High-Speed Board Layout Guidelines

September 2003, ver. 1.1 Application Note 224

Introduction

As both device pin density and system frequency increase, printed circuit board (PCB) layout becomes more complex. A successful high-speed board must effectively integrate the devices and other elements while avoiding signal transmission problems associated with high-speed I/O standards. Because Altera® devices feature fast I/O pins, a wide variety of high-speed features, and edge rates less than a hundred picoseconds, it is imperative that an effective design successfully:

- Reduces system noise by filtering and evenly distributing power to all devices
- Terminates the signal line to diminish signal reflection
- Minimizes crosstalk between parallel traces
- Reduces the effects of ground bounce
- Matches impedance

This application note discusses the following issues and provides guidelines for successful, effective board design using Altera devices:

- Material selection
- Transmission line
- Routing scheme for minimizing crosstalk and maintaining signal integrity
- Termination schemes
- Simultaneous switching noise (SSN)
- Additional FPGA-specific board design/signal integrity tips

Material Selection

Fast edge rates contribute to noise and crosstalk, depending on the PCB dielectric construction material. Dielectric material can be assigned a dielectric constant (ε_r) that is related to the force (i.e., Equation 1) of attraction between two opposite charges separated by a distance in a uniform medium.

Equation 1:

$$F = \frac{Q_1 Q_2}{4\pi\varepsilon r^2}$$

Where Q_1, Q_2 = charges, r = distance between the charges (m), F = force (N), ε = permittivity of dielectric (F/m)

Each PCB substrate has a different relative dielectric constant. The dielectric constant is the permittivity of a relative to that of empty space (i.e., Equation 2).

Equation 2:

$$\varepsilon_r = \frac{\varepsilon}{\varepsilon_0}$$

Where ε_r = dielectric constant, ε_0 = permittivity of empty space (F/m), ε = permittivity (F/m).

The dielectric constant compares the effect of an insulator on the capacitance of a conductor pair, to the capacitance of the conductor pair in a vacuum. The dielectric constant affects the impedance of a transmission line, and signals can propagate faster in materials that have a lower ε_r.

A high-frequency signal that propagates through a long line on the PCB from driver to receiver is severely affected by the loss tangent of the dielectric material. A large loss tangent means higher dielectric absorption. Material with a high loss tangent value affects the high-frequency signal on a long line. Dielectric absorption increases attenuation at higher frequencies. Table 1 shows the loss tangent value for FR-4 and GETEK materials.

The most widely used dielectric material for PCBs is FR-4, a glass laminate with epoxy resin that meets a wide variety of processing conditions. The ε_r for FR-4 is between 4.1 and 4.5. GETEK is another material that can be used in high-speed boards. GETEK is composed of epoxy and resin (polyphenylene oxide) and has an ε_r between 3.6 and 4.2.

Table 1. Loss Tangent Value of FR4 & GETEK Materials

Manufacturer	Material	Loss Tangent Value
GE Electromaterials	GETEK	0.010 @ 1 MHz
Isola Laminate Systems	FR-4	0.019 @ 1 MHz

Transmission Line

The transmission line is a trace, and has a distributed mixture of resistance (R), inductance (L), and capacitance (C). There are two types of transmission line layouts:

- Microstrip
- Stripline

Figure 1 shows a microstrip layout, which refers to a trace routed as the top or bottom layer of a PCB and has only one voltage-reference plane (i.e., power or GND). Figure 2 shows a stripline layout, which uses a trace routed on the inside layer of a PCB and has two voltage-reference planes (i.e., power and/or GND).

Figure 1. Microstrip Transmission Line Layout *Note (1)*

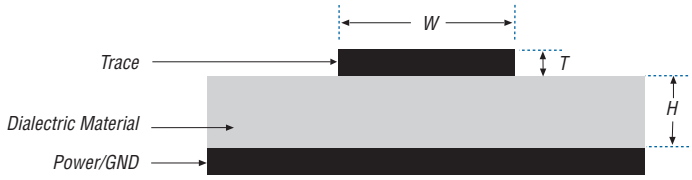

Figure 2. Stripline Transmission Line Layout *Note (2)*

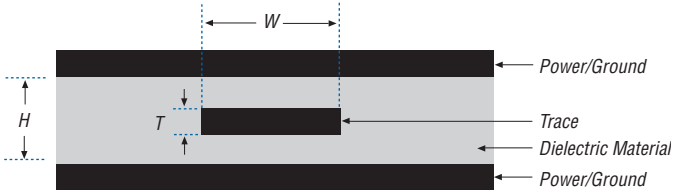

Notes to Figures 1 & 2:

(1) W = width of trace, T = thickness of trace, and H = height between trace and reference plane.

(2) W = width of trace, T = thickness of trace, and H = height between trace and two reference planes.

Impedance Calculation

Any circuit trace on the PCB has characteristic impedance associated with it. This impedance is dependent on width (W) of the trace, thickness (T) of the trace, dielectric constant (ε_r) of the material used, and height (H) between the trace and reference plane.

Microstrip Impedance

A circuit trace routed on an outside layer of the PCB with a reference plane (i.e., GND or V_{CC}) below it, constitutes a microstrip layout. Use Equation 3 to calculate the impedance of a microstrip trace layout.

Equation 3: .

$$Z_0 = \frac{87}{\sqrt{\varepsilon_r + 1.41}} \ln\left(\frac{5.98 \times H}{0.8W + T}\right) \Omega$$

Using typical values of W = 8 mil, H = 5 mil, T = 1.4 mil, ε_r and (FR-4) = 4.1 with Equation 3 and solving for microstrip impedance (Z_o) yields:

$$Z_0 = \frac{87}{\sqrt{4.1 + 1.41}} \ln\left(\frac{5.98 \times (5)}{0.8(8) + 1.4}\right) \Omega$$

$$Z_0 \sim 50\ \Omega$$

The measurement unit in Equation 3 is mils (i.e., 1 mil = .001 inches). Also, copper (Cu) trace thickness (T) is usually measured in ounces (i.e., 1 oz = 1.4 mil).

Figure 3 shows microstrip trace impedance vs. trace width (W) using the values in Equation 3, keeping dielectric height and trace thickness constant.

