High-Speed-Board Design Techniques

Advanced Micro Devices

Application Note

INTRODUCTION

The most important factor in the design of many systems today is speed. 66-MHz processors are common; 90- and 120-MHz processors are becoming readily available. The demand for high speed results from: a) the requirement that systems perform complex tasks in a time frame considered comfortable by humans; and b) the ability of component manufacturers to produce high-speed devices. An example of a) is the large amount of information that must be processed to perform even the most rudimentary computer animation. Currently, Programmable Array Logic (PAL) devices are available with propagation delays of 4.5 ns, and complex PLDs such as MACH have propagation delays of 5 ns. While this might seem fast, it is not the propagation delay that creates the potential for problems, but rather the fast edge rates needed to obtain the fast propagation delays. In the future, much faster devices will become available, with correspondingly faster edge rates.

Designing high-speed systems requires not only fast components, but also intelligent and careful design. The analog aspect of the devices is as important as the digital. In high-speed systems, noise generation is a prime concern. The high frequencies can radiate and cause interference. The corresponding fast edge rates can result in ringing, reflections, and crosstalk. If unchecked, this noise can seriously degrade system performance.

This application note presents an overview of the design of high-speed systems using a PC-board layout. It covers:

- the power distribution system and its effect on board-noise generation,
- transmission lines and their associated design rules,
- crosstalk and its elimination, and
- electromagnetic interference.

1. POWER DISTRIBUTION

The most important consideration in high-speed board design is the power distribution network. For a noise-free board, it is necessary to have a noise-free power distribution network. Note that it is just as important to develop a clean V_{CC} as it is to get a clean ground. For AC purposes, which is what this application note mainly discusses, V_{CC} *is* ground.

The power distribution network also must provide a return path for all signals generated or received on the board. This is often overlooked because the effect of the return path is less apparent at lower frequencies. Many designs work even when the nature of the return path is ignored.

1.1 Power Distribution Network as a Power Source

1.1.1 The Effect of Impedance

Consider a 5" x 5" board with digital ICs and a power supply of +5.0 V. The goal is to deliver exactly +5.0 V to the power pins of every device on the board, regardless of its position relative to the power source. Furthermore, the voltage at the pins should be free of line noise.

A power source with these characteristics would be schematically represented as an ideal voltage source (Figure 1a), which has zero impedance. Zero impedance would ensure that the load and source voltages would be the same. It also would mean that noise signals would be absorbed because the noise generators have finite source impedance. Unfortunately, this is only an ideal.

Figure 1b illustrates a real power source with associated impedances in the form of resistance, inductance, and capacitance. These are distributed over the power distribution network. Because of the network's impedance, noise signals can add to the voltage.

The design goal is to reduce the power distribution network impedances as much as possible. There are two approaches: power buses and power planes. Power planes generally have better impedance characteristics than power buses; however, practical considerations might favor buses.

Publication# **16356** Rev. **B** Amendment **/0**
Issue Date: **February 1996**

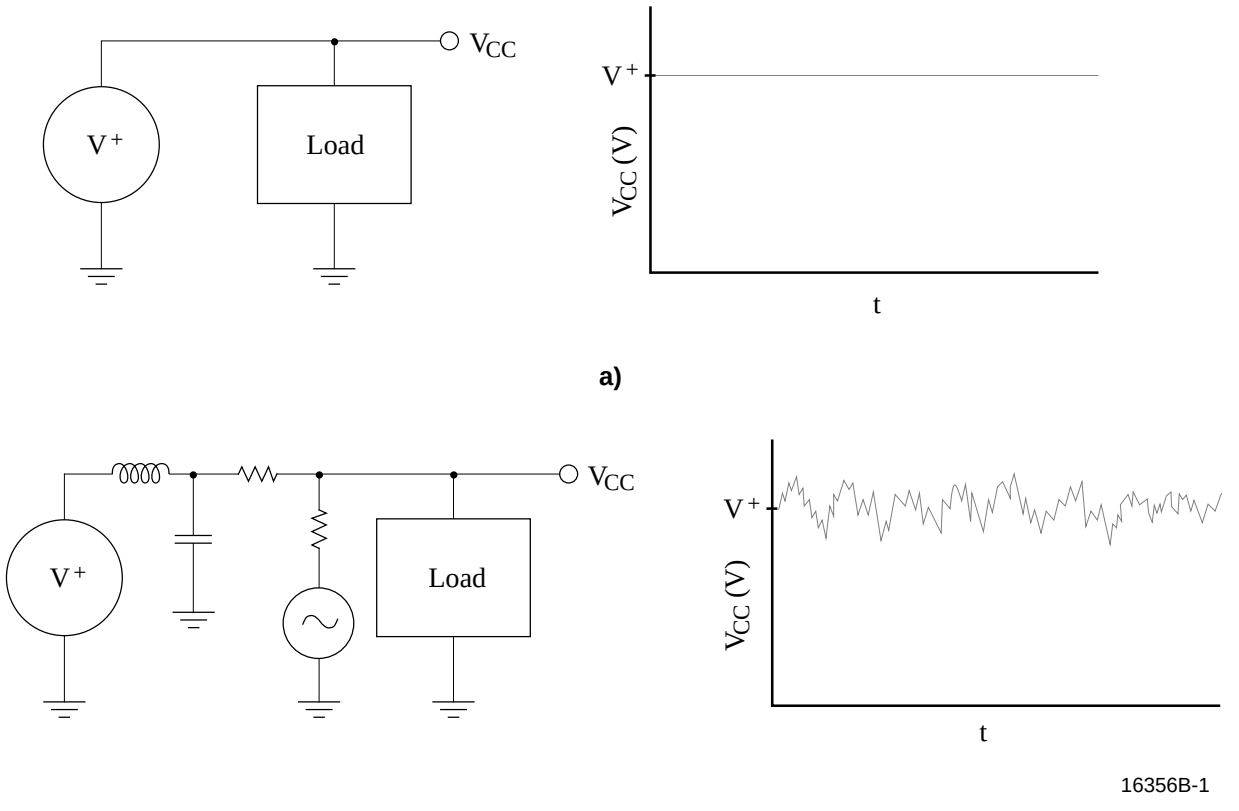

Figure 1. The Power Source. a) Ideal Representation; b) More Realistic Representation

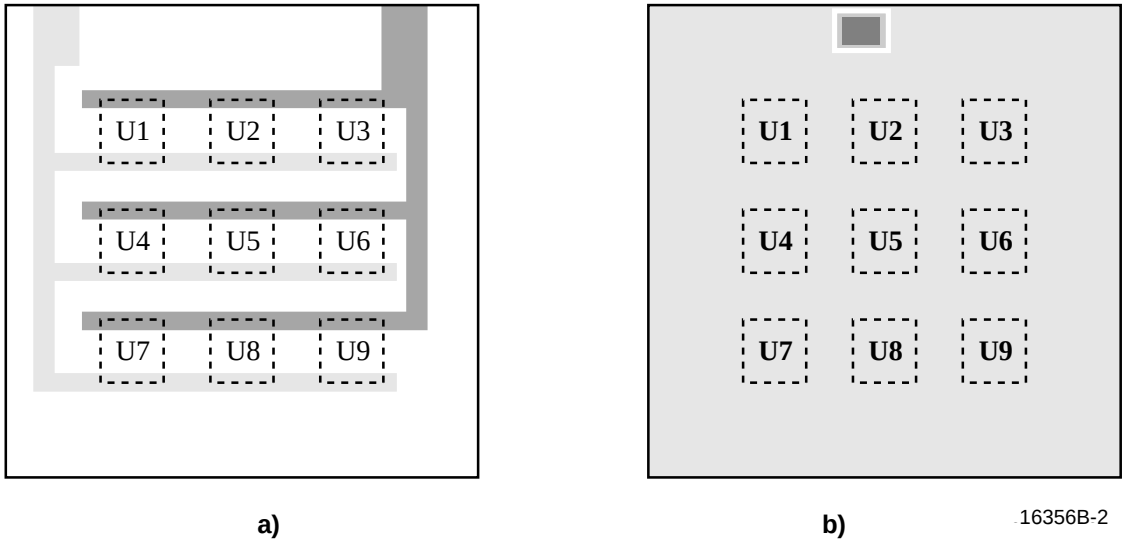

Figure 2. Power Distributions System. a) Power Buses; b) Power Planes

1.1.2 Power Buses vs Power Planes

Two power-distribution schemes are shown in Figure 2. A bus system (Figure 2a) is composed of a group of traces with the various voltage levels required by the system devices. For logic, these are typically +5 V and ground. The number of traces required for each voltage level varies from system to system. A power-plane system (Figure 2b) is composed of entire layers (or sections of layers) covered with metal. Each voltage level requires a separate layer. The only gaps in the metal are those needed for placing pins and signal feed-throughs.

Early designs favored buses because of the expense of devoting entire levels to power distribution. The power

bus shares layers with the signal lines. The bus must supply power to all devices, while leaving room for the signal traces; therefore, buses tend to be long, narrow ribbons. This results in a relatively small cross-sectional area with a small resistance.

Although the resistance is small, it is significant. Even a small board can have 20 to 30 devices on it. If each device on a 20-device board sinks 200 mA, the total current would be 4 A. A bus resistance of only 0.125 Ω has a 0.5 V drop. With a 5 V power supply, the last device on the bus might receive only 4.5 V.

Because the power plane fills an entire layer, the only area constraints are the dimensions of the board. The resistance of a power plane is a small fraction of that of a power bus supplying the same number of devices. Thus a power plane is more likely than a bus to supply full power to all the devices.

On a bus, currents are restricted to paths defined by the bus. Any line noise generated by a high-speed device is introduced to other devices on that power bus. On the board in Figure 2a, noise generated by U9 is sent to U7 by the bus.

On the power plane, the noise currents are distributed because the current path is not restricted. This, along with lower impedance, makes power planes quieter than power buses.

1.1.3 Line Noise Filtering

The power plane alone does not eliminate line noise. Since all systems generate enough noise to cause problems, regardless of the power distribution scheme, extra filtering is required. This is done with bypass capacitors. Generally, a 1 μF to 10 μF capacitor is placed across the power input to the board, and 0.01 μF to 0.1 μF capacitors are placed across the power and ground pins of every active device on the board.

The bypass capacitor acts as a filter. The larger capacitor (≈ 10 μF) is placed across the power input of the board to filter lower frequencies (like the 60-Hz line frequency) that usually are generated off the board. Noise generated on the board by the active devices have harmonics in the range of 100 MHz and higher. A bypass capacitor is placed across each chip and generally is much smaller (≈ 0.1 μF) than the capacitor across the board.

Since the goal is to filter out any AC component on the power supply, it might seem initially that the largest possible capacitor is the best, minimizing the impedance as much as possible. However, this does not take into account that real capacitors do not have ideal characteristics.

A capacitor, which is ideally represented in Figure 3a, is more realistically represented by Figure 3b. Resistance and inductance are the result of the construction of the plates and the leads necessary to build the capacitor. Because the parasitic components are effectively in series with the capacitance, they are called equivalent-series resistance (ESR) and equivalent-series inductance (ESL).

Thus the capacitor is a series resonant circuit for which

$$f_R = \frac{1}{\sqrt{LC}}$$

As shown in Figure 4a, it is capacitive at frequencies below f_R, and inductive at frequencies above f_R. As as result, the capacitor is more a band-reject filter than a high-frequency-reject filter.

As an example, the 10 μF capacitors used for the board-power connections are typically made with rolls of metal foils separated by an insulating material (Figure 5). This results in large ESLs and ESRs. Because of the large ESLs, f_R is generally less than 1 MHz. They are good filters for 60-Hz noise, but not good for the expected 100-MHz and higher switching noise.

The ESL and ESR result from the construction of the capacitor and dielectric material used, rather than from capacitance value. The high-frequency reject capabilities cannot be improved by replacing a capacitor with a larger one of the same type. The impedance of a large capacitor is smaller than that of a small capacitor at frequencies below the f_R of the small capacitor. But at frequencies above f_R, the ESL dominates and there is no difference between the impedance of the two capacitors (Figure 4b). This is because only the capacitance has changed; unless the construction is changed, the ESL remains essentially unchanged. To improve high-frequency filtering, one must replace the capacitor with a type that has a lower ESL.

Various types of capacitors are available for specific frequencies and applications. Table 1 gives a small overview of some available device types.

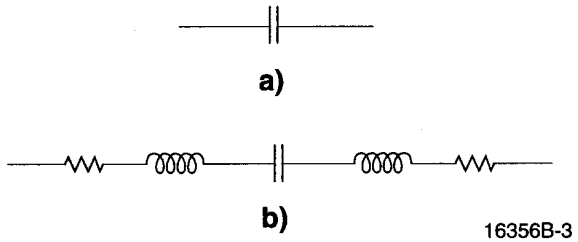

Figure 3. a) Ideal Representation of a Capacitor b) Parasitic Components added to Emulate Real Conditions

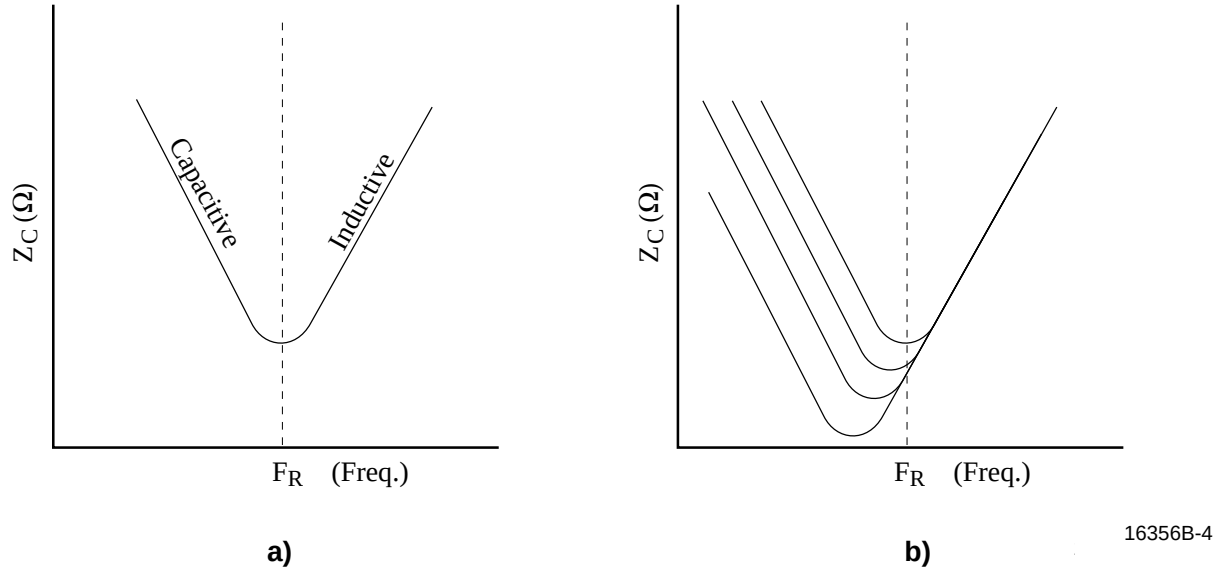

Figure 4. a) Capacitor Impedance Versus Frequency; b) the Effect of Lowering Capacitance While Using the Same Type of Construction (Constant ESL)

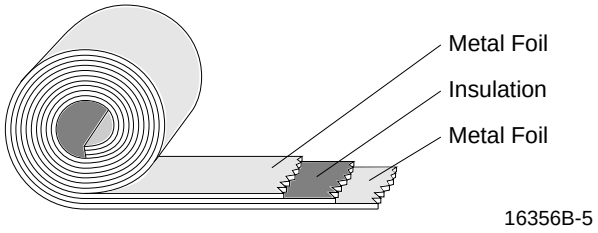

Figure 5. Internal Construction of a Large (> μF) Capacitor

The lowest ESL capacitors often are made with non-ferromagnetic materials, which have a low voltage-capacitance product. Thus it is difficult to make large capacitors with practical breakdown voltages to prevent board failure. However, because of better filtering characteristics, larger values might not be needed. Figure 6 compares a 0.01 μF capacitor of type C0G (non-ferromagnetic) to a 0.1 μF capacitor of another type. Note that the 0.01 μF capacitor gives better filtering at higher frequencies.

