

Line thickness¹⁶ (in.), t
 Effective relative permittivity, ϵ_r
 (Equals the relative permittivity of the surrounding medium)
 For narrow striplines:
 Use these formulas when $w/b < 0.35$ and also when $t/b < 0.25$.
 Impedance (Ω):

$$\frac{60}{\sqrt{\epsilon_r}} \ln \left(\frac{1.9b}{0.8w + t} \right) \quad [4.92]$$

Propagation delay (ps/in.):

$$85\sqrt{\epsilon_r} \quad [4.93]$$

POINTS TO REMEMBER:

For printed circuit board traces, the most critical dimension is the ratio of trace width to height above ground.
 Double the reflection budget to find the allowed mismatch between the characteristic impedance and the terminating resistors.
 Large variations in physical dimensions make a small impact on the resulting impedance.
 The slope of any function plotted on log-log paper is equal to the sensitivity of the function to changes in its argument.
 All formulas for transmission velocity are inversely proportional to the effective square root of electric permittivity.

Ground Planes and Layer Stacking

Ground and power planes in high-speed digital systems perform three critical functions:

- Provide stable reference voltages for exchanging digital signals
- Distribute power to all logic devices
- Control crosstalk between signals.

This chapter focuses on signal crosstalk. Sections 5.1–5.6 assume relatively short traces, for which a lumped analysis of mutual inductance is appropriate. Section 5.7 treats the case of long lines, where we separate coupling into its forward and reverse parts.

Section 5.8 summarizes the rules for designing good printed circuit board layer stacks that control crosstalk.

The formulas in this chapter are accurate only to within a factor of 2. For better accuracy use the formulas as a guide, build a prototype, and then measure its actual performance. Many transmission line structures are easily built out of copper tape and blank circuit board material.

The formulas in this chapter are excellent for showing how electromagnetic effects adjust in response to physical changes. For example, if crosstalk is 30% too high, the formulas can show just how much further to separate your traces. They are less useful for predicting the absolute value of any particular effect.

5.1 HIGH-SPEED CURRENT FOLLOWS THE PATH OF LEAST INDUCTANCE

At low speeds, current follows the path of least resistance. In reference to Figure 5.1, low-speed current transmitted from A to B returns to the driver along the ground plane. This return current flows in wide arcs on its way back to the driver. The current density along each arc corresponds to the conductance of that path.

¹⁶Sometimes reported as plating weight. Each ounce of plating weight is 0.00135 in. thick.




Figure 5.1 At low frequencies current follows the path of least resistance.

At high speeds, the inductance of a given return-current path is far more significant than its resistance. High-speed return currents follow the path of least inductance, not the path of least resistance.

The lowest inductance return path lies directly under a signal conductor, minimizing the total loop area between the outgoing and returning current paths.

Returning signal currents tend to follow this direct path, close underneath a signal conductor. Figure 5.2 shows a typical high-frequency return-current path.




Figure 5.2 At high frequencies current follows the path of least inductance.

Figure 5.3 presents the cross section of a typical printed circuit board trace along with its return-current distribution. The peak current density lies directly under the trace. The current density then falls off sharply to each side.

An approximate relation for the return-current density at a point D inches away from a signal trace is

$$i(D) = \frac{I_0}{\pi H} \cdot \frac{1}{1 + (D/H)^2} \quad [5.1]$$

where I_0 = total signal current, A

H = height of trace above circuit board, in.

D = perpendicular distance from signal trace, in.

$i(D)$ = signal current density, A/in.

Current density at point D is proportional to

$$\frac{1}{1 + (D/H)^2}$$




Figure 5.3 Distribution of high-frequency current density underneath a signal trace.

The current distribution (Equation 5.1) balances two opposing forces. Were the current more tightly drawn together, it would have higher inductance (a skinny wire has more inductance than a broad, flat one). Were the current spread farther apart from the signal trace, the total loop area between the outgoing and returning signal paths would increase, raising the inductance (inductance is proportional to loop area). Equation 5.1 specifies the optimum distribution minimizing the total loop inductance of both outgoing and returning current paths.

The current distribution (Equation 5.1) also minimizes the total energy stored in the magnetic field surrounding the signal wire.

POINTS TO REMEMBER:

- High-speed current follows the path of least inductance.
- Returning signal currents tend to stay near their signal conductors, falling off in intensity with the square of increasing distance.

5.2 CROSSTALK IN SOLID GROUND PLANES

Crosstalk between two conductors depends on their mutual inductance and their mutual capacitance. Usually, in digital problems the inductive crosstalk is as big or larger than capacitive crosstalk, so we will henceforth discuss mainly the inductive coupling mechanism.




Figure 5.5 Mutual coupling experiment




Figure 5.6 Step response of a mutual coupling experiment

The theory behind mutual inductive coupling for lumped circuits appears in Section 1.10. That theory postulates that returning signal currents will generate magnetic fields. Those magnetic fields in turn induce voltages in other circuit traces.

The induced voltages are proportional to the derivative of the driving signal. They become markedly worse as rise times get shorter.

Because the returning current density and its associated local magnetic field strength drop off according to Equation 5.1, we suspect that mutual inductive crosstalk also drops off as we move the two traces in Figure 5.4 away from each other:

$$\text{Crosstalk} \approx \frac{K}{1 + (D/H)^2} \quad [5.2]$$




Figure 5.4 Cross section of two traces showing crosstalk.

Here we are expressing crosstalk as a ratio of measured noise voltage to the driving step size. The constant K depends on the circuit rise time and the length of the interfering traces. It is always less than 1.

We can check this hypothesis with a simple experiment. The traces configured in Figure 5.5 are 26 in. long and separated by 0.080 in. center to center. They lie on a single-layer main circuit board. The ground plane for this experiment is a solid copper sheet fastened below the main circuit board. Together, the main board and the ground layer sandwich a pile of dielectric spacers of known thickness. By this arrangement we can simultaneously vary the effective height above ground of the driven and receiving traces.

As in problems involving characteristic impedance, the ratio of sizes is more important than the absolute dimensions. In this case it is the ratio D/H that determines the crosstalk. By varying the trace height above ground, we can control the ratio D/H .

Figure 5.6 plots the resulting step responses, as measured at point D , when driving the input with a 3.5-V step. Figure 5.6 uses ground separations of 0.010, 0.020, 0.030, 0.040, and 0.050 in. The last trace (biggest noise pulse) was taken with no ground plane at all.