Figure 3. Microstrip Trace Impedance with Changing Trace Width

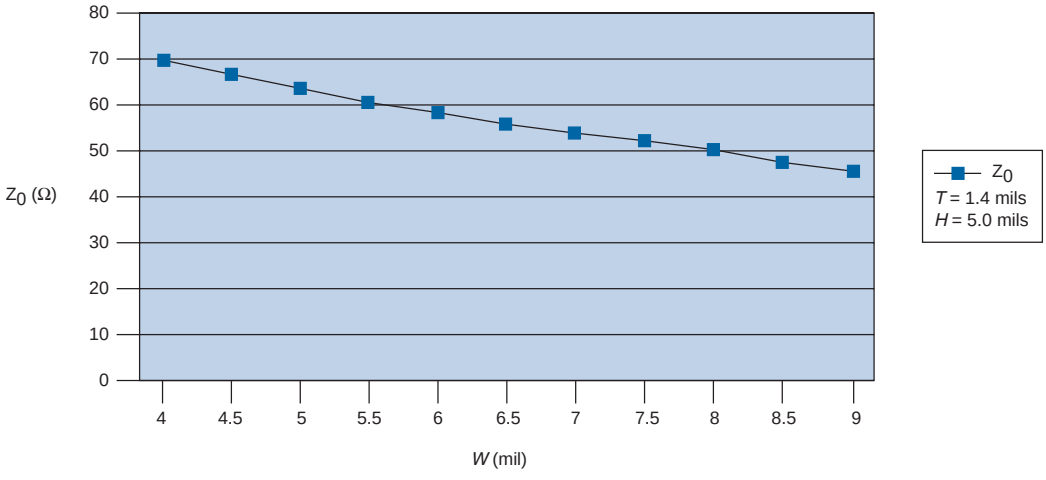

Figure 4 shows microstrip trace impedance vs. height (*H*), using the values in Equation 3, keeping trace width and trace thickness constant.

Figure 4. Microstrip Trace Impedance with Changing Height

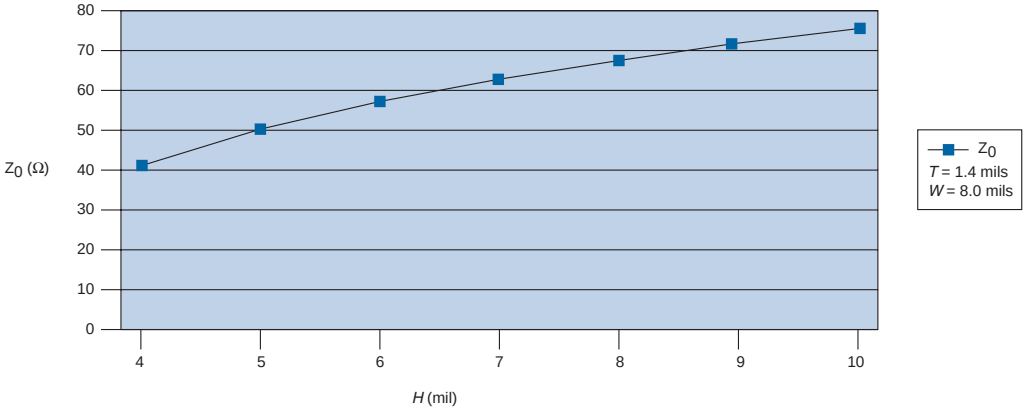

The impedance graphs show that the change in impedance is inversely proportional to trace width and directly proportional to trace height above the ground plane.

Figure 5 plots microstrip trace impedance vs. trace thickness (*T*) using the values in Equation 3, keeping trace width and dielectric height constant. Figure 5 shows that as trace thickness increases, trace impedance decreases.

Figure 5. Microstrip Trace Impedance with Changing Trace Thickness

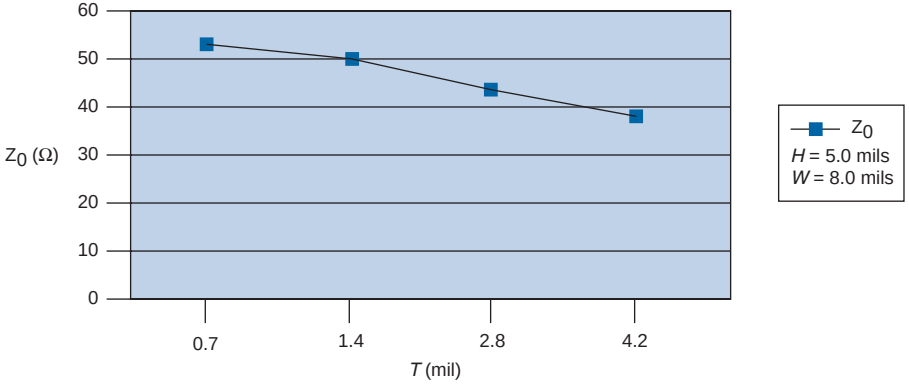

Stripline Impedance

A circuit trace routed on the inside layer of the PCB with two low-voltage reference planes (i.e., power and/or GND) constitutes a stripline layout. You can use Equation 4 to calculate the impedance of a stripline trace layout.

Equation 4:

$$Z_0 = \frac{60}{\sqrt{\varepsilon_r}} \ln\left(\frac{4H}{0.67\pi(T + 0.8W)}\right) \Omega$$

Using typical values of W = 9 mil, H = 24 mil, T = 1.4 mil, ε_r and (FR-4) = 4.1 with Equation 4 and solving for stripline impedance (Z_0) yields:

$$Z_0 = \frac{60}{\sqrt{4.1}} \ln\left(\frac{4\,(24)}{0.67\pi\,(1.4) + 0.8(9)}\right) \Omega$$

$$Z_0 \sim 50\ \Omega$$

Figure 6 shows impedance vs. trace width using Equation 4, keeping height and thickness constant for stripline trace.