The capacitor graphs imply that any one capacitor has a limited effective frequency operating range. Because systems have both high- and low-frequency noise, it is desirable to extend this range. This can be done by putting a high-capacitance, low-ESL device in parallel with a lower-capacitance, very-low-ESL device. Figure 7 shows that this can significantly increase the effective filtering frequency range.

1.1.4 Bypass Capacitor Placement

After the filter capacitors have been chosen, they must be placed on the board. Figure 8a shows the standard placement for boards with slow device speeds. The capacitor is placed near the top of the device to help ensure its accessibility. While simple for layout, this does not give the best high-speed performance.

Note that the V_{CC} capacitor connection is quite close to the chip's V_{CC} connection, but the ground connection is far away. Because noise is not uniform on a power plane, the capacitor is not filtering noise at the chip leads; it is only filtering noise near the chip.

Table 1. Bypass Capacitor Groups

Type	Range of Interest	Application
Electrolytic	1 μF to > 20 μF	Commonly used at power-supply connection on board.
Glass-Encapsulated Ceramic	0.01 μF to 0.1 μF	Used as bypass capacitor at the chip. Also often placed in parallel with electrolytic to widen the filter bandwidth and increase the rejection band.
Ceramic-Chip	0.01 μF to 0.1 μF	Primarily used at the chip. Also useful where low profile is important.
C0G	< 0.1 μF	Bypass for noise-sensitive devices. Often used in parallel with another ceramic chip to increase rejection band.

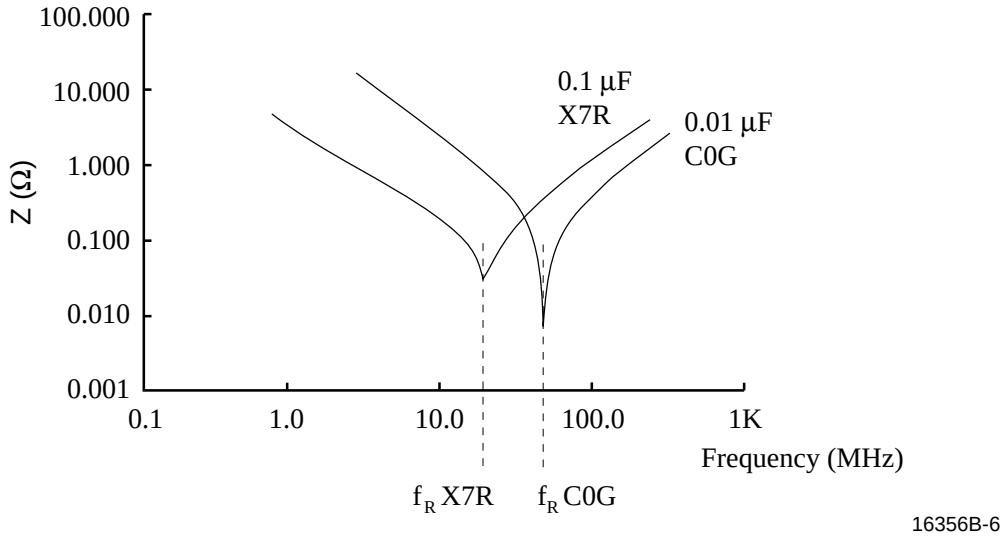

Figure 6. Frequency Response of X7R and C0G Type Construction

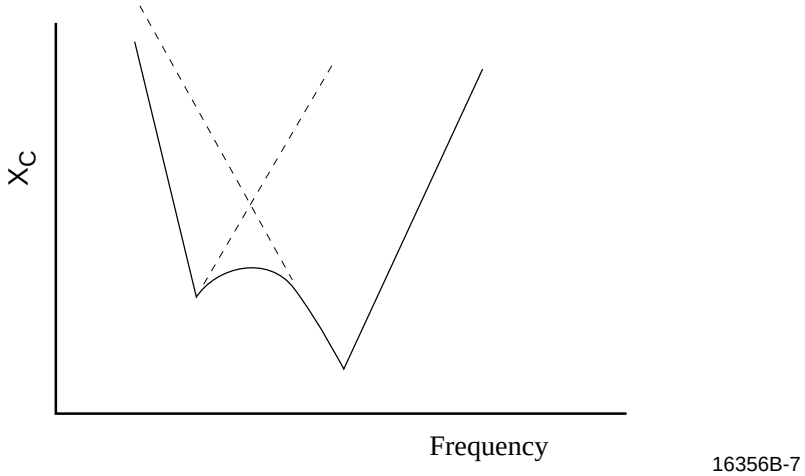

Figure 7. Frequency Response of Two Capacitors in Parallel

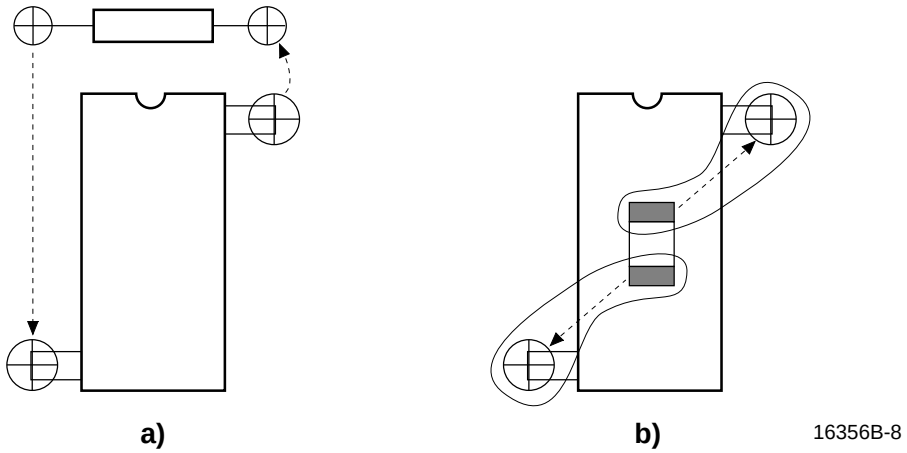

Figure 8. a) Typical Placement of Bypass Capacitors; b) Preferred Placement of Bypass Capacitors

Better performance can be obtained by ensuring that the chip and the capacitor contact the V_{CC} and ground planes at the same point. Because the capacitor size is different from that of the chip, it is necessary to run two traces from the V_{CC} and ground plane contact points to the capacitor, as shown in Figure 8b. These "lead extensions" are placed on a non-power plane and should be kept as short as possible. It is generally best to place the capacitor on the opposite side of the board, directly under the chip. A surface-mount chip capacitor works well here.

Note that the "lead extension" traces from the capacitor to the power pins take up space that could have been used for signal-line routing. However, putting extra effort into routing the signal lines now could prevent much noise-reduction work later on.

For devices with multiple V_{CC} and ground pins, how best to bypass depends on the device. In particular, it depends on whether the power pins are connected internally. On some devices, such as the PAL16R8-4 series, the ground pins are connected by a common ground bus. On these devices, it is only necessary to bypass one ground pin to one V_{CC} pin. If the power is kept separated internally, the separate V_{CC} pins must be decoupled individually. In general, it is best to contact the device's manufacturer for specific recommendations.

1.2 Power Distribution Network as a Signal Return Path

One of the more surprising functions of the power network is the provision of a return path for all signals in the system, whether generated on or off the board. Designs that accommodate this aspect of the power distribution system eliminate many high-speed noise problems.

1.2.1 The Natural Path of the Signal-Return Line

Of greatest concern in high-speed design is the energy generated at the signal switching edges. Each time a signal switches, AC current is generated. Current requires a closed loop. As illustrated schematically in Figures 9a and 9b, the return path needed to complete the loop can be supplied by the ground or V_{CC}. The loop can be represented by Figure 9c.

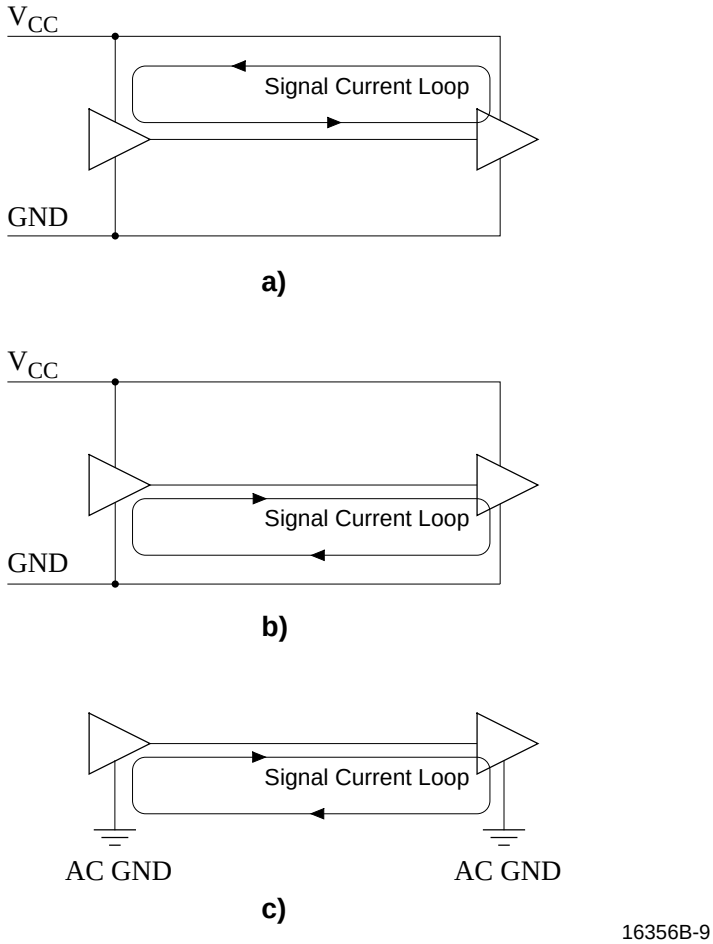

Figure 9. Current Loop of a Signal on the Board.
a) Through V_{CC}; b) Through Ground; c) The Equivalent AC Path

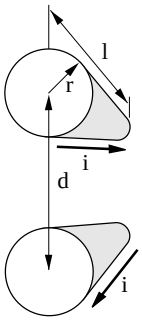

$$L = Kl\,ln\,\frac{d - r}{r}$$

16356B-10

Figure 10. Inductance Increases as the Signal and Return Path are Separated

Current loops have inductance and can be thought of as single-turn coils. They can aggravate ringing, crosstalk, and radiation. The current-loop inductance and associated problems increase with loop size. Minimizing the size of the loop minimizes these problems.

AC return signals have an entire plane in which to choose a path, but they take the path of least impedance (not necessarily least resistance) to the current. Impedance also includes inductance and capacitance. Metal has very little resistance; therefore, the impedance is primarily inductive. Because impedance increases with inductance, the path of least impedance is the path with the smallest inductance.

If the signal line goes from A to B on a random path, the natural return path is not necessarily a straight line, as would be dictated for least resistance. As noted in Figure 10, the inductance of a signal line and its return line increases with the separation of the two paths. The path of least impedance is the path bringing the signal-return line closest to the signal line. If it can, the signal return follows the signal line as closely as possible, resulting in the smallest loop. In multiple layer boards, "as closely as possible" usually means in a ground or V_{CC} plane above or below the signal trace. In a two-layer board, this means the closest ground or V_{CC} trace.

1.2.2 Bus vs Planes for a Signal Return Path

Figure 2a shows that a power bus has a fixed path. The return signal must follow this path, whether optimal or not. Unless the signal lines are purposely laid out near the power buses and oriented to minimize loop size, there will probably be large loops. If the layout of a board using buses for power distribution is not thought out carefully, it can result in a configuration that generates much noise.

The power plane imposes no natural restrictions on current flow. Thus the return signal can follow the path of least impedance, which is the path closest to the signal line. This results in the smallest possible current loops, which makes it the preferred solution for high-speed systems.

Although power planes have an advantage over buses, the benefits they provide can be defeated by the designer. Any break in the natural path of the return signal forces it to go around the break, increasing the loop size (Figure 11). Be careful about cuts in the ground and power planes.

1.3 Layout Rules With Power Distribution Considerations

The following layout rules will help you take advantage of power planes and avoid pitfalls.

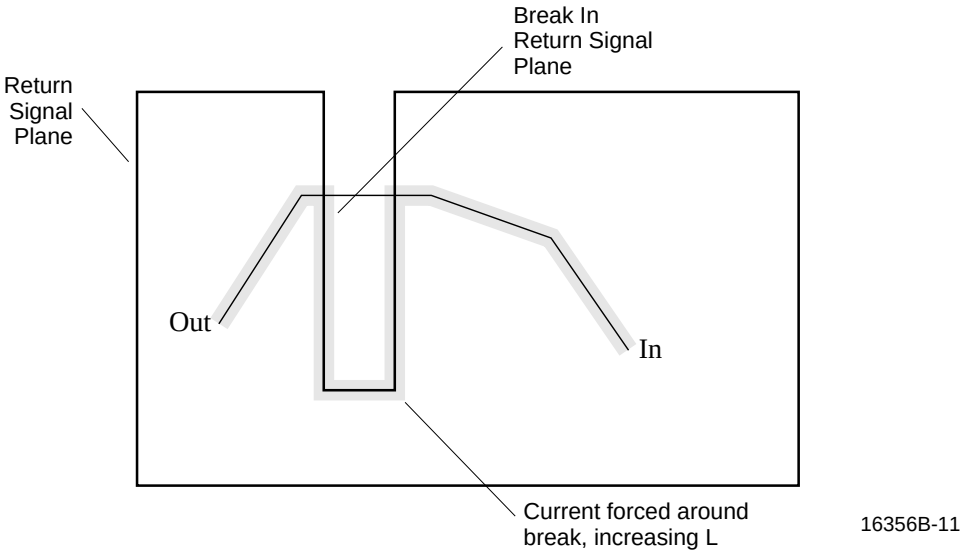

16356B-11

Figure 11. The Increase in Loop Size Due to a Break in the Power Plane

a. Be Careful with Feedthroughs

Cuts in the power plane tend to show up at feedthroughs or vias. These are necessary for traces to cross sides of the board and to connect components and connectors to the board. They are surrounded by small gaps where the power planes are etched away to avoid shorts in the signal lines. If the vias are close and the etchings wide, they might touch and form a barrier to any return path. This can occur with backplane connectors and device sockets.

