CHAPTER 6

DIFFERENTIAL SIGNALING

This chapter explores the essential elements of differential signaling. It illustrates the many benefits of differential operation and lays out the conditions under which those benefits may be obtained. It also gives some practical advice about how to best manage differential connections.

I'll begin with a detailed, start-from-the-basics look at *single-ended signaling* architectures, then extend the analysis to cover *differential signaling*. The point of this presentation is to break down differential signaling into its component parts to show how precisely the various benefits are obtained.

6.1 SINGLE-ENDED CIRCUITS

The scissors in Figure 6.1 have cut the circuit, interrupting communications. With the wire cut, regardless of the position of the switch, no current flows through the bulb. To light the bulb, this elementary circuit needs the top switch closed *and* the bottom wire connected.

It takes two wires to efficiently convey electrical power.

This easy example illustrates an important principle of lumped-element circuit operation—namely, that current flows only in a *complete* loop. Another way of expressing the same loop idea is to simply say that if current comes out of the battery, it must return to the battery. That is the way circuits operate in practice. Electrons can't just come out of the battery, pass through the bulb, and then pile up somewhere. The movement of electrons requires a complete loop path.

The current-loop principle may be stated in general mathematical terms:

The total sum of all current going in and out of the battery is zero.

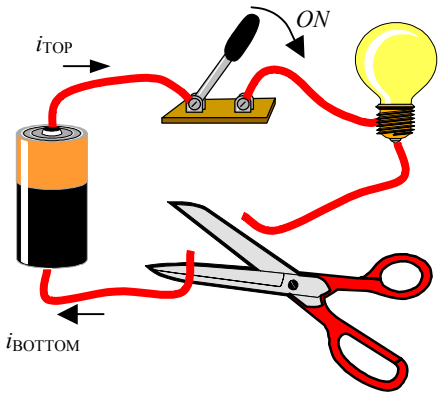

Figure 6.1—Whether cut by the switch or the scissors, this circuit is open and no current flows.

This principle is equivalent to Kirchoff's current law, named after its discoverer, Gustav Robert Kirchoff, 1824–1887, a German physicist. It applies to all lumped circuit elements (not just batteries), including nonlinear circuits, time-varying circuits, passive circuits, active circuits, and circuits with more than two wires. It remains true even in the presence of distributed parasitic capacitances provided one is careful to include the displacement currents.

Applying this principle to Figure 6.1, cutting the bottom wire prevents current from entering the bottom of the battery; therefore, current is also prevented from leaving at the top.

I hope this discussion makes it clear that you cannot propagate electrical current on just one wire. The propagation of electrical current requires a system of at least *two* conductors.

Now let's apply this same reasoning to high-speed digital systems. I'll restrict my attention to *electrical* logic families such as TTL, CMOS, and ECL that have electrical inputs (as opposed to optical or telepathic inputs). First please note that any electric input requires *current* to operate. Specification sheets may emphasize the performance in terms of voltage specifications, but current is still required. Even on CMOS parts, for which the input current is practically zero, it takes a fair amount of current on every rising edge just to charge the parasitic input capacitance.[43] Every electrical input requires current. As a result, the propagation of every (electrical) logic signal requires a system of at least two conductors. Even though the "second wire" does not appear on a schematic circuit diagram, its presence is required.

As you probably know, the two-conductor requirement is rather inconvenient. Most digital logic designers would rather use a single wire for each signal, or at least *make believe that they are doing so*. Toward that end, we designers have adopted a certain convention for the handling of the second wire necessary in all circuits. What we do is hook one side of all transmitters and all receivers to a common reference voltage.[44] In high-speed designs this voltage is usually distributed throughout the system as a ground plane or a pair of power and

[43] It takes 1 mA to charge one picofarad to 1 volt in 1 nanosecond.

[44] A global reference voltage may be distributed by means of a solid ground plane, a solid V_{CC} plane, or the Earth itself. Anything can work as long as it ties everything together through a reasonably low impedance.

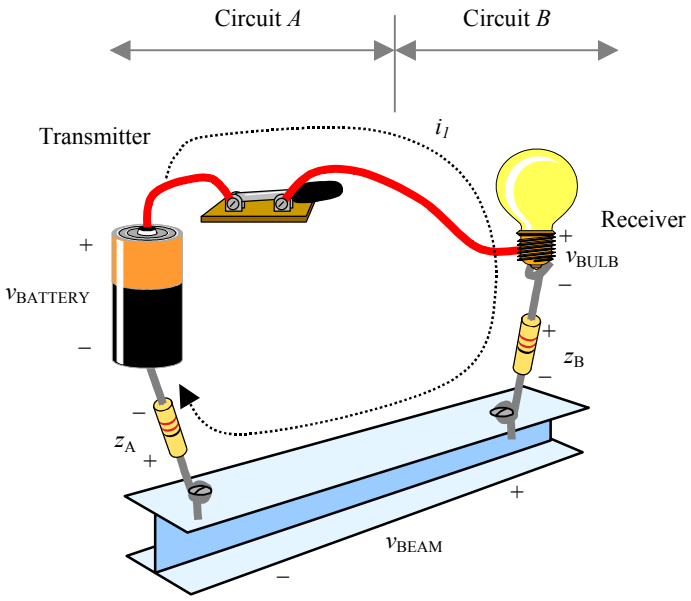

Figure 6.2—Digital systems use a common ground reference for all returning signal currents. Components z_A and z_B represent the ground connection impedances of a semiconductor package.

ground planes. All circuits share this same "second conductor" for the conveyance of returning signal currents. Figure 6.2 illustrates the arrangement for a single transmitter (key switch), receiver (light bulb), and reference system (I-beam). The use of a shared reference voltage for all circuits is called *single-ended signaling*. Single-ended systems require only one apparent wire for each signal. What one needs to keep in mind about this arrangement is that the second wire is still physically there and that it still physically carries the returning signal currents for each transmitter; it's just implemented as a big, common shared connection.

Here's how the circuit in Figure 6.2 works. The transmitter (circuit *A*) includes a battery and a switch. One side of the battery connects to a big, solid, shared-reference system (I-beam). This connection between the transmitter circuit and the reference system passes through impedance z_A, representing the finite impedance of the package pins or balls used to accomplish this connection.

The receiver (circuit *B*) is merely a light bulb connected to the beam through impedance z_B. The bulb brightness indicates the condition of the transmitter (switch *ON* or *OFF*). As with the transmitter, the bulb's connection to the common reference beams has a finite impedance.

Provided that the reference-connection impedances are sufficiently low and that there is no significant voltage drop across the steel I-beam, the circuit functions well. When the switch is depressed, the bulb lights.

Let's take a moment to investigate how the finite-impedance reference connections affect the voltage sensed by the bulb. This effect is controlled by the following physical principle:

The battery voltage equals the sum of the voltages across series-connected loads.

This principle is equivalent to Kirchoff's voltage law. It applies to all voltage sources (not just batteries), including nonlinear sources and time-varying sources.

Applied to the circuit in Figure 6.2, Kirchoff's voltage law predicts the following relation:

$$v_{\text{BATTERY}} = v_{\text{BULB}} + z_{\text{A}} i_{\text{l}} + v_{\text{BEAM}} + z_{\text{B}} i_{\text{l}} \qquad [6.1]$$

Here I have substituted the expression $z_{\text{A}} i_{\text{l}}$ and $z_{\text{B}} i_{\text{l}}$ for the voltages across the two resistors, with the variable i representing the current flowing around the loop. Presuming that the battery voltage remains constant, if either of the voltages $z_{\text{A}} i_{\text{l}}$ or $z_{\text{B}} i_{\text{l}}$ go up, the voltage across the bulb must go down. Let me state that again in a slightly different way: Anything that affects the voltage drop across z_{A}, z_{B}, or the beam also affects the voltage received at the bulb. This statement uncovers a great weakness of single-ended signaling: The reference voltages at transmitter and receiver must match. Unfortunately, as anyone with an oscilloscope knows, noise exists between every two points on a ground plane (or power plane). If the local reference voltages at the receiver and at the transmitter differ by too great an amount, single-ended signaling can't work.

The difficulty here is that the light bulb has no way of knowing what is the *true* voltage coming out of the battery. It has no magic connection to the center of the earth with which to measure the *true* earth potential. It sees only what remains of the transmitted battery voltage after subtraction of the various voltage drops around the signaling loop. In this sense the receiver is a *differential receiver*. Is responds only to the *difference* between the voltage on its input terminal and the voltage on its reference terminal (the reference terminal is the ground pin in TTL logic or the most positive power supply pin in ECL logic). All digital receivers operate the same way. Adding 100 mV of noise to the reference pins of a single-ended IC accomplishes precisely the same thing as adding 100 mV of noise directly to every receiver input. This extra noise detracts directly from the available *noise margin*[45] in your logic family and must be incorporated into any voltage margin analysis for the system.

There is no way around this limitation. If you wish to use single-ended signaling, you must limit the voltage differences within the reference system to a small fraction of the signal amplitude. The impedance of the reference system itself must be low enough to absorb all returning signal currents from all sources without producing objectionable voltage drops at any point.

Figure 6.3 depicts another problem inherent to single-ended signaling. This figure incorporates a second transmitter. Assume the second transmitter is part of circuit *B*. It resides within the same physical package as the receiver; therefore, they share a common ground connection internal to the package. In the figure this shared connection appears as a solid wire connecting the bottom of the light bulb to the bottom of the second battery. This shared internal ground then connects through the impedance of the ground pins or balls on the package (represented by impedance z_{B}) to the circuit board ground (the I-beam). The

[45] The *noise margin* for a logic family is the difference between V_{OL} and V_{IL}, or V_{OH} and V_{IH}, whichever is less.

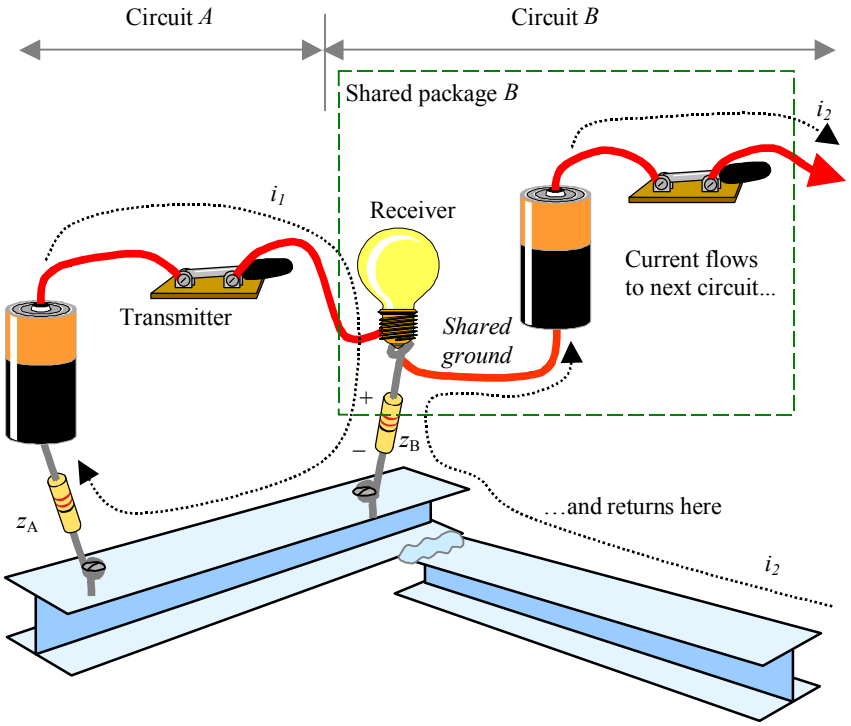

Figure 6.3—*Ground Bounce:* In circuit *B*, a transmitter and receiver share the same ground connection z_B. As current i_2 passes through component z_B, it creates voltage $-i_2 z_B$, which interferes with reception.

boundary of the package containing both the first receiver and the second transmitter is depicted in Figure 6.3 as a box drawn with dashed lines.

Consider what happens when current i_2 flows in the second transmitter circuit. First current i_2 exits the shared package (dotted region). After circulating through its load, this same current must therefore find its way back inside the shared package. There are only two paths available for current returning to the shared package—either through impedance z_B or through the wire leading to the input port of the receiver (light bulb). In a digital system the impedance z_B is much lower (hopefully) than the impedance of the path through the input port, so most of the current reenters the package through z_B. So far, so good. Next let's examine the implications of this returning current.

As returning current i_2 reenters the package, it induces voltage $-z_B i_2$ across the impedance z_B. Because impedance z_B is shared with the receiving circuit, any voltages appearing across it disturb the apparent voltage received by the light bulb (equation [6.1]). This form of interference goes by the name of *common impedance coupling* [60]. It happens any time there is any overlap between the paths of current flow for a transmitter and a receiver.

In a high-speed digital application the inductance of the ground connection z_B usually causes more difficulties that its resistance. This inductance, multiplied times the *di/dt* of the returning current, can generate voltages easily sufficient to disturb normal receiver

operation. The problem of noise generated by returning signal current acting across the finite impedance of a common ground connection within an IC package is called *ground bounce* [61], or more generally *simultaneous switching noise* [59].

The same general problem also happens wherever the reference system is weakened or necked-down. Returning signal currents from many drivers, flowing through the finite inductance of the necked-down region, generate differences in the reference potential (with respect to true earth ground) at different points in the system. These voltages can disturb receiver operation. This problem is often called a *ground shift*, or *noisy ground*. Noticeable ground shifts often happen across places, like connectors and cables, where the integrity of the solid ground plane has been violated.

Power and ground distributions networks are both susceptible to returning signal currents in the same way. Currents flowing through a power distribution net can perturb the power-supply voltages at various points within the system in the same way that currents in the ground network perturb the ground voltages. Whether you care more about power-supply noise or ground noise depends on whether your single-ended circuits use the power or the ground rail as their internal reference for the discrimination of logic signals. TTL integrated circuits and most high-speed digital CMOS circuits use the ground terminal as the designated reference voltage. ECL circuits powered from ground and –5.2 volts use the ground terminal as the designated reference voltage. ECL circuits powered from a positive supply and ground (sometimes called PECL) use their positive supply terminal as the designated reference voltage.

POINTS TO REMEMBER

- The big advantage of single-ended signaling is that it requires only one wire per signal.
- Single-ended signaling falls prey to disturbances in the reference voltage.
- Single-ended signaling is susceptible to ground bounce.
- Single-ended signaling requires a low-impedance common reference connection.

6.2 TWO-WIRE CIRCUITS

Two-wire signaling cures many noise problems at the cost of a second signal trace. As shown in Figure 6.4, a two-wire transmitter sends current on two wires: a first wire, which carries the main signal, and a second wire, which is provided for the flow of returning signal current. As drawn, the *currents* on the two wires will be equal and opposite, but the voltages will not be. This architecture provides three important benefits.

First, it frees the receiver from requiring a global reference voltage.[46] In effect, the second wire serves as a reference for the first. The receiver need merely look at the difference between the two incoming wires. Two-wire signaling renders a system immune to

[46] The reference voltage for TTL, most high-speed CMOS, and ECL is ground; for PECL (positively-biased ECL), it is the power voltage.

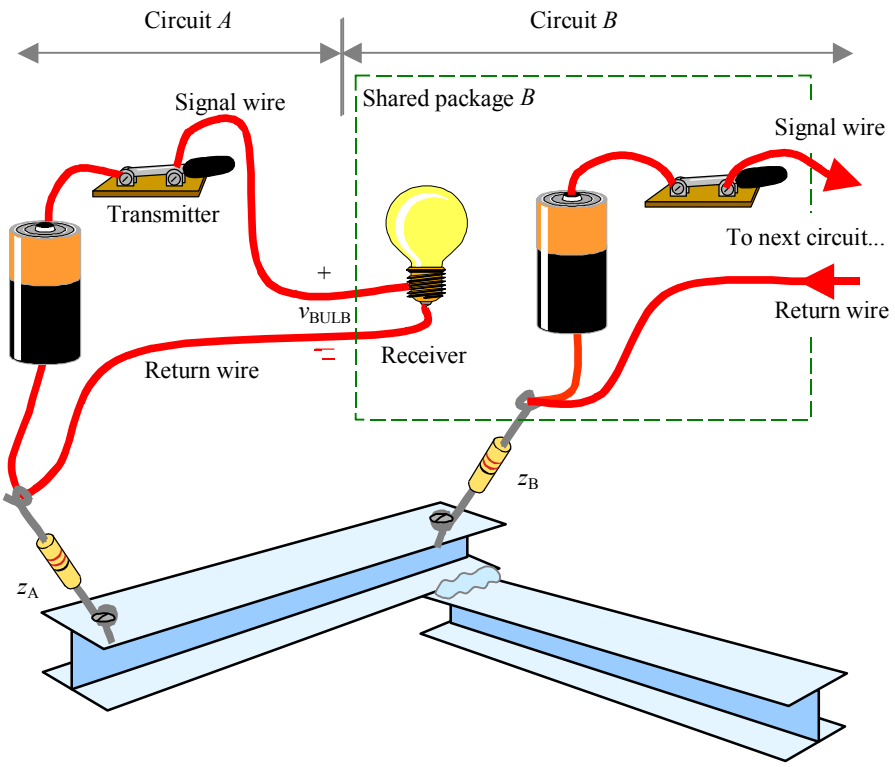

Figure 6.4—*Two-wire* transmission provides a signal wire and return wire for each signal.

disturbances in distribution of global reference voltages, provided the disturbances do not exceed the power-supply noise tolerance of the logic family or the common-mode input range of the receivers.

Second, the two-wire architecture eliminates shared-impedance coupling between a receiver and transmitter in the same package. In Figure 6.4, returning associated with transmitter *B* flows through the return wire back to the battery at *B* without traversing z_B, and therefore *without disturbing the receiver*. By eliminating the shared-impedance coupling between circuits *A* and *B*, two-wire signaling conquers ground bounce locally generated within the package.

Third, two-wire signaling counteracts any type of interfering noise that affects both wires equally. A good example would be the *ground shifts* encountered in a high-speed connector. When two systems are mated by a connector, the net flow of signal current between the systems returns to its source through the ground (or power) pins of the connector. As it does so, tiny voltages are induced across the inductance of the connector's ground (or power) pins. These tiny voltages appear as a difference between the ground (or power) voltage on one side of the connector and the ground (or power) voltage on the other side. This problem is called a *ground shift,* and it is yet another form of common impedance coupling. Two-wire signaling fixes this problem.

These three benefits do not depend on the use of any *changing* voltage on the second wire. As shown in Figure 6.4, the return wire merely carries the local reference voltage (ground, in this case) from the transmitter to the receiver, where it may be observed. This simple circuit renders the system immune to local disturbances in the power and ground voltages, ground bounce generated within a package, and ground bounce generated within a connector. That's pretty good.

The performance of a two-wire signaling circuit hinges on the assumption that no current flows through impedances z_A and z_B. Under this assumption the receiver at *B* can directly observe (on the return wire) the local reference voltage at transmitter *A,* and the next receiver *C* can observe the local reference voltage at transmitter *B*. Any currents flowing through z_A or z_B change the references voltages on the return wires, interfering with reception. The two-wire circuit must be arranged so that it limits the current through z_A and z_B to innocuous levels.

Unfortunately, in a high-speed system all wires couple to the surrounding chassis and other metallic objects, whether you want them to or not. In Figure 6.4 you can model this coupling as a collection of parasitic lumped-element connections connected from each wire to the reference beam. Current transmitted on the signal wire therefore has a choice of returning pathways. It can return to the source along the return wire (the intended path), or it can flow through the parasitic connection to the reference beam and from there return to the transmitter through impedance z_A. The current that flows through the parasitic pathway is called *stray returning signal current*. At high speeds the stray returning signal current is often significant enough to impair the effectiveness of a two-wire signaling system.

Does this impairment defeat the utility of two-wire signaling for high-speed circuits? Not necessarily, provided that you pick a particular, unique signal for the second wire. The second wire must carry a signal equal in amplitude to the first, but opposite in polarity (an *antipodal,* or *complementary,* signal). If you do that, everything still works.

POINTS TO REMEMBER

- Two-wire signaling renders a system immune to disturbances in distribution of global reference voltages.
- Two-wire signaling counteracts any type of interfering noise that affects both wires equally.
- Two-wire signaling counteracts ground bounce (also called simultaneous switching noise) within a receiver.
- Two-wire signaling counteracts ground shifts in connectors.
- Two-wire signaling works when there is no significant stray returning signal current.

6.3 DIFFERENTIAL SIGNALING

The transmission of two complementary signals over identical, matched traces is a special case of two-wire signaling. It is called *differential signaling* (Figure 6.5). A differential

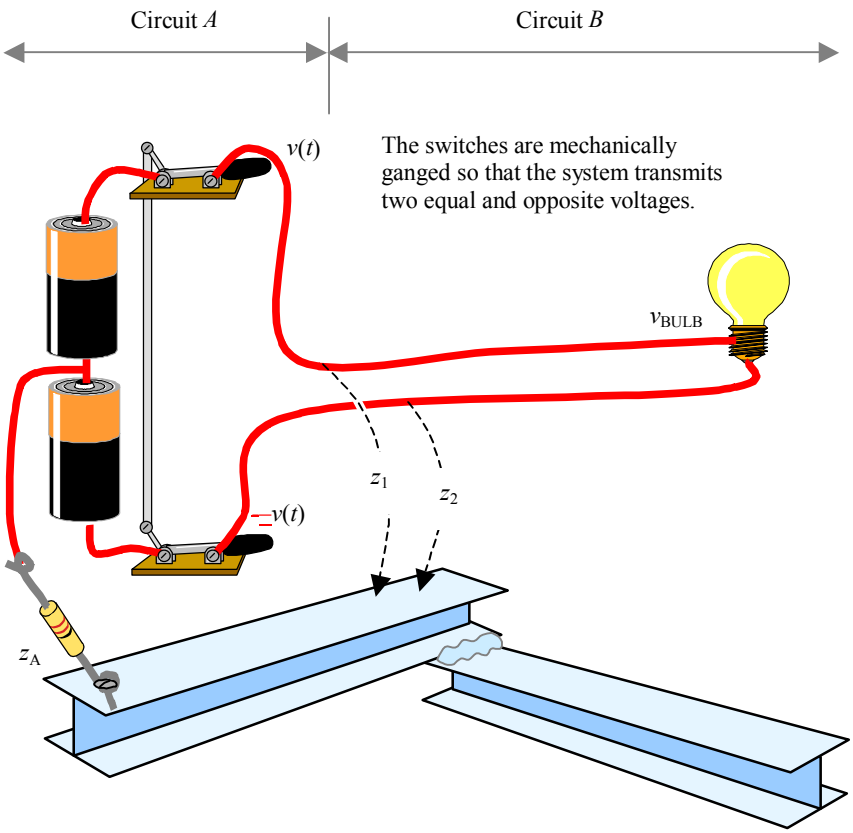

Figure 6.5—*Differential transmission:* As long as parasitic impedances z_1 and z_2 remain well balanced, no net current flows into the reference system.

signaling system delivers equal but opposite AC voltages *and currents* on the two wires. Under these conditions, and assuming the layout is symmetrical so that both wires have equal coupling to the reference system, any AC currents induced in the reference system by one wire are counteracted by equal and opposite signals induced by the complementary wire.