Figure 6. Stripline Trace Impedance with Changing Trace Width

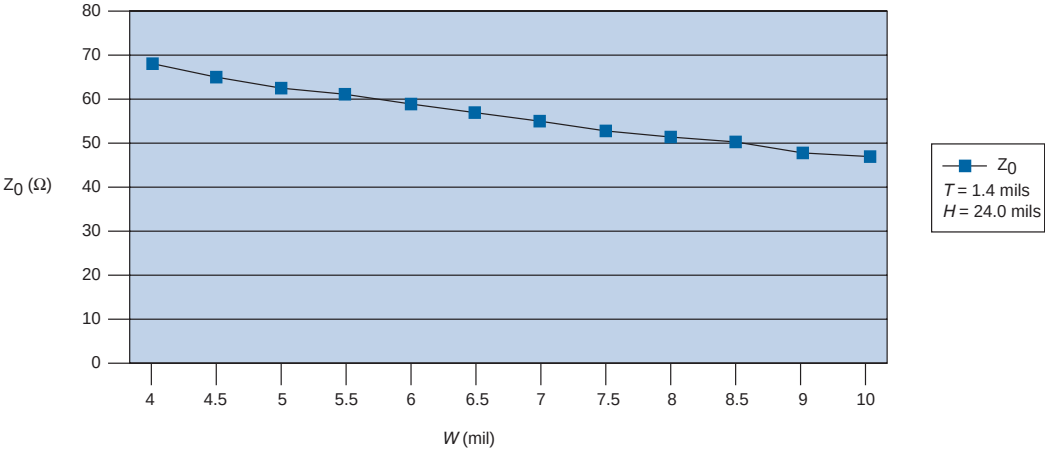

Figure 7 shows stripline trace impedance vs. dielectric height (*H*) using Equation 4, keeping trace width and trace thickness constant.

Figure 7. Stripline Trace Impedance with Changing Dielectric Height

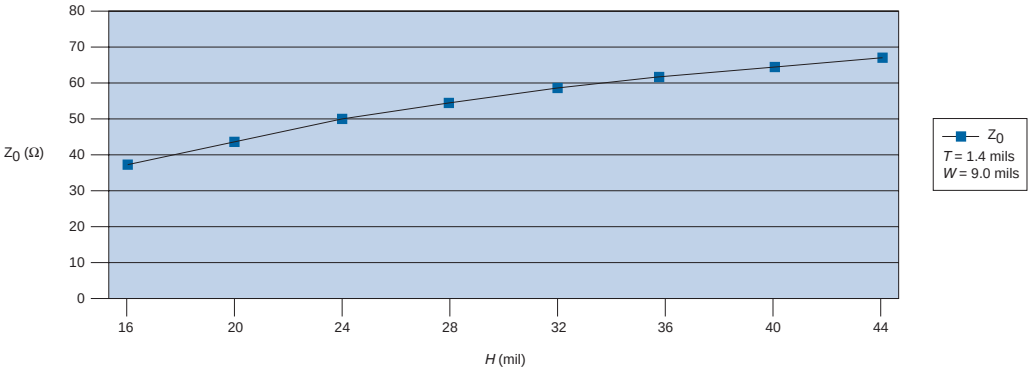

As with microstrip layout, the stripline layout impedance also changes inversely proportional to line width and directly proportional to height. However, the rate of change with trace height above GND is much slower in a stripline layout compared with a microstrip layout. A stripline layout has a signal sandwiched by FR-4 material, whereas a microstrip layout has one conductor open to air. This exposure causes a higher, effective dielectric constant stripline layout compared to microstrip layouts. Thus, to achieve the same impedance, the dielectric span must be greater in stripline layouts compared with microstrip layouts. Therefore, stripline-layout PCBs with controlled impedance lines are thicker than microstrip-layout PCBs.

Figure 8 shows stripline trace impedance vs. trace thickness using Equation 4, keeping trace width and dielectric height constant. Figure 8 shows that the characteristic impedance decreases as the trace thickness increases.

Figure 8. Stripline Trace Impedance with Changing Trace Thickness

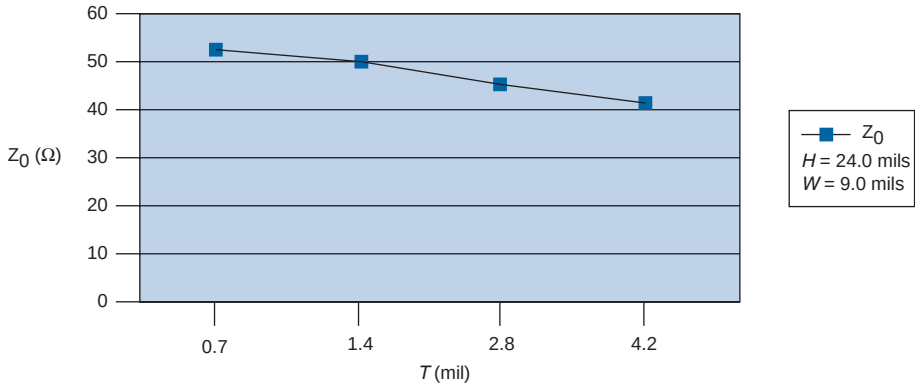

Propagation Delay

Propagation delay (t_{PD}) is the time required for a signal to travel from one point to another. Transmission line propagation delay is a function of the dielectric constant of the material.

Microstrip Layout Propagation Delay

You can use Equation 5 to calculate the microstrip trace layout propagation delay.

Equation 5:

$$t_{PD}\ (\text{microstrip}) = 85\sqrt{0.475\varepsilon_r + 0.67}$$

Stripline Layout Propagation Delay

You can use Equation 6 to calculate the stripline trace layout propagation delay.

Equation 6:

$$t_{PD}\ (\text{stripline}) = 85\sqrt{\varepsilon_r}$$

Figure 9 shows the propagation delay vs. the dielectric constant for microstrip and stripline traces. As the ε_r increases, the propagation delay (t_{PD}) also increases.

Figure 9. Propagation Delay vs. Dielectric Constant for Microstrip & Stripline Traces

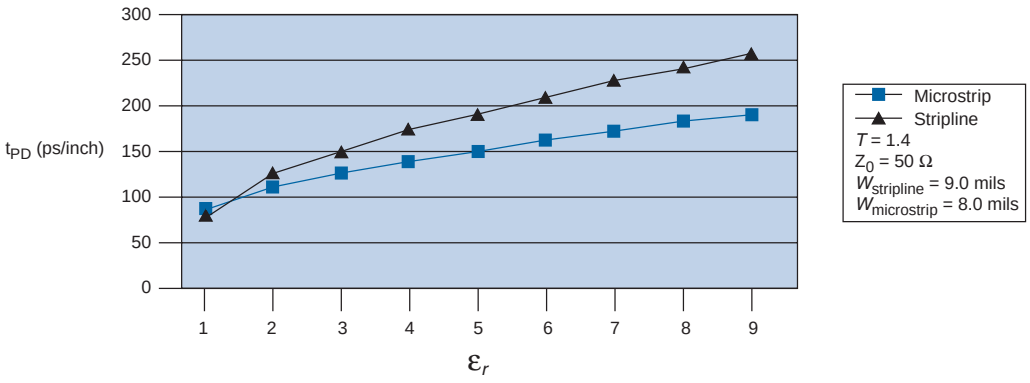

Pre-Emphasis

Typical transmission media like copper trace and coaxial cable have low pass characteristics, so they attenuate higher frequencies more than lower frequencies. A typical digital signal that approximates a square wave contains high frequencies near the switching region and low frequencies in the constant region. When this signal travels through low pass media, its higher frequencies are attenuated more than the lower frequencies, resulting in increased signal rise times. Consequently, the eye opening narrows and the probability of error increases.