For example, this can occur on the connectors on VME backplanes. The 104-pin connector has vias that can block the signal return. All the return signals are forced to the edge of the board. Not only are the loops longer, but the edge is shared by all the return signals; as we will see, this can result in crosstalk (Figure 12).

b. Ground Cables Sufficiently

Current loop considerations are also applicable for cables going off the board. Every signal should be a two-wire pair: one for the signal, and one for the return. The two lines should be kept next to each other to minimize the loop size. Figures 13a and 13b illustrate poorer configurations. Figure 13c illustrates the proper configuration.

c. Separate Analog and Digital Power Planes

High-speed analog devices tend to be sensitive to digital noise. For example, amplifiers can amplify switching noise, making it appear as spikes. Thus on boards with analog and digital functions, the power planes are commonly separated; the planes are tied together at the power source. This causes a problem for devices using both types of signals (such as DACs or voltage comparators). The signal lines must cross the plane boundaries. These boundaries force the return path to the power source before returning to the driver.

The solution is to place jumpers across the ground planes where signals cross (Figure 14). The jumper provides a bridge across the break for the return signal; this helps minimize the current loop.

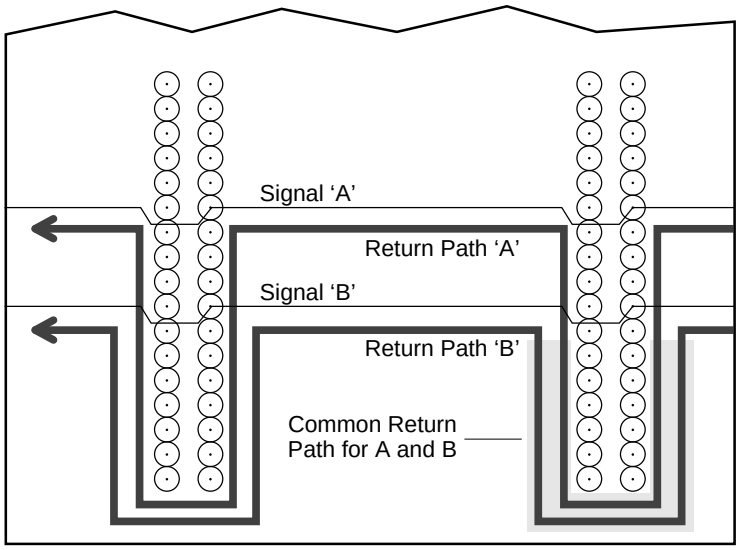

Figure 12. Common Paths of Signal Return Due to Vias

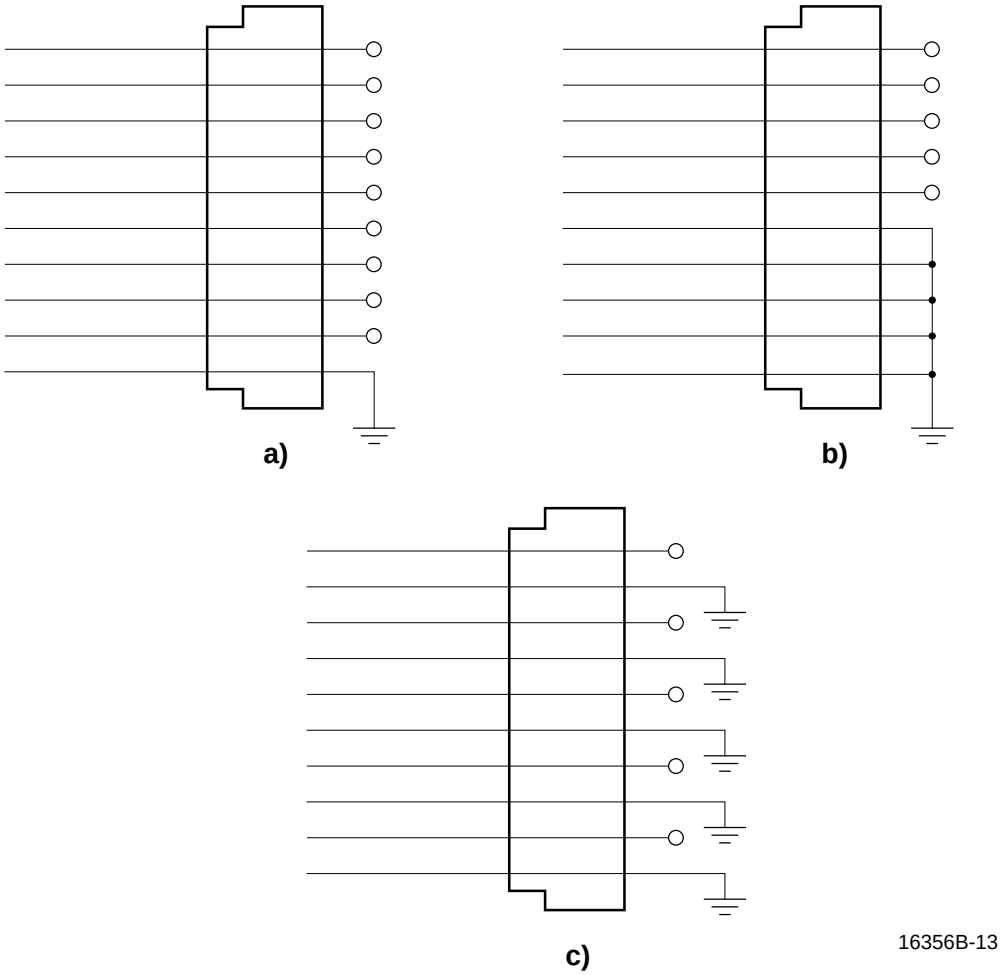

Figure 13. Connector Configuration. a) Insufficient Grounds; b) Enough Grounds but Grounds lumped Together Resulting in Larger Current Loops; c) Grounds Evenly Distributed Among Signal Lines

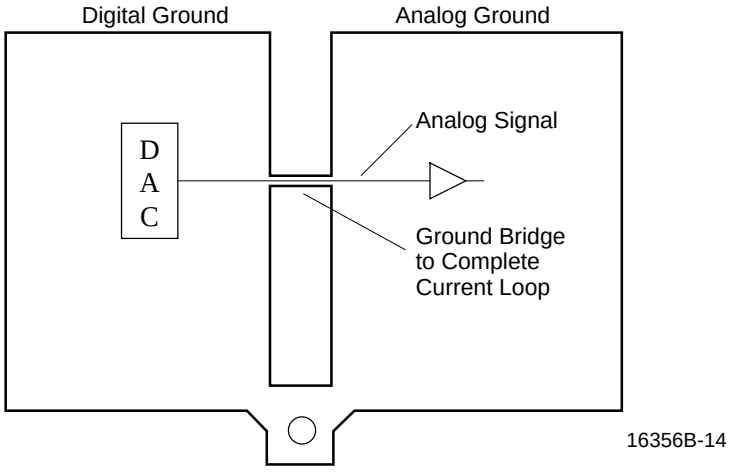

Figure 14. Jumper Between Analog- and Digital-Power Planes for Signal-Return Path

d. Avoid Overlapping Separated Planes

When separate power-planes are used, do not overlap the power plane of the digital circuitry and the power plane of the analog circuitry. The analog and digital power planes are separated to isolate the currents from each other. If the planes overlap, there is capacitive coupling, which defeats isolation.

To ensure separation, take a board and cut between the separated planes. Then inspect the newly-exposed edges of the board. No metal should be showing, except where traces or connections are specifically designed to cross the boundary.

e. Isolate Sensitive Components

Certain devices, such as phase-locked loops, are particularly sensitive to noise interference. They require a higher degree of isolation.

Good isolation can be achieved by etching a horseshoe in the power planes around the device (Figure 15). All signals used by the device enter and leave through the narrow gap at the end of the horseshoe. Noise currents on the power plane must go around the gap and do not come close to the sensitive part.

When using this technique, ensure that all other signals are routed away from the isolated section. The noise signals generated by these lines can cause the interference this technique was designed to avoid.

f. Place Power Buses Near Signal Lines

Sometimes, the designer must use two-layer boards and is forced to use power buses instead of planes. Even then it is possible to control loop size by placing the buses as close as possible to the signal lines. The ground bus could follow the most sensitive signals on the other side of the board (Figure 16). The loop for that signal is the same as it would be if the load used power planes.

2. Signal Lines as Transmission Lines

Controlling the relationship between the signal line and AC ground takes advantage of the return signal's tendency to take the path of least impedance. Another advantage is the constant impedance along the signal line. Such signal lines are called controlled-impedance lines, and they provide the best medium for signal transmission on the board.

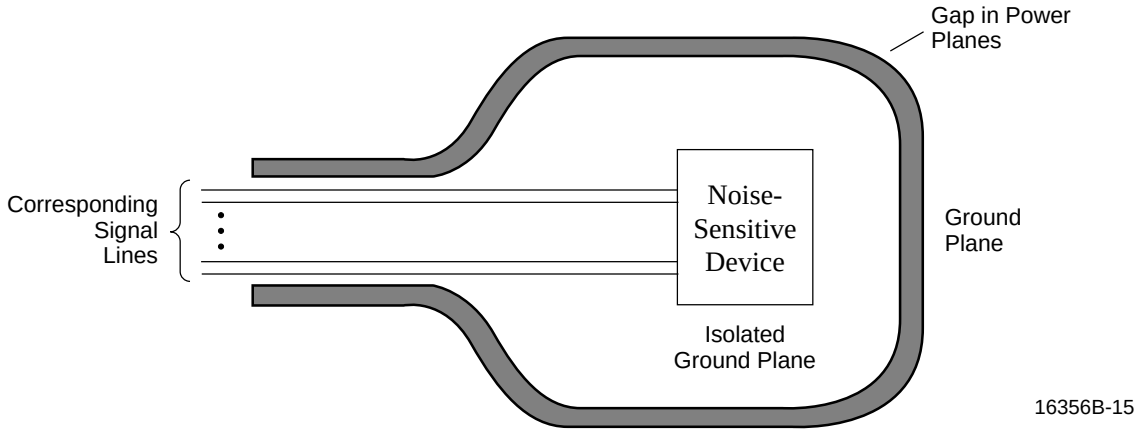

Figure 15. Isolation of Noise Sensitive Components

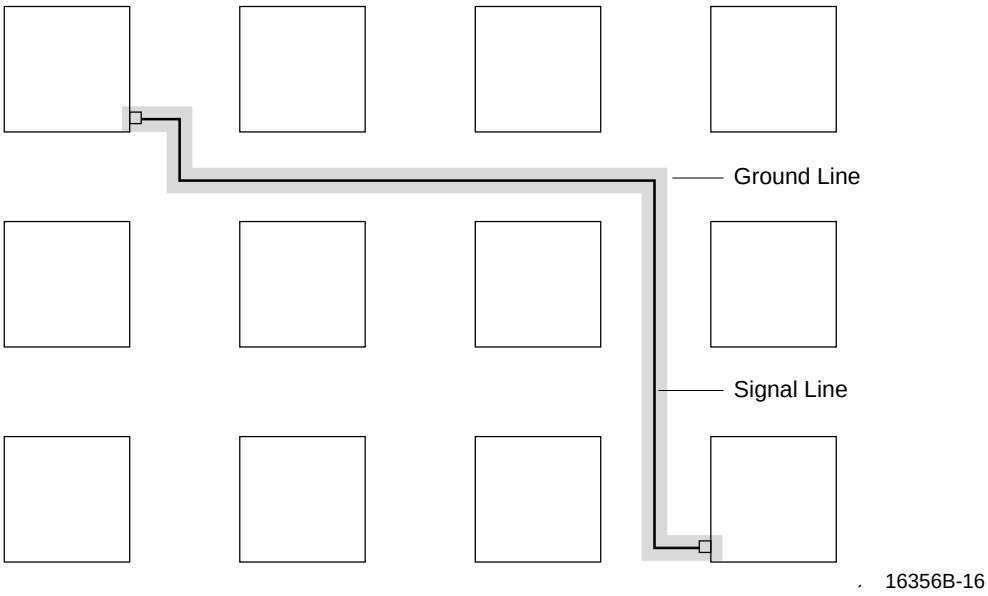

Figure 16. Providing the Optimum Signal-Return Path with a Bus-Power Distribution System

However, when the signal delay is greater than a significant portion of the transition time, the signal line must be treated as a transmission line. An improperly terminated transmission line is subject to reflections, which distort the signal. The signal at the load end of the line can resemble ringing (Figure 17), slowing down the system. It can also cause false clocking, destroying system functionality.

A controlled-impedance signal line can be modeled as shown in Figure 18. Inductance and capacitance are evenly distributed along the length of the line. Their units are henrys per unit length and farads per unit length, respectively.

From the model, we can derive two important parameters: impedance (Z_0), and propagation delay (t_{PD}). On a lossless signal line, Z_0 is an AC resistance; i.e., Z_0 appears to the driver as a pure resistor. Its units are ohms (Ω), and it is equal to

$$Z_0 = \sqrt{\frac{L_0}{C_0}}$$

where

L_0 = Signal Line Inductance in henrys per unit length

C_0 = Signal Line Capacitance in farads per unit length

The propagation delay also depends on L_0 and C_0. It has units of time per unit length, and it is equal to

$$t_{PD} = \sqrt{L_0 C_0}$$

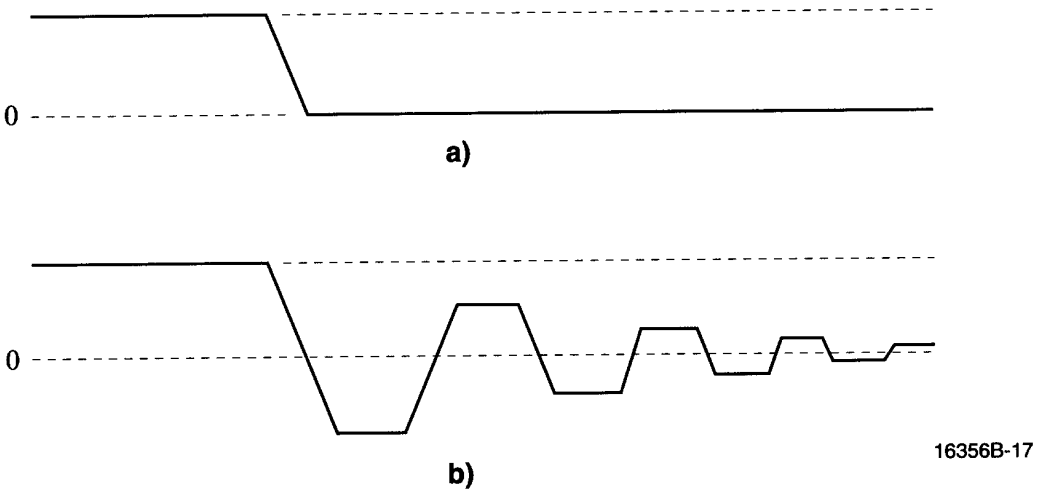

Figure 17. Reflections on a Signal Line. a) at the Driver; b) at the Load

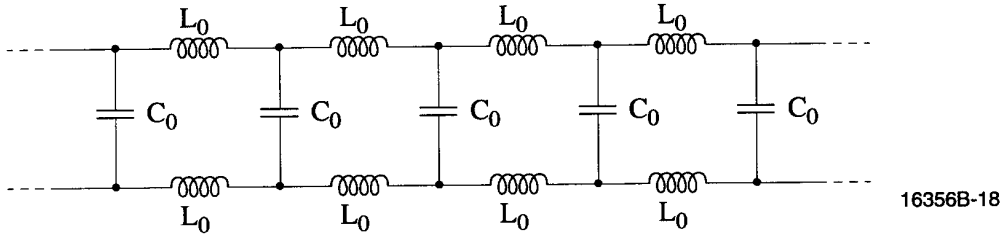

Figure 18. Transmission Line

Transmission Line Categories

Given that the designs discussed in this paper are for printed circuit boards, the possible types of signal lines fall into one of two categories: stripline and microstrip (Figure 19). The stripline has the signal line sandwiched between two power planes. This technique theoretically offers the cleanest signals because the signal line is shielded on both sides. However, the lines are hidden; there is no easy access to the signal lines. Microstrip has the signal line on an outer layer. The ground plane is to one side of the signal line. This technique allows easy access to the signal line.