This effect is illustrated in Figure 6.5, where parasitic impedances z_1 and z_2 represent the impedance from one wire to the surrounding reference system (chassis or other bits of metal) and from the other wire to the reference system respectively. In a well-balanced differential system these two impedances are equal. As long as the AC voltages on the two wires are complementary, the stray currents through z_1 and z_2 will cancel, resulting in *no net flow of stray returning signal current in the reference system*.

The performance of such a system does not hinge on the particular value of impedance z_A (since no current flows through it). Conversely, the system enjoys a measure of immunity from other circuits that do induce currents in the reference structure, because the differential receiver does need the voltages on the reference system to be the same everywhere. A differential receiver requires only that the disparity between reference voltages at either end

of a link not cause the received signal to exceed the common-mode operating range of the receiver.

If the impedances z_1 and z_2 are not well balanced, or if the transmitted voltages are not precisely complementary, some amount of current will flow in the reference system. This current is called *common-mode current*. Assuming that the transmitted voltages are equal, I'd like to point out two physical means for reducing the magnitude of common-mode current: the *weak-coupling* approach and the *precise-balance* approach.

High-performance twisted-pair data cabling (Figure 6.6) uses the *weak-coupling* approach, whereby the parasitic coupling through z_1 and z_2 to Earth is generally weakened. The weakening is accomplished by thickening the plastic jacket on the cable to keep other wires and objects outside the jacket relatively far away from the signal conductors. At the same time, the cable holds the wires of each pair closely together in a 100-ohm differential configuration. This geometry increases the magnitude of z_1 and z_2 relative to the impedance of the load. As a result, regardless of where the cable is laid, the fraction of current that can possibly be conducted through z_1 or z_2 cannot be very great. Further reductions in the common-mode current are obtained by tightly twisting the wires so that any nearby objects are approached an equal number of times by each wire. The tight twists produce a better balance between z_1 and z_2. The combination of a weakened coupling magnitude and good balance is used in category 5, 5e, 6, and 7 unshielded twisted-pair cabling to attain spectacular common-mode rejection.

Lest you doubt that impedances z_1 or z_2 could carry enough signal current to represent a significant radiation problem, let me suggest the following experiment. Beginning with a working LAN adapter, strip back a section of the outer jacket of the twisted-pair cable. Untwist the wires of one of the active pairs, and tape one wire of that pair against a solid metal part of the product chassis for a distance of about two inches. Run in this configuration, the product will not pass its FCC or EN mandated radiation tests. The amount of excess common-mode current flowing through the parasitic capacitance of the one wire taped against the chassis causes enough radiation to exceed FCC or EN mandated limits. Unbalanced stray current to ground makes a big difference in twisted-pair transmision systems and must be strictly limited.

Differential traces on a pcb must take a different approach to the control of common-

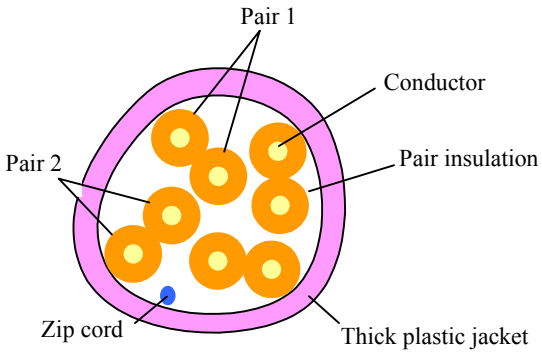

Figure 6.6—Construction of high-performance 4-pair unshielded twisted-pair data cable. The zip cord shown in this figure is a tough string that technicians pull to rip open long sections of the outer insulation.

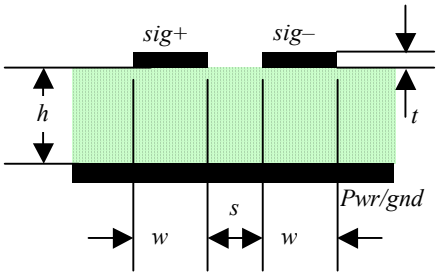

Figure 6.7—Cross section of pcb layout showing edge-coupled microstrip differential pair. Both lines are symmetrically loaded with respect to the underlying solid reference plane.

mode currents. In a pcb with solid plane layers (Figure 6.7) you cannot remove the traces far enough from the solid planes to significantly reduce the magnitude of z_1 or z_2. Common-mode rejection in solid-plane pcb applications is achieved not by weakening the coupling, but by *precisely balancing* it. This is done by using traces of equal height, width, thickness, and length.

In regards to a differential pcb application, it is not particularly important that the traces be coupled tightly to each other. In fact, you couldn't achieve very tight coupling even if you tried. The coupling ratio for typical differential pair on a pcb lies in the 20% to 50% range.

In contrast, for a well-balanced differential twisted pair such as used in the LAN industry, the coupling is excellent. In such a cable if you transmit a fast signal $v(t)$ down one wire of the pair,[47] at the far end you will receive two complementary signals, $v(t)/2$, and $-v(t)/2$. This effect illustrates the good *common-mode rejection* of UTP.

Good twisted-pair cables couple the wires of each pair to each other much better than to the surrounding earth. The crosstalk between the two wires is almost 100%. The pairs affect each other very strongly. The only successful mode of propagation is the differential mode.

In a solid-plane pcb environment, even when the two traces of a differential pair are pressed close together, the coupling between traces still rarely approaches the level of coupling to the surrounding planes. Try it. Working at the near end of a pair of differential traces, impress a signal on trace *A*, with trace *B* grounded. Look at the signals at the far end. You'll see a fairly large amount of crosstalk at the far end, but nothing approaching 100%. Differential traces on a pcb won't do that.

Fortunately, differential traces need not be tightly coupled in order to accrue all the main benefits of a differential architecture. Any reasonable spacing will do, provided that the complementary outputs of the transmitter are well balanced and the traces have symmetrical impedances to the nearby power and ground planes. That's all it takes to get immunity to power and ground noise, immunity to ground bounce, and a marked reduction in emissions.

---

[47] At the near end of the cable, transmit voltage $v(t)$ between wire *A* and earth ground, while holding wire *B* at earth potential. You may do this with a single-ended pulse generator. At the far end of the cable, resistively terminate wires *A* and *B* to a common point, and ground that point to your scope. Look at the signals at the far end of wires *A* and *B*.

Occasionally you will encounter engineers who insist on placing their differential traces very close together. This arrangement does save space, and it marginally improves the crosstalk performance vis-à-vis other nearby signals. What it does not do is magically bullet-proof the system against crosstalk (see Section 6.11.6 "Reducing Local Crosstalk").

With the use of two differential signals, we gain one last important benefit: a reduction in radiated emissions (see Section 6.11.3, "Reducing EMI with Differential Signaling").

POINTS TO REMEMBER

- Differential signaling delivers equal but opposite AC voltages *and currents* on two wires.
- Assuming the layout is symmetrical, any AC currents induced in the reference system by one wire are counteracted by equal and opposite signals induced by the complementary wire.
- Differential pcb traces need not be tightly coupled to be effective.
- Differential signaling markedly reduces radiated emissions.

6.4 DIFFERENTIAL AND COMMON-MODE VOLTAGES AND CURRENTS

On a two-wire transmission circuit, the difference between the instantaneous voltages a and b on the two wires is defined as the differential voltage d:

$$d \triangleq a - b \qquad [6.2]$$

where a and b are each measured with respect to a common arbitrary reference.

On a two-wire transmission circuit, the average of the instantaneous voltages a and b on the two wires is called the common-mode voltage c:

$$c \triangleq \frac{a+b}{2} \qquad [6.3]$$

where a and b are each measured with respect to a common reference, usually a local earth ground, but sometimes a local reference plane or other local reference point.

The differential and common-mode voltages comprise an alternate representation of the original signal, often called a *decomposition* of the original signal. Given the common-mode and differential voltages, you can reconstruct a and b. (The same decomposition applies to currents.)

$$a = c + d/2 \qquad [6.4]$$

$$b = c - d/2 \quad [6.5]$$

In a good differential system one usually strives to limit the AC component of the common-mode signal. This is done because the common-mode portion of the transmitted signal does not enjoy any of the noise-canceling or radiation-preventing benefits of differential transmission. The common-mode and differential signals also propagate differently in most cabling systems, which can lead to peculiar skew or ringing problems if the common-mode component is an appreciable fraction of the overall signal amplitude, especially if those common-mode currents are accidentally converted into differential signals (see Section 6.8, "Differential to Common-Mode Conversion"). Intercabinet cabling, particularly, is extremely sensitive to the presence of high-frequency common-mode currents, which radiate quite efficiently from unshielded cabling.

Another decomposition of the two-wire transmission problem defines odd-mode and even-mode voltages and currents. These are similar to, but slightly different from, differential and common-mode voltages and currents.

An odd-mode signal is one that has amplitude $x(t)$ on one wire and the opposite signal $-x(t)$ on the other wire. A signal with an odd-mode amplitude of $x(t)$ has a differential amplitude of $2x(t)$. If the signal $x(t)$ takes on a peak-to-peak range of y, then the peak-to-peak odd-mode range is simply y, but the *peak-to-peak differential amplitude* is $2y$.

An even-mode signal is the same on both wires. An even-mode signal with a peak-to-peak range of y also has a peak-to-peak common-mode range of y. The even-mode amplitude and common-mode amplitude are one and the same thing.

Two-wire transmission systems sometimes send a signal voltage on one wire, but nothing on the other. In this case the differential-mode amplitude equals the signal amplitude on the first wire. The common-mode amplitude is half that value. In this case the odd-mode and even-mode amplitudes are the same and both equal to half the signal amplitude on the first wire.

Here are the translations between odd-mode and even-mode quantities. The same decomposition applies to currents.

Odd-mode signal: $$o \triangleq \frac{a-b}{2} \quad [6.6]$$

Even-mode signal: $$e \triangleq \frac{a+b}{2} \quad [6.7]$$

Signal on first wire: $$a = e + o \quad [6.8]$$

Signal on second wire: $$b = e - o \quad [6.9]$$

where a and b represent the voltages on the two wires with respect to a common reference.

The differential-and-common-mode decomposition and the even-and-odd mode decomposition share very similar definitions. The discrepancy between the two models has

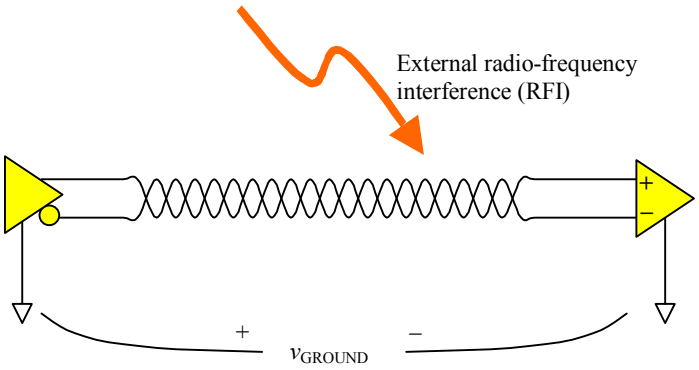

Figure 6.8—A good differential receiver cancels any noise that affects both wires equally, such as external RFI.

to do with the definition of the differential mode. The differential voltage is what you read with an electrical instrument when you put it across two wires. The odd-mode voltage is a mathematical construct that simplifies the bookkeeping in certain situations.

$$o = d/2 \qquad [6.10]$$

$$e = c \qquad [6.11]$$

Any noise like external RF interference that equally affects both wires of a differential pair will induce a common-mode (even-mode) signal, but not a differential-mode (odd-mode) signal (Figure 6.8). A good differential receiver senses only the differential signal and is therefore immune to this type of noise.

POINTS TO REMEMBER

- Differential and common-mode signals are used to describe the voltages and currents on a two-wire transmission system.
- Odd-mode and even-mode signals are yet another way to describe the voltages and currents on a two-wire transmission system.
- Differential receivers cancel common-mode noise.

6.5 DIFFERENTIAL AND COMMON-MODE VELOCITY

In a configuration with a homogeneous dielectric (like a stripline) the propagation velocities for the differential and common-mode (also read odd-mode and even-mode) signals are equal.

Configurations with inhomogeneous dielectrics (like microstrips), however, support slightly different propagation velocities for the two modes. The impact of this difference is

not very great as long as there is not much coupling between the modes (see Section 6.8) and as long as the receiver remains sensitive only to the differential component of the signal. It is theoretically possible to observe deleterious effects from the difference in velocities under the following circumstances:

- A differential signal is created.
- Part of the signal is inadvertently converted to a common-mode signal.
- The differential-mode and common-mode signals propagate independently, and with slightly different velocities, to the far end of a long transmission line.
- The common-mode signal is inadvertently converted back to a differential-mode signal.

In this case the receiver will perceive a superposition of two incident waveforms with slightly different timing and amplitudes (the double-converted signal presumably being smaller).

POINT TO REMEMBER

- Microstrips support slightly different propagation velocities for the differential and common modes. The impact of this difference is not very great.

6.6 COMMON-MODE BALANCE

Common-mode balance is a term that applies to a differential transmission system. Common-mode balance is the ratio of common-mode to differential-mode signal amplitudes within the system. A perfectly differential system with no common-mode component is *perfectly balanced* (i.e., the signals on the two wires are exactly antipodal, or opposite).

When calculating the common-mode balance ratio, it is common to refer only to the AC component of the common-mode signal. For example, a perfectly complementary pair of TTL 3.3-V signals may share a common-mode DC offset of 1.65 volts, but the common-mode AC voltage (the changing part) might still be very small.

The common-mode balance ratio is often expressed in decibels. A common-mode balance of 1 part in 10,000 (–80dB) is exceptionally good. Digital logic parts sold as differential transmitters may exhibit common-mode balance of perhaps only –30 dB or even –20 dB.

Some authors express the common-mode balance in terms of the even-mode to odd-mode signal amplitude. The even-to-odd mode amplitude ratio is half the common-to-differential mode amplitude ratio, making it 6dB less impressive-sounding. An older term used to describe the common-mode to differential-mode amplitude ratio is *longitudinal balance*.

POINT TO REMEMBER

- Common-mode balance is the ratio of common-mode to differential-mode signal amplitudes.

6.7 COMMON-MODE RANGE

Every digital receiver comes with a specification for its common-mode input range. You are expected to ensure that input signals remain within this range at all times. As long as both inputs stay within the common-mode operating range, the component will meet or exceed its specification for the input switching threshold. Beyond that, manufacturers give few clues as to how the component will operate. It may operate normally. On the other hand, it may reverse its outputs, it may saturate and take a long time to recover, it may lock up into a brain-damaged state until power-cycled, or it may permanently fail. You never know which [55]. Don't violate this specification (even for a brief period).

The common-mode range specification is definitely useful, but I'd like to know more about a receiver. For example, let's say you are receiving a 500-mV differential clock input. Add to each signal a common-mode noise voltage of 1-V p-p. How much jitter will come out of the clock? You can't figure that out from the specifications.

Another source of clock jitter is noise from the power supply. Suppose there is a 100-mV AC ripple on V_{CC}. How much jitter will you get? You can't figure that out from the specifications, either.

Both examples deal with the issue of common-mode rejection, which is one measure of how much the input switching threshold changes in response to a defined noise input. For linear amplifiers, it's common to see a common-mode rejection ratio (CMRR) specification for common-mode noise at the input terminals and also for noise at the power terminals. For example, a CMRR of –50 dB for V_{CC} means that a 100-mV ripple on V_{CC} will have the same effect as an equivalent differential noise source of 0.3 mV (that's 100-mV less 50 dB). The CMRR translates each type of noise into an equivalent differential noise level at the input. You can then add up all the equivalent input noise figures to determine the overall signal-to-noise ratio, or jitter performance, of your system.

The common-mode range specification used in digital comparators and receivers doesn't break down the sources of noise, so there's no way to do jitter analysis.

POINT TO REMEMBER

- Don't violate the common-mode input range specification for a receiver (not even briefly).

6.8 DIFFERENTIAL TO COMMON-MODE CONVERSION

Any imbalanced circuit element within an otherwise well-balanced transmission channel creates a region of partial coupling between the differential and common modes of

transmission at that point. The coupling can translate part of a perfectly good differential signal into a common-mode signal, or vice versa.

Such differential-to-common mode conversion problems frequently arise in the design of LAN adapters. For example, assume the output winding of the transformer in Figure 6.9 has equal capacitances C_1 connected from point *a* to ground and from *b* to ground. If the capacitances are exactly equal (and the cable and transformer perfectly symmetrical), the differential signal present on the cable forces equal but opposite currents through these two capacitances. In the product chassis, the two currents perfectly cancel. The perfect cancellation implies that no current circulates between the twisted-pair cable and the surrounding chassis. In actual practice, however, one capacitance is always a little larger than the other.

Let capacitor C_2 in Figure 6.9 represent the small amount of physical imbalance (2 pF) between the parasitic capacitances associated with circuit nodes *a* and *b*. Let's calculate the current flowing through this capacitor, see where it flows, and then decide if it causes any problems.

Using Ethernet 10BASE-T for this example, the drive amplitude is approximately 2V p-p on each wire, at a switching time of 25 ns. The current forced through capacitor C_2 is

$$i(t)_{\text{PEAK}} = C_2 \frac{dv}{dt} = (2\,\text{pF})\frac{2\text{V}}{25\text{ ns}} = 160\,\mu\text{A} \qquad [6.12]$$

This current flows *through* capacitor C_2 to the product chassis. It couples from the product chassis to the Earth (either through the green-wire ground or through the capacitance between the product chassis and the Earth). From the Earth it couples capacitively to the cabling, along which it travels as a common-mode signal riding on the twisted-pair cable

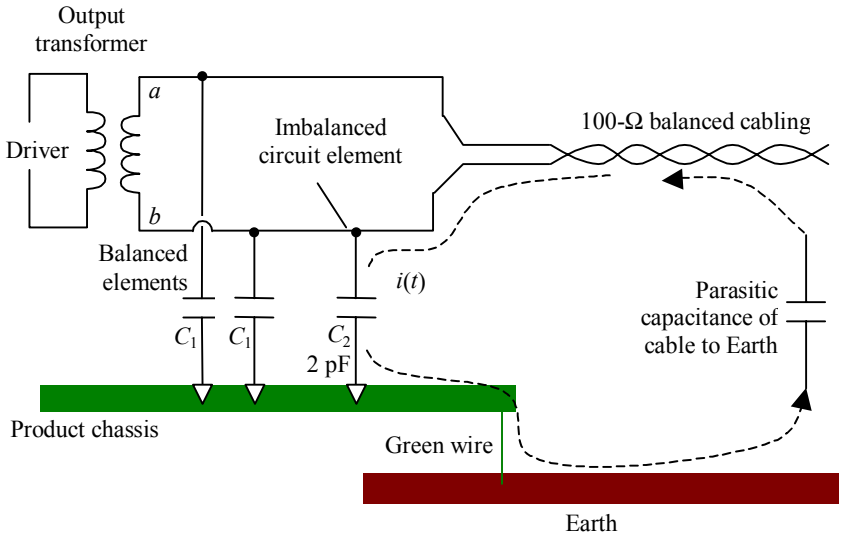

Figure 6.9—An imbalanced circuit element within the transmitter causes current to circulate through the external cabling and the product chassis.

back to the transformer, completing the loop.

A *balanced* load composed of equal-valued capacitors from *a* to ground and from *b* to ground would *not* generate any common-mode currents, because the currents through the two capacitors cancel, leaving nothing to exit the system in common-mode format. In this example the *imbalance* in capacitive loading generates the common-mode current.

A capacitive imbalance even as small as 2 pF causes a big problem in this example, because 160 μA of high-frequency common-mode current on an exposed cable easily violates U.S. and international emissions regulations.

POINT TO REMEMBER

- An imbalanced circuit can translate part of a perfectly good differential signal into a common-mode signal, or vice versa.

6.9 DIFFERENTIAL IMPEDANCE

What is differential impedance? Differential impedance is the ratio of voltage to current on a pair of transmission lines when driven in the differential mode (one signal positive and the other negative).

For example, the circuit in Figure 6.10 drives a signal $x(t)$ differentially into a pair of uncoupled transmission lines.[48] Because the lines are symmetrical, the voltage splits evenly and you see voltage $\frac{1}{2}x(t)$ on the top line and $-\frac{1}{2}x(t)$ on the bottom. The current through each load must therefore equal $\frac{1}{2}x(t)/Z_0$. This current flows through both lines, and through both loads, in the direction shown.

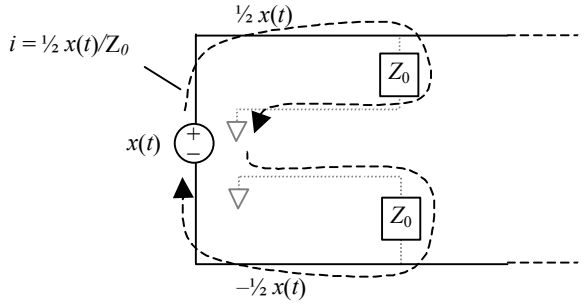

Figure 6.10—The differential input impedance of a pair of noninteracting transmission lines is $2Z_0$.

[48] Traces separated by more than four times the trace height *h* interact so slightly that for purposes of impedance analysis, one usually ignores the effect of one on the other. Such traces are said to be *uncoupled*. The exact degree

The differential impedance (ratio of differential voltage to current) is

$$Z_{\text{DIFF}} = \frac{x(t)}{\left[\tfrac{1}{2}x(t)/Z_0\right]} = 2Z_0 \qquad [6.13]$$

where Z_{DIFF} is the differential impedance of a pair of uncoupled transmission lines (Ω),

Z_0 is the characteristic impedance of either line alone, also called the uncoupled impedance (Ω),

$x(t)$ is the differential voltage applied across both lines (V), and

$\tfrac{1}{2}x(t)/Z_0$ is the current driven into each line by the source (A).

The differential impedance of two matched, noninteracting transmission lines is double the impedance of either line alone.

If the lines are *coupled,* the situation changes. As an example of *coupling,* think of two parallel pcb traces. These traces will always exhibit some (perhaps very small) level of crosstalk. In other words, the voltages and currents on one line affect the voltages and currents on the other. In that way the two transmission lines are coupled together.

With a pair of coupled traces, the current on either trace depends in part on crosstalk from the other. For example, when both traces carry the same signal, the polarity of crosstalk is positive, which reduces the current required on either trace. When the traces carry complementary signals, the crosstalk is negative, which increases the current required on either trace. Figure 6.11 illustrates the situation. The characteristic impedance of each line to ground (including consideration of the adjacent trace) is represented as a single lumped-element component Z_1. The impedance coupling the two transmission lines is depicted as a single lumped-element device Z_2.

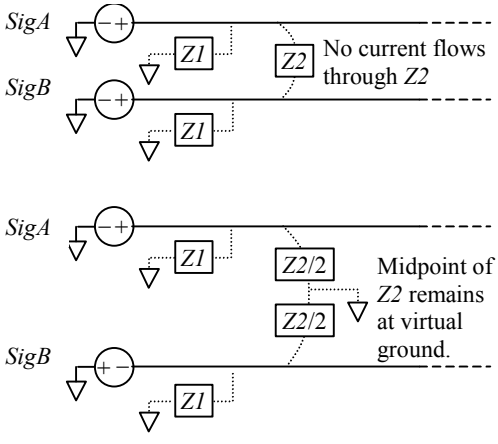

Scenario I
SigB = *SigA*. Impedance seen by *SigA* driver is *Z1*.