Figure 5.7 compiles this measured data into a chart showing the mutual inductance between traces as a function of the ratio D/H . Area is used as a measure of mutual coupling, as explained in Section 1.8. By measuring the area we can factor out the tendency




Figure 5.7 Measured data on mutual coupling.

of the driving waveform to slow down when confronted with a high loop inductance. This effect shows in the resulting noise plots as a lengthening of the noise pulse at high coupling factors.

POINTS TO REMEMBER:

Returning signal currents generate magnetic fields, which in turn induce voltages in other circuit traces.

The induced noise coupled into adjacent traces falls off with the square of increasing distance.

CROSSTALK IN SLOTTED GROUND PLANES

The crosstalk situation depicted in Figure 5.8 is a classic layout mistake called a *ground slot*.

Ground slots happen when a layout engineer runs out of room on the regular routing layers and decides to cram in a trace on the ground plane layer. This is done by cutting a long slot in the ground plane and laying the trace in the slot. Ground slots add inductance to traces passing perpendicularly over the slot, and increase crosstalk. Do not tolerate this practice.




Figure 5.8 Crosstalk in a slotted ground plane.




Figure 5.9 Ground slot caused by improper connector layout.

Ground slots also happen on dense backplanes which pass through fields of connector pins. Always make sure the ground clear-outs around each pin have ground continuity between all pins (Figure 5.9).

As shown in Figure 5.8, the return current from driver at A cannot flow directly under trace $A-B$. Instead, it diverts around the ends of the ground slot.

The diverted current makes a large loop, dramatically increasing the inductance of signal path $A-B$, which slows the rise time of signals received at B .

The diverted current also overlaps heavily with the current loop formed by trace $C-D$ and its return-current path. This overlap leads to a large mutual inductance between the signal trace $A-B$ and trace $C-D$.

The effective inductance in series with trace $A-B$ is

$$L \approx 5D \ln\left(\frac{D}{W}\right) \quad [5.3]$$

where L = inductance, nH

D = slot length (perpendicular extent of current diversion away from signal trace), in.

W = trace width, in.

The slot width (length of exposed signal trace passing across the slot) has almost no bearing on the signal trace inductance. A slot of any width, no matter how thin, causes current diversions around the ends of the slot. Any slot width from zero up to as big as the slot length will have the same effect.

If the trace lies offset toward one end of the slot, the inductance is less. Slots as small or smaller than the trace width have almost no effect. Slots near but not overlapping a trace have little effect.

The amount of rise-time degradation caused by the slot inductance varies, depending on the termination conditions used. The worst case is with a long line, for which the apparent source resistance on either side of the inductance is Z_0 . The resulting 10–90% rise time of the L/R filter thus formed is

$$T_{0-90\%L/R} = 2.2 \frac{L}{2Z_0} \quad [5.4]$$

Combine this rise time with the natural signal rise time using the square root of sum of squares rule:

$$T_{\text{composite}} = \left[(T_{0-90\%L/R})^2 + (T_{0-90\% \text{ signal}})^2 \right]^{1/2} \quad [5.5]$$

For a short line driving a heavy capacitive load C , the 10–90% rise time (assuming critical damping) is

$$T_{10-90} = 3.4(LC)^{1/2} \quad [5.6]$$

Such a circuit might ring. The Q of this circuit depends on R_S , the source resistance of the driver:

$$Q = \frac{(L/C)^{1/2}}{R_S} \quad [5.7]$$

When Q is greater than 1, the circuit rings. When Q is near 1, the circuit rise time follows Equation 5.6. For Q less than 1, the circuit rise time is slower than Equation 5.6.

If a second trace close to the first trace also intersects the same slot, the two traces couple tightly together. The mutual inductance L_M between it and the first trace is the same as L in Equation 5.3.

If the second trace lies offset toward one end of the slot, its mutual inductance with the first trace decreases linearly with the distance between it and the slot's end.

The cross-coupling voltage between the two traces can be figured from knowledge of their mutual inductance and the time rate of change in current in the driving circuit:

$$V_{\text{crosstalk}} = \frac{\Delta V}{T_{10-90}} L_M \quad [5.8]$$

For a long line, the ΔV is just equal to the drive voltage divided by the characteristic impedance:

$$V_{\text{crosstalk}} = \frac{\Delta V}{T_{10-90} Z_0} L_M \quad [5.9]$$

For a short line driving a heavy capacitive load C , the time rate of change in current is a second derivative of voltage:

$$V_{\text{crosstalk}} = \frac{1.52 \Delta VC}{(T_{10-90})^2} L_M \quad [5.10]$$

Equations 5.4–5.10 apply equally well for inductance caused by any interruption in a ground plane.

POINTS TO REMEMBER:

Slots in a ground plane create unwanted inductance.

Slot inductance slows down rising edges.

Slot inductance creates mutual inductive crosstalk.

5.4 CROSSTALK IN CROSS-HATCHED GROUND PLANES

The *power and ground grid* diagrammed in Figure 5.10 saves board area, but at the expense of increased mutual inductance. This technique does not require separate power and ground plane layers. You may connect ordinary signals on the same layers as the power and ground connections. It is appropriate for small low-speed CMOS and ordinary TTL designs but provides inadequate grounding for high-speed logic.




Figure 5.10 Power and ground grid on two layers.

In the ground plane grid scheme, ground lines lie in a horizontal pattern on the bottom of a board, while power traces lie in a vertical pattern on the top. Connecting the two sets of lines at every intersection with a bypass capacitor makes a cross-hatched pattern. Current returns equally well to its source along either the ground or power wiring.

The bypass capacitors used in this system must be particularly good, as some returning current traverses several bypass capacitors on its way back to a driving gate.

The open pattern of power and ground traces leaves plenty of room for routing other signals on the power and ground layers. After completing power and ground connections, horizontal routing channels remain on the ground side and vertical routing channels remain on the power side of the board. If you must use a two layer board, this is a good way to do it.

A related layout pattern is called *cross-hatched ground plane*. This pattern, entirely routed on one layer, consists of vertical and horizontal traces covering the board. The cross-hatched ground plane connects only to ground. No other signals may be routed on this layer.