The high-frequency content of a signal is also degraded by what is called "skin effect." The cause of skin effect is the high-frequency current that flows primarily on the surface (skin) of a conductor. The changing current distribution causes the resistance to increase as a function of frequency.

You can use pre-emphasis to compensate for the skin effect. By Fourier analysis, a square wave signal contains an infinite number of frequencies. The high frequencies are located in the low-to-high and high-to-low transition regions and the low frequencies are located in the flat (constant) regions. Increasing the signal's amplitude near the transition region emphasizes higher frequencies more than the lower frequencies. When this pre-emphasized signal passes through low pass media, it will come out with minimal distortion, if you apply the correct amount of pre-emphasis. See Figure 10 for a graphical illustration of this concept.

Figure 10. Input & Output Signals with & without Pre-Emphasis

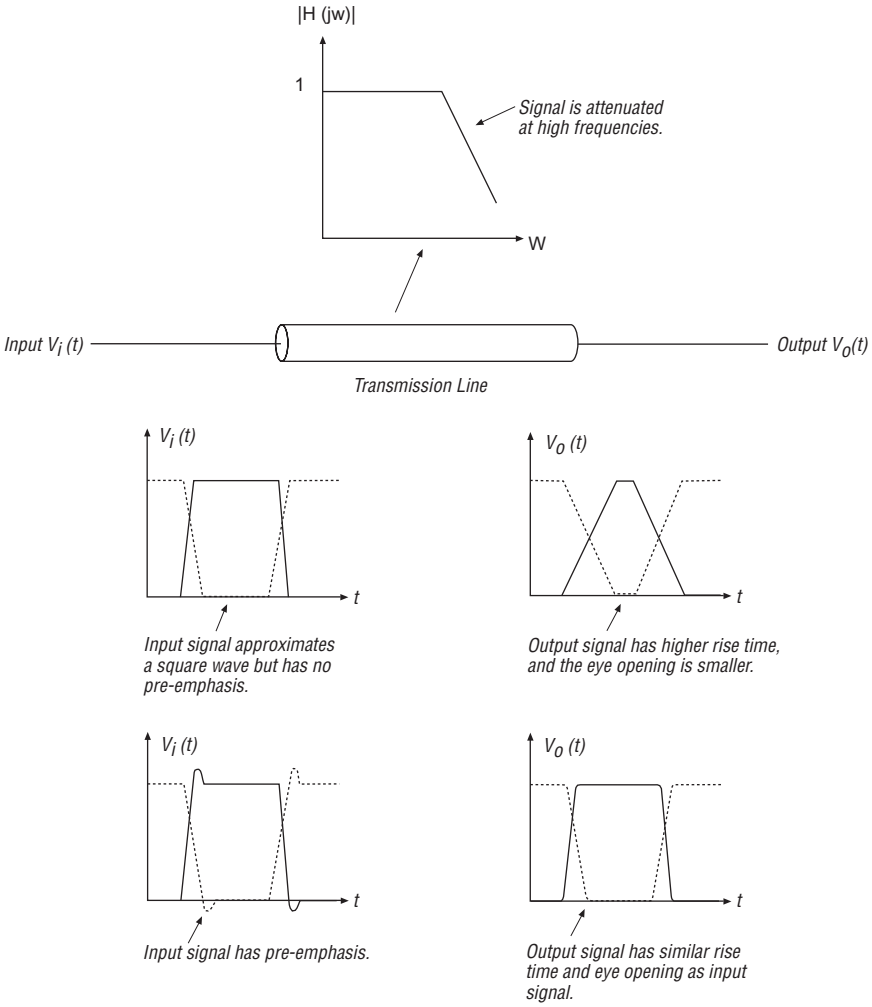

Stratix™ GX devices provide programmable pre-emphasis to compensate for variable lengths of transmission media. You can set the pre-emphasis to be between 5% and 25%, depending on the value of the output differential voltage (V_{OD}). Table 2 shows the available Stratix GX programmable pre-emphasis settings.

Table 2. Programmable Pre-Emphasis with Stratix GX Devices

V_{OD}	Pre-emphasis Setting				
	5%	10%	15%	20%	25%
400	420	440	460	480	500
480	504	528	552	576	600
600	630	660	690	720	750
800	840	880	920	960	1,000
960	1,008	1,056	1,104	1,152	1,200
1,000	1,050	1,100	1,150	1,200	1,250
1,200	1,260	1,320	1,380	1,440	1,500
1,400	1,470	1,540	-	-	-
1,440	1,512	1,584	-	-	-
1,500	1,575	-	-	-	-
1,600	-	-	-	-	-

Routing Scheme For Minimizing Crosstalk & Maintaining Signal Integrity

Crosstalk is the unwanted coupling of signals between parallel traces. Proper routing and layer stack-up through microstrip and stripline layouts can minimize crosstalk.

To reduce crosstalk in dual-stripline layouts, which have two signal layers next to each other (see Figure 11), route all traces perpendicular, increase the distance between the two signal layers, and minimize the distance between the signal layer and the adjacent reference plane.