Scenario II
SigB = –*SigA* (complementary signaling). Impedance seen by *SigA* driver is *Z1* in parallel with (*Z2*/2).

Figure 6.11—Driving point impedance of two different two-wire scenarios.

of separation required to produce an inconsequential interaction of course depends upon what you consider to be inconsequential.

In the first scenario *SigB* is driven identically with *SigA*. The voltages on the two wires are therefore the same at all points in time, so (conceptually) no current flows through the coupling impedance Z_2. The impedance (ratio of voltage to current) measured for each wire under this condition equals Z_1. This is called the *even-mode impedance* of the transmission structure.

In the second scenario *SigA* and *SigB* are driven with complementary signals. This is called differential signaling, or sometimes antipodal signaling. Under these conditions the midpoint of impedance Z_2 remains at a virtual ground. An AC analysis of this circuit reveals that the impedance measured for each wire under this condition equals Z_1 in parallel with half of Z_2. This value of impedance is called the odd-mode impedance. *The odd-mode impedance is always less than the even-mode impedance.*

The terms even-mode impedance and odd-mode impedance are closely related to the terms *common-mode impedance* and *differential-mode impedance*. Common-mode impedance is measured with the two wires driven in parallel from a common source. Common-mode impedance is by definition half the even-mode impedance.

Differential-mode impedance is twice the odd-mode impedance. It is measured using a well-balanced source, computing the ratio of the differential voltage (twice the odd-mode voltage) to the current on either line.

Combining your knowledge about common and differential mode impedances, you should be able to prove that the common-mode impedance always exceeds one-fourth of the differential-mode impedance.

Just as an example, suppose we have two 50-ohm, uncoupled transmission lines. As long as the lines are far enough apart to remain uncoupled, the even-mode and odd-mode impedances for these two lines will be the same, and equal to 50 ohms. Should you connect these two lines in parallel, the common-mode impedance would be 25 ohms. Should you connect these two lines to a differential source, the differential input impedance of the pair of lines would be 100 ohms.

POINTS TO REMEMBER

- Differential impedance is the impedance measured *between* two conductors when they are driven in the differential mode.
- Odd-mode impedance is the impedance measured *on either of* two conductors when they are driven with opposite signals in the differential mode.
- The value of differential-mode impedance is twice the value of odd-mode impedance.
- The differential impedance of two matched, uncoupled transmission lines is double the impedance of either line alone.
- The odd-mode impedance of two matched, uncoupled transmission lines equals the impedance of either line alone.
- Coupling between two parallel pcb traces decreases both differential and odd-mode impedances.

- Common-mode impedance is the impedance measured on two wires in parallel when they are driven together.
- Even-mode impedance is the impedance measured *on either of* two wires when they are driven with identical signals in the common mode.
- The value of common-mode impedance is half the value of even-mode impedance.

6.9.1 Relation Between Odd-Mode and Uncoupled Impedance

I should now like to discuss the relation between the odd-mode impedance and the *uncoupled impedance*. The uncoupled impedance Z_C is what you would measure if the same transmission lines were widely separated, so they couldn't interact. What you need to know is simple: The odd-mode impedance of a coupled transmission line is always *less* than the uncoupled impedance. The even-mode impedance is always *greater*. The closer you place the line, the more coupling you will induce, and the greater a discrepancy you will see between the odd-mode, uncoupled-mode, and even-mode impedances.

Let's codify this into a differential impedance principle:

*Coupling between parallel pcb traces **decreases** their differential (or odd-mode) impedance.*

When implementing tightly coupled differential traces on a pcb, one normally reduces the width of the lines within the coupled region in order to compensate for the expected drop in differential impedance.

6.9.2 Why the Odd-Mode Impedance Is Always Less Than the Uncoupled Impedance

The proof relies on the construction of a thing called an equipotential plane midway between two differential traces. Due to symmetry, all the electric fields in the odd-mode situation will lie perpendicular to this plane. Therefore, the potential everywhere along the equipotential plane will be zero. If the potential everywhere along the plane is zero, I could replace the imaginary equipotential plane with a real, solid copper wall and it wouldn't make any difference. The odd-mode characteristic impedance is not affected by the wall. What's really neat about this construction is that *once the wall is in place,* the problem is partitioned into two noninteracting zones. This gives us a way to evaluate the odd-mode impedance using a single trace and a solid copper wall instead of two traces. If you think about the impact of the wall on the characteristic impedance of a single trace, it's pretty obvious that the wall will add capacitance and decrease the impedance. The net result of this argument is that the odd-mode impedance of a coupled structure is always less than the uncoupled impedance Z_0.

6.9.3 Differential Reflections

***High-Speed Digital Design Online Newsletter*, Vol. 2, Issue 21**

John Lehew writes

In the *High-Speed Digital Design* book and in a few other places it states the fractional reflection Γ coefficient caused by a mismatch in impedance is

$$\Gamma = \frac{Z_2 - Z_0}{Z_2 + Z_0} \qquad [6.14]$$

where Z_0 is the characteristic impedance of a primary transmission line, and

Z_2 is the characteristic impedance of a mismatched section to which the primary line is coupled.

This formula is typically used to calculate reflections that happen in a single transmission line that is referenced to a ground plane. Does this formula also apply to differential or balanced lines?

Reply

Thanks for your interest in *High-Speed Digital Design*.

Aside from the complications introduced by unbalanced modes, differential transmission lines behave pretty much like single-ended ones. Equation [6.14] applies to both.

Assume I have a section of differential transmission line with a differential impedance of Z_1. Assume I couple that into a load with differential impedance Z_2 (it doesn't matter whether Z_2 is a lumped-element load or another section of differential transmission line with characteristic impedance Z_2). The size of the signal that bounces off the joint, in comparison to the size of the incoming signal, is given by your equation for the reflection coefficient Γ.

Let's do an example using unshielded twisted-pair cabling (UTP). Suppose I couple a section of category 5, 100-Ω (nominal) UTP cabling to another section of category 4 120-Ω (nominal) UTP cabling (such cable is available only in France). The reflection off the joint will be of (nominal) size:

$$\Gamma = \frac{120 - 100}{120 + 100} = 0.09 \qquad [6.15]$$

Now, what could go wrong with this simple example? If the cable is inherently *un*balanced (i.e., more capacitance from one side to ground than on the other), then you have a more complicated situation. In general, there are four modes of propagation involved in the problem, one differential mode and one common mode

for each of the two cables. The complete problem is described by a 4x4 coupling matrix whose entries vary with frequency.

Imperfections in the balancing of the cable result in cross-coupling between the differential modes and the common modes at the joint, which is one of the things that creates EMI headaches.

POINT TO REMEMBER

- Aside from the complications introduced by unbalanced modes, differential transmission lines behave pretty much like single-ended ones.

6.10 PCB CONFIGURATIONS

The requirements for high-speed differential traces in a solid-plane pcb are these:

1. The two traces carry complementary voltages,
2. The two traces carry complementary currents—in conjunction with point 1 this implies that their characteristic impedances be the same,
3. The two traces have equal impedances to the surrounding reference system—ground planes, V_{CC} planes, or both, and
4. The two wires have equal propagation delay.

These requirements may be satisfied by many trace configurations. The most popular cases are the differential microstrip, the edge-coupled stripline, and the broadside-coupled stripline (Figure 6.12).

The author assumes you have access to a good 2-D field solver that can predict the differential impedance for various combinations of six variables: trace width, trace height, trace separation, trace thickness, configuration, and dielectric constant. If you don't have a 2-D field solver, then get one (it comes with any signal integrity package). Tell your

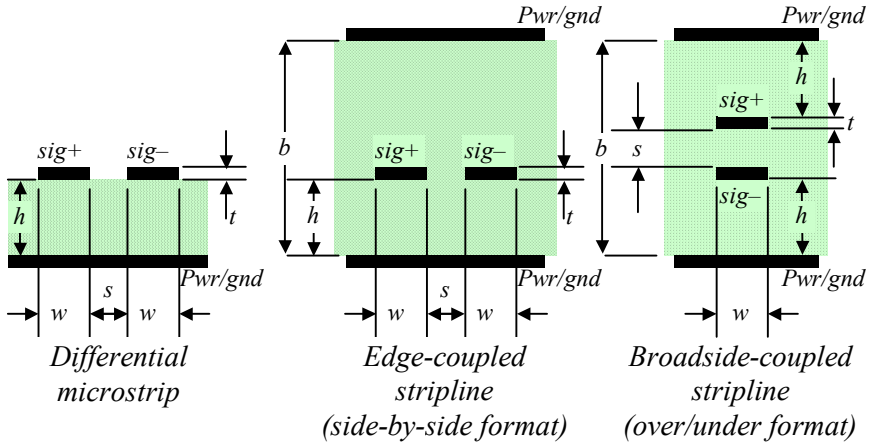

Figure 6.12—Differential pcb traces can be arranged in many different ways.

manager I said you have to have it. A 2-D field solver is absolutely the best way to compute the impedance of any pcb transmission structure.

The following sections present some limited data on differential trace impedance, but don't expect to find all possible combinations of the six primary variables. I will concentrate instead on helping you understand the meaning and purpose of the various adjustments you can make to trace geometry.

6.10.1 Differential (Microstrip) Trace Impedance

***High-Speed Digital Design Online Newsletter*, Vol. 5, Issue 2**

Mitch Morey of San Diego writes

I'm working on a board with 100-ohm differential signaling that I would like to design for microstrip routing. I've used the Polar Instrument calc, the ADS LineCalc software, and have got two additional stack-up constructions from our fabrication houses, and have talked to numerous people on this.

Here are the recommendations I have gathered so far. All configurations represent 100-Ω differential microstrips operating at 2.4Ghz speeds using a 5-mil FR-4 dielectric.

- .005" lines with .005" edge to edge (fab shop 1)
- .004" lines with .008" edge to edge (fab shop 2)
- .005" lines with .008" edge to edge (ADS LineCalc)
- .006" lines with .0065" edge to edge (ADS LineCalc #2)
- .016" lines with .016" edge to edge (engineer #1)

Why the discrepancies?

Reply

Thanks for your interest in *High-Speed Digital Design*.

What you need is a piece of software called a 2-D E&M field solver. This program calculates the magnetic and electric fields surrounding your traces, and from that data extracts the impedance and delay. This is the best way to do impedance calculations. The good field solvers allow you to specify the trace width, height, spacing, thickness, dielectric constant, *and* they allow you to overlay the trace with a soldermask layer.

I'm not sure what ADS LineCalc uses, but if it's not a 2-D field solver, you shouldn't trust its results. I have reason to distrust the accuracy of the examples you have provided.

First let me give you some general principles to help you understand what's happening, and then I'll rule out a couple of the solutions below.

The first thing you need to know is that the patterns of electric field lines in a dielectric medium follow the same shapes as patterns of current flow in salt water. This sounds pretty obtuse, but it's going to help you in a major way, because it will help you *see* what is happening when you change the trace geometry.

Follow me for a minute on this mental experiment. Imagine a microstrip of length *x*. I want you to mentally replace the dielectric medium surrounding this trace

with a slightly conductive salt-water mixture. Now imagine that you connect an ordinary ohmmeter between the trace and the ground plane. The value of DC resistance you measure in this experiment will be exactly proportional to the *impedance* of the trace. I hope you can now imagine what would happen if you press the trace closer to the ground plane. Can you see that the impedance must go *down,* because there is now less water between the trace and the reference plane? If the trace is pressed down to the point where it nearly touches the reference plane, its resistance to ground (i.e., impedance) approaches zero.

What about doubling the width? This adjustment doubles the surface area of the trace, substantially lowering its resistance to the reference plane (i.e., impedance). I like this DC analogy because most engineers find it a lot easier to imagine simple patterns of DC current flow than they do high-frequency electromagnetic fields. The constant of proportionality isn't important—I just want you to see what's going to happen as you make various adjustments.

So far I've shown two things that decrease the impedance in microstrips:

1. Moving a trace closer to the reference plane decreases its impedance.
2. Fattening a trace (i.e., increasing its width) decreases its impedance.

The converse statements are also true:

1. Moving a microstrip further away from the reference plane increases its impedance.
2. Shaving down the trace width increases its impedance.

Stripline traces are a little more complicated in that you must account for the distance from your trace to both top and bottom reference planes. The general result for offset striplines is that whichever plane lies closest to the trace has the most influence on the impedance. Smack in the middle, the planes are both equally important.

Let's now imagine a differential configuration with two traces, still embedded within your slightly salty water. Connect the ohmmeter between the two traces (from one to the other). The resistance you read now will be proportional to the *differential impedance* of two-trace configuration. [Note: One-half the differential impedance is defined as the odd-mode impedance.]

If your two traces are set far apart, and they have the same dimensions as in the first experiment, your new differential measurement will be exactly *twice* as great as before. If you draw out the patterns of DC current flow, you can see why. For *widely* separated traces, the current flows mostly from one trace straight down to the nearest reference plane, then it shoots across the plane to a position underneath the second trace, and from there it leaks back up to the second trace. As this current flows, it encounters a resistance R when leaving the first trace, practically zero resistance flowing across the plane, and then another amount R as it flows back up to the second trace. The total resistance encountered is $2R$.

The differential impedance of two widely separated traces equals twice the impedance to ground of either trace alone.

Now let's see what happens to the differential impedance as you slide the two traces towards each other. When they get close enough, significant amounts of current begin to flow directly between the traces. You still get the same old currents going to and from the reference plane, but in addition to those currents you have now developed a new pathway for current, directly from trace to trace. This additional current pathway acts like a new resistance in *parallel* with the original, widely spaced current pathways. The new parallel pathway lowers the differential impedance of the configuration. You may conclude that

The differential impedance of a tightly spaced pair is less than twice the impedance to ground of either trace alone.

If the traces are moved so close that they nearly touch, the differential resistance (impedance) approaches zero. In general, the differential impedance is a monotonic function of the trace separation.

All other factors being equal, the tighter the intertrace spacing, the less the differential impedance.

I view any decrease in impedance as an annoying side effect of close spacing. If I could redesign the universe, I'd try to make it not happen. Fortunately, you can counteract the annoying drop in impedance by shaving down the width of your traces. If you shave off just enough width, you can push the impedance back up to where it belongs. In this way, the trace separation and trace width become somewhat interchangeable.

To maintain constant impedance, a reduction in spacing must be accompanied by a reduction in trace width (or an increase in trace height).

With these eight rules in mind, let's now look at the specific recommendations you have been given.

With your 5-mil dielectric, the individual impedance of a 16-mil trace on FR-4 already falls below 50 ohms, so the differential impedance will be less than 100 ohms regardless of what spacing you use. You can therefore rule out the 16-mil configuration. I suspect your engineer #1 may have been thinking about using a thicker dielectric than what you propose.

The two ADS LineCalc results conflict with each other. Starting from a pair 5-mil wide with an 8-mil space, *increasing* the trace width to 6 mils will lower the impedance, and *decreasing* the spacing to 6.5 mils will lower it even further. Therefore, one of these results must be wrong. They cannot both be 100-ohm solutions. Therefore, I suspect something is wrong with either your copy of ADS LineCalc or (dare I say it) your use of the tool.

Table 6.1 presents some data from another commercial 2-D field solver (HyperLynx). All these combinations should give you a 100-ohm differential microstrip impedance under the conditions listed in the table:

Each row lists the trace height h, the finished, plated trace width w, the finished, plated edge-to-edge separation s, the proximity factor k_p, the skin-effect resistance R_{AC}, the skin-effect loss coefficient α_r, and the effective permittivity ϵ_{re}. All

Table 6.1—AC Resistance and Skin-Effect Loss (at 1 GHz) for Selected 100-Ω Differential Edge-Coupled Microstrips

h mil	w mil	s mil	k_p	R_{AC} Ω/in.	R_{AC} Ω/m	α_r dB/in.	α_r dB/m	ϵ_{re}
5	8	30	3.48	1.54	60.7	0.067	2.63	3.17
5	7	11	3.17	1.57	61.7	0.068	2.68	2.97
5	6	7	3.01	1.69	66.5	0.073	2.89	2.85
5	5	5	2.91	1.89	74.4	0.082	3.23	2.78

NOTE (1)—AC parameters R_{AC}, k_p, α_r, ϵ_{re}, and FR-4 dielectric $\epsilon_r = 4.3$ are specified at 1 GHz.
NOTE (2)—These microstrip examples assume copper traces of 1-oz thickness (including plating) with $\sigma = 5.98 \cdot 10^7$ S/m, plus a conformal coating (soldermask) consisting of a 12.7-μm (0.5-mil) layer having a dielectric constant of 3.3.
NOTE (3)—The propagation velocity v_0 (m/s) and delay t_p (s/m) are found from ϵ_{re}, where $1/t_p \triangleq v_0 \triangleq c/\sqrt{\epsilon_{re}}$ and $c = 2.998 \cdot 10^8$ m/s.

AC parameters are specified at 1 GHz. The trace construction is assumed to be 1/2-oz etch plus 1/2-oz of plating, yielding a total 1-oz thickness, with a conformal coating (soldermask) consisting of a 12.7-μm (0.5-mil) layer having a dielectric constant of 3.3. If you select a different type of soldermask, your pcb vendor will adjust the trace width to compensate for the thickness of the dielectric material above the traces. A thicker soldermask will slightly reduce the finished propagation velocity.

The resistance data was developed using a method-of-moments magnetic field simulator with 120 segments equally spaced around each pcb trace, with the current linearly interpolated across each segment. The author estimates the accuracy of the data generated by this simulator at approximately ±2%.

Whatever you choose to do, insist that your board fabrication shop place differential impedance test coupons on your panels and test each one to verify that you are getting the correct impedance.

6.10.2 Edge-Coupled Stripline

In a pcb application, the differential impedance of closely spaced traces varies with their spacing. As the gap between traces narrows, the differential impedance goes down. In extreme cases, the widths of the traces may require adjustment in order to keep the differential line impedance within a specified target range. This pesky adjustment in width is the one key disadvantage of closely spaced differential lines.

Figure 6.13 plots the impedance of edge-coupled differential traces versus spacing and trace width. The data for this plot were generated using a method-of-moments magnetic field simulator with 120 segments equally spaced around each pcb trace, with the current linearly interpolated across each segment.

At small separations (less than 9 mils, as shown in Figure 6.13) the traces couple significantly, so the impedance varies with both trace width and spacing. When the intrapair

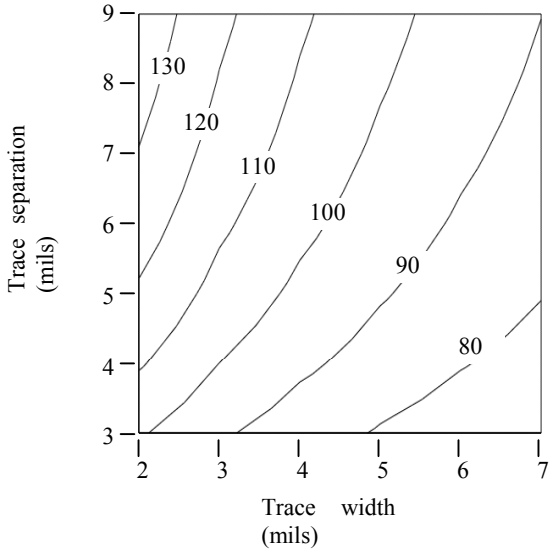

Per Figure 6.12
b = 24 mil
h = 6 mil
t = 0.68 mil
ϵ_r = 4.3

Figure 6.13—Differential impedance Z_{DIFF} (ohms) of edge-coupled differential striplines.

separation exceeds about four times the trace height h (a separation of 24 mils in this example), the coupling between the two traces of the pair usually becomes so weak that the traces hardly affect each other. The differential impedance then depends mostly on just the trace width. At great separations the contours become completely vertical.

Very widely separated traces comprise an *uncoupled* or *loosely coupled* differential pair. For an uncoupled differential pair, the even-mode and odd-mode impedances are the same, and the differential impedance is twice the impedance of either line alone.

Differential traces *can* be pushed really, really close together. Squishing them together saves board area. If you do so, you will need to compute a new trace width to compensate for the fact that the differential impedance goes down for closely spaced pairs. Widely spaced pairs are not subject to this picky, difficult-to-implement requirement.

Another disadvantage of closely spaced pairs has to do with trace routing. Once the signals are closely paired, they *cannot* be separated, or else you will mess up their impedance (unless you readjust their widths). This effect imposes a routing penalty on edge-coupled traces, because it is difficult to get closely spaced, edge-coupled pairs to go around obstacles without temporarily separating them (Figure 6.14).

Finally, the use of tight coupling decreases the trace width, exacerbating the skin-effect loss.

What do I do? Unless absolutely pressed for space, I normally set the trace separation at about four times the trace height h. This setting usually yields a less-than-6% reduction in impedance, a small enough value to simply ignore. All stripline traces, differential or not, then have the same width. I instruct my layout professional to keep differential pairs near each other, but allow them to separate from time to time as needed to go around obstacles. I also insist that the elements of each pair be equal in length, to within 1/20 of a risetime, which limits the common-mode signal contributed by trace skew to less than 2.5% of the single-ended signal amplitude.

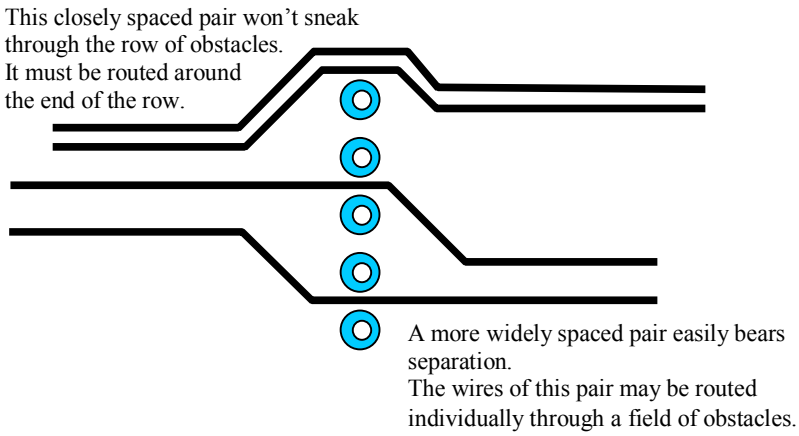

Figure 6.14—Closely spaced differential pairs can be difficult to route.

Figure 6.15 illustrates the magnetic lines of force surrounding an edge-coupled differential stripline pair. It shows intense concentrations of magnetic flux near the corners of the traces, indicating a substantial peaking of the current density at the corners. The proximity factor k_p for these traces takes into account the crowding of current at the corners of the traces, plus a slight current concentration on the inside-facing surfaces of the pair, plus an allotment for the current induced on the top and bottom reference planes of the stripline cavity. The value of the proximity factor k_p for the traces illustrated in Figure 6.15 is 3.08. Differential pcb traces with impedances between 100 to 150 ohms possess a proximity factor typically within the range of 2.5 to 3.5. Figure 6.16 illustrates the distribution of current around the periphery of the signal conductors, and also the lower reference plane. Current flows also on the upper reference plane, although that plane lies above the limits of the vertical scale in the drawing, so it is not shown.