The cross-hatched ground helps implement high-impedance transmission structures on a thin board. Sometimes on a thin dielectric the narrow width required to implement a satisfactory impedance is too small to fabricate reliably. In this case, a cross-hatched ground plane pattern etched into the ground layer adds series inductance and reduces

shunt capacitance, thus raising the line's characteristic impedance. Don't try to implement controlled impedance lines on a cross-hatched ground plane unless the lines run at 45 degrees to the hatch direction. The hatches must be much smaller than the length of a rising edge for this approach to work.

Both the power and ground grid and cross-hatched ground plane layouts induce a lot more mutual inductance between traces than solid ground planes. The question is, will your circuit work with that much mutual inductance?

First let's estimate the self-inductance of a single trace running across a cross-hatched ground plane. This estimation applies equally well to a power and ground grid layout:

$$L \approx 5Y \ln \left(\frac{X}{W} \right) \quad [5.11]$$

where L = inductance, nH
 X = hatch width, in.
 W = trace width, in.
 Y = trace length, in.

If the trace lies offset near one cross-hatched line, the inductance is a little less. Cross-hatched patterns that are as small or smaller than the width of a trace have almost no effect.

If a second trace, close to the first trace, runs between the same two cross-hatched members, the two traces couple tightly together. The mutual inductance L_M between it and the first trace is the same as L in Equation 5.11.

If the second trace is offset by a good distance D , its mutual inductance with the first trace decreases with a denominator similar to Equation 5.2, but with the cross-hatched dimension X replacing the term H :

$$L_M \approx \frac{5Y \ln(X/W)}{1 + (D/X)^2} \quad [5.12]$$

Use the formulas from Section 5.3 for computing rise-time degradation and crosstalk voltage from this self-inductance and mutual inductance.

POINT TO REMEMBER:

If you must work with only two layers, use the power and ground grid system.

5.5 CROSSTALK WITH POWER AND GROUND FINGERS

The *power and ground fingers layout*, diagrammed in Figure 5.11, like the power and ground grid, allows some mutual inductive coupling but saves even more board area.




Figure 5.11 Ground fingers layout.

This old layout appears in old computer equipment (like the PDP-8) manufactured before the era of FCC-mandated radiation guidelines. The power and ground fingers layout is also used on cheap wire wrap frames. Don't use it.

The power and ground fingers technique works only with very-low-speed logic implemented on small circuit cards. It's main benefit is that the power and ground wiring can be implemented on a single layer. Signal traces require a second layer.

In the power and ground fingers scheme, one ground trace lies on the right side of the board and one power trace lies on the left. These traces extend from left to right when needed, like long fingers or rungs of a ladder.

Integrated-circuit packages straddle the rungs, attached with short connections to ground or power traces. Bypass capacitors between adjacent power and ground wires complete the picture.

The problem with this layout is that most returning signal currents must go all the way around the edge of the board to get back to their driver. This diversion introduces massive amounts of self-inductance and mutual inductance.

If you must work with only two layers, use the power and ground grid approach in Section 5.4. If some unknown force compels you to use the ground fingers layout, first build up a sample board and measure the mutual inductance between traces. Then calculate whether your circuit has a chance of working. It might work with low-speed CMOS

logic or with the old LS-TTL series, but not with any fast logic families. In addition to the danger of the product simply not functioning, electromagnetic radiation from the open current loops will almost certainly fail FCC radiation tests.

Here is an approximation for loop inductance on a power and ground fingers layout:

$$L \approx 5Y \ln \left(\frac{X}{W} \right) \quad [5.13]$$

where L = inductance, nH

X = board width, in.

W = trace width, in.

Y = trace length, in.

Notice that doubling the trace width has almost no effect on the overall inductance. Fat ground traces do not help; what is needed is a web of smaller ground traces covering the surface of the board.

If a trace lies offset to one side, the inductance is a little less.

Because returning currents divert around the outside edge of the board, magnetic fields are everywhere. Any second trace intersects these fields, coupling the traces tightly together. The mutual inductance L_M between any two traces is practically the same as L in Equation 5.13. There is not much of a decrease in mutual coupling with distance.

Use the formulas from Section 5.3 for computing rise-time degradation and crosstalk voltage from this self-inductance and mutual inductance.

POINT TO REMEMBER:

For high-speed-logic, avoid the ground fingers layout.

5.6 GUARD TRACES

Guard traces appear extensively in analog design. At audio frequencies, on a two layer board having no solid ground plane, a pair of grounded traces running parallel to a sensitive input circuit can reduce crosstalk by an order of magnitude.

In the digital world, a solid ground plane provides most of the benefits of grounded guard traces. Beyond the solid ground plane, guard traces provide little additional benefit.

As a rule of thumb, the coupling between microstrips is halved by inserting a third line, grounded at both ends, between them. Their coupling is halved yet again if the third line connects through vias to the local ground plane at frequent intervals.¹ If you have

¹J. A. Coekin, *High-Speed Pulse Techniques*, Pergamon Press, Oxford, 1975, pp. 203-205.

more than one ground plane layer, then ground the guard trace at each end, but not in the middle. In digital problems, if two traces lie separated far enough to permit introduction of a guard trace, the coupling is usually already low enough that guarding is unnecessary. See Example 5.1.

EXAMPLE 5.1: Guard Trace Calculations

Two traces, shown in Figure 5.12, lie separated by three trace widths. That is just enough room to fit in a guard trace.




Figure 5.12 Guard trace positioning.

What is the estimated crosstalk? Using Equation 5.1, the crosstalk fraction can't be any worse than

$$\text{Crosstalk} < \frac{1}{1 + (D/H)^2} \tag{5.14}$$

The centerline separation is 0.040, and the trace height is 0.005, so the ratio D/H is 8.

$$\text{Crosstalk} < \frac{1}{1 + (8)^2} = 0.015 \tag{5.15}$$

This is not enough crosstalk to worry about in a digital system.

How much crosstalk is too much? In analog systems, high-powered signals which cross over to low-level inputs require very high crosstalk immunity. Heterogeneous digital systems that mix logic families are sensitive to crosstalk when higher voltage signals (like TTL) get near lower-voltage parts (like ECL).

For ordinary homogeneous digital systems, a crosstalk level of 1–3% between adjacent wires is fine. This assumes there exists a solid ground plane such that each wire interacts only with its nearest neighbors. Cross-coupling from other more remote wires is negligible. Using a hatched or fingers ground system, where many wire pairs interact, we must sum all the crosstalk contributions before arriving at the composite crosstalk level on a given signal.

