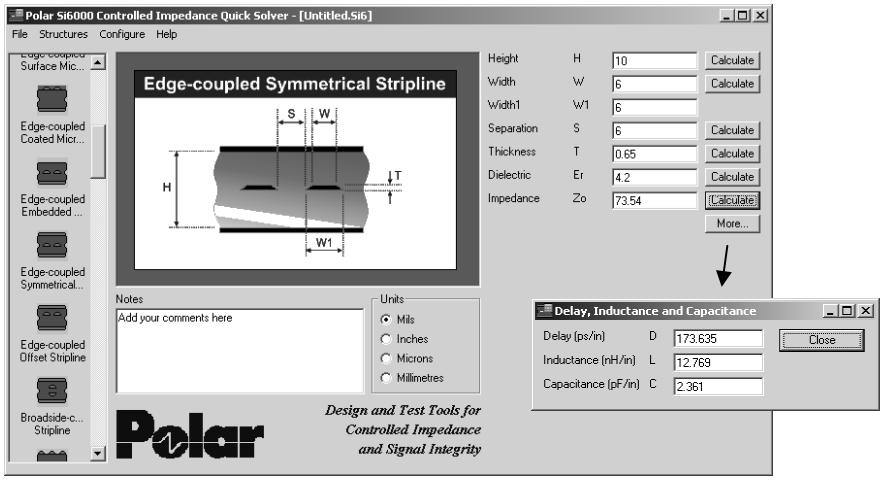

Figure 14-15 Polar Instruments' "QuickSolver" impedance calculator (see at www.polarinstruments.com).

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.