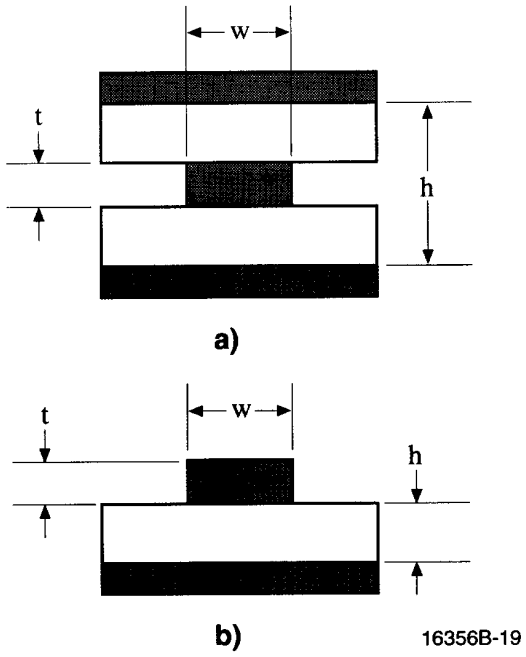

Figure 19. Signal Line Construction on a Circuit Board. a) Stripline; b) Microstrip

The parameters C_0, L_0, Z_0, and t_{PD} can be determined from the physical dimensions of the signal line and the dielectric properties of the board material. They are discussed below.

For Stripline

$$Z_0 = \frac{60}{\sqrt{\varepsilon_R}} \ln \frac{4h}{0.67\pi w\ (0.8 + \frac{t}{w})}\ \Omega$$

$$t_{PD} = 1.017\sqrt{\varepsilon_R}\ ns/ft$$

$$C_0 = 1000\ \frac{t_{PD}}{Z_0}\ pF/ft$$

$$L_0 = Z_0^2 C_0\ pH/ft$$

For Microstrip

$$Z_0 = \frac{87}{\sqrt{\varepsilon_R + 1.41}} \ln \frac{5.98h}{0.8w + t}\ \Omega$$

$$t_{PD} = 1.017\ \sqrt{0.457\varepsilon_R + 0.67}\ ns/ft$$

$$C_0 = 1000\ \frac{t_{PD}}{Z_0}\ pF/ft$$

$$L_0 = Z_0^2 C_0\ pH/ft$$

e_R is the relative dielectric constant of the board material. A common material is epoxy-laminated fiberglass, which has an average e_R of 5.

Example

The dimensions of the trace and board are restricted by certain rules. Generally, the vendor sells the board with 1 oz of copper, so the metal thickness is about 1 mil. The trace width should be between 8 and 15 mils. Signal lines thinner than 8 mils tend to be harder to control. Signal lines thicker than 15 mils tend to have excess capacitance. A typical value is 10 mils. The layer separation is determined by the required board thickness and the number of layers to be used. For this example, 30 mils is adequate.

Based on these assumptions, it is possible to calculate the parameters for a typical signal line: width = 10 mils, thickness = 1 mil, separation = 30 mils, and $e_R = 5$.

$$Z_0 = \frac{87}{\sqrt{5} + 1.41} \ln \frac{5.98 * .03}{0.8 * .001 + .01}\ \Omega$$

$$= 67.0$$

$$t_{PD} = 1.017\ \sqrt{0.456 * 5 + 0.67}\ ns/ft$$

$$= 1.75\ n$$

$$C_0 = 1000 * \frac{1.75}{67.05}\ pF/ft$$

$$= 26.1\ pF/ft$$

$$L_0 = 67.05^2 * 26.1\ pH/ft$$

$$= 117\ nH/ft$$

Distributed Load Calculations

The calculations above are for a signal line with a lumped load at the end of the trace (Figure 20). If the load is distributed along the signal line (Figure 21), the capacitance of the load devices is also distributed along the line and adds to the line capacitance. This changes the signal-line parameters Z_0 and t_{PD}. The new parameters are derived from the original values based on the added capacitance, C_L, in farads per unit length:

$$Z_0 = \frac{Z_0}{\sqrt{1 + \frac{C_L}{C_0}}}\ \Omega$$

$$t_{PD} = t_{PD} * \sqrt{1 + \frac{C_L}{C_0}}\ ns/ft$$

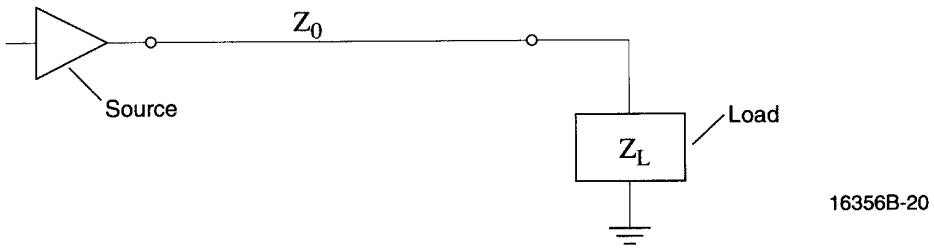

Figure 20. Transmission Line with a Lumped Load

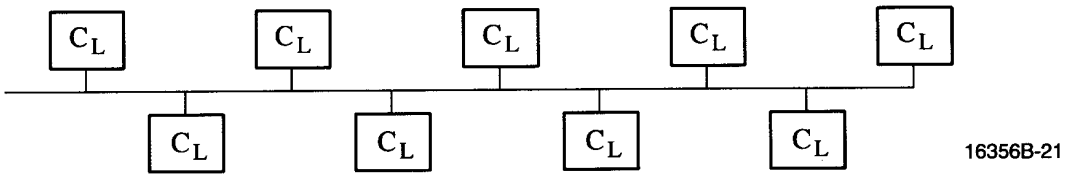

Figure 21. Transmission Line with a Distributed Load

Distributed loading is common in memory banks. The input capacitance on these devices can range from 4 pF to 12 pF. The following example uses 5 pF. The physical size of memory devices usually permits placing two of them per inch. The distributed added capacitance is then:

$$C_L = \frac{5\ pF}{0.5\ in * \frac{1\ ft}{12\ in}}$$

$$= 120\ pF/ft$$

The new values of Z_0 and t_{PD}, based on distributed loading, are:

$$Z_0 = \frac{67.05\ \Omega}{\sqrt{1 + \frac{120\ pF/ft}{26.1\ pF/ft}}}$$

$$= 28.34\ \Omega$$

$$t_{PD} = 1.75\ ns/ft * \sqrt{1 + \frac{120\ pF/ft}{26.1\ pF/ft}}$$

$$= 4.14\ ns/ft$$

With this distributed load, the impedance has been greatly reduced, and the signal is now much slower.

Reflections

The source generates a signal with an energy content determined by $Z_0\ \Omega$. Even though the line is seen as a resistance, the signal line does not dissipate energy. The energy in the signal must be dissipated by the load impedance (Z_L), as shown in Figure 20.

The maximum transfer of energy from source to load requires that the load impedance equal the source impedance. For the entire signal to be transferred to Z_L, Z_L must equal Z_0. If they are not equal, some of the signal energy is dissipated, and the rest is reflected back toward the source. The source generator output then adjusts to compensate for the "new" load.

The waveform of the signal at the load can be thought of as the sum of the originally generated signal and the reflection from the load. The appearance of the waveform depends on the mismatch of the load and line impedances and the ratio of the signal-transition time (t_R) to the propagation delay of the line (τ), t_R/τ. If the transition time is significantly longer than the propagation delay of the line, the reflection reaches the source when the original signal has changed only a small amount. The generator compensates for the "new" load and transmits the corrected signal with little signal disturbance. The signal at the load then has a small overshoot.

If the propagation delay of the line is long enough for the reflection to reach the source after the signal has changed a significant percentage, the generator must change significantly to compensate for the load. The load reflects the new transition, which results in the ringing shown in Figure 17.

The amount of overshoot usually varies proportionally with the signal-line length until the signal-line delay is equal to the transition time. At this point, the overshoot can be as much as the original transition, effectively doubling the swing of the transition.

A signal line long enough to produce significant reflections acts like a transmission line. The point at which the signal line is considered a transmission line depends on the amount of tolerable distortion. A liberal rule of thumb is to consider a signal line a transmission line when the transition time of the original signal is less than four times the propagation delay of the signal (Figure 22); that is, when $t_R/\tau \geq 4$.

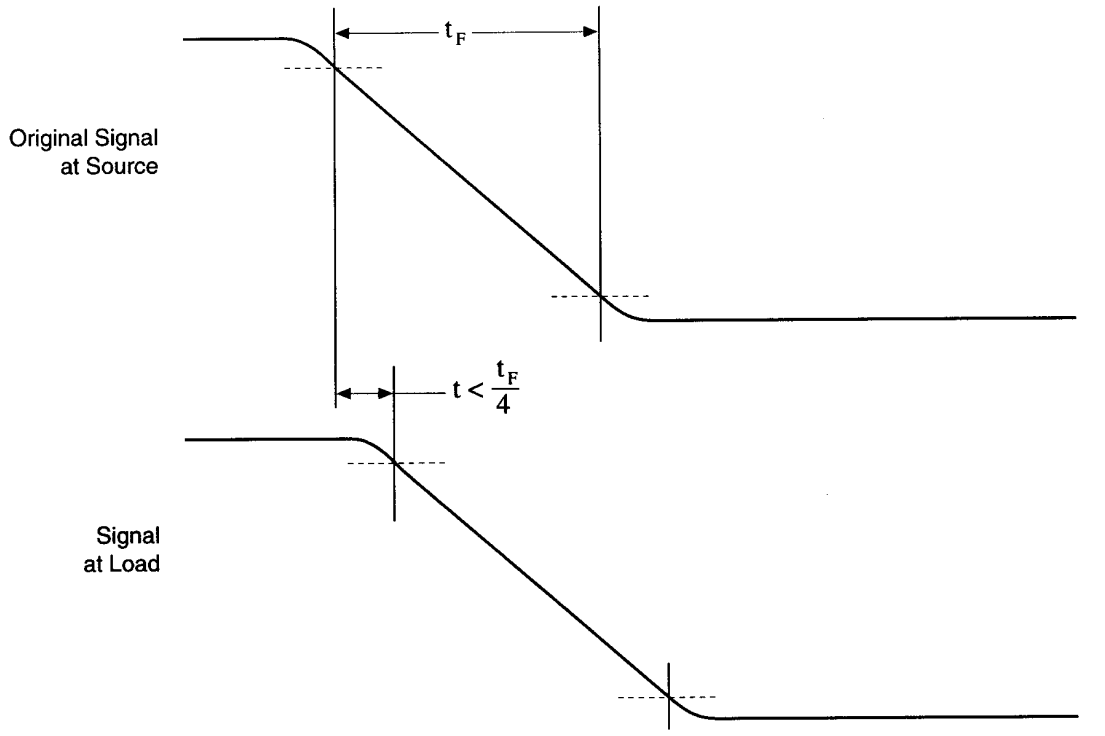

16356B-22

Figure 22. Minimum Delay Between Original and Reflected Signal Which Results in a Transmission Line

A more conservative rule is to consider the signal line a transmission line when t_R/τ is less than eight times the propagation delay. Generally, the larger the transition time is in relation to the propagation delay of the signal line, the cleaner the resultant signal.

From this it is possible to determine what length of the microstrip line discussed above must be treated as a transmission line. On available devices, t_R ranges from 5 ns (especially those using bipolar technology) to 1 ns (newer bipolar and CMOS devices). The rise times and corresponding signal-line lengths are shown in Table 2 for the example given above.

Table 2
Example: t_R and Corresponding Transmission-Line Length for $\frac{t_R}{\tau} = 4$

t_R (ns)	Line Length (inch)
5	8.6
4	6.9
3	5.1
2	3.4
1	1.7

For older devices with 5 ns transition times, signal lines shorter than 8.6" do not have to be treated as transmission lines. For newer, high-speed devices, even a two-inch line is a transmission line. Practically all signal lines are transmission lines on boards with high-speed devices.

If the transmission line has the distributed load in the example above, then the minimum transmission-line length must be reconsidered. As shown in Table 3, a four-inch line is a transmission line when $t_R = 5$ ns. If $t_R = 1$ ns, a signal line smaller than one inch is a transmission line.

Table 3
Example: t_R and Corresponding Transmission-Line Length with Lumped and Distributed Loads for $\frac{t_R}{\tau} = 4$

	Line Length (inch)	
t_R (ns)	Lumped Load	Distributed Load
5	8.6	3.6
3	5.1	2.17
2	3.4	1.4
1	1.7	0.75

Quantifying Reflections

Given that the signal line is long enough to be considered a transmission line, the size of the reflected signal depends on the difference between Z_0 and Z_L. The numerical indicator of the percentage, or the original signal that is reflected, is called the reflection coefficient (K_R). K_R is equal to:

$$K_R = \frac{Z_L - Z_0}{Z_L + Z_0}$$

The percentage of the original signal reflected back is 100 * K_R.

Referring back to the open load:

$$K_R = \frac{\infty - Z_0}{\infty + Z_0}$$

$$= 1$$

For an shorted load:

$$K_R = \frac{0 - Z_0}{0 + Z_0}$$

$$= -1$$

For open and shorted loads, the entire signal is reflected without attenuation. K_R is negative for the shorted load. This indicates that the reflected signal is inverted from the original.

With a printed-circuit board, it is possible to estimate the expected type of mismatch. Z_0 typically ranges from 30 Ω to 150 Ω. Input impedances range from 10 kΩ (for bipolar devices) to over 100 kΩ (for CMOS devices). Output impedances can be very low. A CMOS PAL device, such as the PALCE16V8, has a typical-output LOW voltage of 0.2 V at 24 mA for about 8 Ω. The output-HIGH impedance is about 50 Ω, which is closer to the expected Z_0.