Figure 11. Dual-Stripline Layout

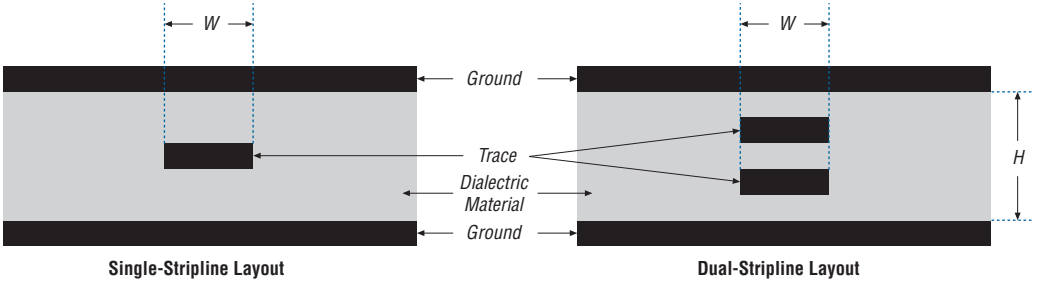

Use the following steps to reduce crosstalk in either microstrip or stripline layouts:

- Widen spacing between signal lines as much as routing restrictions will allow. Try not to bring traces closer than three times the dielectric height.
- Design the transmission line so that the conductor is as close to the ground plane as possible. This technique will couple the transmission line tightly to the ground plane and help decouple it from adjacent signals.
- Use differential routing techniques where possible, especially for critical nets (i.e., match the lengths as well as the gyrations that each trace goes through).
- If there is significant coupling, route single-ended signals on different layers orthogonal to each other.
- Minimize parallel run lengths between single-ended signals. Route with short parallel sections and minimize long, coupled sections between nets.

Crosstalk also increases when two or more single-ended traces run parallel and are not spaced far enough apart. The distance between the centers of two adjacent traces should be at least four times the trace width, as shown in Figure 12. To improve design performance, lower the distance between the trace and the ground plane to under 10 mils without changing the separation between two traces.

Figure 12. Separating Traces for Crosstalk

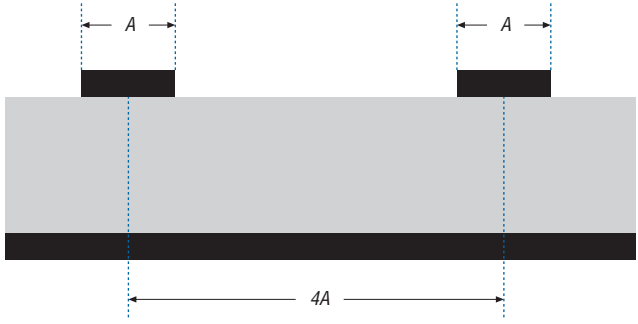

Compared with high dielectric materials, low dielectric materials help reduce the thickness between the trace and ground plane while maintaining signal integrity. Figure 13 plots the relationship of height vs. dielectric constant using Equations 3 and 4, keeping impedance, width, and thickness constant.

Figure 13. Height vs. Dielectric Constant

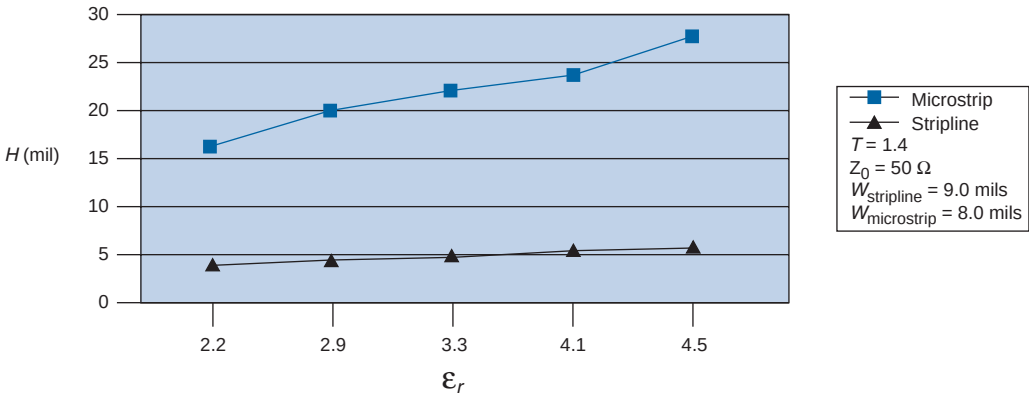

Signal Trace Routing

Proper routing helps to maintain signal integrity. To route a clean trace, you should perform simulation with good signal integrity (SI) tools. The following section describes the two different types of signal traces available for routing:

- Single-ended trace
- Differential pair trace

Single-Ended Trace Routing

A single-ended trace connects the source and the load/receiver. Single-ended traces are used in general point-to-point routing, clock routing, low-speed, and non-critical I/O routing. This section discusses different routing schemes for clock signals. You can use the following types of routing to drive multiple devices with the same clock.

- Daisy chain routing
  - With stub
  - Without stub
- Star routing
- Serpentine routing

Use the following guidelines to improve the clock transmission line's signal integrity:

- Keep clock traces as straight as possible. Use arc-shaped traces instead of right-angle bends.
- Do not use multiple signal layers for clock signals.
- Do not use vias in clock transmission lines. Vias can cause impedance change and reflection.
- Place a ground plane next to the outer layer to minimize noise. If you use an inner layer to route the clock trace, sandwich the layer between reference planes.
- Terminate clock signals to minimize reflection.
- Use point-to-point clock traces as much as possible.

Daisy Chain Routing With Stubs

Daisy chain routing is a common practice in designing PCBs. One disadvantage of daisy chain routing is that stubs, or short traces, are usually necessary to connect devices to the main bus (see Figure 14). If a stub is too long, it will induce transmission line reflections and degrade signal quality. Therefore, the stub length should not exceed the following conditions:

$$TD_{stub} < (T_{10\% \text{ to } 90\%})/3$$

Where TD_{stub} = Electrical delay of the stub

$T_{10\% \text{ to } 90\%}$ = Rise or fall time of signal edge

For a 1-ns rise-time edge, the stub length should be less than 0.5 inches (see "References" on page 33). If your design uses multiple devices, all stub lengths should be equal to minimize clock skew. Figure 14 shows stub routing. If possible, you should avoid using stubs in your PCB design. For high-speed designs, even very short stubs can create signal integrity problems.