Figure 5.13 illustrates a typical guard trace application. The driver sends a known voltage step down trace A. Crosstalk from this signal can be received on trace B or trace C. The traces are 26 in. long, with a characteristic impedance of 50 Ω.




Figure 5.13 Guard trace demonstration.

The various step responses for the system of microstrips appear in Figure 5.14. The large impulse is the crosstalk between wires A and B, with C left disconnected. The middle impulse is crosstalk from A to C, with B left disconnected. It is four times smaller than the A-to-B interference, as predicted by Equation 5.2. With trace B shorted to ground




Figure 5.14 Example showing guard trace effect on coupling.

at both ends, we get the smallest coupling from A to C. This coupling is about one-half as big as the middle trace. This halving of coupling is the guard trace effect.

POINT TO REMEMBER:

A solid ground plane provides most of the benefit of grounded guard traces.

ARE-END AND FAR-END CROSSTALK

The crosstalk examples in Sections 5.1–5.6 use lumped-circuit analysis. This simple mutual inductive coupling model works well for many coupling problems but not for long lines.

This section treats coupling between two long distributed transmission lines. Such coupling involves both mutual inductance and mutual capacitance.

5.7.1 Inductive Coupling Mechanism

In this section we consider only the mutual inductive coupling. Section 5.7.2 discusses mutual capacitive coupling. B. L. Hart presents a somewhat more mathematical description of this same material.²

Figure 5.15 illustrates a typical crosstalk situation. The ends of this system are marked “near” and “far,” as is common in the language of long-wire crosstalk.

Wire A-B carries a signal whose magnetic fields induce voltages in wire C-D. Magnetic coupling (mutual inductance) normally acts like a transformer. Because the mutual inductance is distributed, it appears as a succession of small transformers connected between the two lines.

Assuming the coupling is small (it had better be), the transformers do not significantly affect the propagation of signals from A to B.

As a voltage step moves from A to B, at each coupling transformer a small blip of interference appears on the adjoining line. Each blip propagates both forward and backward along line C-D.

For the moment, let's just consider the blip caused by transformer k . When a positive step arrives from A, the changing current induces a momentary voltage across transformer k , as shown in Figure 5.15. This voltage blip is the reaction of inductor k to a change in its current:

$$L_M = L \frac{di}{dt} \quad [5.16]$$

²B. L. Hart, *Digital Signal Transmission Line Circuit Technology*, Van Nostrand Reinhold, New York, 1988.




Figure 5.15 Mutual coupling between two long transmission lines.

The transformer reproduces this voltage blip on line C-D, with the polarity as marked.

The interesting thing about mutual inductive coupling is that the polarity at each end of transformer k is different. A positive blip of voltage goes down line C-D to the left (reverse coupling), while a negative blip goes to the right (forward coupling).

The reflection diagram in Figure 5.16 shows the blips from all transformers adding in a curious pattern. The negative (forward) blips all arrive together at the far end. The positive (reverse) blips arrive at different times, requiring a total of $2T_p$ to get to the near end.

Let's work with just the forward crosstalk coefficients for a moment. Each of the forward blips is proportional to the derivative of the input signal and to each mutual inductance L_M . Since all the forward blips arrive at the far end simultaneously, the total forward-blip size is proportional to the total mutual inductance between the two lines. If the line were stretched longer, the total mutual inductance would increase and so would the forward mutual inductive crosstalk.

Reverse mutual inductive coupling is different. The total amount of coupling (total area) is the same as for forward coupling, but it spreads out over a period of $2T_p$. In actual practice, all the reverse blips smooth out into one continuous blob of reverse coupling. The ideal reverse step response due to mutual inductive coupling is a rectangular function as shown in Figure 5.17.

If the line were stretched longer, the total mutual inductance would increase. The reverse coupling would respond by increasing its duration but not its height.

5.7.2 Capacitive Coupling Mechanism

Distributed mutual capacitance works almost the same way as distributed mutual inductance. The difference lies in the polarity of couplings.

When a voltage step passes one of the mutual capacitors in Figure 5.15, a small blip of interference appears on the adjoining line. Each blip propagates both forward and backward along the line C-D.




Figure 5.16 Reflection diagram showing mutual inductive coupling from the four transformers shown in Figure 5.15.

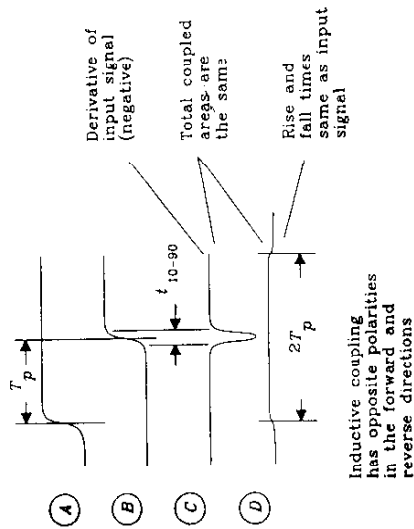



Figure 5.17 Forward and reverse mutual inductance coupling (distributed).
Inductive reverse coupling = $K[V(t) - V(t - 2T_p)]$

The polarities of mutual capacitive coupling are positive for forward and backward blips. Other than that, they behave the same as mutual inductive interference blips.

The mutual capacitance forward coupling looks like the derivative of the input signal. It gets larger for longer lines. Its polarity is positive (opposite of forward mutual inductive coupling).

The mutual capacitance reverse coupling area is the same as for forward coupling, but it spreads out over a period of $2T_p$. The ideal reverse step response due to mutual capacitive coupling is a rectangular function as shown in Figure 5.18.



Capacitive coupling has the same polarities in the forward and reverse directions.

Figure 5.18 Forward and reverse mutual capacitive coupling (distributed).
Capacitive reverse coupling = $K[V(t) - V(t - 2T_p)]$

5.7.3 Combining Mutual Inductive and Mutual Capacitive Coupling

Under normal conditions, over a solid ground plane, the inductive and capacitive crosstalk voltages are of roughly equal size. The forward crosstalk components (voltage at D) cancel, while backward crosstalk components (voltage at C) reinforce.

Strip-line circuits show particularly good balance between inductive and capacitive coupling and have tiny forward-coupling coefficients. Microstrips, for which the electric field lines responsible for crosstalk travel mostly through air instead of through the dielectric, have somewhat less capacitive crosstalk than inductive, leading to a small negative forward-coupling coefficient.