Consider the microstrip line derived earlier, with a CMOS device as its load. The following discussion shows what happens on the HIGH to LOW transition.

The driver's output impedance (Z_S) is:

$$Z_S \approx \frac{V_{OL}}{I_{OL}}$$

$$= \frac{0.2\ V}{24\ mA} \approx 8.3\ \Omega$$

A more accurate number can be obtained from an actual I/V curve of the output.

The input impedance of the load is greater than 100 kΩ. This is so much greater than Z_0 (67 Ω), that K_R at the load is practically equal to one. K_R at the source is:

$$K_R = \frac{8.3 - 67}{8.3 + 67}$$

$$= -0.78$$

The driver generates a signal switching from 3.5 V to 0.2 V. Since the driver-output impedance and Z0 make up a voltage divider, the generated signal is:

$$\Delta V = \frac{(0.2\ V - 3.5\ V)\ Z_0}{Z_0 + Z_S}$$

$$= \frac{(0.2\ V - 3.5\ V)\ 50}{50 + 8}$$

$$= 2.84\ V$$

The resultant signal at the source is:

$$V_S = 3.5\ V - \Delta V = 3.5\ V - 2.84\ V$$

$$= .066\ V$$

When the signal reaches the load, V_L changes by –2.84 V from the original transmission and a further –2.84 V from the reflection. Since V_L originally was 3.5 V, it is now –2.19 V.

At the start, V_S = 0.66 V. The reflected signal returns to the source. Some of it is reflected per the source KR. VS is equal to the sum of the original signal, the reflected signal, and the second reflected signal. The second reflection is equal to:

$$V_R = -.78 * -2.84$$

$$= 2.21\ V$$

$$V_S = 0.66\ V + -2.84\ V + 2.21\ V$$

$$= -0.035\ V$$

The second reflection goes to the load. When it arrives,

$$V_L = -2.19 + 2.21 + 2.21$$

$$= 2.24$$

The signal continues like this, bouncing back and forth, getting smaller each time. This is illustrated in the lattice diagram in Figure 23. The lines at the left and right are the voltage at the source and load, respectively. The angled lines show the value of the transmitted signal and the reflections.

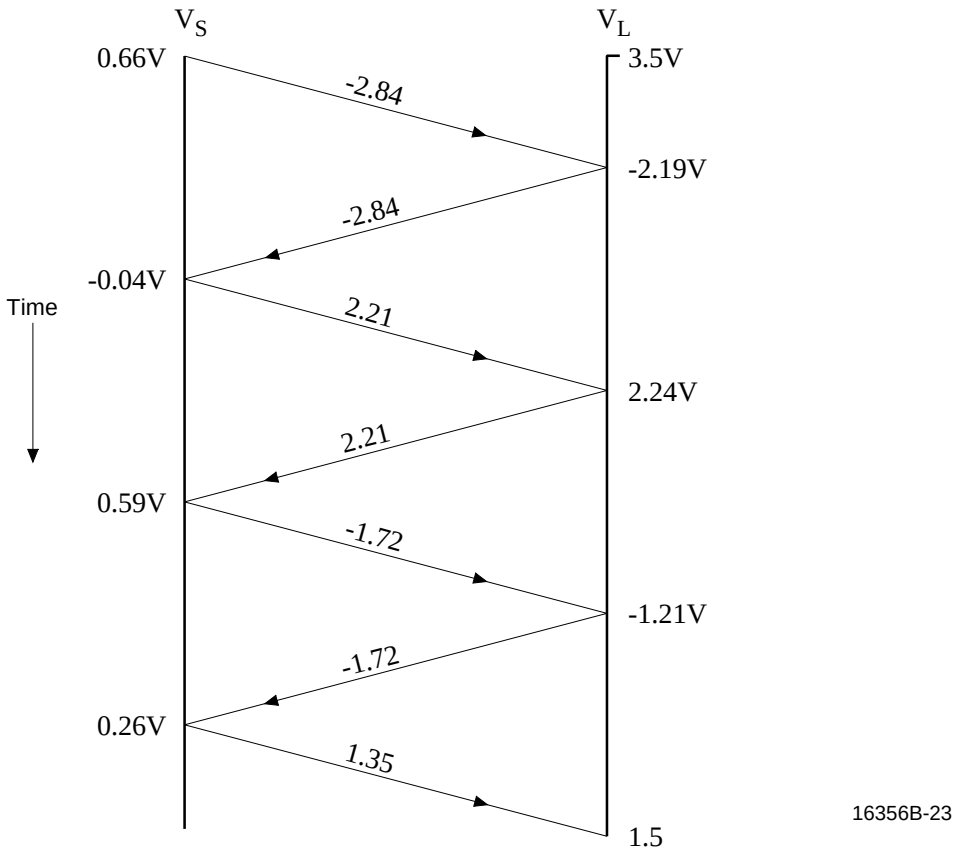

Figure 23. Lattice Diagram Representation of a Reflected Signal

The same information in the time domain is shown in Figure 24. The top part of the Figure shows the source; the bottom shows the load signal. Note that it takes five complete cycles for the signal strength to drop below the input threshold. Propagation delays are typically from 2 ns/ft to 5 ns/ft. With t_{PD} = 3 ns/ft and a 6-inch line, the delay across the line is about 1.5 ns. The signal can be safely considered valid at about 13.5 ns after the original transition.

TERMINATION

The amount of reflections shown in the last example would be too much for most systems. A technique is needed to eliminated, or at least reduce, the reflections. Since the reflections are eliminated when $Z_L = Z_0$, it is necessary to change Z_L to equal Z_0.

To understand this, look at the nature of the input and output impedances of the PAL devices. As noted above, input impedances tend to be high. Bipolar is in the 10 kΩ range, while CMOS is in the 100 kΩ range. Output drivers tend to have low impedance.

There are two schemes for termination: reduce Z_L to Z_0 to eliminate load reflections, or increase Z_S to Z_0 to eliminate secondary reflections at the source. Z_L can be reduced by placing a resistor in parallel with the load—parallel termination; Z_S can be increased by placing a resistor in series with the source and the line—series termination.

Parallel termination is shown in Figure 25a. Because of the extremely high input resistance of most devices, R_L can be made equal to Z_0.

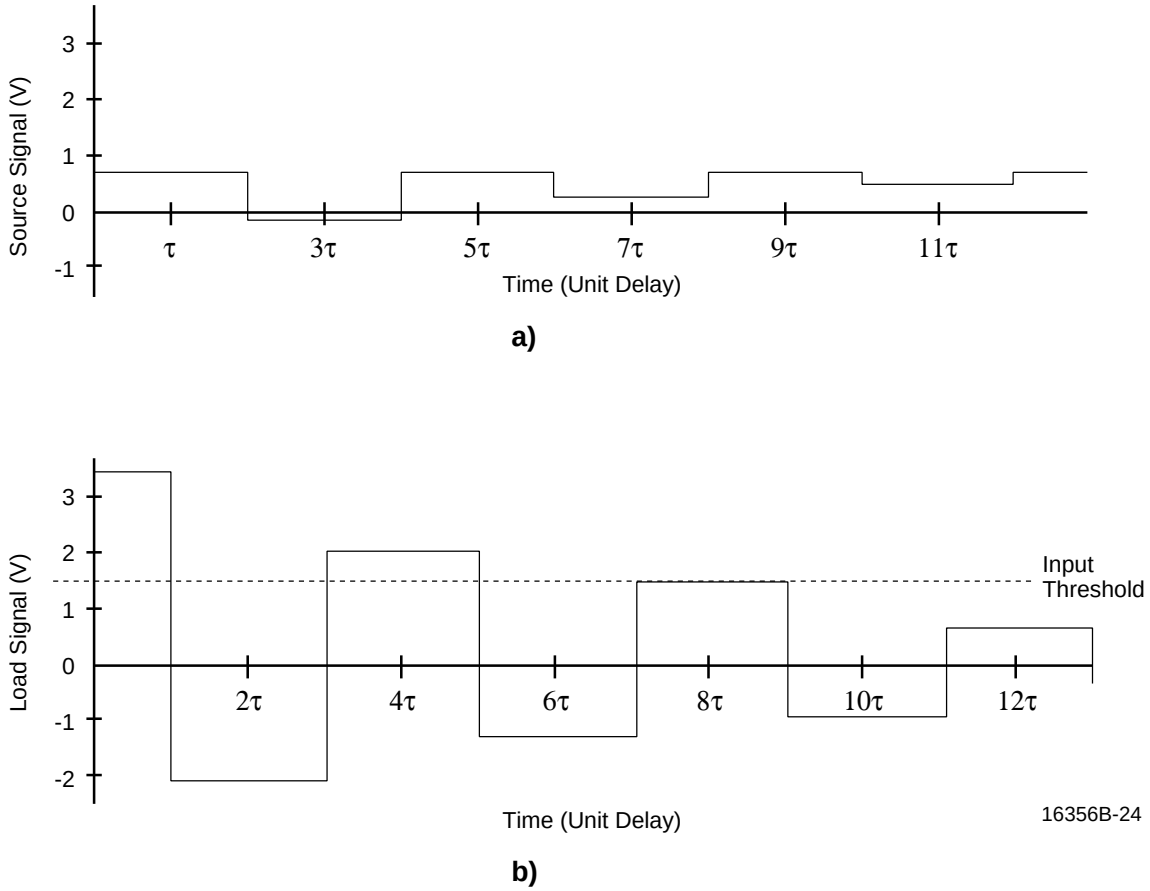

Figure 24. Time Representation of a Reflected Signal; a) at the Source b) at the Load

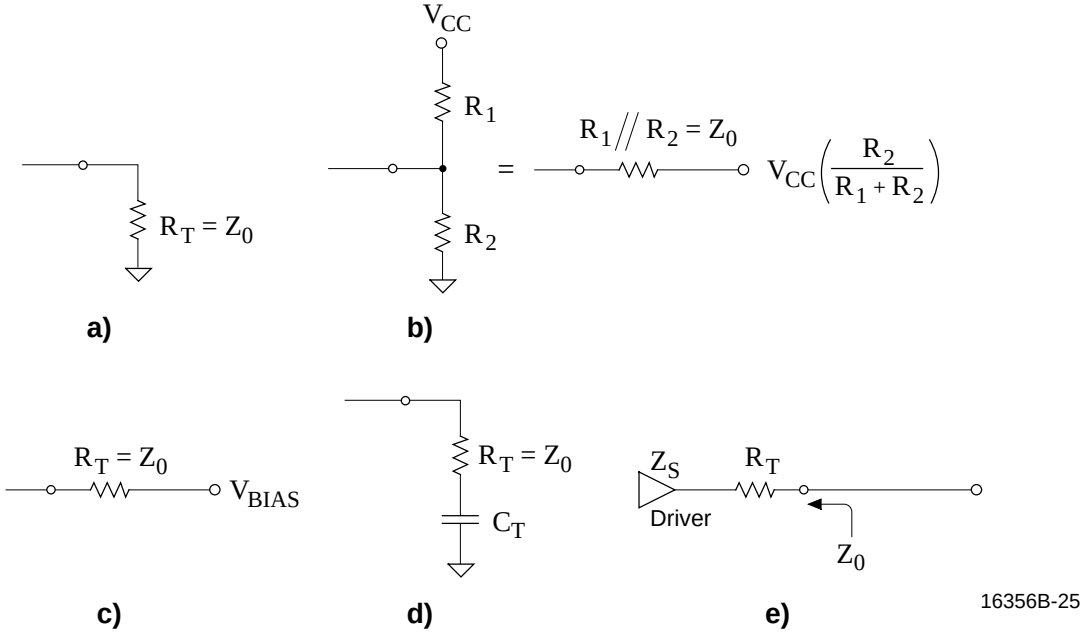

Figure 25. a) Parallel Termination; b) Thévenin Equivalent; c) Active Termination; d) Series Capacitor; e) Series Termination

This scheme has one disadvantage: the current drain is high for the HIGH-output state. For a 50-Ω termination, it can be as much as 48 mA. Most drivers are rated for an I_{OH} of 3.2 mA. This is clearly above the level that the device can support and still maintain an adequate V_{OH}.

Terminating to V_{CC} can help, since I_{OL} is usually higher than I_{OH}. However, most CMOS devices designed for board-level applications have drivers rated for an I_{OL} of 24 mA or less. This is still below the level that can support and maintain an adequate V_{OL} for a low-impedance transmission line.

The current can be reduced considerably by using two resisters, as shown in Figure 25b. The resistors form a voltage divider with the Thévenin voltage equal to:

$$V_{TH} = \frac{V_{CC} * R_2}{R_1 + R_2}$$

The Thévenin resistance is equal to:

$$R_{TH} = \frac{R_1 * R_2}{R_1 + R_2}$$

Although this is a good solution, there is higher power-supply current because the resistors are between V_{CC} and ground.

Another approach to reducing load current is to reference the resistor to a positive voltage between V_{OH} and V_{OL} (Figure 25c). The current flow from 3 V to 2.5 V through a 50-Ω resistor is considerably less than the flow from 3 V to ground through the same resistor. This does not present any signal problems, because the DC voltage reference is AC ground. However, it is difficult to find a terminating voltage source that can switch from sinking current to sourcing current fast enough to respond to the transitions.

Another technique is to replace the original terminating resistor with a resistor and capacitor series-RC network (Figure 25d). The resistor is equal to Z_0. The capacitor can be on the order of 100 pF; the exact value is not important. At these frequencies, the capacitor is an AC short but it blocks DC. Thus the driver does not see the DC loading effect of R_L. This technique is referred to as AC termination.

Techniques that terminate at the load are designed to eliminate the first reflection. An alternate approach is to increase Z_S to equal Z_0 by placing a resistor in series with the source (Figure 25e). When added to Z_S, this resistor makes the new source impedance look like Z_0.

This type of termination works best with a lumped load because the voltage divider formed by the Z_S and Z_0 attenuates the signal (Figure 26 a and b). The original transition is cut in half by this voltage divider, since $Z_S + R_L = Z_0$. This half-transition tracks down the transmission line until it is reflected at the load, which is unterminated. Since the reflection causes the original half-transition to double, it brings the signal at the load to its final value (Figure 27a). The reflection then travels back up the line, completing the transition all along the line (Figure 27b).