Figure 14. Daisy Chain Routing with Stubs

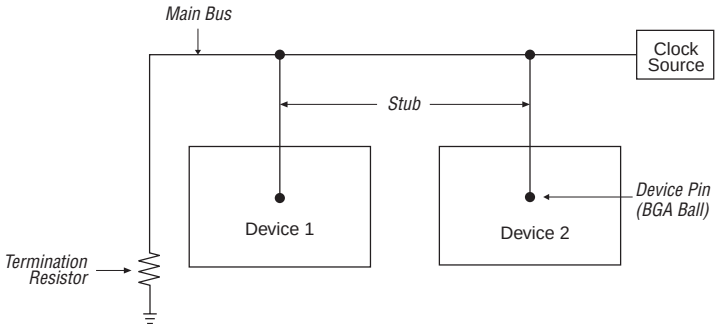

Figures 15 through 17 show the SPICE simulation with different stub length. As the stub length decreases, there is less reflection noise, and the eye opening increases due to less reflection noise.

Figure 15. Stub Length = 0.5 Inch

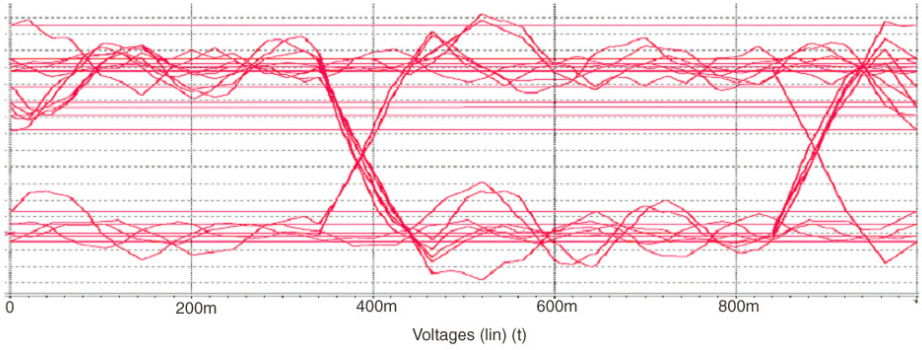

Figure 16. Stub Length = 0.25 Inch

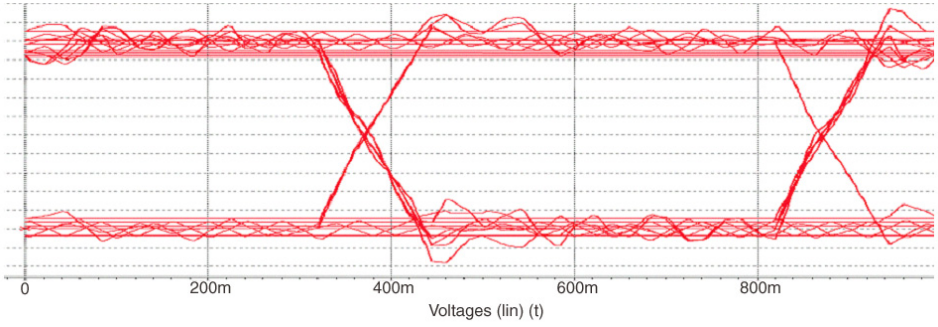

Figure 17. Stub Length = Zero Inches

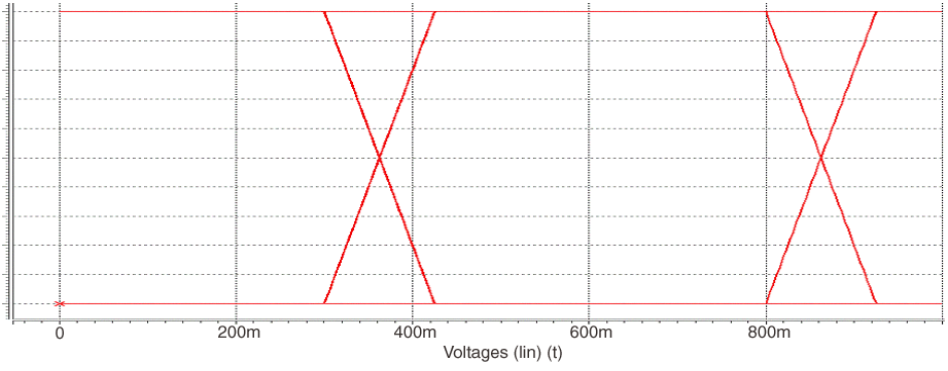

Daisy Chain Routing without Stub

Figure 18 shows daisy chain routing with the main bus running through the device pins, eliminating stubs. This layout removes the risk of impedance mismatch between the main bus and the stubs, minimizing signal integrity problems.

Figure 18. Daisy Chain Routing without Stubs

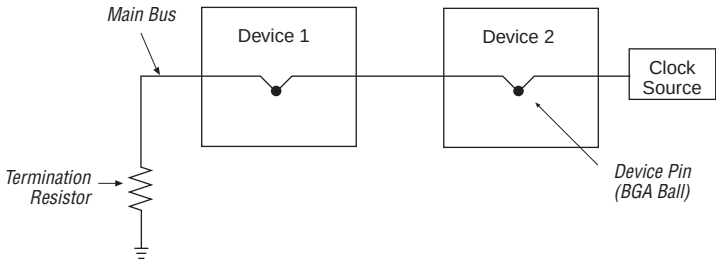

Star Routing

In star routing, the clock signal travels to all the devices at the same time (see Figure 19). Therefore, all trace lengths between the clock source and devices must be matched to minimize the clock skew. Each load should be identical to minimize signal integrity problems. In star routing, you must match the impedance of the main bus with the impedance of the long trace that connects to multiple devices.

Figure 19. Star Routing

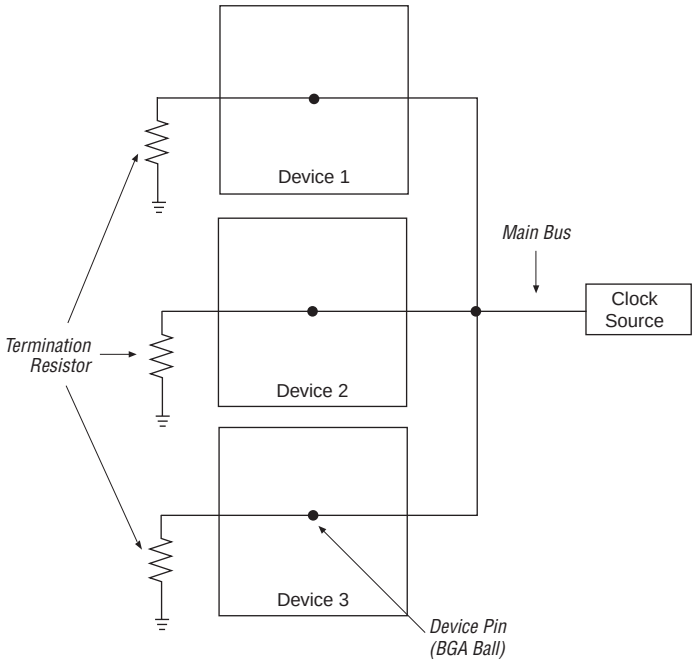

Serpentine Routing

When a design requires equal-length traces between the source and multiple loads, you can bend some traces to match trace lengths (see Figure 20). Improper trace bending affects signal integrity and propagation delay. To minimize crosstalk, ensure that $S \geq 3 \times H$, where S is the spacing between the parallel sections and H is the height of the signal trace above the reference ground plane. See Figure 21.