Figures 6.15 and 6.16 were generated using a 2-D field solver. This same program also calculates skin-effect resistance and skin-effect loss. A listing of such calculations for selected edge-coupled stripline geometries appears in Table 6.2. These values provide a good starting point for planning a high-speed interconnection. The table shows that the most important determiner of skin-effect loss is simply the trace width. All other factors are

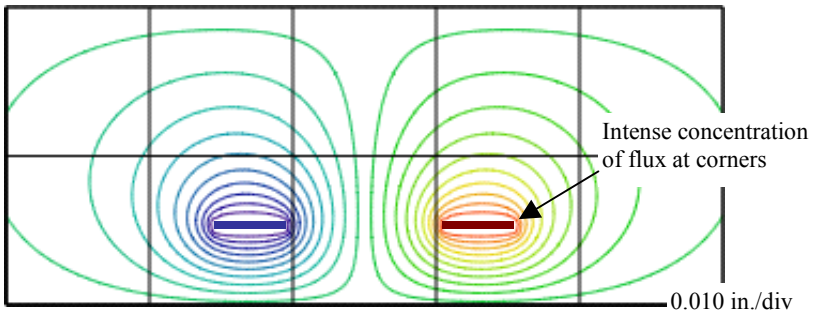

Figure 6.15—This cross-sectional view of the magnetic field in the vicinity of a differential edge-coupled stripline shows intense field concentrations near the corners of the traces.

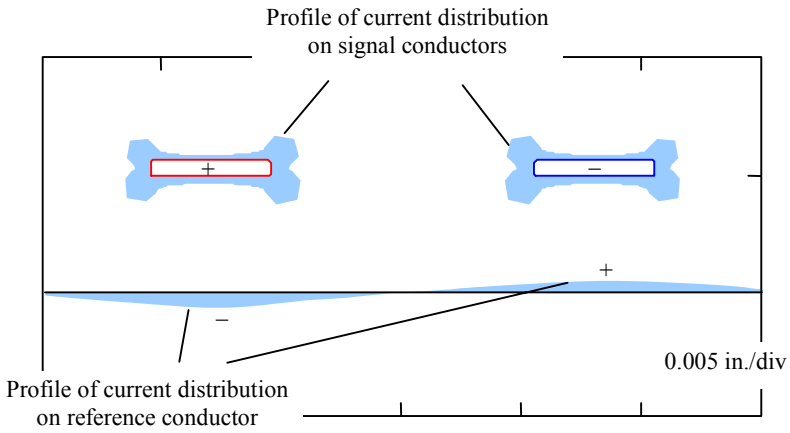

Figure 6.16—The distribution of current around the perimeter of a differential edge-coupled stripline takes on a dog-bone appearance, with the current peaking at the corners of the trace. The inside-facing corners of the pair have a slightly higher peak current than the outside-facing corners.

secondary.

If you don't find the configuration you need in the table, you may interpolate values of r_{SKIN} and α_R from the entries listed in the table. The interpolation process begins by first correcting the listed table values to account for your actual value of b. This is accomplished by scaling the trace geometry as follows:

- Select any row from the table with an interplane spacing b_2 that is reasonably close to your target b_1.
- Compute the scaling constant $k = b_1/b_2$.
- Multiply h, w, s and also the trace thickness t by the constant k, producing a new row for the table.
- The impedance of the new row will be identical to the original Z_{DIFF}.
- The skin-effect resistance (at 1 GHz) for the new row will equal the original r_{SKIN} divided by k.
- The skin-effect attenuation constant for the new row will equal the original α_r divided by k.

This scaling process is exact; however, it produces an unusual value of t (the thickness is usually standardized to either 1/2 or 1 oz of copper). Because t has only a second-order effect on characteristic impedance and resistance, you can usually ignore this change. Using the scaled values of b, h, w, and s while leaving t fixed will introduce some small inaccuracies into the scaling process, but still produce a new configuration with a characteristic impedance workably close to Z_{DIFF}.

Once you have corrected a few of the table entries to account for your actual value of b, you can then interpolate along the h and w dimensions to estimate the skin-effect resistance and skin-effect loss.

This principle of scaling and interpolation works with the results from any field solver.

Example showing scaling of edge-coupled differential stripline

Suppose you have implemented a differential layout with $b = 20$, $h = 7$, $w = 5$, and $s = 7$. The skin-effect loss (Table 6.2) is 0.0929 dB/in versus a hypothetical budget of 0.1 dB/in.

How much can you squeeze the trace width without exceeding the skin-effect loss budget?

First let's try changing w and s without changing b. The next smallest trace size listed in Table 6.2 shows $w = 4$, $s = 5.2$, yielding a resistive loss of 0.1089 dB/in. You can adjust w and s 58% of the way towards this next smallest size, where the fraction 58% is determined by the ratio of skin-effect loss numbers:

$$58\% = \frac{0.1 - 0.0929}{0.1089 - 0.0929} \qquad [6.16]$$

The interpolated values of w and s are

$$\begin{aligned} w &= 5 + 0.58(4 - 5) = 4.42 \\ s &= 7 + 0.58(5.2 - 7) = 5.96 \end{aligned} \qquad [6.17]$$

Next let's start the problem over, this time changing b. You have plenty of excess budget for skin effect loss in your initial configuration, so the new b can be smaller than the original, where the ratio of 93% is determined by the ratio of skin-effect losses:

$$93\% = \frac{0.0929}{0.1} \qquad [6.18]$$

Scaling all values by the factor 93% produces

$$\begin{aligned} b &= 0.93 \cdot 20 = 18.6 \\ h &= 0.93 \cdot 7 = 6.51 \\ w &= 0.93 \cdot 5 = 4.65 \\ s &= 0.93 \cdot 7 = 6.51 \end{aligned} \qquad [6.19]$$

This example has produced two new approximate configurations. The next step is to tweak the exact spacing s for each configuration to zero in on the exact impedance. This is done using your 2-D field solver. Then choose the exact interpair pitch to satisfy your crosstalk constraint. After that you will be in a position to determine which approach gives you the best layout density: just shrinking w and s while keeping b constant or shrinking b and everything else along with it.

POINTS TO REMEMBER

- Differential traces *can* be pushed really, really close together. If you do so, compute a new trace width to compensate for the fact that the differential impedance goes down for closely spaced pairs.
- Widely spaced (i.e., loosely-coupled) pairs are not subject to picky, difficult-to-implement spacing and width requirements.
- The most important determiner of skin-effect loss is the trace width.
- An interpair trace separation of four times *h* yields about a 6% effect on impedance, a small enough value in many cases to simply ignore.
- Matching the elements of each pair to within 1/20 of a risetime limits the common-mode signal contributed by trace skew to less than 2.5% of the single-ended signal amplitude.

Table 6.2—AC Resistance and Skin-Effect Loss (at 1 GHz) for Selected 100-Ω Differential Edge-Coupled Striplines

b	*h*	*w*	*s*	R_{AC} Ω/in.	R_{AC} Ω/m	α_r dB/in.	α_r dB/m	Z_{DIFF} (Ω)
10	3	3	40.0	3.50	137.8	0.152	5.99	99.0
10	4	3	7.0	3.24	127.6	0.139	5.47	100.4
10	5	3	7.0	3.22	126.8	0.137	5.40	101.2
10	5	4	40.0	2.76	108.7	0.126	4.95	94.6
14	4	3	5.5	3.19	125.6	0.136	5.36	101.0
14	4	4	12.0	2.74	107.9	0.118	4.63	100.3
14	5	3	4.5	3.14	123.6	0.135	5.31	100.1
14	5	4	7.5	2.60	102.4	0.112	4.39	100.5
14	5	5	40.0	2.33	91.7	0.101	3.98	99.5
14	7	3	4.5	3.11	122.4	0.132	5.19	101.8
14	7	4	6.5	2.56	100.8	0.110	4.31	100.6
14	7	5	13.0	2.24	88.2	0.096	3.77	100.9
14	7	6	40.0	2.01	79.1	0.091	3.59	95.0

NOTE—All values *b*, *h*, *w*, and *s* in mils.
NOTE—AC parameters R_{AC}, α_r, and Z_{DIFF} specified at 1 GHz with FR-4 dielectric $\epsilon_r = 4.3$.
NOTE—These stripline examples assume copper traces of 1/2-oz thickness with $\sigma = 5.98 \cdot 10^7$ S/m.

Table 6.2 (continued)—AC Resistance and Skin-Effect Loss (at 1 GHz) for Selected 100-Ω Differential Edge-Coupled Striplines

b	*h*	*w*	*s*	R_{AC} Ω/in.	R_{AC} Ω/m	α_r dB/in.	α_r dB/m	Z_{DIFF} (Ω)
20	5	3	4.4	3.14	123.6	0.134	5.28	101.0
20	5	4	6.5	2.59	102.0	0.111	4.37	100.7
20	5	5	11.0	2.27	89.4	0.097	3.84	100.6
20	5	6	40.0	2.09	82.3	0.092	3.61	98.4
20	7	3	3.9	3.13	123.2	0.134	5.26	101.1
20	7	4	5.2	2.55	100.4	0.109	4.29	100.9
20	7	5	7.0	2.17	85.4	0.093	3.66	100.7
20	7	6	10.0	1.91	75.2	0.082	3.24	100.3
20	7	7	19.0	1.75	68.9	0.075	2.96	100.7
20	7	8	40.0	1.61	63.4	0.072	2.85	96.3
20	10	3	3.7	3.14	123.6	0.135	5.30	100.6
20	10	4	5.0	2.54	100.0	0.108	4.25	101.6
20	10	5	6.5	2.15	84.6	0.092	3.60	101.4
20	10	6	8.5	1.88	74.0	0.081	3.17	100.7
20	10	7	12.0	1.68	66.1	0.072	2.85	100.5
20	10	8	25.0	1.56	61.4	0.067	2.65	100.3
30	5	3	4.3	3.15	124.0	0.135	5.30	100.7
30	5	4	6.3	2.60	102.4	0.111	4.38	100.7
30	5	5	10.0	2.27	89.4	0.097	3.83	100.7
30	5	6	22.0	2.09	82.3	0.090	3.54	100.2
30	6	3	4.0	3.14	123.6	0.134	5.27	101.0
30	6	4	5.4	2.56	100.8	0.110	4.33	100.7
30	6	5	7.5	2.20	86.6	0.094	3.71	100.7
30	6	6	11.2	1.96	77.2	0.084	3.31	100.6
30	6	7	20.0	1.81	71.3	0.078	3.07	100.4

NOTE—All values *b*, *h*, *w*, and *s* in mils.
NOTE—AC parameters R_{AC}, α_r, and Z_{DIFF} specified at 1 GHz with FR-4 dielectric $\epsilon_r = 4.3$.
NOTE—These stripline examples assume copper traces of 1/2-oz thickness with $\sigma = 5.98 \cdot 10^7$ S/m.

Table 6.2 (continued)—AC Resistance and Skin-Effect Loss (at 1 GHz) for Selected 100-Ω Differential Edge-Coupled Striplines

b	*h*	*w*	*s*	R_{AC} Ω/in.	R_{AC} Ω/m	α_r dB/in.	α_r dB/m	Z_{DIFF} (Ω)
30	7	3	3.8	3.14	123.6	0.134	5.29	100.9
30	7	4	5.0	2.56	100.8	0.109	4.31	100.9
30	7	4.5	5.7	2.35	92.5	0.101	3.96	100.7
30	7	5	6.5	2.17	85.4	0.093	3.67	100.6
30	7	6	8.8	1.91	75.2	0.082	3.22	100.7
30	7	7	12.5	1.73	68.1	0.074	2.92	100.8
30	7	8	21.0	1.60	63.0	0.069	2.72	100.5
30	8	3	3.7	3.15	124.0	0.134	5.29	101.0
30	8	4	4.7	2.56	100.8	0.110	4.33	100.6
30	8	5	6.0	2.17	85.4	0.093	3.66	100.7
30	8	6	7.7	1.89	74.4	0.081	3.20	100.6
30	8	7	10.2	1.69	66.5	0.073	2.85	100.9
30	8	8	14.0	1.55	61.0	0.066	2.62	100.6
30	8	9	23.0	1.45	57.1	0.062	2.45	100.5
30	8	10	40.0	1.36	53.5	0.060	2.37	97.8
30	10	3	3.6	3.16	124.4	0.135	5.30	101.2
30	10	4	4.5	2.57	101.2	0.110	4.32	101.0
30	10	5	5.5	2.17	85.4	0.093	3.67	100.7
30	10	6	6.8	1.89	74.4	0.081	3.19	100.7
30	10	7	8.4	1.67	65.7	0.072	2.82	100.9
30	10	8	10.5	1.51	59.4	0.065	2.55	100.8
30	10	9	13.5	1.38	54.3	0.059	2.34	100.6
30	10	10	19.0	1.29	50.8	0.055	2.18	100.6
30	10	11	34.0	1.22	48.0	0.053	2.08	100.2
30	10	12	40.0	1.15	45.3	0.052	2.03	96.4
30	15	3	3.5	3.17	124.8	0.135	5.33	101.0
30	15	4	4.3	2.58	101.6	0.111	4.36	100.7
30	15	5	5.3	2.18	85.8	0.093	3.65	101.2
30	15	6	6.3	1.89	74.4	0.081	3.19	100.9
30	15	7	7.5	1.67	65.7	0.072	2.82	101.0
30	15	8	9.0	1.50	59.1	0.064	2.53	100.9
30	15	9	11.0	1.36	53.5	0.058	2.30	101.0
30	15	10	13.0	1.25	49.2	0.054	2.14	100.0
30	15	11	17.0	1.17	46.1	0.050	1.98	100.3
30	15	12	25.0	1.10	43.3	0.047	1.87	100.5

NOTE (1)—All values *b*, *h*, *w*, and *s* in mils.
NOTE (2)—AC parameters R_{AC}, α_r, and Z_{DIFF} specified at 1 GHz with FR-4 dielectric $\epsilon_r = 4.3$.
NOTE (3)—These stripline examples assume copper traces of 1/2-oz thickness with $\sigma = 5.98 \cdot 10^7$ S/m.

6.10.3 Breaking Up a Pair

Article first published in *EDN Magazine*, November 9, 2000

The two traces comprising a differential pair, when routed close together, share a certain amount of cross-coupling. This coupling lowers the differential impedance between the traces. For example, when two traces with $Z_0 = 50\ \Omega$ are well separated (uncoupled), the differential impedance between them should be precisely $Z_{DIFF} = 2Z_0 = 100\ \Omega$. When you jam the same two traces close together (tightly coupled), as you might do to improve routing density, the differential impedance will be a lower value, perhaps something in the range of 70 to $90\ \Omega$.

If the coupling effect lowers the differential impedance too much for your taste, you can fix it. Just reduce the trace widths, thus raising their impedance. The traces remain coupled, but you can (theoretically) always push the differential impedance back up to $100\ \Omega$ by making the traces skinnier.[49]

Model the mismatched region as a short transmission line of impedance Z_{DIFF} plus a lumped inductance L_{EXCESS}.

What happens to a tightly coupled pair when it traverses an obstacle, such as a via (Figure 6.17)? If you have room to route both traces on the same side of the obstacle, maintaining their constant separation, no special problems arise. If, on the other hand, you separate the traces to pass by, then the differential impedance in the separated region reverts to its original, uncoupled value of $2Z_0$. If you have thinned the trace widths to produce exactly $100\ \Omega$ in the coupled state, then the reverted, uncoupled impedance with skinny traces will *exceed* $100\ \Omega$.

To calculate (approximately) the effects of such a mismatch, let's assume you have a long, uniform differential transmission configuration with differential impedance Z_{DIFF}. Insert into the middle of this line a short section with differential impedance Z_2, having length (in time) t_d. Further assume that t_d is much less than

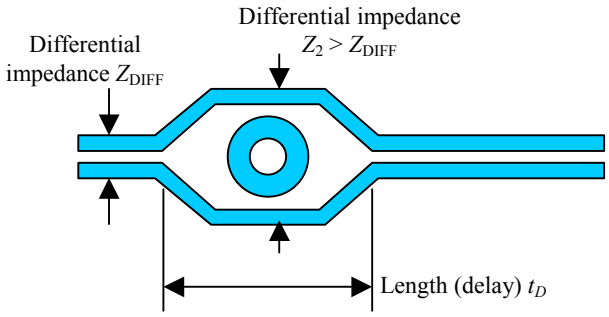

Figure 6.17—The impedance of a tightly-coupled differential pair changes when you separate the wires.

[49] You must avoid placing the traces so close together that the required trace width becomes unmanufacturable.

the signal risetime (or fall time) t_r, so the mismatched section acts as a simple lumped-element circuit.

Analysis begins by computing the values L_2 and C_2 corresponding to the mismatched section: $C_2 \triangleq t_d / Z_{DIFF}$ and $L_2 \triangleq t_d Z_{DIFF}$. Next, mentally break L_2 into two pieces, $L_2 \triangleq L_N + L_{EXCESS}$, with $L_N \triangleq Z_{DIFF}^2 C_2$.

The value L_N is the *natural inductance* you would expect to accompany capacitance C_2 in a differential transmission line with differential impedance Z_{DIFF}, as evidenced by the ratio $\sqrt{L_N / C_2} = Z_{DIFF}$.

Inductor L_{EXCESS} represents the *excess inductance* in the mismatch region *above and beyond* L_N. In other words, this procedure models the mismatch region as a short transmission line of impedance Z_{DIFF} (comprising L_N and C_2) plus a series inductance L_{EXCESS}. You will next model the reflection generated by component L_{EXCESS}.

When a fast step input hits any small inductive discontinuity, you get a reflected pulse. The pulse duration equals the rise (or fall) time of the incoming step. You can approximate the reflection coefficient Γ (ratio of reflected pulse height to the incoming step size) generated by L_{EXCESS} as follows:

$$\Gamma \approx \frac{1}{2t_r} \frac{L_{EXCESS}}{Z_{DIFF}} \qquad [6.20]$$

Substituting the definition of L_{EXCESS} yields

$$\Gamma \approx \frac{1}{2t_r} \frac{L_2 - Z_{DIFF}^2 C_2}{Z_{DIFF}} \qquad [6.21]$$

Further substituting your basic expressions for L_2 and C_2 gives

$$\Gamma \approx \frac{1}{2t_r} \frac{t_d Z_2 - Z_{DIFF}^2 \left(t_d / Z_2 \right)}{Z_{DIFF}} \qquad [6.22]$$

And consolidating the terms provides

$$\Gamma \approx \frac{t_d}{2t_r} \left(\frac{Z_2}{Z_{DIFF}} - \frac{Z_{DIFF}}{Z_2} \right) \qquad [6.23]$$

That's about the best approximation you will find for the case of a short separation between the elements of a differential pair. If you want better accuracy, use a time-domain simulator.

In the example of Figure 6.17, supposing the ratio of Z_2 / Z_{DIFF} to be $(122\,\Omega)/(100\,\Omega) = 1.22$, equation [6.23] reduces to

$$\Gamma \approx \frac{t_d}{2t_r}\left(1.22 - \frac{1}{1.22}\right) = 0.200\frac{t_d}{t_r} \qquad [6.24]$$

If this amount of signal degradation is troublesome, try thickening the traces in the separated region to match the impedance of the thinner, more highly coupled traces elsewhere.

If you've kept t_d/t_r less than 1/6, you can expect a proportional accuracy of a couple of percentage points from approximation [6.23]. If you try to stretch the approximation to a ratio of t_d/t_r as big as 1/3, expect the approximation to be good only to about 20%. Beyond that, at t_d/t_r = 1/2 it falls completely apart, delivering totally erroneous answers.

The same approximation works for BGA layouts, where signals escaping from the inner rows neck down to pass through the BGA ball field. The neck-down region raises the local trace impedance in a small region.

In the event your separated traces pass particularly close to the edges of a via pad, the parasitic capacitance between your trace and the via may add to the value of C_2 in a noticeable way, modifying the values of L_N, L_{EXCESS}, and the reflection coefficient Γ.

If the parasitic capacitance is large enough to produce a negative value of L_{EXCESS} (meaning that the effective impedance Z_2 within the mismatched region is less than Z_{DIFF}), then your analysis must begin with the known value of L_2 and then compute the *excess capacitance* above and beyond that necessary to balance out the inductance L_2. It so happens that the reflection coefficient for the case of $Z_2 < Z_{DIFF}$ is also given by equation [6.23]. The negative amplitude associated with the result in that case indicates that a positive-going step edge produces a negative reflected pulse.

POINT TO REMEMBER

- If you separate elements of a tightly-coupled pair the differential impedance reverts to twice the uncoupled value of Z_0.

6.10.4 Broadside-Coupled Stripline

Figure 6.18 plots the impedance of broadside-coupled differential traces versus spacing and trace width. The data for this plot were generated using a method-of-moments magnetic field simulator with 120 segments equally spaced around each pcb trace, with the current linearly interpolated across each segment.

As you can see in the figure, widening the traces always decreases the differential impedance. Making the traces skinnier always raises the impedance.

Changing the height invokes a more complicated behavior. The impedance is maximized by a trace height (to the centerline of the trace) of 25% of the interplane separation (6 mils, as drawn in the figure).

From the 25% maximal point, a reduction in *h* moves the traces nearer to the planes, increasing the trace capacitance to the planes. This maneuver reduces the characteristic impedance. The more tightly you press the traces against the planes, the less direct coupling exists between the traces themselves.

Going in the other direction from the 25% maximal point, an increase in *h* moves the traces nearer to each other (*s* gets smaller). This increases the direct capacitance between the two traces. This maneuver also reduces the characteristic impedance. The closer you bring the traces towards each other, the greater the coupling between traces.

At the 25% maximum-impedance-point the trace impedance is least sensitive to changes in width or height. Also, given a fixed interplane spacing, the 25%-point also maximizes the trace width, thus minimizing the skin-effect losses.

The only *disadvantage* to the 25%-point is that it maximizes crosstalk between adjacent broadside-coupled pairs; however, the advantages to be gained in trace pitch don't seem to this author to justify the penalties associated with using any trace height much different from 25%.

When converting from an edge-coupled pair on the surface layer of a pcb to a broadside-coupled pair on the inner layers, there is a subtle asymmetry built into the conversion. The asymmetry is illustrated in Figure 6.19.

In the figure, signal **A** starts on layer 1 and then proceeds through blind via **Y** to layer 4. The signal current has no difficulty changing layers at this position; it just passes through **Y**. The return current associated with signal **A** (dotted line) has a more difficult time managing the layer transition. At the left of the figure, the return current for **A** is shown flowing on the top surface of solid plane layer 2. In the center of the figure, most of the return current for this trace flows on solid plane layer 5 (returning signal current always

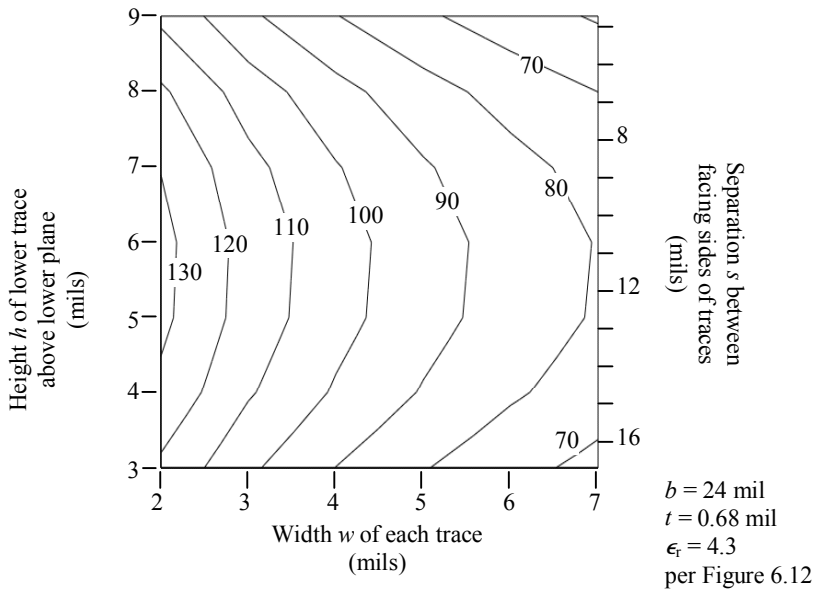

Figure 6.18—Differential impedance Z_{DIFF} (ohms) of broadside-coupled differential striplines.

flows mostly on the nearest solid plane). In the vicinity of the blind via, the return current must find some path (other than the capacitance of the planes, which is not terribly significant) through which it can hop from layer 2 to layer 5. Wherever a signal changes reference planes, the return current must always find a way to follow along.