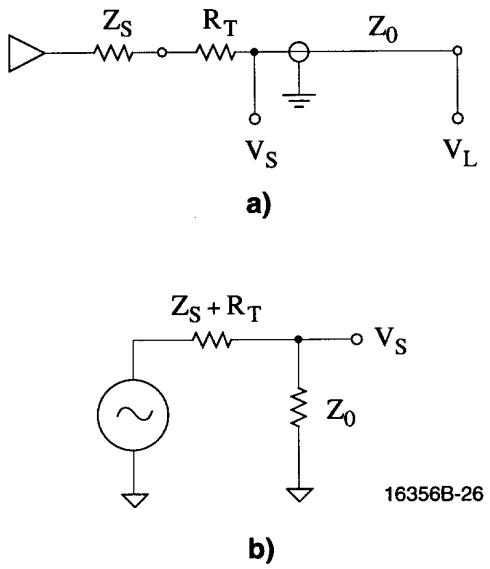

Figure 26. a) Series Termination; b) Voltage Divider formed by Series Termination

This can be illustrated by putting a series terminating resistor on the unterminated microstrip example considered earlier. A 59-Ω resistor (68Ω – 9Ω) is placed in series with the driver. For a LOW to HIGH transition, the signal at the source is:

$$\Delta V = \frac{(0.2\ V - 3.5\ V)\ Z_0}{Z_S + Z_0 + 59\ \Omega}$$

$$= \frac{(0.2\ V - 3.5\ V) * 67\ \Omega}{8\ \Omega + 67\ \Omega + 59\ \Omega}$$

$$= -1.65\ V$$

$$V_S = 3.5\ V + \Delta V$$

$$= 3.5\ V - 1.65\ V$$

$$= 1.85\ V$$

If the load is effectively an open circuit, then a –1.65 V reflection returns. When the reflected signal reaches the source, no new reflections occur because Z_S is matched to Z_0 by R_T. V_S is 1.85 V – 1.65 V = 0.2 V.

The reflection at the load causes V_L to equal 0.2 V when the original signal arrives. V_S does not equal 0.2 V until the reflected signal returns, in this example, 3 ns later (Figure 27).

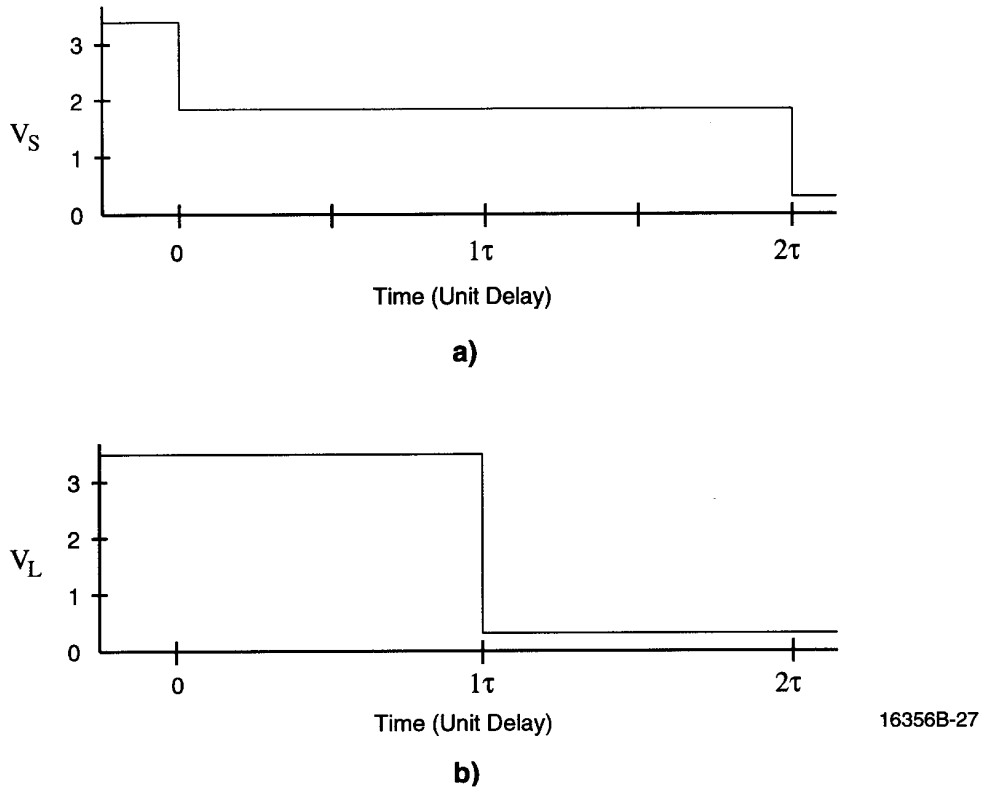

Figure 27. a) Signal at Source; b) Signal at Load end

This can be a risky approach if the load is distributed along the line, since those loads not at the end of the line will see some intermediate voltage until the reflection cleans them up on its return to the source. In addition, this technique adds the delay of the return trip because the signal cannot be considered valid until the device closest to the driver has a valid input. The input to the device closest to the driver becomes valid upon the return of the reflection. The delay is longer than indicated in the last example because the added capacitance of the distributed load reduces Z_0 and increases t_{PD}.

Despite this drawback, series termination is successfully used with DRAM drivers, even when the DRAMs are distributed along the signal line. The risk of the signal spending time near threshold and the extra delay are reduced by choosing R_T so that the resultant Z_S is slightly less than Z_0. The voltage swing at the line is larger, and the voltage level is closer to V_{OL}, below the input threshold. If the line is terminated with 20 Ω, V_S becomes:

$$V_S = 3.5\ V + \frac{(0.2\ V - 3.5\ V)\ Z_0}{Z_S + Z_0 + 20\ \Omega}$$

$$= 3.5\ V + \frac{(0.2\ V - 3.5\ V) * 67\ \Omega}{8\ \Omega + 67\ \Omega + 20\ \Omega}$$

$$= 1.17\ V$$

Because the termination is not an exact match, some ringing occurs. However, if the ringing is below a tolerable level, it can be used successfully. The designer must decide on the compromise. Furthermore, the high capacitance of memory lines often swamps out the ringing.

Often, an exact match is not possible because of the differences between the HIGH- and LOW-output impedances. The output impedance of TTL-compatible devices is different for HIGH and LOW levels. For example, the PALCE16V8 is 8 Ω when LOW, and about 50 Ω when HIGH. This complicates the choice of a terminating resistor because no single value is ideal for both cases. A compromise value must be chosen that results in acceptable results in both transition directions.

Layout Rules for Transmission Lines

The controlled impedance signal line is the best practical medium for signal transfer on a board, and proper termination helps ensure proper noise-free operation. However, it is still possible to generate noise with an inefficient layout. The following layout rules further enhance board operation.

1. Avoid Discontinuities

Discontinuities are points where the impedance of the signal line changes abruptly; they cause reflections. The formula for K_R is as valid here as it as at the end of the line. Because they cause reflections, they should be avoided. Discontinuities can be at sharp bends on the trace or at vias through the board.

At bends on the trace, the cross-sectional area increases, and Z_0 decreases. It is possible to compensate for the bend by cutting the trace as shown in Figure 28. The cut is chosen so that the resulting diagonal is equal to the trace width. This minimizes the delta in cross-sectional area, as well as the discontinuity. Using two 45° bends makes use of the same concept and is a common way of smoothing out bends. A smooth circular arc would be ideal but is harder to generate with many tools.

Vias take signals through the board to the other side (Figure 29). The vertical run of metal between layers is an uncontrolled impedance, and the more of these there are, the greater is the overall amount of uncontrolled impedance in the line. This contributes to reflections. Also, the 90° bend from horizontal to vertical is a discontinuity that generates reflections. If vias cannot be avoided, use as few as possible.

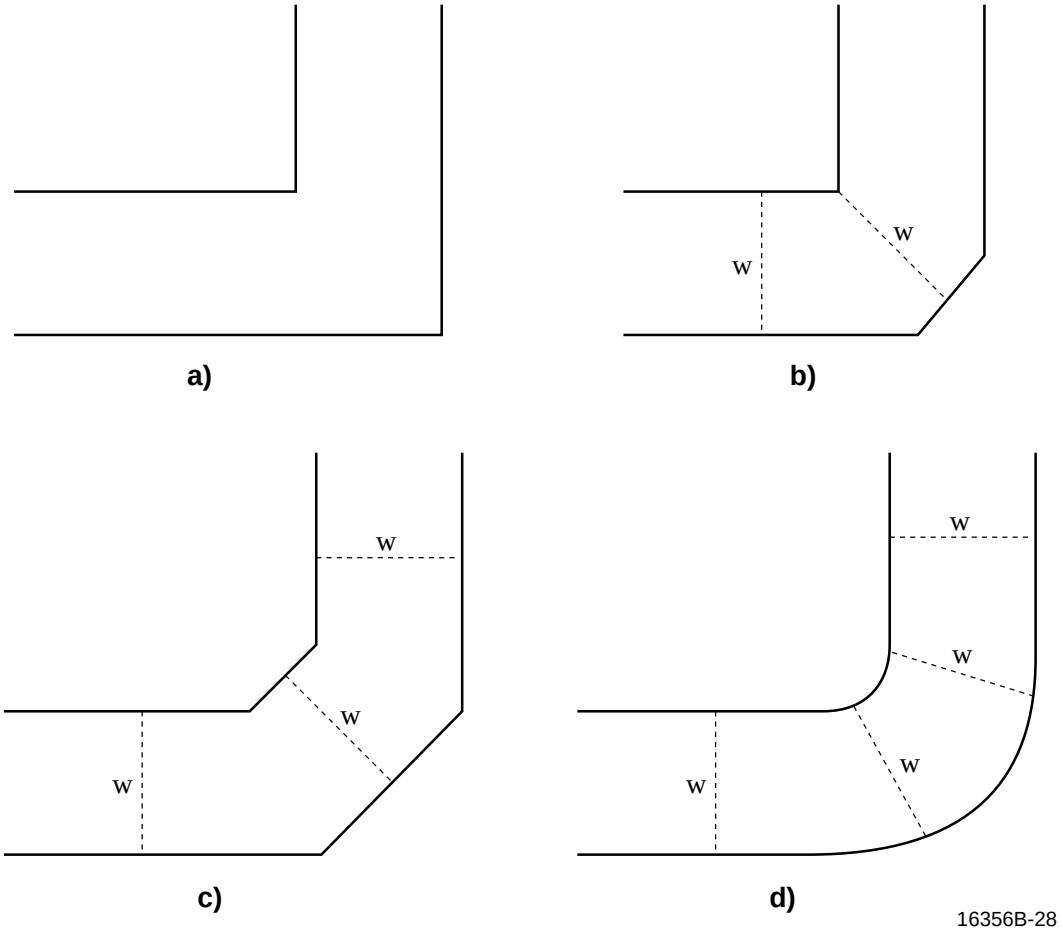

Figure 28. Reducing Discontinuity. a) Corner on PC Board Trace which Causes Discontinuity; Solved: b) by Shaving the Edge; c) by 45° Corner; d) by Using Curves

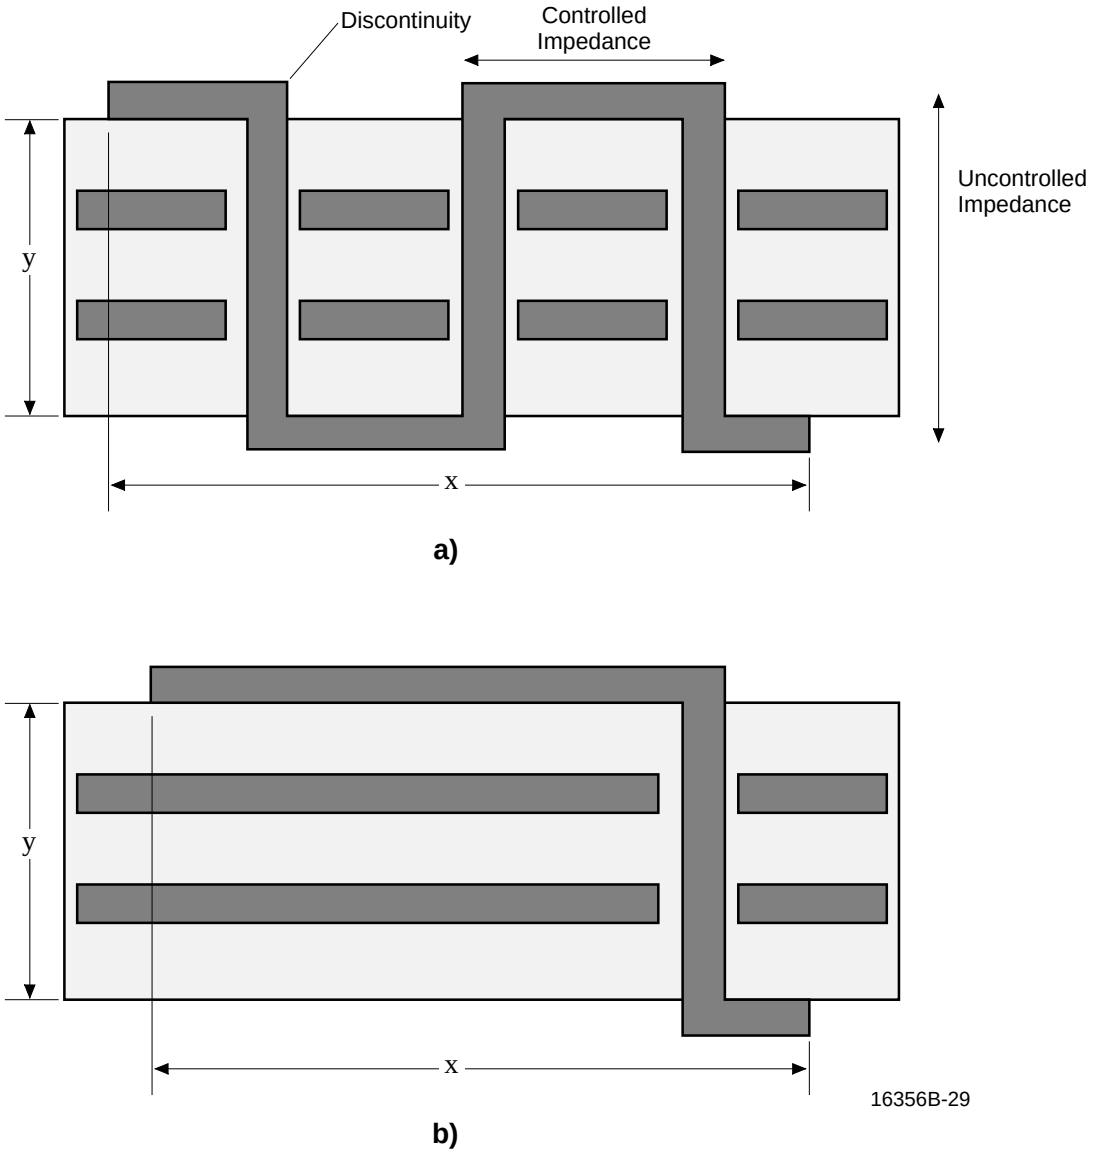

Figure 29. a) Excessive Number of Vias; b) Preferred Solution

Note that changing from an outer layer to an inside layer (or vice versa) generates an impedance change, since the design effectively is changed from stripline to microstrip (or vice versa). While it is theoretically possible to change geometries to compensate and keep impedances the same, it is very difficult to do so in production. The best results are obtained if outside signals remain outside, and inside signals remain inside.

2. Do Not Use Stubs or Ts

When laying out the signal lines, it is often convenient to run stubs or Ts to the devices, similar to Figure 30a. Stubs and Ts can be noise sources. If long enough, they are transmission lines with the main line as the source and are subject to the same type of reflections.

The signal lines should avoid long stubs and Ts. As long as the stubs are very short, a single line can be used with a single termination at the end, although Z_0 must then be derated to account for the distributed load. Given the example in Figure 30a, if the stubs are too long, the signal line could be made into two signal lines, as shown in Figure 30b. Both are transmission lines and require terminating; however, this is preferable to terminating each long stub individually.