Figure 20. Serpentine Routing

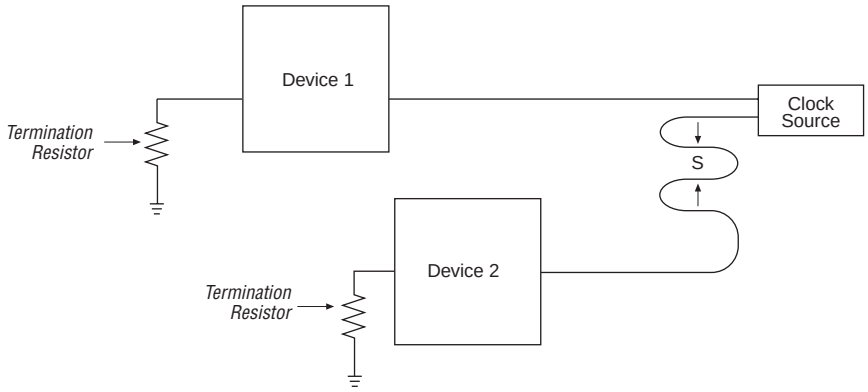

Altera recommends avoiding serpentine routing if possible. Instead, use arcs to create equal-length traces.

Differential Trace Routing

To maximize signal integrity, proper routing techniques for differential signals are important for high-speed designs. Figure 21 shows a differential pair using the microstrip layout.

Figure 21. Differential Trace Routing *Note (1)*

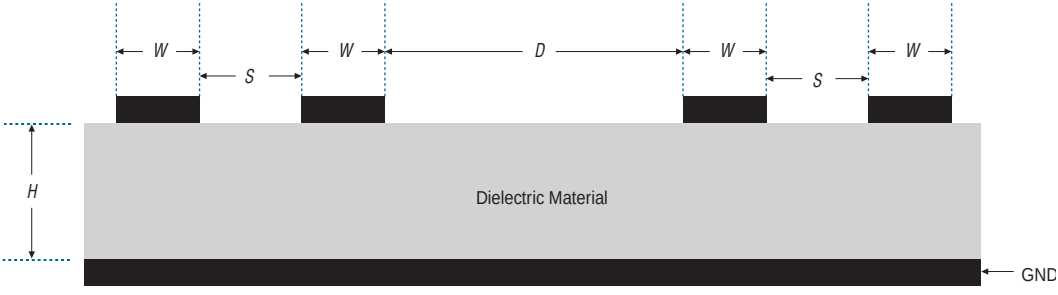

Note to Figure 21:

(1) *D* = Distance between two differential pair signals; *W* = Width of a trace in a differential pair; *S* = Distance between the trace in a differential pair; and *H* = Dielectric height above the group plane.

Use the following guidelines when using two differential pairs:

- Make sure $D > 2S$ to minimize the crosstalk between the two differential pairs.
- To minimize reflection noise, place the differential traces $S = 3H$ as they leave the device.
- Keep the distance between the differential traces (S) constant over the entire trace length.
- Keep the length of the two differential traces the same to minimize the skew and phase difference.
- Avoid using multiple vias, because they can cause impedance mismatch and inductance.

Termination Schemes

Mismatched impedance causes signals to reflect back and forth along the lines, which causes ringing at the load receiver. The ringing reduces the dynamic range of the receiver and can cause false triggering. To eliminate reflections, the impedance of the source (Z_S) must equal the impedance of the trace (Z_o), as well as the impedance of the load (Z_L). Stratix devices feature support for on-chip implementation of the resistor. This section discusses the following signal termination schemes:

- Simple parallel termination
- Thevenin parallel termination
- Active parallel termination
- Series-RC parallel termination
- Series termination
- Differential pair termination
- On-chip termination

Simple Parallel Termination

In a simple parallel termination scheme, the terminating resistor (R_T) is equal to the line impedance. Place the termination resistor as close to the load as possible to be efficient. See Figure 22.

Figure 22. Simple Parallel Termination

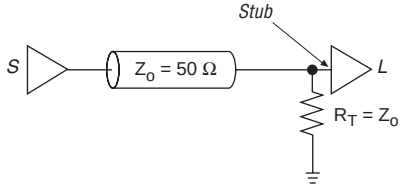

The stub length from R_T to the receiver pin and pads should be as small as possible. A long stub length causes reflections from the receiver pads, resulting in signal degradation. If your design requires a long termination line between the terminator and receiver, the placement of the resistor becomes important. For long termination line lengths, use fly-by termination (see Figure 23).

Figure 23. Simple Parallel Fly-By Termination

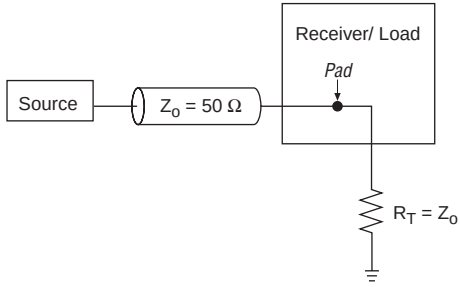

Thevenin Parallel Termination

An alternative parallel termination scheme uses a Thevenin voltage divider (see Figure 24). The terminating resistor is split between R_1 and R_2, which equals the line impedance when combined. Although this scheme reduces the current drawn from the source device, it adds current drawn from the power supply because the resistors are tied between V_{CC} and GND.