In the figure, because planes 2 and 5 happen to carry the same potential, they are connected with via **X**, which forms the path for returning signal current. If the planes carry different voltages, the return-current path must traverse a bypass capacitor. In either case the returning signal current temporarily diverts away from the signal current.

The return current associated with signal **B** displays no such difficulty. At the position of blind via **Z** the returning signal current (dotted line) merely needs to change from the top surface of solid plane layer 2 to the bottom surface of that same layer.[50] As you can see, there is a clearance hole located at blind via position **Z** on layer 2. The returning signal current easily pops through this hole from one side of the plane to the other. There is no significant diversion of signal and return current at this location.

The effect of the dissimilarity in return paths is that signal **A** experiences an extra delay as the return current finds its way from plane to plane. To minimize this additional delay, make sure you put a number of plane-to-plane connections near any points where the signals dive into an over/under configuration, where they change layers within the board, and again where they emerge.

Blind vias are used in Figure 6.19 for illustrative purposes; the same effect applies to through-hole vias.

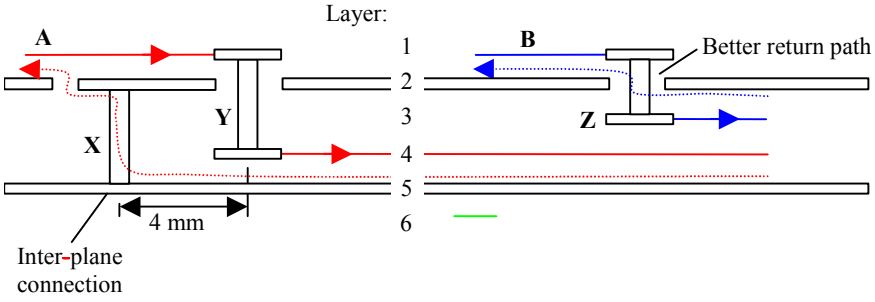

Figure 6.19—The diversion through X in the path of returning signal current for signal A creates more delay than for trace B (cutaway view).

Example showing asymmetry in broadside configuration

The broadside-coupled layout illustrated in Figure 6.19 includes two traces, **A** and **B**, that begin on the surface layer of the pcb and then pass through layer 2 into a broadside-coupled stripline configuration on layers 3 and 4. The plane-stitching via **X** is located 4 mm away from signal via **Y**. The return current for trace **A** therefore has to divert 4 mm out of its way to find the plane-stitching via, plus another 4 mm to get back, giving a total estimated additional delay of (4 mm)·2·(4 ps/mm) = 56 ps. If this same asymmetry exists at both ends of the broadside configuration, the total additional delay on trace **A** equals 112 ps.

[50] The reference planes in a pcb are many skin depths thick; high-speed currents do not penetrate the planes, but rather flow only on one surface or the other of the plane.

This crude estimate doesn't perfectly model what happens, as both the capacitance and the inductance of the via configuration are involved, but if an intrapair skew number anywhere near 100 ps matters to you, then either place the plane-stitching vias closer to the signal vias or don't use a broadside configuration.

In addition to possible asymmetry caused by imbalances in the return paths, the broadside configuration falls prey to any differences between the AC voltages on the two planes. Because the top trace is coupled more heavily to the top plane, and the bottom trace to the bottom plane, any differences in the voltages on these two planes induce a differential signal on the two traces. When using the broadside configuration, it pays to use the same power-supply voltage on both planes and nail them together with numerous vias on a tight grid. I like to use whatever plane voltage delivers the best common-mode noise rejection at the receiver (usually the ground plane).

If you use different power-supply voltages on the two planes, all the power-supply noise between them couples directly into the differential broadside-coupled configuration. In the side-by-side configuration, both traces naturally couple equally to the same nearby plane, so differential pickup of power-supply noise doesn't happen.

Broadside-coupled traces suffer from the dielectric layer-thickness tolerances on the layers separating the traces from their respective solid planes. For example, in a design with 5-mil separation from each trace to its respective plane, a layer-thickness tolerance of +/– 1 mil might result in one trace being 4 mils and the other 6 mils away from the reference planes. This arrangement impairs your ability to achieve good symmetry between the traces, which is after all the whole purpose of using a differential configuration. Edge-coupled traces, because they are etched at the same time under the same conditions on the same layer with the same layer thickness, are generally more symmetric.

The impedance of a broadside-coupled trace is affected by the mechanical registration tolerance of the two signal layers (3 and 4 in Figure 6.19).

The one possible area where broadside-coupled traces have an advantage over edge-coupled traces is routing density. For example, if you need to interleave a large bus through a succession of connector pin fields on a large backplane, the broadside configuration requires only single-track routing between pins, whereas an edge-coupled configuration might require double-track routing between pins to achieve the same layout density. The broadside configuration is also somewhat easier to lay out by hand, because both traces go everywhere together (except at the launch and recovery sites). I avoid broadside-coupled traces unless they are made necessary by routing considerations.

Figure 6.20 illustrates the magnetic lines of force surrounding a broadside-coupled differential pair. It shows intense concentrations of magnetic flux near the corners of the traces, indicating a substantial peaking of the current density at the corners. The proximity factor k_p for these traces takes into account the crowding of current at the corners of the traces, plus any current concentrations on flat surfaces of the pair, plus an allotment for the current induced on the top and bottom reference planes of the stripline cavity. The proximity factor for the traces illustrated in Figure 6.20 is 2.73.[51] Differential broadside-coupled pcb traces with impedances between 75 to 135 ohms possess a proximity factor typically within the range of 2.5 to 3.5.

[51] The proximity factor is the ratio of the actual AC resistance to the resistance one would compute assuming current distributed uniformly around the periphery of one signal conductor only, and taking into account the skin depth.

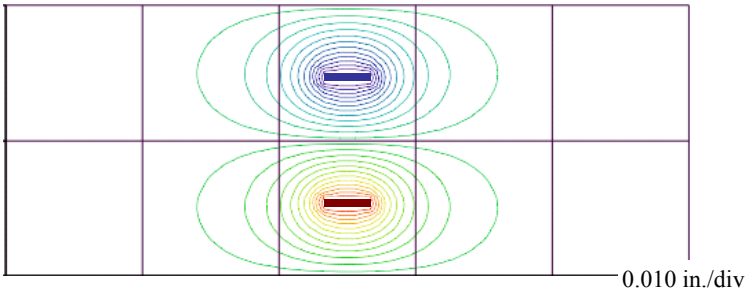

Figure 6.20—This cross-sectional view of the magnetic field in the vicinity of a broadside-coupled stripline shows intense field concentrations near the corners of the traces.

Table 6.3 presents the skin-effect resistance and skin-effect attenuation for selected broadside-coupled stripline configurations.

If you need results for some value b_1 that does not appear in the table, you can linearly interpolate the values in the table to accommodate an intermediate value of b, using the b-column as the x-axis and any other column as the y-axis in your interpolation.

POINTS TO REMEMBER

- Broadside differential trace impedance is maximized by a trace height equal to 25% of the interplane separation.
- The bottom trace of a broadside-coupled differential pair has some extra delay built in at the endpoints.
- Avoid broadside-coupled traces unless they are made necessary by routing considerations.

Table 6.3—AC Resistance and Skin-Effect Loss (at 1 GHz) for Selected 100-Ω Differential Broadside-Coupled Striplines

b mil	*h* mil	*w* mil	R_{AC} Ω/in. @1GHz	R_{AC} Ω/m @1GHz	α_r dB/in. @1GHz	α_r dB/m @1GHz
14	4	1.9	4.04	159.1	0.175	6.89
20	5	3.5	2.71	106.5	0.117	4.61
30	7	5.3	2.01	78.9	0.087	3.43
45	10	9.1	1.32	52.0	0.057	2.24

NOTE (1)—All values *b, h,* and *w* in mils.
NOTE (2)—AC parameters R_{AC}, and α_r specified at 1 GHz with FR-4 dielectric $\epsilon_r = 4.3$.
NOTE (3)—These stripline examples assume copper traces of 1/2-oz thickness with $\sigma = 5.98 \cdot 10^7$ S/m.

6.11 PCB APPLICATIONS

The following sections describe the main applications for differential signaling on pcbs.

6.11.1 Matching to an External, Balanced Differential Transmission Medium

Differential traces are often used to connect to balanced cabling. For this purpose, the tightness of coupling between the two traces making up the differential pair is irrelevant. What matters is that the differential characteristic impedance of trace configuration matches the differential characteristic impedance of the balanced cabling.

The most popular types of balanced cabling are 100-Ω twisted-pair cabling (ISO 11801 categories 3, 5, 5e, 6, and 7), and the old 150-Ω shielded twisted-pair cabling (IBM Type I). When connecting directly to these cabling types, one normally uses two 50-ohm traces (or two 75-ohm traces for 150-Ω cabling) to couple into the cable.

Figure 6.21 depicts a typical LAN coupling situation. The target cable is a 100-Ω unshielded twisted-pair cable. Your objective in this application is to generate a purely differential transmitted signal of standard size with as little high-frequency power and as low a common-mode content as practicable.

The low-pass filter formed by L1-C1 (and L2-C2) truncates any unnecessary power in the frequency range above the bandwidth of the digital signal. The natural balance of the transformer combined with the additional balancing properties of the common-mode choke together serve to limit the common-mode content of the transmitted waveform. The reason common-mode balance is so important is that common-mode radiation from an unshielded twisted-pair cable is many orders of magnitude more efficient than differential-mode radiation. Minimizing the common-mode current minimizes the cable emissions.

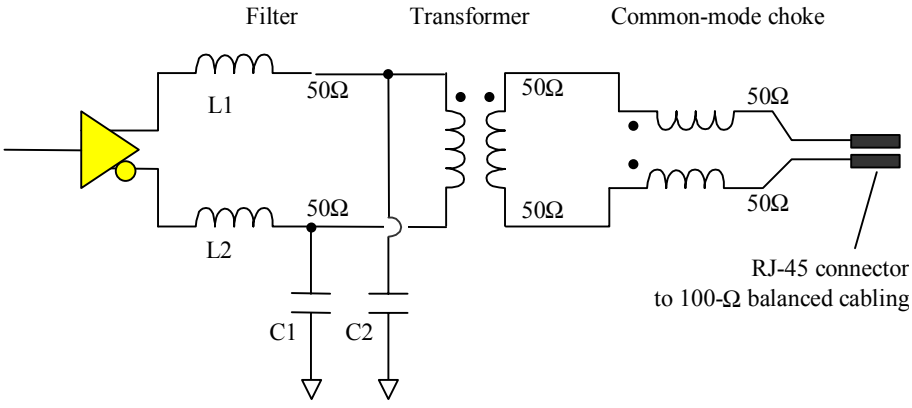

Figure 6.21—Ethernet 10/100BASE-T interface, showing use of 50-Ω transmission lines to match 100-Ω balanced load.

After passing through the common-mode choke you should make the two pcb traces as symmetrical as possible, with equal impedances to ground. The traces must be symmetrically positioned with respect to all nearby grounded objects, but they do not necessarily need to be tightly coupled to each other.

POINTS TO REMEMBER

- Match the differential characteristic impedance of two pcb traces to the differential characteristic impedance of a balanced cable.
- Make the two pcb traces as symmetrical as possible, with equal impedances to ground.

6.11.2 Defeating Ground Bounce

Differential signals arrive naturally at a receiver with a built-in reference voltage. The receiver of a differential signal need not rely on its own internal reference, which could be corrupted by ground bounce or other disturbances in the reference supply. Differential signaling defeats ground bounce.

For ground-bounce cancellation to work, the receiver must see two complementary signals with equal delays from the driver. Any ground shifts or disturbances along the way that affect both wires equally will be cancelled at the receiver.

Note that the two halves of a differential signal must arrive synchronously so that they will be equally influenced by noise, but they do not necessarily need to be tightly coupled to each other.

POINT TO REMEMBER

- Differential signaling defeats ground bounce.

6.11.3 Reducing EMI with Differential Signaling

Differential signals radiate less than single-ended signals. That's one of the benefits of differential logic. If the two complementary signals of a differential pair are perfectly balanced, the degree of field cancellation is determined entirely by the separation between traces.

If, however, the two complementary signals are not perfectly balanced, then the degree of attainable field cancellation will be limited to a minimum value determined not by the trace spacing, but by the common-mode balance of the differential pair. Because the common-mode balance of most digital drivers is not particularly good, it often happens that differential pairs radiate far more power in the common-mode than in the differential mode. In such a situation no radiation benefit remains to be gained from squeezing the differential traces more closely together.

Figure 6.22 plots the theoretical radiation gain attained by a differential microstrip pair as a function of trace separation. The figure assumes the measurement antenna is located in

the plane of the board, removed in a broadside direction a distance $r = 10$ m away from the traces (worst case). The radiation from one trace of the differential pair is *supposed to be* cancelled by the equal and opposite radiation from the adjacent trace, resulting in a marked reduction in emissions. The depth of cancellation is related to the ratio $2\pi s/\lambda$, where λ is the free-space wavelength of the highest frequency of interest and s is the separation between traces. This ratio controls the relative phase relationship of the two near-complementary waves as they leave your board. The cancellation is also related to the ratio $r/(r + s)$, which speaks to the relative intensities of the two near-complementary waves as they reach the antenna. The formula in Figure 6.22 shows that differential cancellation improves as you reduce s.

The common-mode radiation from the two traces of a differential microstrip reinforces rather than cancels, so that common-mode radiation does not vary strongly with trace separation. You can adjust the differential-mode radiation by adjusting the trace spacing, but you can't do much about the common-mode radiation (except to install a driver with better common-mode balance).

Under FCC class B measurement conditions, the differential-mode radiation from a differential microstrip pair with 0.5-mm (0.020-in.) separation should theoretically yield a 40-dB radiation improvement at 1 GHz, compared to the radiation measured if the same signal were implemented as a single ended layout. Smaller separations should yield even more improvement. While that theory sounds appealing, in practice you will rarely if ever achieve as much as a 40-dB improvement in overall radiation because your gains will be limited by the degree of balance available on the two outputs of your differential transmitter. Unless the outputs are balanced to better than 1 part in 100, a *common-mode radiation component* of at least 1% of the differential amplitude will emanate from your differential pair anyway. Given a 1% common-mode imbalance, even a differential spacing of zero would not improve the total radiation by more than 40 dB.

Taking an example from the LVDS differential driver family, which prescribes a differential balance no better than 1 part in 16, even the most Herculean efforts at trace

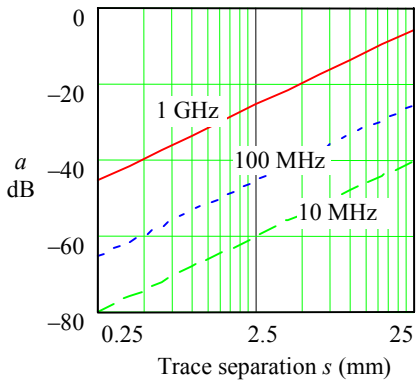

$$a = 20\log\left|1 - \frac{r}{r+s}e^{\frac{-j2\pi s}{\lambda}}\right|$$

where r is the distance to the receiver,
s is the trace separation, and
λ is the free-space wavelength of the highest frequency of operation.

Figure 6.22—Theoretical radiation improvement *a* for the differential portion of the far-field radiation from a microstrip, as a function of trace separation *s*.

balancing will never improve the overall radiation for that logic family by more than a factor of 16 (24 dB).

In plain terms, a differential trace spacing of 0.5 mm is close enough to deliver all the EMI benefit you are likely to ever achieve. Because radiation problems on digital pcbs are usually dominated by the common-mode radiation, you need not struggle to place ordinary differential digital traces any closer than 0.5 mm for any EMI purpose.

POINT TO REMEMBER

- You need not struggle to place ordinary differential digital traces any closer than 0.5 mm (0.020 in.) for any EMI purpose.

6.11.4 Punching Through a Noisy Connector

When two systems are mated by a connector, the net flow of signal current between the systems returns to its source through the ground (or power) pins of the connector. As it does so, tiny voltages are induced across the inductance of the connector's ground (or power) pins. These tiny voltages appear as a difference between the ground (or power) voltage on one side of the connector and the ground (or power) voltage on the other side. This problem is called a *ground shift,* and it is yet another form of common impedance coupling. In a single-ended communications system the ground shift voltages detract directly from the available noise margin for your logic family. In a differential signaling system, the ground-shift voltages don't matter, because they affect both wires of the differential pair equally. Subject to the limits of common-mode rejection, ground shifts generated within a connector are totally cancelled within the receiver.

Differential signaling usually reduces crosstalk generated by either mutual inductance or mutual capacitance within the connector itself. The exact gains available depend on the relative spacing of the differential signals as they pass through the connector and the distance to the nearest aggressive source. If the aggressive source is closer to one element of the pair, the crosstalk will not affect both pairs equally, and it will therefore not be cancelled in the receiver.

The cancellation of ground shifts generated by a connector and the cancellation of nearby aggressors within a connector both have a lot to do with the position of the signal pins within the connector, but very little to do with the pcb trace layout, or intertrace coupling.

POINT TO REMEMBER

- Subject to the limits of common-mode rejection, ground shifts generated within a connector are totally cancelled within a differential receiver.

6.11.4.1 Differential Signaling (Through Connectors)

***High-Speed Digital Design Online Newsletter*, Vol. 3, Issue 12**
Sal Aguinaga writes

I have 16 differential pairs that go through a connector and terminate on a daughter card. What is the best signal-to-ground ratio and pattern I should consider?

In this case the connector is a high-density pin connector. If the differential impedance is 100 Ω, do I need a special ground pattern as the signals go through the connector to maintain the differential impedance close to 100 Ω?

Reply

Thanks for your interest in *High-Speed Digital Design*.

Regarding your correspondence, there is no general formula for the number of grounds required, as it depends on the spacing and sizes of the connector pins, and how they are bent.

Here are a few thoughts for you to consider.

- In a connector with an open field of pins, put the two elements of each differential pair on the same row of the connector. That will ensure that they get the same pin lengths and go through the same pattern of elbow bends.[52]
- On a synchronous bus, if you have enough time to wait for the crosstalk to settle, you may not need to isolate the differential pairs from each other.
- If isolation between pairs is required (for low crosstalk between nonsynchronized bus signals or for a clock or other asynchronous signal), place the pairs so that no signal wire from one pair falls adjacent to any wire from another pair. This implies that you will be using at least as many grounds as signal pins to separate the pairs, and probably more.

The differential impedance of most open-pin-field connectors is probably going to be a little higher than you want. You can measure this. You will need a pair of test boards on which you can mate the connector halves. The boards don't use any traces. They can be solid copper with holes drilled for the connector pins. Ground all the pins that will be grounded (or tied to a power plane) in your application. Use two RG-174 50-Ω coaxial cables to route a differential, 100-Ω signal into the designated signal pin pair. Let this signal go through the connector to the far side. On the far side of the connector, terminate the signal differentially with 100 Ω.

For any signal speed that will work with an open-pin-field connector, you will find that a 1/8-watt axial, 100-Ω resistor works fine as a terminator.

Blast in a differential signal from your 100-Ω source. Make a record of the resulting waveform as measured at the source. This should show the source waveform going out and a first reflection coming back (you've made a crude TDR

[52] Connectors with solid ground shields between columns of pins are usually designed to accommodate differential pairs collocated within each shielded cavity. This implies that the pin lengths may differ for the two elements of each pair. A perfect differential shielded connector would be designed to match the skew on each element of a differential pair as it traverses the connector, even if the elements were located on different rows. If your connector produces a known skew, it's up to you to cancel it somewhere else in the layout.

instrument). Use a step risetime commensurate with what you are going to be using in the real system.

Don't mess around with fancy 35-ps step edges on this type of connector. They will just show you a bunch of fine-structure detail that isn't going to matter in the real system.

Now disconnect the coaxial cables from the connector. Place the 100-Ω termination directly across the coaxial cable outputs, with the coaxial grounds tied together. Repeat the measurement. You should (ideally) see no reflection.

Looking at the difference between the first measurement and the second, if the reflected waveform bumps up in the positive direction (same polarity as the step input), the connector impedance is a little too high. If it bumps negative, the connector impedance is too low. If you don't see a bump then the impedance is just right.

Adding more ground pins around the signal pair lowers the impedance.

Spacing the signal pins further away from ground raises the impedance.

Bonus idea: Adding a little lumped-element capacitance from signal to ground on each side will lower the effective impedance. This may be implemented in the pcb layout by just using larger-than-normal via pads. Experimentation and remeasurement is required to get this idea to work. The "big-pad" concept works when the connector through-delay is less than 1/6 of the signal risetime and the connector is acting like a lumped-element inductor (too high an impedance).

6.11.5 Reducing Clock Skew

When a digital component receives a clock, the precise moment at which the clock is recognized depends upon the *switching threshold* for that component. For 5-V TTL logic, the switching threshold is defined to be somewhere between 0.8V (V_{IL}) and 2.4V (V_{IH}). The spread between V_{IL} and V_{IH} defines a window shown in Figure 6.23 within which the actual clock transition takes place:

$$t_{\text{UNCERTAINTY}} = \frac{V_{IH} - V_{IL}}{\left(dv/dt\right)} \qquad [6.25]$$

where $t_{\text{UNCERTAINTY}}$ is the uncertainty (skew) in the clock switching moment,

V_{IH} and V_{IL} are the worst-case guaranteed high and low logic thresholds respectively, and

dv/dt is the rate of change of voltage on the clock input (roughly equal to the logic swing ΔV divided by the 10% to 90% risetime of the driver).[53]

[53] The risetime emanating from an unreasonably fast driver will be modified by the package parasitics of the receiver. The risetime of the signal measured at the input pads of the receiver die is therefore generally not quite as fast as the risetime measured external to the package. It is the risetime of the internal signal combined with the receiver switching thresholds that determines the actual amount of switching uncertainty. If the driver and receiver are implemented in similar technology, and the package is not deemed a significant impediment to reception, you may simply use the driver risetime for the calculation of *dv/dt* in equation [6.25].

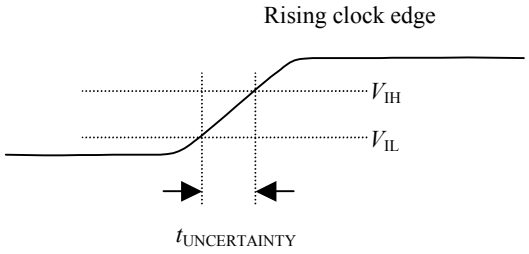

Figure 6.23—The input thresholds and the risetime of the input signal combine to create an uncertainty in the precise switching time for a clock.