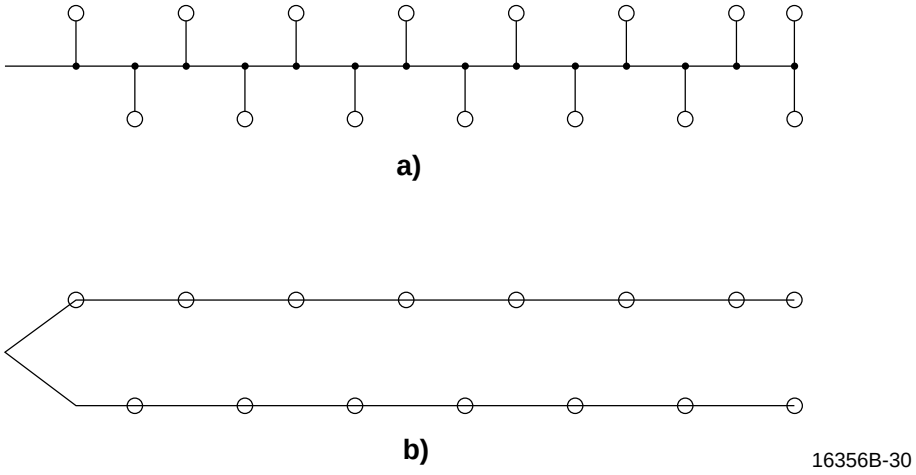

Figure 30. a) Stubs off of Transmission Line; b) Preferred Solution

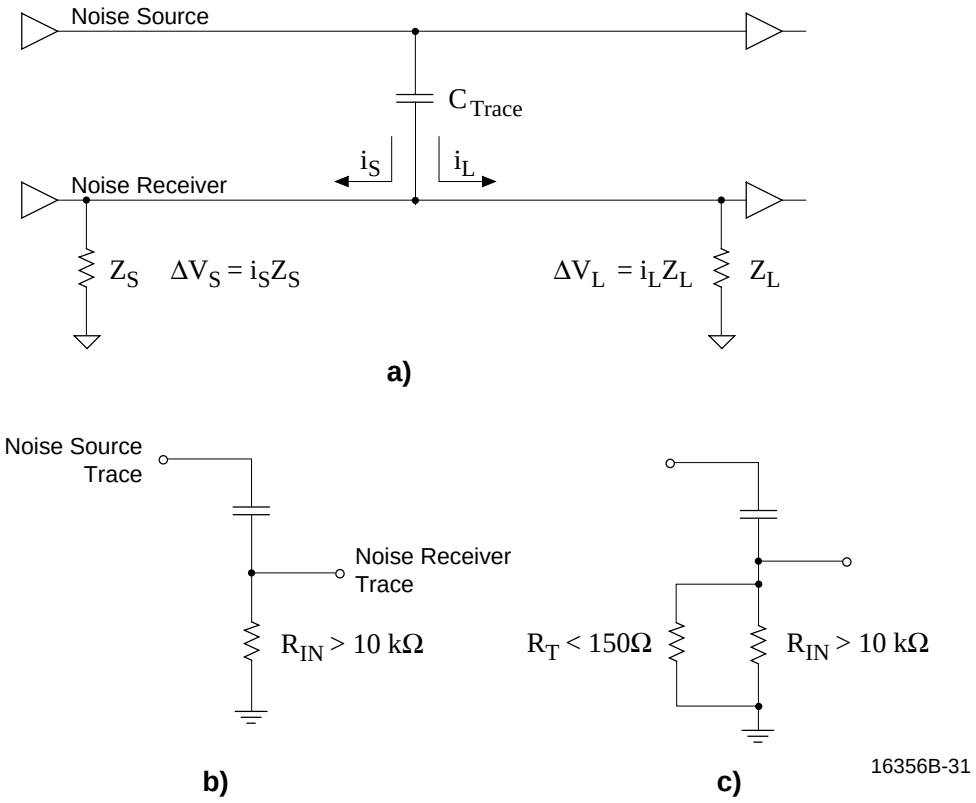

Figure 31. a) Capacitive Crosstalk; b) Equivalent Circuit; c) Solution

3. Crosstalk

Crosstalk is the unwanted coupling of signals between traces. It is either capacitive or inductive. Crosstalk can be handled effectively by following a few simple rules.

3.1 Capacitive Crosstalk

Capacitive crosstalk refers to the capacitive coupling of signals between signal lines. It occurs when the lines are close to each other for some distance.

The circuit representation in Figure 31 shows two signal lines, called the noise source and the noise receiver. Because of capacitance between the lines, noise on the source can be coupled onto the receiving line. This occurs in the form of current injected into the receiving line. In a transmission line, the current sees Z_0 in both directions, and propagates both ways, until it can be dissipated across the source and load. The voltage spike this causes on the line is determined by Z_0. When the current

pulse gets to Z_S and Z_L, it dissipates across these resistors with a voltage proportional to the impedance. If there is an impedance mismatch at the source or load, reflections occur. In the case of an unterminated load, the voltage spike across Z_L can be very large. Terminating the load can significantly reduce the voltage noise seen at the input of the next device.

Capacitive crosstalk can also be reduced by separating the traces. The farther apart the signal traces are, the less the capacitance, and the smaller the crosstalk.

Space constraints on the board may put limits on how far apart the signal lines can be placed. An alternate approach is to put a ground trace between adjacent-signal lines, as shown in Figure 32. The signal is now coupled to ground, not to the adjacent-signal line.

Note that the ground trace must be a solid ground. If it is only connected to the ground plane at the trace ends, the trace has a relatively high impedance. For good grounding, the ground trace should be connected to the ground plane with taps separated a quarter wavelength ($\lambda/4$) of the highest frequency component of the signal.

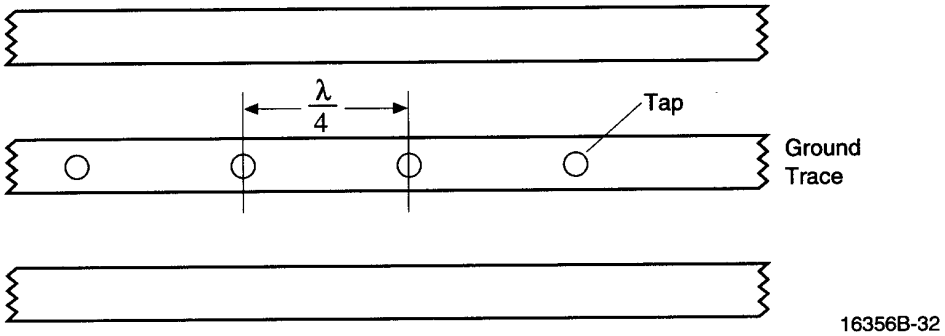

Figure 32. Isolating Traces with a Ground Trace

The wavelength is the distance the signal travels in a single period or:

$$\lambda = vel * Period$$

$$= \frac{1}{t_{PD}} * \frac{1}{freq}$$

With digital signals, the highest significant frequency harmonic of interest is usually assumed to be $1/\pi t_R$. Consider an example where t_R = 1.25 ns (possible for PAL16R8-4 devices). The upper-frequency component is:

$$f_{MAX} = \frac{1}{1.25\ ns * \pi}$$

$$= 255\ MHz$$

The distributed load delay for our example in section 2 was 4.14 ns/ft. λ is equal to the period divided by t_{PD}.

$$\lambda = \frac{1}{255\ MHz} * \frac{1}{4.14\ \frac{ns}{ft}} * \frac{12\ in}{ft}$$

$$= 11.4\ in$$

$$\lambda/4 = \frac{11.4}{4}$$

$$= 2.8\ in$$

For maximum isolation, the ground trace must have taps to the ground plane no more than 2.8 inches apart.

3.2 Inductive Crosstalk

Inductive crosstalk can be thought of as the coupling of signals between the primary and secondary coils of an unwanted transformer (Figure 33). The transformer windings are the current loops on the board (or system). These can be either artificial loops inadvertently created by inefficient layout (Figure 34a) or natural loops resulting from the combination of the signal path and the signal return path (Figure 34b). Artificial loops are sometimes hard to locate, but can be eliminated as shown in Figure 34c.

The amount of unwanted signal coupled to the load depends on the proximity and size of the loops, as well as the impedance of the affected load. The amount of energy transferred increases as the loops become larger and get closer together. The size of the signal seen at the load, on the secondary loop, increases with the load impedance.

3.2.1 Loop Size and Proximity

The inductance of a loop (L) increases with loop size. When two loops interact, one will have a primary inductance (LP) and the other will have a secondary inductance (LS), as shown in Figure 33b. Because the signal lines are not purposely designed to be transformers, the coupling is loose; however, it can create interference on the secondary loop.

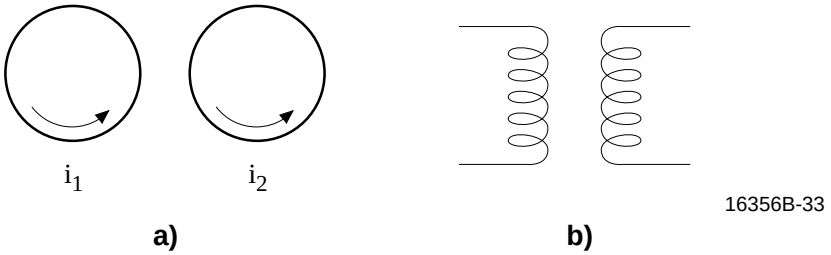

Figure 33. a) Inductive Crosstalk; b) Transformer Equivalent

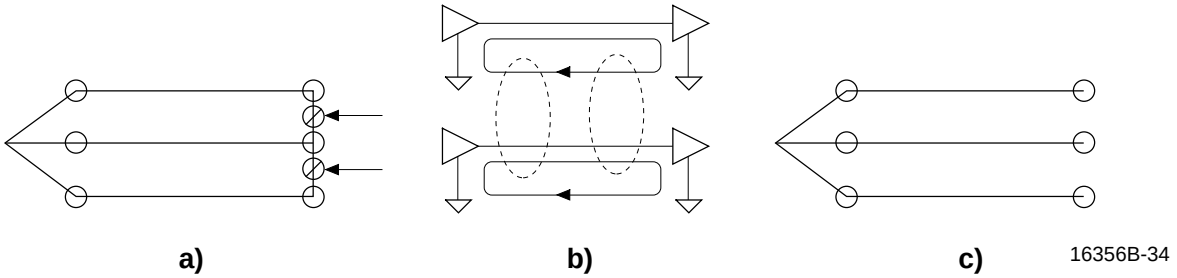

Figure 34. a) Artificial Loops; b) Schematic Equivalent; c) Solution

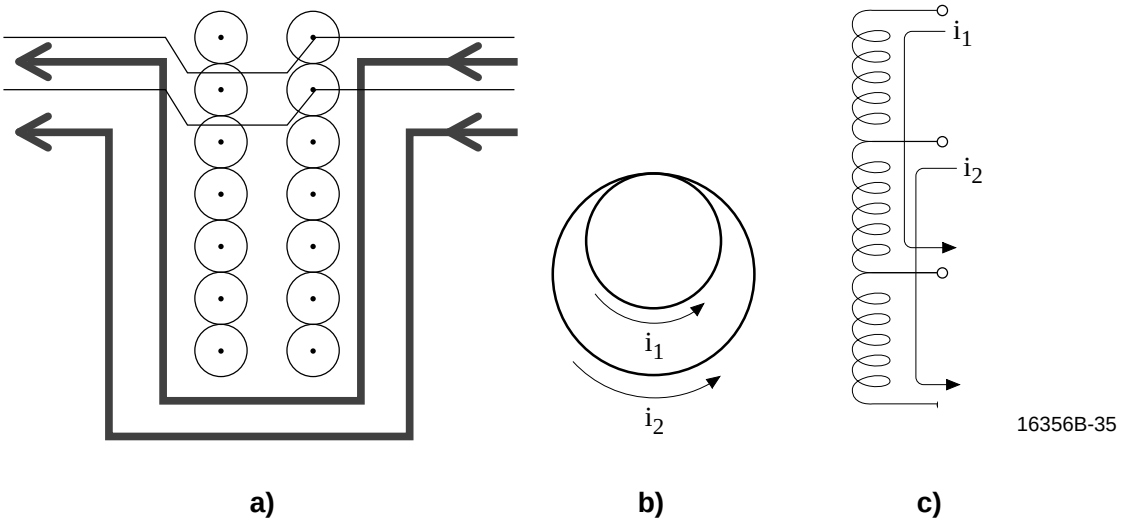

Figure 35. a) Common Return Path; b) Loop; c) Autotransformer Equivalent Circuit

If portions of the return paths of two signal lines coincide, the resulting loops might form an auto-transformer (Figure 35 a and c). An example of this is the VME-backplane example discussed above. Ensuring that each signal has its own return path can eliminate this source of crosstalk.

3.2.2 Load Impedance

If inductive crosstalk comes about due to artificial loops, the solution is to open the loops. Unfortunately, locating the loops can often be a challenge. If the crosstalk is generated by natural signal / return-signal loops, then clearly the loop cannot broken. But by keeping the load impedance low, the effect of the crosstalk can be minimized. Figure 36 shows a simplified schematic representation of a secondary "natural" loop with a load. Here Z_S is the intrinsic impedance of the secondary loop. Note the series current (i_S). Because the impedances are in series, i_S is the same everywhere in the loop. With a constant i_S, the voltage drop is largest across the largest impedance. On an unterminated line, this usually is the load at the end of the line; i.e., at the input of the receiving device.

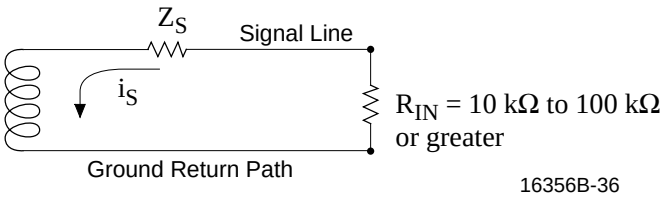

Figure 36. Series Inductive Loop

Large noise signals are most unwanted at the inputs, where noise signals should be minimal. If the maximum signal is developed across the largest impedance, the signal developed at the input can be reduced by terminating the signal line at the receiver end, which reduces R_{IN} to R_T.

R_T is usually in the 30 Ω to 150 Ω range. This reduction in R_{IN} is at least two orders of magnitude. The voltage drop across R_{IN} is reduced accordingly. The exact drop is difficult to predict because it depends on the value of Z_S , which is difficult to determine. But reducing R_{IN} by orders of magnitude should have a significant effect.

3.3 Crosstalk Solutions Summary

The following steps summarize the ways in which the effects of crosstalk can be minimized.

1. The effect of both capacitive and inductive crosstalk increases with load impedance. Thus all lines susceptible to interference due to crosstalk should be terminated at the line impedance.
2. Keeping the signal lines separated reduces the energy that can be capacitively coupled between signal lines.
3. Capacitive coupling can be reduced by separating the signal lines by a ground line. To be effective, the ground trace should be connected to the ground plane every λ/4 inches.
4. For inductive crosstalk, the loop size should be reduced as much as possible. Where possible, loops should be eliminated.
5. For inductive crosstalk, avoid situations where signal return lines share a common path.