Over a slotted, hatched, or otherwise imperfect ground plane, the inductive crosstalk component is much larger than capacitive coupling and the forward crosstalk is large and negative. The forward crosstalk is never larger than the reverse crosstalk.

5.7.4 How Near-end Crosstalk Becomes a Far-end Problem

In Figure 5.15, the forward and reverse coupling signals are different. Each signal propagates to its respective end of cable C-D and extinguishes at the terminator.

Practical applications are often different from this model. In digital applications without source terminations, the device connected to the left end in Figure 5.19 is a low-impedance driver. Just like any other signal, reverse crosstalk, when it hits this driver, reflects. The reflection coefficient for a low-impedance driver is almost -1 . This changes the reverse coupling from positive polarity to negative polarity and sends it back toward the far end.

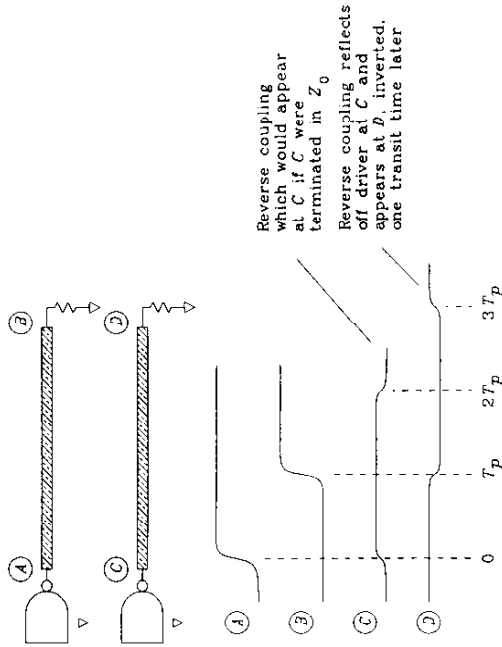


Figure 5.19 Reverse coupling reflecting from a low-impedance driver.

The signal seen at the far end D is a copy of the reverse coupling signal at C, delayed by one propagation time and inverted.

Because the mutual inductive and mutual capacitive portions of the forward coupling nearly cancel each other, the forward coupling is not visible when juxtaposed with this much larger reflected reverse coupling. When we measure crosstalk as defined in Figure 5.20, we are mostly measuring reflected reverse coupling.

EXAMPLE 5.2: REFLECTED REVERSE CROSSTALK

Figure 5.20 shows the measurement setup which produced the reflected reverse crosstalk waveforms in Figure 5.21.

The pulse generator drives trace A-B with a 2.5-V step input having a rise time of 880 ps. The waveform monitored at A appears in the top of Figure 5.21 at a scale of 1 V/division.




Figure 5.20 Setup for reflected reverse crosstalk measurement.




Figure 5.21 Reflected reverse crosstalk measurement.

Crosstalk signals at points D and F connect to regular coax probes and display in Figure 5.21 at 50 mV/division.

All probes are the same length. All probes terminate at the scope in 50 Ω .

The near ends of both pickup traces connect directly to ground, simulating a low-impedance driver.

Both crosstalk signals start together 4.5 ns after the initial rising edge:

$$T_p = 4.5 \text{ ns} \quad [5.17]$$

Both signals have a duration of 9 ns, and are of negative polarity.

$$\text{Crosstalk duration} = 2T_p = 9 \text{ ns} \quad [5.18]$$

The crosstalk signals measured at points D and F are

$$D = (2.5 \text{ divisions})(50 \text{ mV/division}) = 125 \text{ mV} \quad [5.19]$$

$$F = (0.7 \text{ division})(50 \text{ mV/division}) = 35 \text{ mV} \quad [5.20]$$

The crosstalk ratios (output divided by input) measured for these two geometries are

$$\frac{D}{A} = \frac{0.125}{2.5} = 0.05 \quad [5.21]$$

$$\frac{F}{A} = \frac{0.035}{2.5} = 0.014 \quad [5.22]$$

The crosstalk ratios predicted by Equation 5.2 ($K = 1$) for these geometries are

$$\frac{D}{A} \approx \frac{1}{1 + (0.040 / 0.010)^2} = 0.059 \quad [5.23]$$

$$\frac{F}{A} \approx \frac{1}{1 + (0.080 / 0.010)^2} = 0.015 \quad [5.24]$$

5.7.5 Characterizing Crosstalk Between Two Lines

Forward crosstalk is proportional to the derivative of the driving signal, and to line length. The constant of proportion depends on the balance between inductive and capacitive coupling. Once we measure this ratio for one known signal, modeling the response for other signals is trivial.

Modeling reverse crosstalk for fast-rising edges is equally simple. The reverse coupling locks like a square pulse, with rise and fall times comparable to the input signal and height proportional to the driving signal amplitude. The reverse coupling percentage is a physical constant determined by the line parameters which is independent of line length. The duration of the pulse is $2T_p$.

Reverse crosstalk for slow-rising edges is a little more involved. Once we have found the reverse coupling coefficient for fast edges, the reverse signal for any input can be found from

$$\text{Reverse coupling}(t) = \alpha_R [V(t) - V(t - 2T_p)] \quad [5.25]$$

where t = time, s

$V(t)$ = driving waveform, V

α_R = reverse coupling coefficient for fast-edged signal

T_p = propagation delay of line, s

For lines longer than half the signal rise time, the reverse coupling has time to build up to its full value. The reverse coupling coefficient for such a line equals approximately

$$\alpha_R \approx \frac{1}{1 + (D/H)^2} \quad [5.26]$$

where D = separation between lines, in.

H = line height above ground, in.

For lines shorter than half the signal rise time, the reverse coupling ramps up and then back down, never reaching its steady-state maximum value.

5.7.6 Using Series Terminations to Reduce Crosstalk

A series terminator extinguishes reverse-coupled crosstalk at the near end. An end terminator attenuates the returning reflection of the main signal, the reverse coupling from which would be again pointed toward the far end. Using both terminators eliminates both sources of reverse coupling noise, improving overall crosstalk considerably.

The reduced coupling gained by combining series and end terminations lets us route parallel bus traces closer than would otherwise be practical.

POINTS TO REMEMBER:

Regarding long transmission lines:

Over solid grounds, inductive and capacitive crosstalk are equal.

Forward-crosstalk components cancel, while reverse crosstalk reinforces. Over a slotted or imperfect ground plane, the inductive coupling exceeds capacitive coupling, making forward crosstalk large and negative.