The smaller you can make this uncertainty, the less clock skew you will have to recognize in your timing budget. In the differential world a similar effect takes place, but the specifications for V_{IH} and V_{OL} are replaced by a specification for the offset voltage of the receiver. The input offset voltage is the differential input voltage at which a particular receiver actually switches. An ideal receiver would switch precisely at zero (when the inputs are equal), regardless of common-mode voltage, temperature, power supply quality, age, and so on. Such a part would have an input offset voltage of zero. Practical receivers always switch at some finite, hopefully small, value. There is no way to predict the polarity of the offset.[54] The spec sheet for a differential receiver usually provides a worst-case upper bound for the magnitude of the offset. The effective spread between V_{IH} and V_{IL} for a differential receiver is twice the maximum offset magnitude.

A relevant figure of merit for comparing differential logic families in this regard is the ratio of the spread in differential input offset voltage to the peak-to-peak differential output voltage swing.[55] For single-ended logic the corresponding figure of merit is the spread between V_{IH} and V_{IL} divided by the peak-to-peak output voltage swing. Differential logic usually fares better on this measure of performance.

In a differential clock distribution system one usually attempts to match the delays of the two complementary signal traces as they traverse a pcb. The two traces need not follow the same path; they just need to have the same delay. If the delays of the two traces are unequal, if affects the switching time. For example, suppose the two complementary halves of a differential signal arrive at successive times t_1 and t_2. Let the separation between t_1 and t_2 be a small fraction (perhaps 1/10) of a rising edge. Under these conditions the receiver will switch very nearly at the average arrival time $(t_1 + t_2)/2$. In the worst case, if the separation is as great as a risetime, the receiver will switch no earlier than t_1 and no later than t_2. I normally match the delay of the two traces in a differential signal to within 1/20 of the signal risetime.

The clock skew contributed by a clock receiver is a function of the input risetime, the switching thresholds, and, for differential signaling, the degree of similarity in the times of arrival of the two complementary signals. Clock skew has little or nothing to do with trace spacing or geometry (other than delay).

---

[54] Unless the device is manufactured with a purposeful offset in one direction or another.

[55] In a differential signaling architecture, the peak-to-peak differential output voltage swing is twice as large as the peak-to-peak voltage swing on either of the two inputs.

POINTS TO REMEMBER

- Differential receivers often have more accurately specified switching thresholds than single-ended receivers.
- Uncoupled differential traces need not follow the same path; they just need to have the same delay.

6.11.6 Reducing Local Crosstalk

Differential traces on a pcb do a relatively poor job of reducing local crosstalk. As illustrated in Figure 6.24, when some local aggressive trace approaches a differential pair, the interference is not balanced. Interference couples much more strongly to the near side of a differential pair than to the far side.

In the figure, clock+ is twice as close to the aggressor as clock–, so you get a 4:1 difference in the crosstalk coupled into the two sides. Imbalanced crosstalk of this sort cannot be cancelled by a differential receiver. The receiver sees almost the full value of crosstalk from the aggressive trace to the clock+, with little or no cancellation from clock–.

The best way to prevent crosstalk onto a differential pair is to design a keep-out zone around the sensitive traces, forcing other traces to stay at a respectable distance. All modern layout systems support separation rules by net class that will allow you to keep big, dirty signal traces away from delicate, sensitive ones.

Cramming the traces of a differential pair closer together does yield marginal improvements in crosstalk reduction, but you don't get the *big* benefit you get from simply moving the whole pair further away from the problem. Figure 6.25 provides data to support this assertion. The data for this plot were generated using a sheet-conductance measurement method [56].

The figure plots crosstalk for three trace configurations. Each plot shows the measured near-end crosstalk coefficient in dB versus the separation x between traces.[56]

The first scenario shows how a single-ended aggressor affects a differential pair. The next two scenarios show how two differential pairs affect each other. The difference between scenarios II and III is that the intrapair separation changes from 8 to 4 mils. In all cases the separation between planes is 24 mil, and the height of all traces above the nearest plane is 6 mil. The widths of all traces have been adjusted to give a characteristic impedance

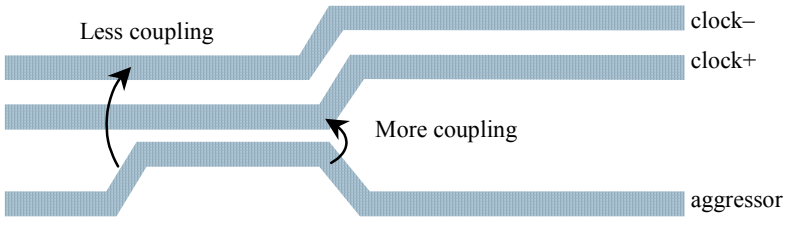

- The aggressor couples more strongly to clock+ than to clock–
- It is a good practice to enforce a bigger-than-normal spacing between clock signals and data signals

Figure 6.24—Differential signaling does *not* offer much help with nearby crosstalk.

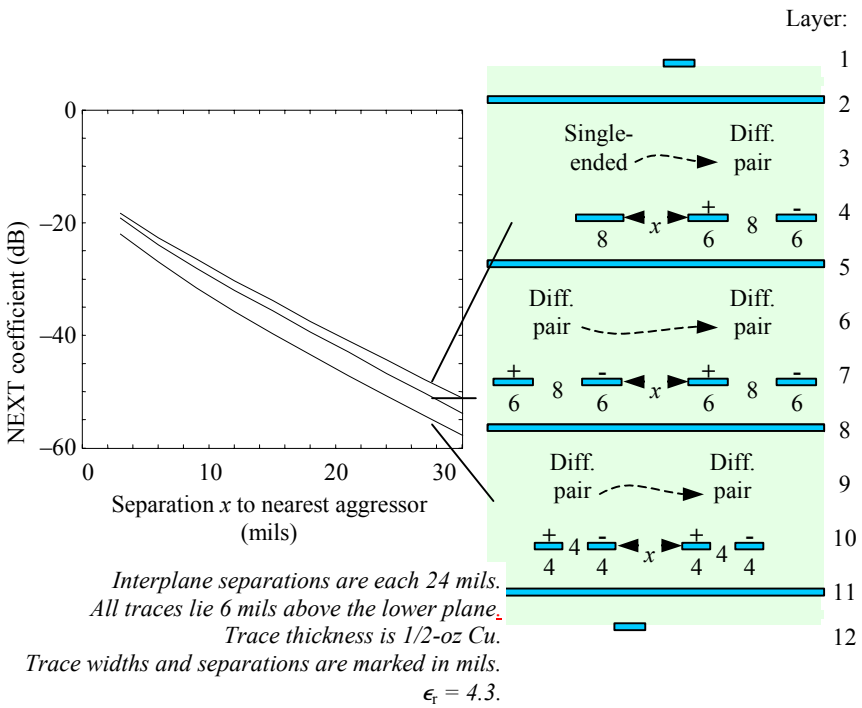

Figure 6.25—The crosstalk immunity of differential pairs is not much improved by tight coupling; mostly what matters is the interpair spacing.

of 100 ohms for all differential transmission-line pairs and 50 ohms for the single-ended aggressor.

In all cases the crosstalk falls off steeply as you increase the separation *x* between the aggressor and victim. Looking at the difference between scenarios I and II, the addition of a complementary signal to the aggressor (changing it into a differential pair) tends to cancel a little bit of the interference, but the complementary signal is too far away to have much of an effect. The net reduction in interference is less than 2 dB. When you go from scenario II to III, reducing the intrapair separation, the complementary signals is brought into play. This reduces crosstalk by an additional 4 dB, which is useful, but still nowhere near as significant as simply moving the aggressor further away.

Tightly coupling a differential pair delivers only a modest improvement in crosstalk, and therefore only a modest improvement in the achievable pair pitch (see also Section 6.10.2, "Edge-Coupled Stripline").

POINT TO REMEMBER

- Tightly coupling a differential pair delivers only a modest improvement in crosstalk.

6.11.7 A Good Reference about Transmission Lines

The *Transmission Line Design Handbook* by Brian C. Wadell [58] compiles handy approximations for transmission-line impedance, delay, skin effect loss, dielectric losses, and radiation losses. It is very comprehensive. The book addresses most of the popular transmission-line formats in use today, including microstrip, buried microstrip, offset stripline, and edge- or broadside-coupled differential striplines.

Wadell heavily references the original research articles and measurements. He doesn't pull any punches when the research is fuzzy or contradictory. He shows you just what is known and indicates what is not known. If you're looking for closed-form approximations, this is the best source.

P.S.: If you have an old copy of Wadell's text, check out his errata list on the Web.

6.11.8 Differential Clocks

***High-Speed Digital Design Online Newsletter*, Vol. 1, Issue 10**

Fabrizio Zanella writes to the SI-List

I understand the benefits of using differential pairs for signals running at 100MHz and above. Can you speak about the impact of using differential clocks in a parallel bus? Do the differential clocks maintain the noise suppression characteristics when daisy-chained in a multidrop environment? Has anyone tried this and had positive experiences versus single-ended multidrop clocks?

Reply

Thanks for your interest in *High-Speed Digital Design*.

In general, I have found differential distribution to be a very effective means of combating ground bounce in the transmitting package, ground bounce in the receiving package, as well as the ground shifts that occur on either side of the connectors in high-speed systems. These benefits accrue in multidrop configurations as well as point-to-point configurations.

I have found differential distribution to be of little value in reducing the impact of crosstalk generated locally by other traces on the same pcb. This is because the crosstalk function from nearby traces falls off very steeply with distance. The impact of this is that differential pairs cannot be placed particularly close to any aggressive signal. For example, imagine a system that has one aggressive trace and a nearby victim trace that is receiving unacceptable amounts of crosstalk. Now I propose to protect the victim trace by splitting it into a differential pair and using a true differential receiver. Assuming that I don't want to affect the layout density, I plan to implement the centerline of the new pair right on top of the original victim trace. In other words, when we split the victim, one member of the pair will have to move closer to the offending source, while the other moves away. Unless the two traces of the pair are extremely close together (less than a third of the original separation between centerlines), the extra crosstalk we pick up from the nearby side of the differential pair overwhelms any "balancing" effect we might have hoped to gain from the far side of the differential system. To mitigate this effect, you have to

separate the victim pair from the aggressive signal. In the final analysis, it's usually the extra separation that is providing most of the crosstalk benefit, not the fact of balanced signal distribution. When you want to battle crosstalk picked up on a pcb (over a solid ground plane), increasing the trace separation will probably result in a more dense design than using differential distribution.

For clocks, I see differential signaling used a lot, both in point-to-point distribution and in multidrop distribution. The multidrop aspect does not diminish the ground-bounce-canceling properties of differential signaling. There are only a few clock nets in a system (compared to the number of data nets) and it isn't that difficult to provide this extra measure of protection.

For parallel bus signals, I rarely see differential distribution used because it doubles the number of wires required. That will cause the phenomenon known as "routing headaches" among your pcb layout staff.

POINT TO REMEMBER

- The benefits of differential signaling apply to multidrop configurations.

6.11.9 Differential Termination

Article first published in *EDN Magazine,* June 8, 2000

I am designing a piece of equipment to interface to a digital tape recorder designed by another company. This recorder uses a differential ECL interface, and the user's manual recommends terminating the clock lines slightly differently from the data lines. Each clock line employs a split terminator (160-Ω to –5.2V and 100-Ω to ground), but the data signals simply use a single 120-Ω resistor between the wires of each pair. Because these two methods are Thevenin equivalents, why does the user's manual recommend different termination schemes? As far as I can tell, the transmitters and receivers for the clock and data lines are electrically equivalent, and all signals have 390-Ω pull-down resistors to –5.2V at the transmitter to properly bias their emitter-follower outputs. I have contacted the company that designed the circuit, but the original designers are unavailable. Does one termination scheme have any advantage over the other?
—Raymond Bullington

Engineers often have a difficult time figuring out why something was done. Sometimes there is no reason, sometimes there is a multitude of good reasons, and sometimes (as is common in the standards world) everyone wants it done the same way but all for conflicting and different reasons.

Anyway, differences do exist between the termination schemes you described. The single-resistor scheme (120 Ω across the two lines) terminates all differential-mode signals into 120 Ω but provides no termination for common-mode signals.

Your four-resistor scheme (independent terminations for each line) terminates all differential signals *and* all common-mode signals. The difference between these

two styles matters only if a common-mode signal is present. And where might a common-mode signal come from? It can come from any skew naturally present in the clock driver output plus any imbalances in the circuit that convert part of the output from the differential mode to the common mode.

Every long, differential link needs a differential termination for signal quality and *also* a common-mode termination to prevent common-mode resonance.

Consider the single-resistor termination shown in Figure 6.26. Say that a positive-going edge $x(t)$ arrives first on trace A, and then, after a tiny skew interval Δt, the opposite signal $-x(t - \Delta t)$ arrives on trace B. During the tiny skew interval Δt, the single 120-Ω resistor R1 creates two tiny artifacts. First, the initial rising edge on line A shoots right through the resistor onto trace B, creating a little blob of crosstalk. The amplitude of the crosstalk compared to the amplitude of the incoming signal is 1/2. Second, coincident with the crosstalk, you get a small signal reflected back onto line A. The amplitude of the reflected signal compared to the amplitude of the incoming signal is also 1/2. Both artifacts have positive polarity, creating what amounts to a common-mode reflection.

After time Δt, the opposite signal $-x(t - \Delta t)$ arrives on line B. At this time you get a second set of crosstalk and reflection artifacts, but with negative polarities this time (because they originated on the negative half of the differential pair). The second set of artifacts partially cancels the first, with the degree of cancellation depending on the exact temporal alignment of the two signals. The two sets of artifacts perfectly cancel only when the signals on A and B arrive in perfect synchronism.

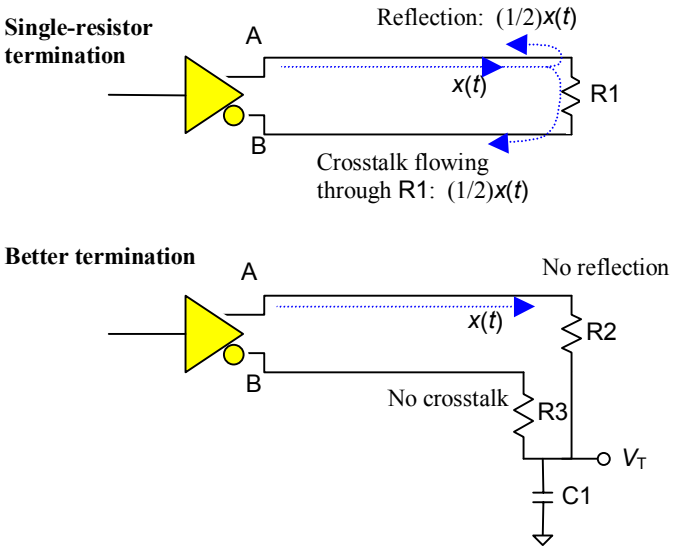

Figure 6.26—A network of two resistors terminates each half of the differential pair independently.

In this example, both the crosstalk and reflection amplitude coefficients equal 1/2. You may express the residual common-mode signal $\gamma(t)$, induced on either trace by the single-resistor end-terminator, as $\gamma(t) = (1/2)(x(t) - x(t - \Delta t))$. If the skew is less than the signal rise or fall time t_{10-90}, and you define the signal step height as ΔV, the peak amplitude of the reflected common-mode noise roughly equals $(1/2)\Delta V(\Delta t/t_{10-90})$.

Once created, the common-mode noise returns to the driver. In your case, the ECL driver presents a very low output impedance to the line, generating a big reflection. The reflected noise then proceeds to the receiver, where it once again encounters the single-resistor termination, but this time as a purely common-mode signal. Because a common-mode signal presents the same voltage on both traces, the single-resistor terminator draws zero current and acts as an open circuit. The open circuit generates another big reflection. After that point, the common-mode noise trapped on the line happily bounces back and forth between the driver (low impedance) and the receiver (open circuit) for a long time.

Terrible things happen to the common-mode noise if your trace delay equals one-quarter of the clock period. In that case, the little common-mode artifacts from each edge build and superimpose, cycle after cycle, magnifying the common-mode noise at the receiver and also magnifying the common-mode radiated emissions. This problem is called a *common-mode resonance*.

To avoid common-mode resonance, every long, differential link needs two terminations: first, a good differential termination at one end or the other to provide good differential signal quality, and second, a reasonable common-mode termination at one end or the other to prevent severe common-mode resonance. An ECL driver does not provide a good common-mode termination at the source; therefore, one is required at the load.

The four-resistor termination that was recommended to you for the end of the clock net independently terminates both lines, damping both differential and common-mode signals at that point.

An even better termination circuit appears in Figure 6.26. This circuit terminates both differential and common-mode signals but requires only two resistors (R2 and R3 are each set to half the differential-line impedance). The capacitor need be only large enough to hold its charge steady during the brief interval of skew Δt. In your case the ECL sources incorporate pull-down resistors, so you don't need to supply a special terminating voltage (V_T) to the capacitor.

Any driver that provides a reasonable common-mode termination at the source relieves you of the responsibility of providing one at the destination. For example, a source-terminated driver works fine with the single-resistor termination.

POINT TO REMEMBER

- Every long, differential link needs at least one good differential termination *and also* a reasonable common-mode termination to prevent severe common-mode resonance.

6.11.10 Differential U-Turn

Article first published in *EDN Magazine,* September 1, 2000

What is the effect of a split in a solid plane on the impedance of a coplanar differential pair? The differential pair passes over a solid plane (logic return) and then crosses a 50-mil void into an I/O area that has a solid plane of its own that is tied to the chassis.
—Boris Shusterman

Significant currents flow on the solid reference plane beneath your differential traces. Cutting the plane interrupts these currents.

Consider first a single-ended pcb trace. When a changing voltage propagates down a *single-ended* transmission structure, currents flow through the distributed capacitance of the trace to all nearby objects, especially the big, solid plane underneath the trace. This capacitive effect generates a returning (reverse) flow of current on the solid reference plane. The return current flows all the way back to the source along the reference plane, staying underneath the signal trace the whole way, making a complete circuit. (Current always makes a loop.)

Counteract the U-turn by shrinking both the reference-plane gap and the spacing between traces.

Now, consider a differential-pcb-trace pair. Differential structures have capacitance from each trace individually to the reference plane and also *between* the traces. The between-trace, or mutual, capacitance of the differential pair induces returning current on the other trace as well as on the solid reference plane. In a differential-pcb pair, most of the returning current from each trace still flows on the solid plane, not the other trace, because each differential trace couples much more strongly to the big, solid nearby plane than it does to its little, skinny differential buddy.

Try to visualize the propagation of a differential signal as a quad of four currents: two currents, i+ and i–, on the two signal traces, and the returning currents, r+ and r–, on the reference plane underneath the traces (Figure 6.27). In most cases the currents are almost as big as i+ and i–.

When a differential signal encounters the gap between reference planes, the two signal currents continue across the gap on the signal wires, but the gap blocks the two return currents r+ and r–. This situation forces the return currents to execute a U-turn maneuver, whereby each return current U-turns into the other position. On the new reference plane, a similar effect takes place with a second U-turn formation creating a pair of currents r′+ and r′– to complete the quad current formation.

In the space between the reference planes, current circulates clockwise on all sides of the U-turn zone. This current behaves like a small current-loop antenna, generating a substantial, fast-changing magnetic field within the U-turn zone.

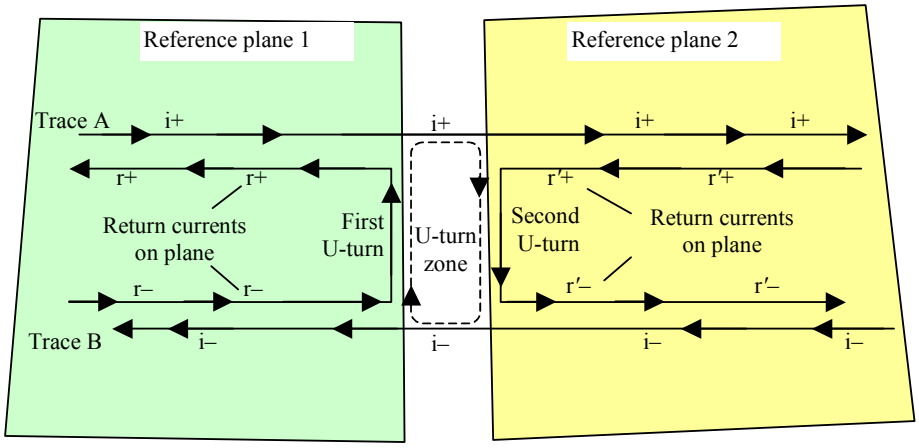

Figure 6.27—The U-turn zone at a reference-plane gap separates the return currents from the primary signal currents on traces A and B.

The magnetic field makes the circuit behave as if an inductor were in series with the signal path. The length and width of the U-turn zone determine the inductance. For example, a trace-to-trace separation of 0.100 in. and a plane-to-plane separation of 0.100 in. would generate an inductance on the order of 10 nH, which in a 100-Ω differential system would introduce a low-pass filter response with a time constant of 100 psec. If risetime exceeds 1 nsec, you probably won't even notice the effect. On the other hand, at very high speeds, a 100-psec risetime could mess up your signals.

The U-turn zone is more than merely a region of increased impedance due to the absence of the reference planes. It is an effect whose physical dimensions span both the gap between planes and the spacing between traces. You counteract the U-turn by shrinking both the reference-plane gap and also the spacing between traces. You can also practically eliminate it by providing continuous pathways adjacent to each signal trace for the conveyance of return currents from plane to plane, eliminating the need for a U-turn zone. If the planes carry different DC voltages, a bypass capacitor next to each trace isn't perfect, but it helps.

The magnetic fields within the U-turn zone induce crosstalk and EMI. The crosstalk couples to all the differential pairs that pass over the same gap. Both EMI and crosstalk vary in proportion to the size of the U-turn zone.

POINT TO REMEMBER

- Visualize the propagation of a differential signal as a quad of four currents.

6.11.11 Your Layout Is Skewed

Article first published in *EDN Magazine*, April 18, 2002

Passing through a sharp turn with an edge-coupled differential pair, the outside trace travels further than the inside trace. The difference in distance traveled contributes a small amount of skew to your differential signals. The skew acts as a mode converter, changing part of your differential signal power into common-mode power.

The pair-turning skew becomes noticeable only when the skew contributed by the turn rises to a level comparable with the natural skew already coming out of your driver. Therefore, before worrying about turns, first determine the skew of your driver. In many cases the driver skew is not specified, in which case you can assume the skew will be *at least* 10% of the signal risetime. Digital differential drivers just aren't balanced very well. Analog transceivers often are, which accounts for the importance of meticulous skew-matching in some analog applications.

Chamfering or rounding of differential corners does not eliminate skew.

For example, let's say an Ethernet 100BASE-TX LAN transceiver with a well-balanced output transformer and common-mode choke puts out differential signals balanced to 1 part in 1000—meaning that the common-mode output is 1000 times smaller than the differential signal. To avoid amplifying the common-mode signal on the wires (and thereby the radiation), the aggregate skew contributed by all components used with this transceiver must remain less than 1/1000 of 1 risetime. The risetime of a 100BASE-TX signal is approximately 8 nS, corresponding to roughly 94 inches of propagation in air, 1/1000 of which works out to 0.094 in., so the skew budget within cable connectors and board layout should be set somewhere around 0.1 inch. Worrying about little bitty skew effects much smaller than 0.1 inch doesn't buy you anything.