4. ELECTRO-MAGNETIC INTERFERENCE (EMI)

EMI is becoming more critical with speed. High-speed devices are naturally more susceptible to interference. They accept fast glitches, which slower devices ignore. Even if the board or system is not susceptible, the FCC in the United States, along with VDE and CCITT in Europe, places severe limitations on the high-frequency noise (both radiated and line noise) that the board can generate.

The designer can reduce EMI through shielding, filtering, eliminating current loops, and reducing device speed where possible. Although shielding is outside the scope of this article, all the other issues are discussed as follows.

4.1 Loops

Current loops are an unavoidable part of every design. They act as antennae. Minimizing the effects of loops on EMI means minimizing the number of loops and the antenna efficiency of the loops. Do not create artificial loops; and keep the natural loops as small as possible.

1. Avoid artificial loops by ensuring that each signal line has only one path between any two points.
2. Use power planes whenever possible. Ground planes automatically result in the smallest natural current loop. When using ground planes, ensure that the signal-return line path is not blocked.

   If power buses are necessary, have the fast-signal lines run either over or next to a power bus.

4.2 Filtering

Filtering is standard for power lines. It can also be used on signal lines, but is recommended only as a last resort, when the source of the signal noise cannot be eliminated.

Three options are available for filtering: bypass capacitors, EMI filters, and ferrite beads. Bypass capacitors are discussed in section 1. EMI filters are commercially available filters; they are available over a wide frequency range. Ferrite beads are ferrite ceramics that add inductance to any wire within their proximity. They are used as high-frequency suppressors.

4.2.1 EMI Filters

EMI filters are commercially manufactured devices designed to attenuate high-frequency noise. They are used primarily to filter out noise in power lines. They act to isolate the power outside the system (referred to as the line) from the power inside the system (referred to as the load). Their effect is bi-directional: they filter out noise going into, and coming out of, the device or board.

EMI filters consist of combinations of inductors and capacitors. In general, the configuration to use depends on the impedance of the nodes to be connected. A capacitor should be connected to a high-impedance node; an inductor should be connected to a low-impedance node. EMI filters are available in variations of the following configurations: feedthrough capacitor, L-Circuit, PI-Circuit, and T-Circuit.

- The feedthrough capacitor's only component is a capacitor (Figure 37a). It is a good choice when the impedances connected to the filter are high. Note that it provides no high-frequency current isolation between nodes.
- The L-Circuit has an inductor on one side of a capacitance (Figure 37b). It works best when the line and load have a large difference in impedance. The inductive element is connected to the lowest impedance.

- The PI-Circuit has an inductor surrounded by two capacitors (Figure 37c). PI filters are best when the line and load impedances are high and when high levels of attenuation are needed.
- The T-filter has inductors on either side of the capacitor in a T fashion (Figure 37d). It is a good choice when both line and load impedances are low.

LC filters are rated according to insertion loss, which is the amount of signal lost due to the insertion of the filter. Insertion loss is usually stated in decibels. Filter manufacturers provide graphs of their filters over prescribed frequency ranges.

4.2.2 Ferrite Noise Suppressors

Ferrite noise suppressors are ferrite ceramics placed in proximity to the conducting material. They are available as beads for single wires and clamps for cables. When using beads, the wire is placed through a hole in the bead (Figure 38a). When using clamps, the ferrite material is clamped around the cable (Figure 38b). Clamps are popular with ribbon cable.

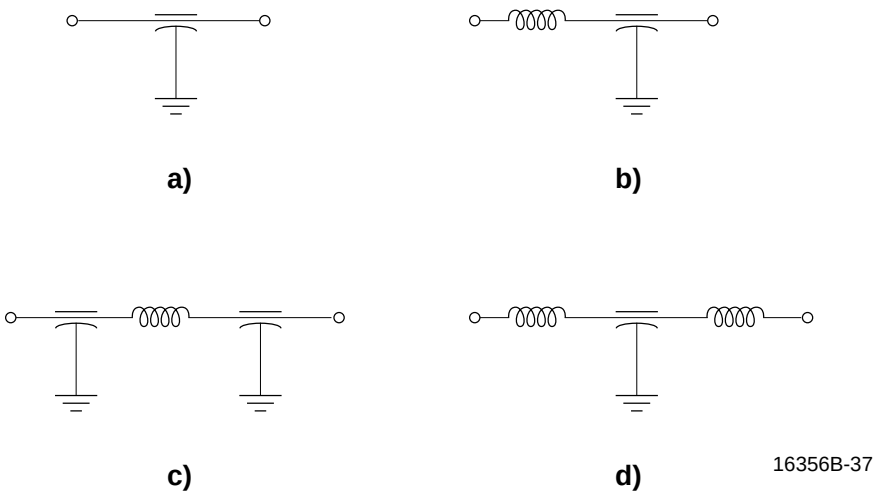

Figure 37. Line-Noise Filters. a) Capacitor; b) LC Filter; c) PI Filter; d) T Filter

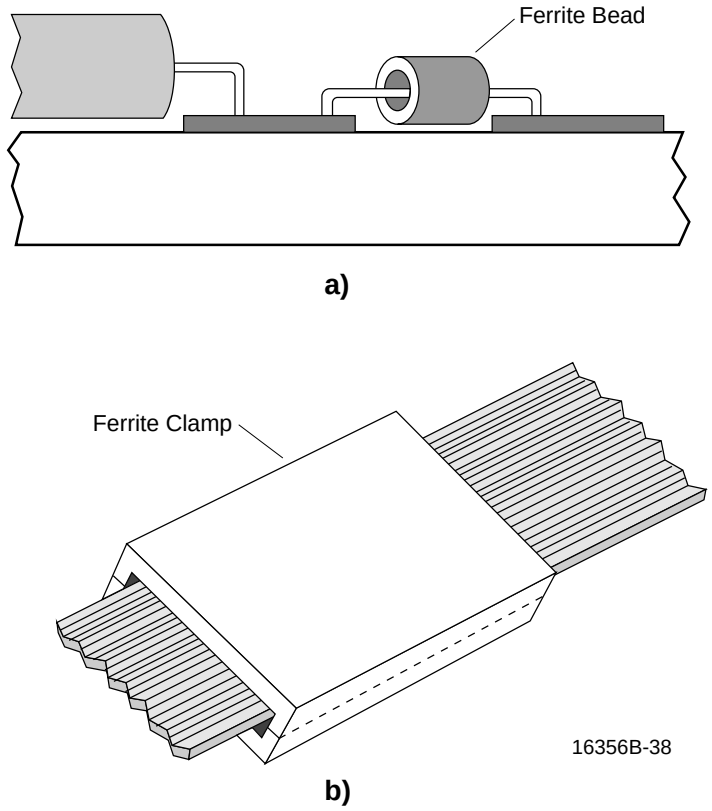

Figure 38. a) Ferrite Bead; b) Ferrite Clamp

Ferrite suppressors work by adding inductance in series with the line (Figure 39). Ferrite manufactures supply graphs similar to those in Figure 40, which shows the added impedance as a function of frequency. The system designer must determine the insertion loss. The formula is:

$$\text{Loss (db)} = 20 \log_{10} \frac{Z_S + Z_L + Z_F}{Z_S + Z_L}$$

where Z_S = Source Impedance
Z_L = Load Impedance
Z_F = Ferrite Impedance

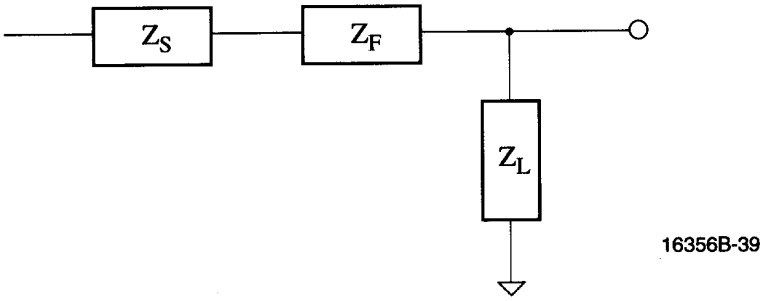

Figure 39. Ferrite Filter Equivalent Circuit

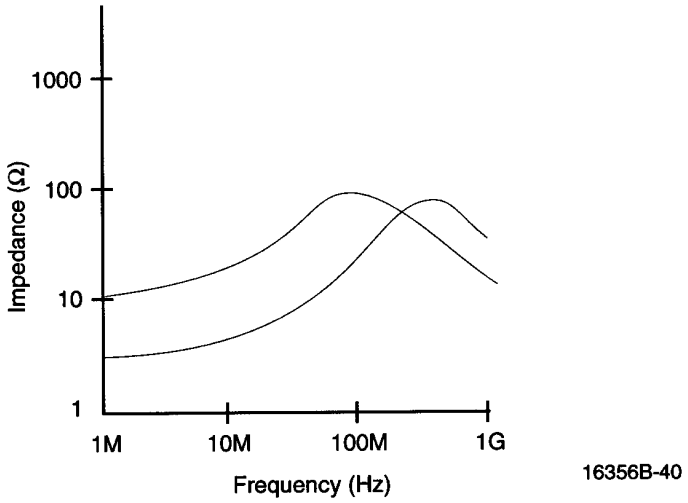

Figure 40. Frequency Response of Ferrite Filter

Ferrite suppressors add inductance to the line without adding DC resistance. This makes an ideal choice for line-noise suppressors on the V_{CC} pins of devices.

Because ferrite beads are small and easy to handle, they can sometimes be used in signal lines to suppress high-frequency noise signals. This is not recommended for two reasons: first, it masks the cause of most problems; second because it might affect the edge rates of the signal. However, when the board is already laid out, ferrite beads can be used on noisy signal lines as a last resort.

4.3 Device Speed

The less energy a device generates in a given frequency range, the less noise can be radiated in that range. Faster devices, by definition, have shorter transition times. Because shorter transition times have more energy in the high-frequency range, faster devices can generate more high-frequency noise.

Figure 41a is an outline of the Fourier transform of a square wave (Figure 41b). There are two corners of interest: $1/\pi tL$(this frequency is determined by the period of the signal) and $1/\pi tf$ (determined by the transition time of the signal; this is also the frequency we used in determining the wavelength in the discussion on capacitive coupling). After $1/\pi tf$, the curve drops off very rapidly. For practical purposes, $1/\pi tf$ is the highest significant frequency component of the signal.

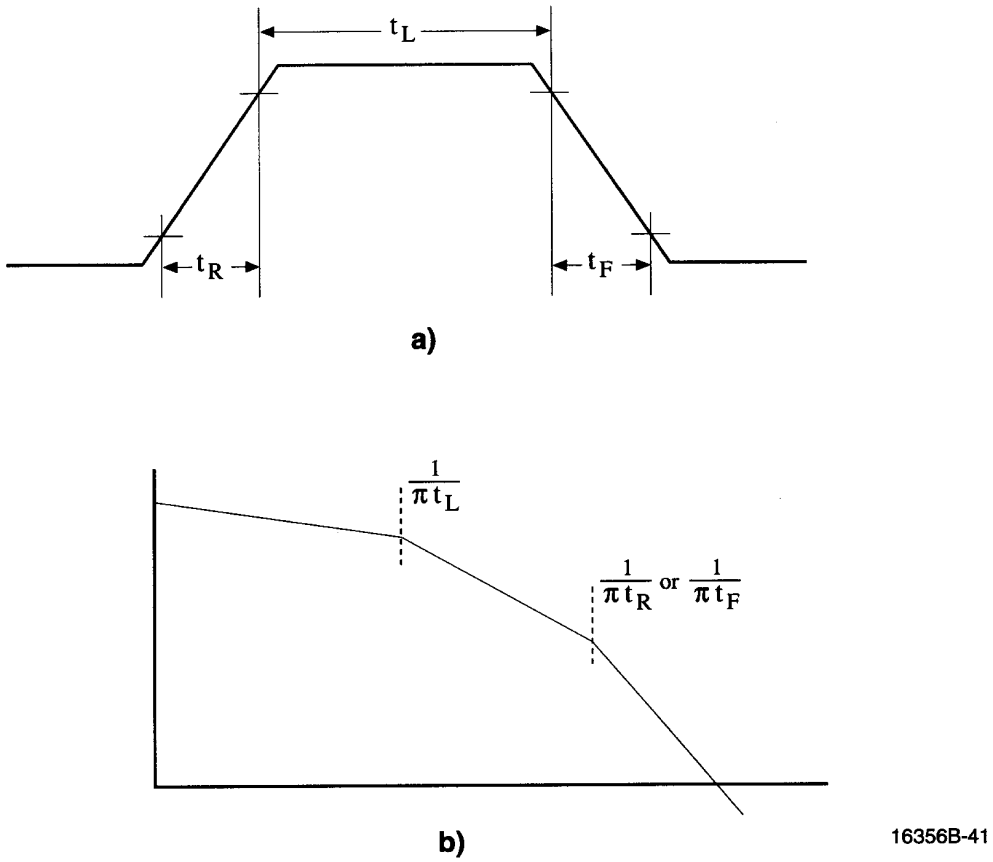

Figure 41. a) Single Pulse; b) Fourier Transform of Pulse

For example, the PAL16R8-4 series has a typical transition time of 2 ns. It can be as short as 1.25 ns. The frequency component of the edge is:

$$f = \frac{1}{\pi * 1.25\ ns}$$

$$= 254\ MHz$$

The output signal has a high-frequency component of 254 MHz, regardless of the clock frequency.

Because of the high-frequency component, the board might require extra filtering and possibly shielding in order to comply with EMI emissions restrictions required by regulatory agencies.

If the system speed requirements are high enough (for example, clock rates over 80 MHz), devices this fast must be used, and the extra effort required to meet compliance is justified. However, if a slower device can meet system requirements, it should be used. By virtue of the longer transition time, the slower device generates less energy at the higher frequencies. In general, try to use devices that are fast enough to meet the system requirements, but no faster.

SUMMARY

While faster technologies provide the theoretical possibility of faster systems, extra care must be taken to turn this possibility into reality. The largest noise components can be eliminated by addressing the following:

- integrity and stability of power and ground;
- termination and careful layout of transmission lines to eliminate reflections;
- termination and careful layout to reduce the effects of capacitive and inductive crosstalk;
- noise suppression for compliance with radiation regulations.

There are many other second-order issues that could be addressed, but that are beyond the scope of the application note. Some references for additional information are listed below.

1. Sherman Lee, Mark McClain, Dave Stoenner. "Am29000 32-Bit Streamlined Instruction Processor Memory Design Handbook," Advanced Micro Devices Inc., Sunnyvale, CA, Appendix A, Memory Array Loading Calculations.
2. William R. Blood Jr. "ECL Systems Design Handbook," Motorola Semiconductor Products Inc., Mesa, AZ, May, 1983 (Fourth Edition) Chapters 3 and 7.
3. Ramo, Whinnery, and Van Duzer, "Fields and Waves in Communications Electronics," John Whilley & Sons, 1965, Chapter 1.