Forward crosstalk is proportional to the derivative of the input signal and to line length.

Reverse coupling looks like a square pulse, with a constant height and duration equal to $2T_p$. For short lines, reverse coupling does not climb to its full value.

Reverse crosstalk, when it hits a low-impedance driver, reflects toward the far end.

5.8.1 HOW TO STACK PRINTED CIRCUIT BOARD LAYERS

A *layer stack* for a printed circuit board specifies the arrangement of circuit board layers. It specifies which layers are solid power and ground planes, the dielectric constant of the substrate, and the spacings between layers. When planning a layer stack, also compute the desired trace dimensions and minimum trace spacing.

Manufacturing constraints heavily influence the layer stack. As a rule, the greater our circuit wiring density, the greater your production costs per square inch. This section sets out some basic rules of thumb for planning layer stacks.

5.8.1 Power and Ground Planning

Design the power and ground layers first. To plan a power and ground system, first establish the signal rise times, the number of signals, and the physical dimensions of the circuit board.

Included among the physical dimensions, make a guess as to the trace width. The trace width assumption is not particularly critical at this stage.

Next estimate the self-inductance and mutual inductance using solid, hatched, and fingers ground plane models. At this point it is usually clear which model suits the design. Remember that for the ground fingers mode, all traces interact. For the hatched model, traces running along the same hatch grid interact. For the ground plane model, only adjacent traces interact.

If you will be using a solid ground plane, plan on using ground and power planes in pairs. The symmetric pairing of solid planes in a layer stack helps prevent warping in the circuit board. A board with a single plane, offset to one side, can warp noticeably.

Power planes may be used as low-inductance signal return-current paths just as ground planes. Assuming adequate bypass capacitors between power and ground (Chapter 8), transmission lines routed over a power plane operate as well as those routed over a ground plane. Transmission striplines routed between one power and one ground layer, or two power layers, also work.


5.8.2 Chassis Layer

Sometimes you will want to run a signal outside your digital system. For this application you may pick a low-speed or controlled rise-time driver. This is a good choice because it reduces external radiation, which helps with FCC problems.

If the ground for the driver connects to ordinary digital logic ground, the effective driver output equals its intended drive voltage plus any noise voltage present on the digital logic ground. See Figure 5.22.

Digital logic grounds are notorious for high-frequency noise voltages. Grounds carry fluctuating voltages caused by the action of many returning signal currents acting across their self-inductance. These high-frequency fluctuations are too small to cause trouble in the digital circuits, but plenty large enough to exceed FCC limits. Any wire connected to the digital logic ground which runs outside the cabinet almost always fails FCC tests.

Without other precautions, the controlled rise-time driver effectively picks up ground noise and broadcasts it outside the chassis.



A ground connection at high frequencies between logic ground and true earth ground may be ineffective.

Figure 5.22 Using a controlled rise-time driver.

One solution to this problem adds a *chassis plane* to the layer stack. This plane stacks directly next to a ground plane, giving a very tight capacitive coupling between the two planes. At high frequencies, the two planes are effectively tied together. The chassis plane then screws, solders, or welds to the external chassis along one continuous axis near the controlled rise-time driver. At high frequencies, we have effectively shorted the digital ground plane to the chassis. This reduces the amount of digital ground noise at that point, also reducing noise carried by the controlled rise-time driver to the outside world.

Ordinary capacitors will not function as a short between chassis and digital logic ground, because they have too much lead inductance. Only the large, wide, parallel surface area between the chassis plane and digital ground plane has a low enough inductance to effectively short the two together.

With the chassis plane approach, the digital logic and external chassis remain electrically isolated at low frequencies. This may be desirable for safety or other reasons. If the isolation does not matter, simply short digital logic ground directly to the chassis without using a separate chassis layer. Make this connection by screwing, soldering, or welding the ground plane to the external chassis along one continuous axis near the controlled rise-time driver.

When using a chassis plane, counterbalance it in the layer stack with some other solid plane layer. For mechanical reasons, always lean toward using symmetric arrangements of planes in your layer stack.

5.8.3 Selecting Trace Dimensions

Squeezing traces tightly together increases the circuit packing density. Very dense designs require fewer circuit board layers. Since printed circuit board cost is proportional to the number of layers, as well as to the board surface area, we are tempted to always design using the fewest number of layers that will do the job.

Smaller, more closely spaced traces also yield more crosstalk and less power-handling capacity. This tradeoff among crosstalk, routing density, and power is critical to low-cost product design.

Let's deal with power-handling capacity first, because it is the simplest of the constraints.

The power-handling capacity of a printed circuit trace depends mostly on its cross-sectional area and the allowable temperature rise. For a given cross-sectional area, a trace's temperature rise above ambient is roughly proportional to the power it dissipates. Large temperature rises are unreliable and heat up nearby digital circuits. A conservative upper limit on trace heating inside digital products is 10° C.

Figure 5.23 relates maximum power-handling capacity to temperature rise. The horizontal axis in Figure 5.23 measures cross-sectional area in units of square inches.³ The vertical axis of Figure 5.23 shows the allowable current for that trace at a given temperature rise.

For example, a 0.010-in.-wide trace of 1-oz copper (0.00135 in. thick) can safely pass 750 mA of current at a temperature rise of 10° C.

Power is rarely a serious constraint except for large power-distribution buses. As thin-film technology, with extremely small trace cross sections, becomes more widely available, heating limitations may become more prevalent.

A second lower bound on trace width results from the manufacturing process. Table 5.1 lists the minimum trace widths attainable in various production processes.




Figure 5.23 Current-carrying capacity of copper printed circuit traces.

³Example: A trace 0.010 in. wide, fabricated in a 1-oz-weight copper (0.00135 thickness), has a cross-sectional area of 1.35×10^{-5} in.²

TABLE 5.1 MINIMUM LINE WIDTHS ATTAINABLE WITH VARIOUS PRODUCTION PROCESSES

Process	Minimum line width (in.)
Gold screened onto thick film substrate	0.010
Etched copper on epoxy board with plating	0.004
Etched copper on epoxy board with no plating	0.003
Gold evaporated onto thin film substrate and then etched	0.001

With any process the manufacturing yield will drop, and cost will go up, as one approaches the minimum attainable trace width. This factor prevents most designers from using the minimum attainable line width.