To take a faster example, a 2.5 Gb/s serial link driver with a risetime of 200 ps has an output skew of probably no better than 20 ps (maybe a lot worse). In this case a skew budget of perhaps 20 ps seems reasonable.

Figure 6.28 illustrates the skew calculations for three alternative corner treatments, each with a trace pitch (centerline to centerline) of p. In each case, the two traces within the pair share the same number and type of sharp corners (diagonal-striped regions). The differences between the inside and outside traces are shaded dark with white lettering. The lengths added to the outside trace are $2p$, $1.65p$, and $1.57p$ respectively for the three corner styles. Apparently, chamfering or rounding of differential corners does not eliminate skew; it only makes at best a modest improvement.

If your trace separation p equals 20 mils, and the propagation delay of your media is 160 ps/in., the skew associated with a distance of p is (0.020 in. x 160 ps/in.) = 3.2 ps. According to Figure 6.28, the skew accumulated when rounding a corner of any of the three types shown ranges in that case from a

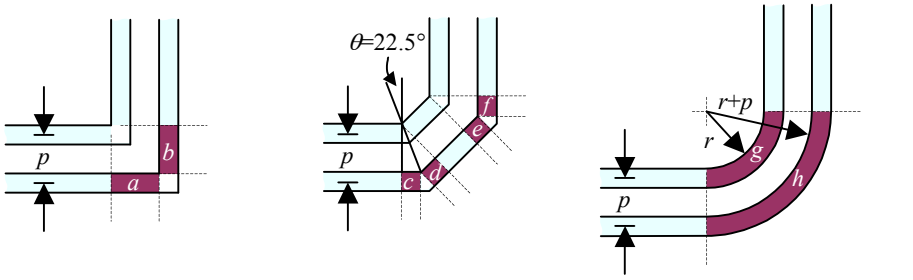

Figure 6.28—These three corner treatments all generate similar amounts of skew.

high of $2p \rightarrow 6.4$ ps to a low of $1.57p \rightarrow 5.0$ ps. If this amount of skew is superceded by the skew from your driver, then don't worry about the turns.

When skew becomes a problem, you can mitigate its impact in two ways. First, use a smaller spacing. The smaller you make *p,* the less skew you will get. This is one of the few benefits of tightly coupled pairs. Second, position your ICs so the traces leave the driver headed in the same direction that they enter the receiver. For example, a differential pair that starts out headed north and ends up headed north has by definition equal numbers of right- and left-hand turns no matter what happens in the middle (unless it makes a spiral), so the net skew accumulated is zero.

POINT TO REMEMBER

- Chamfering or rounding of differential corners does not eliminate skew.

6.11.12 Buying Time

Article first published in *EDN Magazine*, May 2, 2002

The previous article "Your Layout Is Skewed" concerned various styles of corners and bends normally used on differential edge-coupled pairs. It pointed out that all corners, whether chamfered or not, add extra length to the outside trace as it rounds the bend. The equivalent trace length added by 90 worth of bending ranges from one and one half to two times the intrapair trace pitch depending on how the corner is chamfered. The extra time added to the outside trace is a form of intrapair skew.

This article considers two strategies for minimizing the intrapair skew accumulated by a differential net. The six BGA chips within the dotted-line region in Figure 6.29 illustrate the first strategy.

Pair **A** exits the bottom chip heading north. It enters the receiver (top chip) also heading north. Along the way, this pair takes one right turn and one left turn. *The skew accumulated in the two successive turns cancels to zero*.

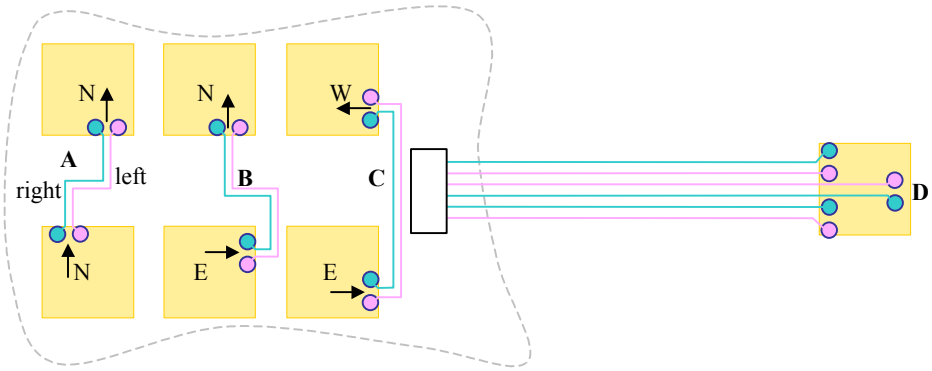

Figure 6.29—The positioning of entrances and exits affects the intra-pair skew.

In a general routing problem, the number and types of turns required depends on the relative orientations of the driver and receiver. Because pair **A** starts and ends going in the same direction, this pair will always make equal numbers of right-hand and left-hand turns no matter what happens in the middle (unless it makes a spiral). The net skew accumulated on any pair with a chip floor plan like **A** is zero.

Pair **B** doesn't fare as well. It exits the bottom chip headed east. It takes one left turn to get going north (the orientation of the receiver), after which the number of left and right turns remain balanced. The total skew accumulated by pair **B** equals the amount generated by one left-hand turn.

Pair **C** is the worst of all. It exits to the east and enters to the west. It therefore requires two extra left turns to achieve the correct orientation. If you are going fast enough so that every turn worth of skew matters, you should carefully plan your chip orientations so the accumulated skew naturally balances to zero.

A pair that starts and ends going north has by definition equal numbers of right-hand and left-hand turns.

The second strategy concerns the precise manner in which your pair enters or exits a ball grid array (or any field of connector pins, package pins, or vias). This strategy works best when the ball pitch exceeds the intrapair trace pitch. It works by offsetting the centerline of the pair as it enters (or exits) the BGA, as shown in the figure at **D**. By offsetting the top pair down half a position, extra time delay accrues to the topmost trace. Offsetting the bottom pair up half a position adds time to the bottommost trace. This strategy "buys some time," which you can use to pay for other floor planning inadequacies. It isn't perfect, but it delivers what you need—balanced skew.

When you have to adjust the skew, I favor doing it near either the driver or the receiver, whichever has the poorest termination. That way the skew adjustment can't possibly affect the quality of the good termination at the other end of the line. If both ends have high-quality terminations, then you place the adjustment at either end. In an imperfect world, that's as well as you can do.

To those who yearn for a perfect layout with zero bends I say, with all due credit to the Rolling Stones, "You can't always get what you waaaannnnnnnt… You

can't always get what you waaaannnnnnnt… but if you buy some time, you just might find, you'll get what you need."

Oooh, yeah!

POINT TO REMEMBER

- A pair that starts and ends going north has by definition equal numbers of right-hand and left-hand turns.

6.12 INTERCABINET APPLICATIONS

The term used to describe the high-speed differential wiring often used between pieces of equipment is *balanced cabling*. This term has been adopted by the ISO building-wiring standards committee to describe any cable that provides one or more pairs of wire, each pair having a defined differential-mode impedance and each pair having a defined immunity to crosstalk from the other pairs within the same jacket.

There are two basic construction techniques used to produce balanced cabling: the *twisted pair* and the *quad configuration* (see Figure 6.30). Both arrangements hold the wires of each pair in a fixed arrangement with a uniform cross section. This stabilizes the differential impedance of the cable. Both guarantee low crosstalk among the pairs.

The twisted-pair cable guarantees low crosstalk by virtue of having a different rate of twist on all the pairs within the same jacket. The different rates of twisting are an essential part of the crosstalk cancellation process. This happens because of the way transmission-line coupling works between two adjacent differential pairs. The basic rule of thumb for pair-wise crosstalk is this: *When you flip one pair, the crosstalk reverses polarity*.

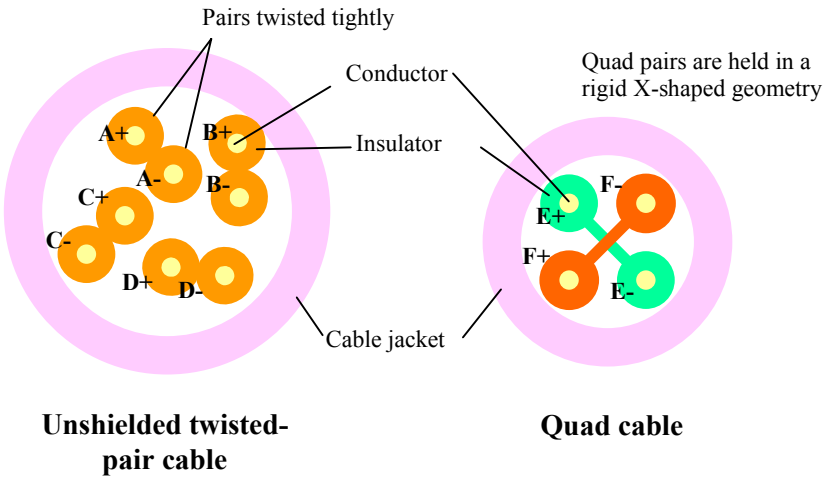

Figure 6.30—Construction of balanced cables.

A corollary to the flipping rule is this: *When you flip both pairs, the crosstalk retains the same polarity*. That might happen often if the twist rates were the same on two adjacent pairs. Every time both pairs flipped over, the crosstalk would remain the same. The crosstalk therefore might never cancel. To avoid this effect, the pairs in a multi-pair twisted cable are usually twisted at different rates. This randomizes the relation of one pair vis-à-vis its partner, nulling out the crosstalk.

On a well-constructed twisted-pair cable, one of the colored pairs will carry a noticeably tighter twist than the others. The crosstalk to and from this pair will be the best in the bunch, a nice property. Do not, however, be deluded into thinking that all cables will have the same hierarchy of twist performance. There are few, if any, standards concerning which of the pairs should carry the tightest twist. Manufacturers are free to change the twist pattern at will, including reassignment of the hierarchy of twist performance. The crosstalk standards for most cables specify only the worst-case crosstalk between any two pairs. They do not designate any particular pairs as having better performance than the others.

The quad cable guarantees low crosstalk by virtue of its unique geometrical alignment. Both capacitive and inductive coupling mechanisms between the pairs are cancelled by this construction technique. With regards to interpair crosstalk within the same jacket, quad cable, if carefully constructed, can exceed the performance of twisted-pair cabling. With regards to crosstalk with other objects and cables outside the jacket, quad cable performs less well than twisted-pair cable. Twisted-pair cable does a better job of canceling electromagnetic radiation from the cable and providing good common-mode rejection.

The following sections describe the main applications for balanced cabling between cabinets.

POINTS TO REMEMBER

- The twisted-pair cable guarantees low crosstalk by virtue of having a different rate of twist on all the pairs within the same jacket.
- Quad cable guarantees low crosstalk by virtue of its unique geometrical alignment.

6.12.1 Ribbon-Style Twisted-Pair Cables

Ribbon-style twisted-pair cables have the same twist pitch on all pairs and yet still deliver reasonable crosstalk performance. It seems counterintuitive that this would work, because when one pair twists (inverting its local field polarity) the adjacent pair twists as well (inverting its sensitivity). The crosstalk would seem to reinforce with the same polarity at every twist. How can it work?

This paradox is solved by looking closely at the exact variations in crosstalk as the wires turn about one another. Imagine an axis run horizontally through the centerline of both pairs. Now imagine you can continuously control the angle of rotation on each pair about their respective axes. Begin with a rotational phase of 0º (left view in Figure 6.31). The coupling for this configuration is dominated by the inside two wires, A– to B+, and so has a negative polarity.

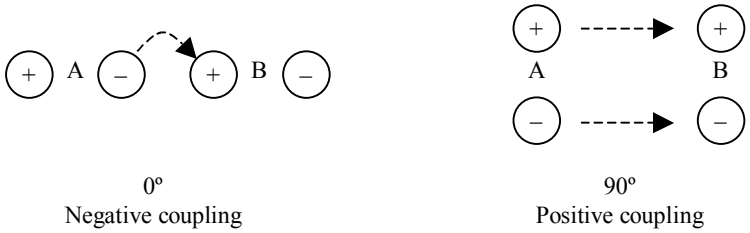

Figure 6.31—In a twisted ribbon cable the crosstalk coupling polarity reverses every 90º.

At a rotational phase of 90º (right view in Figure 6.31), the coupling changes dramatically. In this case the A+ wire couples mostly to B+, and A– to B–, yielding a coupling amplitude almost exactly the same as in the previous case, but with opposite (positive) polarity.

The coupling reverts to the original (negative) polarity at 180º and inverts back to positive once again at 270º. As long as the rotational axes of each pair is held in a fixed position, the coupling averaged throughout the entire rotational cycle nulls to near zero. The wires of a twisted ribbon cable are varnished into place to ensure they maintain the correct geometry. Different arrangements of the starting phases and rotational directions are possible.

In a practical, multipair, jacketed cable it is not generally possible to hold all the wires in fixed positions. The wires in the 90º case are likely to slump towards each other, upsetting the cancellation. To circumvent this difficulty the manufacturers of multi-pair twisted cables resort to the ruse of varying the twist rate on each pair.

POINT TO REMEMBER

- Ribbon cables can use the same twist pitch on every pair because the wires are held in a rigid geometry.

6.12.2 Immunity to Large Ground Shifts

Differential signaling with unshielded twisted-pair cables does not require a direct ground connection between the two ends of the link. As long as the potential difference between the transmitter and receiver remains within the common-mode input range of the receiver, the system will function. In most cases, the existing green-wire ground connection implemented on most computer equipment keeps the product chassis at either ends of the link within an acceptable voltage range (see box). No additional grounding needs be added to the system.

High-frequency single-ended signaling, on the other hand, *does* require a direct ground connection between the two ends of the link. Because this ground connection carries high-frequency returning signal currents, it must follow closely along with the signal wires in a low-inductance, controlled-impedance structure. The green-wire ground is woefully inadequate for this purpose. If single-ended signaling is to be used between cabinets, additional grounding means (such as a coaxial cable shield) must be implemented.

These additional grounding means may violate one of the most sacred AC power safety principles:

Never introduce a metallic connection between any two frames powered by different AC power sources.

As explained in the box "Earth Potential," violation of this rule may draw significant currents through the green-wire connection. This is a problem because it upsets the sensitive green-wire current detectors built into the main electrical panel of most modern buildings. These detectors look for early warning signs of electrical malfunction. For example, a partial short between any *hot* wire and a product chassis will transmit green-wire currents back to the electrical panel where they may be detected. The circuit that detects these currents is called a ground-fault interrupter, or *GFI,* circuit breaker. When the detected current exceeds a critical threshold, power may be removed from that section of the building. Messing with the green wire is serious stuff. Don't do it.

If you must electrically connect the metallic frames of two systems, make sure that both systems are served by a green-wire ground connected to the same Earth potential. There are multiple ways to do this. For systems located within the same rack-mounted chassis, just screw all the frames to the same rack. For boxes located in the same room, but not in the same rack, provide a way to plug the AC power cord of one system into a convenience outlet on the other system. This arrangement daisy chains the green-wire connections, so you know they are all at the same potential. If daisy chaining is not possible, try to plug all the systems into the same outlet or power strip. For boxes located within different rooms, use differential signaling, fiber, or RF connections that don't require a metallic connection between frames.

Earth Potential

The potential across the surface of the Earth is not constant. Various mechanisms, including spurious power distribution currents, magnetic-field interactions, and lightening, induce large currents in the surface of the Earth. These currents, working across the surface resistance of the soil layers, produce noticeable potential differences. Between the ends of a typical building, one may observe several volts of potential difference.

Large buildings are typically divided into several grounding domains. Each domain is powered by a local transformer, which typically sits near the center of the domain. At that location, the neutral wire of the transformer secondary, the green-wire ground, and a copper ground stake are all bonded together (see Figure 6.32). The green-wire grounds between domains do not touch. This arrangement limits the voltages between the machinery in your office (whose outer metallic skin is connected to the green-wire safety ground) and the actual local ground where you are standing to something just below the level of human perception (a few volts).

If you connect together the metallic skin of a box in one domain with the skin of another box in another domain, several things will happen. First, you may see a noticeable spark. After that, several amps of current that used to be flowing in the Earth will now begin flowing through the connection. This current flows from the ground stake at position *A,* through a green wire to chassis *B,* through connection *C* to chassis *D,* and from there back to ground stake *E*.

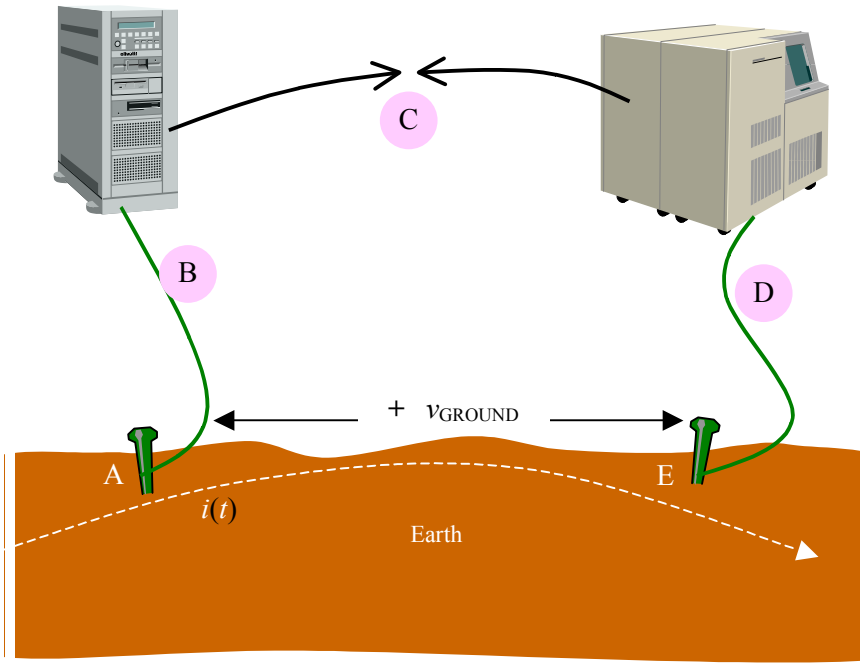

Figure 6.32—Huge currents *i*(*t*) circulating through the Earth's crust cause measurable differences in the electric potential at the Earth's surface.

POINTS TO REMEMBER

- Never introduce a metallic connection between any two frames powered by different AC power sources.
- If you must electrically connect two boxes, make sure that both boxes are served by green-wire grounds connected to the same Earth potential.
- Differential signaling with unshielded cables does not require a direct ground connection between the two ends of the link.

6.12.3 Rejection of External Radio-Frequency Interference (RFI)

External RF fields impinging on a twisted-pair cable tend to affect both wires equally. Any interference mostly appears as a common-mode signal on the cable, which is cancelled at the receiver. I say mostly because, as usual, a number of things can go wrong. As good as twisted-pair cables are for rejecting RF interference, here's what happens in the real world:

1. Part of the RF field energy is absorbed by the cable. This has to do with the efficiency of the complete cable structure as an electrical antenna, a subject outside the bounds of this book. For more information about antenna

efficiency as it relates to RFI problems, see the excellent text by Clayton Paul [57].

2. The RF field energy absorbed is converted to a common-mode current. This happens according to the relation $i_{\text{COMMON}} = \sqrt{P/Z_{\text{COMMON}}}$, where P is the power received and Z_{COMMON} is the common-mode impedance of the cable, with respect to true Earth ground, including its common-mode terminations.
3. If the common mode of propagation is unterminated at both ends of the cable (as it would be using ordinary transformer-coupling at both ends), a significant resonance may occur that amplifies the received common-mode current.
4. Some fraction of the common-mode current flowing in the cable is converted to a differential-mode current. This conversion may take place due to a natural imbalance in the construction of the cable itself, an imbalance in the connectors, or an imbalance in the transmitter or receiver circuitry.
5. The receiver interprets the differential-mode current as a true differential signal.

POINTS TO REMEMBER

To get the best RF-rejection performance from your cabling,

- Use a tightly twisted, well-balanced cable. Twisted cables work better than quad cables in this respect.
- Don't scrimp on connectors. Buy and use connectors designed to go with the cable.
- Use well-balanced circuitry for both transmitter and receiver.

6.12.4 Differential Receivers Have Superior Tolerance to Skin Effect and Other High-Frequency Losses

Let's say you need to communicate one digital signal from box A to box B located 15.2 m (50 feet) away. You choose a single-ended 3.3-V 50-ohm line driver, running on RG-58 coax at 1000 Mbaud (one nanosecond per bit), with a rise/fall time of 250 ps.

The response of this system is shown in Figure 6.33. The figure shows the actual eye pattern, as predicted by simulation, using solid lines. The ideal transmitted waveform, assuming no skin-effect distortion or attenuation, is depicted with a dashed line. The transmitted data pattern is …1111010001111….

The worst-case high and low receiver thresholds for single-ended 3.3-V JEDEC LVTTL logic are drawn in place at 2.0 and 0.8 volts respectively. Notice that the low-side threshold fails to catch the first negative-going excursion—causing a bit error. Even when the LVTTL receiver does properly interpret the data, you can expect a fair amount of jitter in the received waveform.

Differential receivers are commonly specified with more accurate switching thresholds than ordinary single-ended logic. Had you selected a differential receiver and a

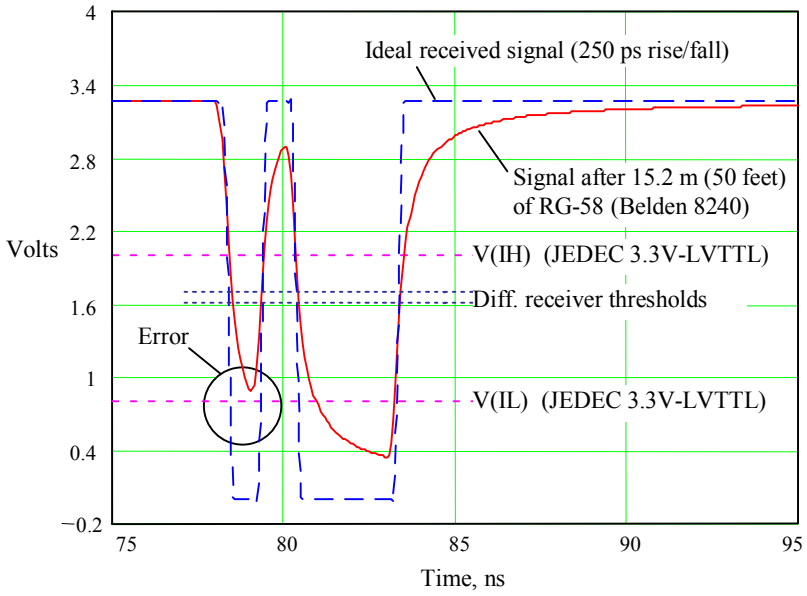

Figure 6.33—Fifty feet of Belden RG-58 distorts this 1-Gb/s signal.

differential cabling system, the effective receiver thresholds would have been more nearly centered in the middle of the data pattern.