Other factors tend to increase line width. Poor control over the etching process can result in large line width variations. At low line widths, the percentage line width variation, which controls the percentage impedance tolerance, may be unacceptable. A need for accurate impedance control can force the use of lines much wider than the minimum attainable trace width.

Use the formulas for trace impedance in Appendix C to find a combination of trace width and height sufficiently large that, over expected variations in width and layer height, the impedance stays within your design range. Remember that you must also allot room in your impedance budget for variations in the electric permittivity of the substrate.

Considerations of power, cost, and impedance tolerance usually drive the selection of a particular trace width. Given the width, the impedance constraint sets the layer height. Next, using the formula for crosstalk (see Section 5.7.5 and Equation 5.2), figure the minimum spacing between adjacent traces (center to center). This number is called the minimum *trace pitch*. The unused distance between traces is called the *trace separation*. The sum of the trace separation and the trace width equals the trace pitch.

POINTS TO REMEMBER:

As a rule, the greater your circuit wiring density, the greater your production costs per square inch.

Printed circuit board cost is proportional to the number of layers and to the board surface area.

Design the power and ground layers first.

For mechanical reasons, lean toward using a symmetric arrangement of ground and power planes in your layer stack.

Smaller, more closely spaced traces yield more crosstalk.

5.8.4 Routing Density Versus Number of Routing Layers

With more layers, we can spread the traces out farther. That makes routing easier and reduces the risk of crosstalk problems. Unfortunately, the cost of multilayer printed circuit boards is proportional to the number of layers times the surface area. Using more layers costs more.

With fewer layers, we must use a narrower trace pitch, which can also cost extra. Not only that, at a fine enough pitch we run the risk of too much crosstalk.

Determining the least cost number of layers for a board is a matter of experience and guesswork. The central issue involves estimating the trace pitch required by N connections routed on a certain-size board, using M layers. Knowing the trace pitch, we can find out how much the board will cost and at the same time model the crosstalk.

Trace pitch is determined from wiring density. One useful model for wiring density is called Rent's rule, after the IBM engineer who popularized it. Rent noticed that most large square boards, when divided into quadrants, reveal half their wiring going between quadrants and half staying within each quadrant. Further subdividing of each quadrant, reveals the same distribution. If, upon traveling between two quadrants, we assume (this is very hypothetical) a wiring length on average equal to the spacing between quadrants, we arrive at a total average wiring length equal to three-eighths of a board side.

Knowing the average wire length and the number of wires we can compute the total board surface area occupied by those wires using any trace pitch. This knowledge, in the form of Equation 5.27, shows the trace pitch required to route N connections on a fixed-size board using M layers.

Of course, if we have some other information about the routing requirements, such as large buses or other structures, we should use it. Given no other information, we can attempt to calculate the required trace spacing on a board using Rent's idea:

$$P_{\text{avg}} = \frac{(XY)^{\frac{1}{2}}}{N} \quad 2.7M \quad [5.27]$$

where N = number of connections (assumed distributed according to Rent's rule)

P_{avg} = average trace pitch, in.

X = board width, in.

Y = board height, in.

M = number of routing layers

For example, on an 8 in. \times 12 in. board having 800 interconnections routed on four layers, we need a trace pitch, on average, of 0.132 in. This means that if the board is pretty much covered with DIP through-holes, we will need one trace routed between almost every pin.

Don't count on filling more than half the spaces between pins. In the above example we should plan on more layers or plan on using double-track routing (two traces between each pin).

For through-hole boards, the average pitch predicted by Equation 5.27 and the minimum pitch required are very different. Use the minimum pitch determined from crosstalk considerations to determine if you need double- or triple-track routing between pins. Use the average pitch predicted by Equation 5.27 to determine how many available routing tracks we need.

Inner layers of surface-mounted boards may have more routing space available than DIP boards. The total number of vias is roughly the same, but the vias are smaller in surface-mount designs since IC pins don't have to stick through them. The average and minimum pitch for inner layers of surface-mount boards may be similar.

Routing up to four tracks between pins is possible on inner layers of epoxy circuit boards, but this may lead to severe crosstalk problems.

Extra routing space on a board can be gained by leaving large spaces between chips, but this takes more total surface area. Most designers choose to go with more layers instead.

If crosstalk is a problem, ensure that the layout squeezes traces together only to fit between pins and then immediately spreads them back out again on their way to the next chip. This requires a lot of hand adjustment, but even modest increases in trace separation reduce crosstalk.

With any luck, we can find some number of circuit board routing layers that yield acceptable crosstalk and won't cost too much.

POINTS TO REMEMBER:

Don't count on filling more than half the spaces between through-hole pins.


When all else fails, use Rent's reasoning to figure the average trace length.

5.8.5 Classic Layer Stacks

Figures 5.24–5.26 illustrate three classic layout stacks for 4, 6, and 10 layers. These stacks are designed for use with the ordinary epoxy multilayer fabrication process described below. Beyond 10 layers, designers usually incorporate additional ground planes to isolate sets of routing layers.

These stacks are appropriate for high-speed computer products embedded within well-shielded card cages. If you are planning systems that must pass FCC, VDE, Tempest, or other mandated rules for electromagnetic emissions without the benefit of a well-shielded card cage, these simple stacks will be inadequate for your purposes.

In each figure, the terms *horizontal routing* and *vertical routing* refer to the orientation of traces on that layer (horizontal or vertical). Traces on each layer traditionally lie



Comments:
(1) Easy to produce due to large trace widths. Good control over impedance.
(2) Ratio D/H for crosstalk is 5.0.

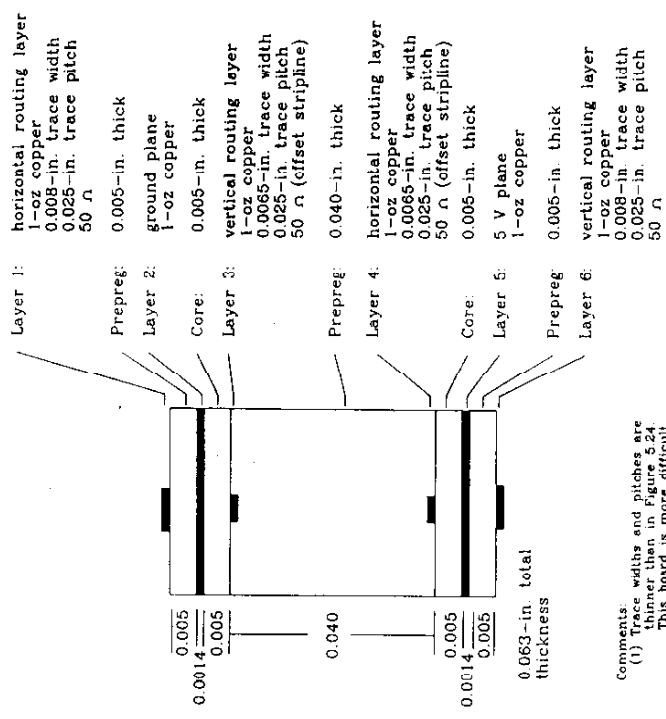
Figure 5.24 Four-layer stack.

parallel to each other, with every other layer built at right angles to the one below it. Very few traces cut diagonally across a layer or make large 90-degree bends. This increases the routing efficiency.

Power and ground layers in Figures 5.24–5.26 are marked with thick solid lines. Trace layers show the proportional line width and trace height. The notations *core* and *prepreg* refer to materials used in the substrate lamination process. The following paragraphs briefly describe one commonly used process for building circuit cards. If you want to tightly control trace-to-ground spacings, you need to know about core and prepreg layers.

The multilayer buildup process starts with a set of raw two-sided laminate layers coated with copper on each side. Surfaces destined to become inner layers are etched in this form. Surfaces destined to become outer layers are left fully copper-plated. These etched laminates are called cores. The thickness between opposing layers on a core depends on the thickness of the original laminate.

The cores are then stacked together, with a sheet of prepreg epoxy material placed between each pair of cores. This sheet melts into an epoxy glue when heated and pressed. The thickness of the prepreg sheet determines the spacing between core layers. The prepreg cures into a hard epoxy layer having the same dielectric constant as the core layers. Core and prepreg layers alternate.



Comments:
(1) Trace widths and pitches are thinner than in Figure 5.24. This board is more difficult to produce.
(2) Thin traces, combined with prepreg, make impedance more difficult to control.
(3) The trace pitch is twice as dense as in Figure 5.24, and this board has twice the number of routing layers. Overall, the routing density is four times that of Figure 5.24.
(4) The D/H ratio for crosstalk is 5.0.

Figure 5.25 Six-layer stack.

Because the prepreg partially melts in the process, traces on opposite sides of the prepreg tend to sink into the melted, gluey prepreg material. Accurate analysis of trace-to-ground separation takes into account the fact that the spacing between facing sides of two traces on either side of a prepreg layer will be reduced by twice the trace thickness as they sink into the prepreg. Ground layers do not sink into the prepreg.

Manufacturers sometimes form the outer layer from one side of a core, and in other cases they form it from a metal foil pressed onto a top layer of prepreg. In either case, the outer layer is a solid copper sheet (no etching yet).

After the prepreg cures, holes are drilled through the assembly. The drilling exposes the various copper layers and pads which it penetrates in the inner layers, but they are at this point not connected electrically.

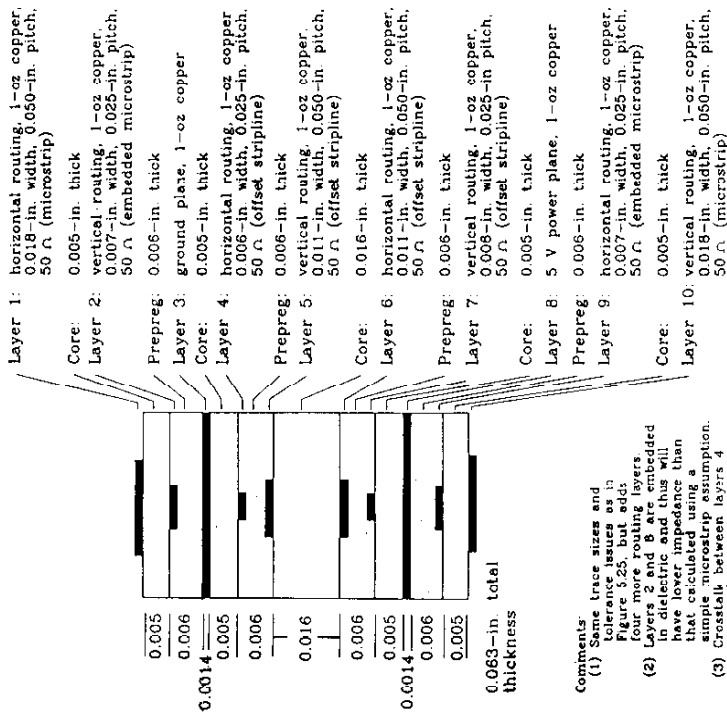


Figure 5.26 Ten-layer stack.

The plating step covers the inside surfaces of the holes and the outer surfaces of the board simultaneously. To save plating material and time, most manufacturers mask off the outer surfaces except around holes and the outer-layer traces. After plating, the outer-layer traces come out a little thicker than the raw copper sheet. It is this additional thickness which causes greater uncertainty in the finished line widths of outer layers compared to inner layers.

The last step etches away unwanted copper on the outer layers, leaving a finished board. The board is then tinned (optional), coated with solder mask, and silk-screened on both sides.

POINTS TO REMEMBER:

- Core and prepreg layers alternate.
- Outer layers, if plated, have greater trace width variation than inner traces.
- Traces on routing layers tend to sink into the prepreg mixture. Their thickness doesn't add to the total board thickness.
- The thickness of solid ground plane layers always adds to the total board thickness.

5.8.6 Extra Hints for High-speed Boards

For the very-highest-speed boards, place power and ground layers directly together. This positioning maximizes their capacitive coupling, reducing power supply noise. Use plenty of extra ground planes (not power planes) to isolate sets of routing layers. Sprinkle around lots of ground vias, connecting together the many ground planes. Returning signal currents will flow through these ground vias as they jump from layer to layer following the contorted path of each signal trace.

Had we used a mix of power and ground planes for isolating routing layers, instead of just ground planes, the return currents, which always flow in the nearest plane, would have had to traverse many bypass capacitors as they jumped between ground and power planes. This is not a good idea, because any currents we send through the bypass capacitors will cause voltages across them. These voltages radiate very effectively from the power and ground planes, adding to our radiated noise problems.

POINTS TO REMEMBER:

- At the highest speeds, keep ground and power planes directly adjacent.
- Use extra ground planes, not power planes, to isolate routing layers.