For example, the chart shows the differential receiver thresholds for LVDS logic. These thresholds still properly discriminate the data even in the face of severe pulse distortion. In general, for the same amount of transmission-line distortion, a differential receiver generates less jitter than a single-ended receiver. This advantage follows from the generally better threshold tolerances available in differential receivers, not of the differential architecture itself.

In the example of Figure 6.33, you could improve the single-ended system performance by using a differential receiver with its negative input terminal tied to a stable and accurate source of 1.65 V. That simple change would create a single-ended receiver with much better control over the input threshold than indicated in the figure.

If the transmission cable is end-terminated, the end termination is best tied to some voltage halfway between V_{IH} and V_{IL} (or tied to a split terminator with a Thevenin equivalent voltage between V_{IH} and V_{IL}). That way the DC attenuation of the cabling symmetrically affects both high and low logic levels, keeping the received signal centered.

POINT TO REMEMBER

- Differential receivers have more accurate switching thresholds than ordinary single-ended logic.

6.13 LVDS SIGNALING

This section is not intended as a promotion of LVDS logic. It is a primer on how to interpret the specifications of any differential logic family.

Low-Voltage Differential Signaling (LVDS) is a good example of a high-speed differential logic family. The LVDS standard was generated in 1995 by the IEEE. [62]

The LVDS standard contemplates both *general-purpose* and *short-range* applications. Here I will confine my remarks to the *general-purpose* version of the standard. It defines a differential data path operating at data transfer rates in the range of 200 to 500 MHz with a source-synchronous clock and various bus widths up to 128 bits.

Table 6.5 lists some of the key performance specifications for LVDS general-purpose transceivers.

6.13.1 Output Levels

LVDS being a differential logic family, there are two (complementary) outputs per logic signal. The nominal steady-state operating conditions for these outputs are 1.0 and 1.4 volts,

Table 6.5—LVDS General-Purpose Link Specs (adapted from ANSI/IEEE P1596-3-1995)

Transmitter specifications					
Signal	Parameter	Conditions	Min	Max	units
V_{oh}	Output voltage high, either wire	R_{load} = 100Ω ± 1%		1475	mV
V_{ol}	Output voltage low, either wire	R_{load} = 100Ω ± 1%	925		mV
$\|V_{od}\|$	Output differential voltage	R_{load} = 100Ω ± 1%	250	400	mV
R_0	Output impedance, single-ended	V_{cm}=1.0V and 1.4V	40	140	Ω
V_{os}	Output offset voltage		1125	1275	mV
ΔV_{os}	Change in VOS between 0 and 1 states (this specification defines the AC common-mode output voltage)	R_{load} = 100Ω ± 1%		25	mV
t_{rise}, t_{fall}	V_{od} rise/fall time 20% to 80%	R_{load} = 100Ω ± 1%	300	500	ps
Receiver specifications					
V_i	Input voltage range, either input	$\|V_{gpd}\|$ < 925 mV	0	2400	mV
V_{idth}	Input differential threshold	$\|V_{gpd}\|$ < 925 mV	–100	+100	mV
V_{hyst}	Input differential hysteresis	V_{idthh}–V_{idthl}	25		mV
R_{in}	Receiver differential input impedance	—	90	110	Ω
C_{in}	Not specified				
Implementation specifications					
	Pcb skew allocation	Worst case		50	ps

for the low and high states respectively. When one wire goes to 1.0 volts, the other goes to 1.4, and vice versa.

You may decompose this situation into a steady-state common-mode component of 1.2 volts, plus a changing differential voltage of ±0.4 volts.[57] The differential voltage varies from –0.4 V to +0.4 V, so we say that the differential peak-to-peak voltage is 800 mV. The peak-to-peak voltage on either wire alone would be 400 mV.

The standard requires that the steady-state differential voltage representing either a zero or one state be at least 250 mV, but no larger than 400 mV *when the outputs are loaded differentially by 100 ohms*. It also requires that no output ever fall below 925 mV or exceed 1475 mV under the same conditions.

Concerning output levels under other loading conditions, the standard provides little guidance. The only hints come in the form of stated constraints on output impedance. From the standard it is impossible to determine, for example, the output levels that would result from loading the outputs with a 75-ohm differential load. Nor is it possible to determine something else that the IBIS community has long sought in these sorts of documents: a precise specification of the shape of the rising and falling edge.

POINT TO REMEMBER

- Normal operating voltages for LVDS logic are 1.2 ± 0.2 V on each wire.

6.13.2 Common-Mode Output

The output offset voltage (DC bias) can change by 25 mV when switching from the one state to the zero state. This is a peak-to-peak change. The AC amplitude range of the common-mode voltage is therefore ±12.5 mV.

The AC amplitude of the differential output voltage lies somewhere in the range of ±250 to ±400 mV.

When considering certain radiation and crosstalk problems, it is handy to know the approximate ratio between common-mode and differential-mode emissions coming out of a driver. For LVDS, the ratio of common to differential amplitudes can be as poor as 12.5/250 = 5%.

POINT TO REMEMBER

- LVDS, like most digital transceivers, is not extraordinarily well balanced.

6.13.3 Common-Mode Noise Tolerance

Looking at the common-mode operating range of the receiver (called V_i in the specification), it is apparent that the receiver can tolerate inputs anywhere in the range of 0 to 2400 mV. From these numbers you can derive how much common-mode noise the

[57] Or an even-mode voltage of 1.2 volts plus an odd-mode voltage of ±0.2 volts.

system will tolerate. First assume a driver is at its lowest permissible level (925 mV). Now figure out how much common-mode noise you can add in the negative direction without the receiver input falling outside its guaranteed operating range. The answer is –925 mV. Next assume the driver is at its highest permissible level (1475 mV). Now figure out how much common-mode noise you can add in the positive direction without the input falling outside its guaranteed operating range. The answer is +925 mV. LVDS therefore tolerates a common-mode difference (V_{gpd}) between the ground potential at the driver and the ground potential at the receiver as great as ±925 mV.

If you exceed the common-mode operating range of the receiver, all bets are off. It could do anything (see Section 6.7, "Common-Mode Range").

POINT TO REMEMBER

- The common-mode noise tolerance for general-purpose LVDS logic is ±925 mV.

6.13.4 Differential-Mode Noise Tolerance

Next let's look at the differential noise margin. In the worst case the transmitter differential output may be as small as 250 mV. At the same time, the receiver threshold may be offset by as much as 100 mV. The difference between these two figures is the differential noise margin, which is 150 mV. As a percentage of the signal swing on either wire (400 mV p-p), the 150 mV figure represents 37%. A transmitter with a larger output swing would enjoy an even bigger percentage noise margin. These are excellent noise margins for digital logic. Most single-ended logic families have noise margin percentages on the order of only 10% to 15%.

POINT TO REMEMBER

- The high noise margin gives LVDS a built-in natural advantage in combating ringing, overshoot, and crosstalk from like devices.

6.13.5 Hysteresis

A receiver with hysteresis has two input switching thresholds, one used for positive-going signals and one used for negative-going signals. The positive-going threshold is always set a little higher than the negative-going threshold. Once the input crosses into positive territory, the receiver automatically switches to the negative-going threshold so that the signal must turn around and descend below the negative threshold before it can cause another switching event. Once the signal crosses below the negative threshold, the receiver flips back to using the positive threshold. A receiver equipped with hysteresis requires a certain amount of signal change before flipping to the other state.

The hysteresis feature is intended to avoid annoying self-oscillation that might happen with slowly changing, but very clean, inputs. The feature is usually implemented as a form

of limited positive feedback from the receiver output back to own its input. With hysteresis, the inputs tend to switch quickly and firmly, once they reach an acceptable level, and then stay there. It's a good feature.

LVDS inputs have a guaranteed amount of hysteresis. You still, however, shouldn't supply any of these parts with a slowly moving input, because in that case as the input slowly sweeps through the transition region, any crosstalk that happens to exceed the hysteresis switching range will cause glitches (and therefore more noise) in the receiver.

POINT TO REMEMBER

- Always provide fast-edged inputs to LVDS logic.

6.13.6 Impedance Control

The LVDS specification goes to a lot of trouble to control ringing and reflections. This is one of the strongest provisions of the specification. Excellent control of ringing and reflections makes it possible to obtain first-incident-wave switching in most LVDS applications.

LVDS uses a both-ends termination strategy to control reflections. Each LVDS transmission line is terminated first at the source and again at the end of the line.[58]

The source impedance of the driver is constrained to the range of 40 to 140 Ω. That is a ratio of only 3.5:1 from highest value of allowed output impedance to the lowest. To a board-level analog designer, this sounds easy, because you can go out and buy very accurate lumped-element resistors. To a chip designer, it's a nightmare. You just can't control the absolute value of $R_{DS}(ON)$ or the absolute value of transconductance very accurately.

The achievement of a 3.5:1 output impedance specification represents a major accomplishment for the chip industry, one which I am sure will pay off in terms of higher volume, given the user-friendly advantages of both-ends termination (see box "Both-Ends Termination").

The worst-case reflection coefficient at the transmitter, assuming it is coupled to a perfect 100-Ω transmission line, will therefore be the worse of these two numbers:

$$\Gamma_{\text{MAX RO}} = \frac{140-100}{140+100} = +0.167 \qquad [6.26]$$

$$\Gamma_{\text{MIN RO}} = \frac{40-100}{40+100} = -0.428 \leftarrow \text{worst case} \qquad [6.27]$$

The input impedance of the receiver is constrained to the range 90 to 110 Ω. The LVDS specification recommends (but does not require) that this be implemented as a built-in terminator, placed inside the integrated chip package right at the die. From a signal integrity perspective, that would definitely be the best place to put it. Initial LVDS implementations,

[58] Variants of LVDS are available in which the drivers can enter a tri-state (high-impedance) mode, and the receivers do not incorporate terminations. These versions are suitable for building multidrop bus structures.

however, did not do this. Due in part to the difficulty of fabricating accurate on-chip resistances, early implementations of LVDS left the 100-ohm termination as an external component.

If you have to design with external terminations, use a 100-Ω ±10% external terminating resistor in a low-inductance package (0805 or smaller package) directly attached to the transmission line at the input terminals of the package, with very small pads (for low parasitic capacitance).[59]

The worst-case reflection at the terminator, assuming a line impedance of precisely 100 Ω, will be the worse of these two numbers:

$$\Gamma_{\text{MAX RIN}} = \frac{110-100}{110+100} = +0.047 \qquad [6.28]$$

$$\Gamma_{\text{MIN RIN}} = \frac{90-100}{90+100} = -0.053 \leftarrow \text{worst case} \qquad [6.29]$$

Multiplying together the worst-case transmitter and receiver reflection coefficients, $\Gamma_{\text{MIN RO}}$ [6.27] and $\Gamma_{\text{MIN RIN}}$ [6.29], shows that the amplitude of any residual reflections in the transmission structure (meaning anything that arrives after the initial step edge) can in no case exceed 2.25% of the initial signal amplitude. Therefore, you can expect solid first-incident wave switching performance from this system, assuming a perfect implementation with perfect 100-Ω differential transmission lines.

You may be wondering how far the line impedance may stray from the ideal value of 100-Ω while still guaranteeing first-incident-wave switching. Figure 6.34 reveals the answer. This figure shows the magnitude of the residual reflections remaining after the

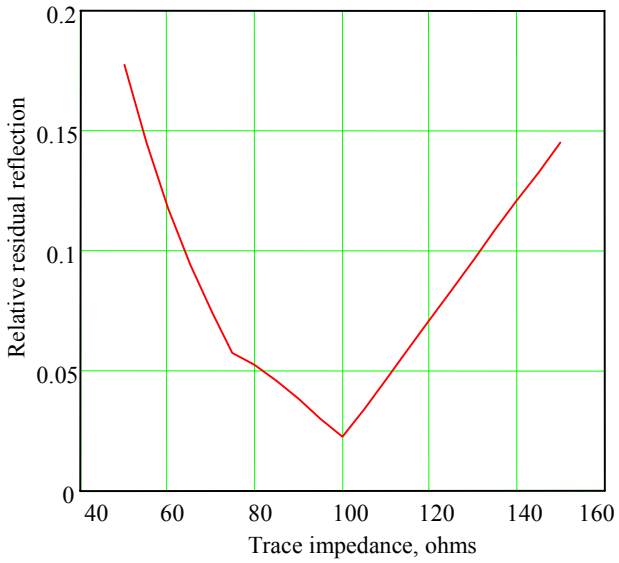

Figure 6.34—Residual reflection after arrival of initial step edge for terminated LVDS logic with worst-case transmitter and receiver impedances, as a function of trace impedance.

arrival of the first-incident waveform. The figure is a compilation of four different constraint lines corresponding to different combinations of worst-case high and low R_o interacting with worst-case high and low R_{in}. At various specific values of trace impedance, one constraint or another takes precedence, which accounts for the segmented appearance of the curve.

The chart indicates that a line impedance of 100Ω ±10 would produce an initial residual reflection no greater than 5% of the incoming step amplitude. A plus-or-minus 20-ohm tolerance would increase the initial residual to no greater than 7%. LVDS logic, because it uses a both-ends style termination, tolerates a fairly wide range of line impedances.

The existence of a significant residual reflection may not by itself endanger the performance of a particular link, depending on the polarity and timing of the arrival of the reflected signal power. Only time-domain simulation can tell.

This figure assumes the transmitter and receiver impedances are purely resistive. Any significant reactances at either the transmitter or receiver will further degrade the performance.

Both-Ends Termination

The both-ends-termination strategy uses a source termination at the transmitter and an end termination at the far end of the line. This is what I call the ax-murderer approach to reflection control. An ax murderer never takes just one whack at a problem. He keeps working and working on the problem until it is completely, finally solved.

That's how both-ends termination works. The main signal, once launched from the transmitter, proceeds toward the far end at full speed. Upon arrival, the end termination metes out one terrible, swift chop that reduces the incoming signal to little more than a small reflected bit of rubble. This surviving reflection retreats post-haste towards the safety of the transmitter, where the source terminator delivers another fatal blow. Only a very small remaining fraction of the original signal amplitude survives this double-chop to crawl back toward the far end a second time.

The second-incident wave amplitude is the product of the original transmitted signal size, the reflection coefficient at the far end, and the reflection coefficient at the near end. In a reasonably well-designed system this amplitude is small enough to simply ignore. That's the big advantage of both-ends-terminated systems, there are hardly any reflections to worry about.

Another advantage of both-ends termination is its high tolerance for obstacles in the middle of the transmission line (like vias). Single-end terminated transmission lines have a lower tolerance for obstacles. In a single-end terminated line, there is always some reflection pattern that can take one bounce off the obstacle, and one additional bounce off some unterminated end, and wind up at the receiver. In a both-ends-terminated design the reflection coefficient at each end of the line is small, so *all* reflection modes are damped, even ones that bounce off of obstacles in the middle of the line.

POINT TO REMEMBER

- LVDS works best with 100-Ω transmission lines.

6.13.7 Trace Radiation

In the LVDS specification the number ΔV_{OS} defines the degree of balance between the two complementary outputs. It calls for a peak-to-peak common-mode (or even-mode) content in the transmitted signal of no more than 25 mV. Compared to the peak-to-peak signal level on either of the two signal wires (400 mV), that's a relative common-mode content of 6.25%.

The common-mode content limits the degree of attainable radiated field cancellation to a value of –24 dB (=20log(0.0625)). You can easily achieve this amount of cancellation at all frequencies up to 1 GHz by placing the differential traces at any separation of 0.5 mm or less (see Section 6.11.3). Unless you need to save the circuit board space, it is not, in this author's opinion, worth the effort trying to cram LVDS traces closer together than 0.5 mm.

In individual circumstances with particularly well-balanced transmitters it is possible to get better cancellation, but you can't depend on always having parts that beat the specification.

POINT TO REMEMBER

- You need not struggle to place ordinary differential digital traces any closer than 0.5 mm (0.020 in.) for any EMI purpose.

6.13.8 Risetime

I'm glad to see a specification for the minimum risetime. That' s a big help when dealing with all manner of high-speed phenomena, especially the calculation of crosstalk. I offer my sincere thanks to all the standards weanies who voted for this provision.

6.13.9 Input Capacitance

The last of the receiver specifications is the input capacitance. Sadly, this specification is lacking. The closest we get in the standard to addressing the input capacitance is a vague statement that the input capacitance "should not limit the high-frequency, 250-MHz operation of the receiver." That's nice, but it's not a specification. Standards like this leave open the possibility of receivers that meet the spec as written, but don't interoperate.

6.13.10 Skew

The LVDS committee did a lot of work on clock-to-data skew. They carved out an overall skew budget, defining a permissible amount of skew for each signal in an LVDS link. Their budget assumes a link architecture that includes two pcbs, each with a connector and each plugged into some sort of backplane media.

In this architecture the specification pcb designers need to worry about is the pcb skew number of 50 ps. If every data and clock signal in an LVDS link is matched to within this amount of delay, the timing for the link as a whole should work.

Keep in mind that skew accumulates as your signal progresses. If your signal must traverse more than two connectors, the skew budget for each is less than in a simpler system.

If you are using an FR-4 dielectric, the 50-ps delay number gives you an allowance for about 1/4 inch of line length imbalance between any two signals in an LVDS link. This figure is definitely achievable, but don't depend entirely on your autorouter—you need to take a close look at the final artwork to make sure you've stayed under the limit.

The LVDS specification does not make any specific reference to the degree of skew imbalance permitted between the two wires of an individual differential signal. My rule of thumb is that the skew imbalance in any differential pair should be kept to less than 1/10 of the risetime.

POINT TO REMEMBER

- Always double-check your final artwork to make sure you've met the specifications for skew.

6.13.11 Fail-Safe

LVDS components from National Semiconductor include a fail-safe circuit in the receivers. This feature shuts off the output in the event the input is disconnected (zero differential input). This feature is permitted by the standard, but not required, so check carefully if you will be mixing different vendors to make sure they all do it in a compatible fashion.

Figure 6.35 illustrates how fail-safe is implemented in the National LVDS logic family.

The figure depicts the mandated differential thresholds, V_{IH} and V_{IL}, for an LVDS receiver. The manufacturer of the receiver in Figure 6.35 has created a part with better control over the input threshold than the mandated minimum. The actual thresholds, $V_{TH}(+)$ and $V_{TH}(-)$, are specified at ±30 mV. The close tolerance of the actual thresholds is exploited to create the fail-safe feature.

The fail-safe feature is created by forward biasing the inputs. When the input is disconnected from any source (the transmitter is turned off or unplugged), biasing resistors R2 and R3 trickle enough current through the external end-termination resistor R1 to forward-bias the input by 50 mV. This level is above the actual component threshold, so the receiver output stays locked at 1.

When the input is connected to a source with a differential output impedance of 100 ohms, the current from resistors R2 and R3 forward-biases the input by only half as much, or only about 25 mV, shifting $V_{TH}(+)$ and $V_{TH}(-)$ to new worst-case values of +55 mV and –5 mV respectively. These values remain well within the mandated limits of ±100 mV.

Some applications require a greater margin of safety for the fail-safe feature. For example, let's say you are making a twisted-pair communication link. When the transmitter is powered off, you may expect more than 25 mV of differential noise. This can be

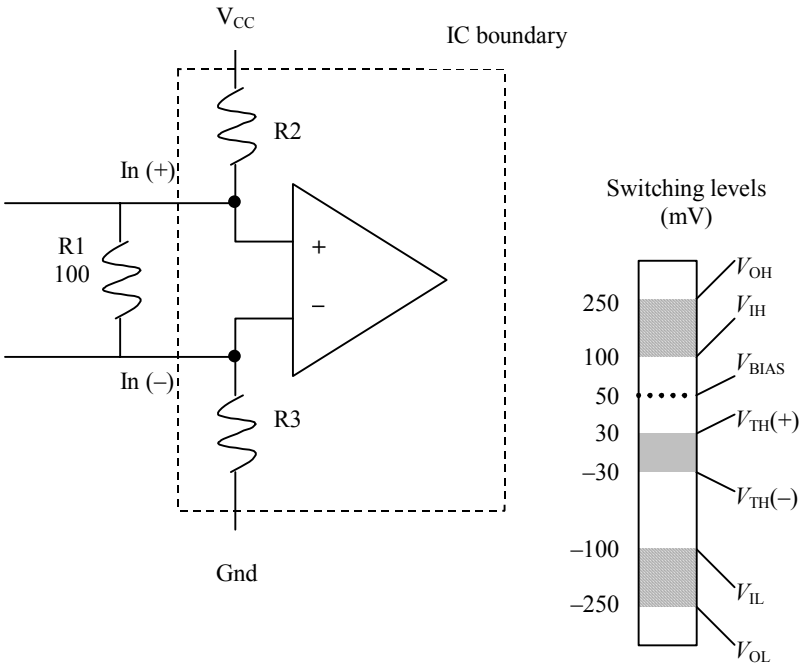

Figure 6.35—Switching levels for National LVDS logic family.

implemented by adding two new resistors, R4 and R5, in parallel with the existing bias resistors R2 and R3, but outboard of the IC package. The new resistors can be sized to enforce an arbitrary amount of offset in the case the transmitter is disconnected. One disadvantage of this technique, if taken to an extreme, is that the fail-safe bias current may be large enough to disturb normal operations.

The circuit in Figure 6.36 fixes this problem. In the event the transmitter is powered off or unplugged, the fail-safe resistors R4 and R5 provide a large amount of bias current. In the event the transmitter is connected and powered on, you can pick values for R6 and R7 in the transmitter that will source an equal and opposite amount of current, canceling the offset. Resistors R2 through R5 appear in parallel with the differential impedance of the termination network and must be taken into account when selecting values for R1 and trace impedance. The same applies for resistors R6 and R7 at the source.

POINT TO REMEMBER

- Fail-safe features are permitted by the LVDS standard, *but not required*.

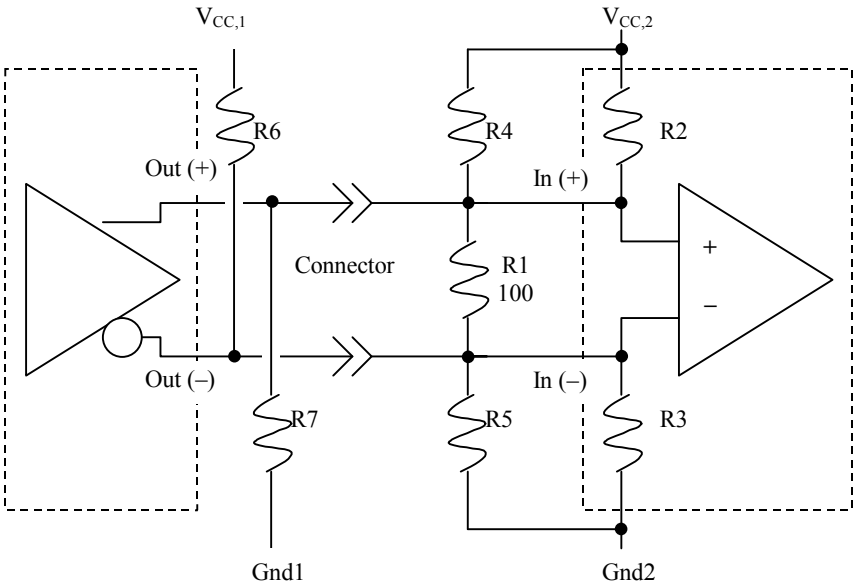

Figure 6.36—Combining external bias resistors at receiver and transmitter creates zero bias during operational mode.