Figure 24. Thevenin Termination

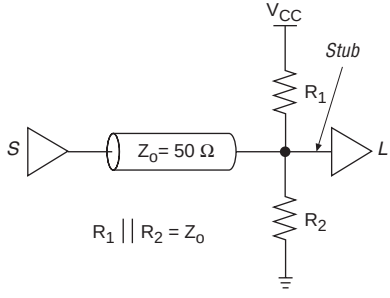

As noted in the previous section, stub length is dependent on signal rise and fall time and should be kept to a minimum. If your design requires a long termination line between the terminator and receiver, use fly-by termination or Thevenin fly-by termination. See Figures 23 and 25.

Figure 25. Thevenin Fly-By Termination

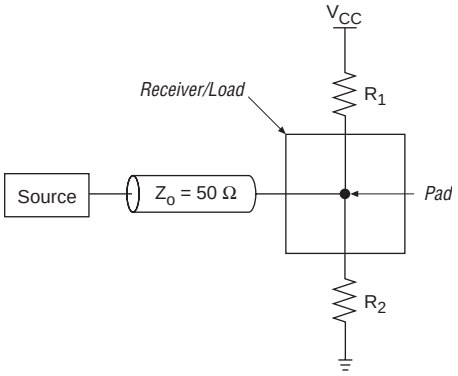

Active Parallel Termination

Figure 26 shows an active parallel termination scheme, where the terminating resistor ($R_T = Z_o$) is tied to a bias voltage (V_{BIAS}). In this scheme, the voltage is selected so that the output drivers can draw current from the high- and low-level signals. However, this scheme requires a separate voltage source that can sink and source currents to match the output transfer rates.

Figure 26. Active Parallel Termination

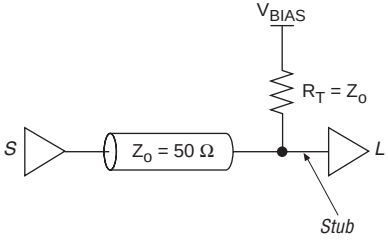

Figure 27 shows the active parallel fly-by termination scheme.

Figure 27. Active Parallel Fly-By Termination

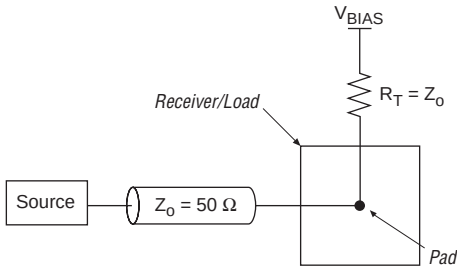

Series-RC Parallel Termination

A series-RC parallel termination scheme uses a resistor and capacitor (i.e., series-RC) network as the terminating impedance. The terminating resistor (R_T) is equal to Z_0. The capacitor must be large enough to filter the constant flow of DC current. However, if the capacitor is too large, it will delay the signal beyond the design threshold.

Capacitors smaller than 100 pF diminish the effectiveness of termination. The capacitor blocks low-frequency signals while passing high-frequency signals. Therefore, the DC loading effect of R_T does not have an impact on the driver, as there is no DC path to ground. The series-RC termination scheme requires balanced DC signaling (i.e., the signals spend half the time on and half the time off). AC termination is typically used if there is more than one load. See Figure 28.

Figure 28. Series-RC Parallel Termination

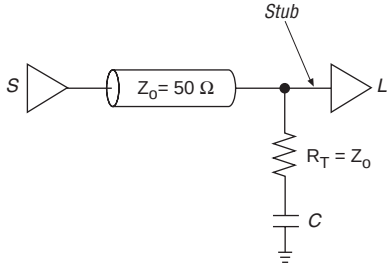

Figure 29 shows series-RC parallel fly-by termination.

Figure 29. Series-RC Parallel Fly-By Termination

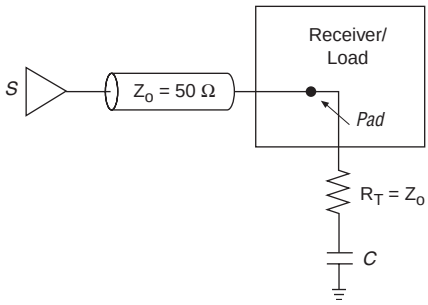

Series Termination

In a series termination scheme, the resistor matches the impedance at the signal source instead of matching the impedance at each load (see Figure 30). The sum of R_T and the impedance of the output driver should be equal to the Z_0. Because Altera device output impedance is low, you should add a series resistor to match the signal source to the line impedance. The advantage of series termination is that it consumes little power. However, the disadvantage is that the rise time degrades due to the increased RC time constant. Therefore, for high-speed designs, you should perform the pre-layout signal integrity simulation with Altera input/output buffer information specification (IBIS) models before using the series termination scheme.

Figure 30. Series Termination

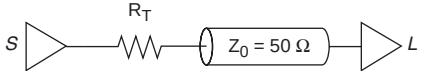

Differential Pair Termination

Differential signal I/O standards require a termination resistor between the signals at the receiving device (see Figure 31). For the LVDS and LVPECL standard, the termination resistor should match the differential load impedance of the bus (i.e., typically 100 Ω). Altera Stratix, Stratix GX, and Mercury™ devices have an on-chip termination option. Using on-chip termination decreases required board space. For more information, see "On-Chip Termination for Stratix GX Transceivers" on page 25.

Figure 31. Differential Pair (LVDS & LVPECL) Termination

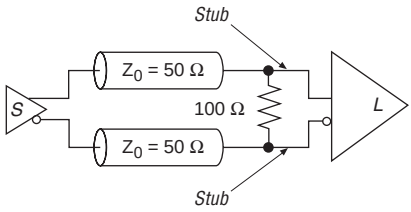

Figure 32 shows the differential pair fly-by termination scheme for the LVDS and LVPECL standard.

Figure 32. Differential Pair (LVDS & LVPECL) Fly-By Termination

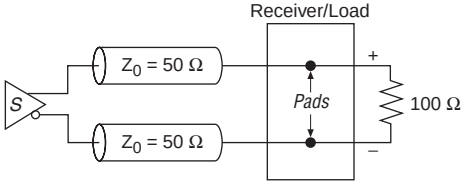

3.3-V PCML uses two parallel 100-Ω termination resistors at the transmitter and two parallel 50-Ω termination resistors at the receiver (see Figure 33). The termination voltage (V_T) is the same as the V_{CCIO} voltage (3.3 V).

Figure 33. Differential Pair (3.3-V PCML) Termination

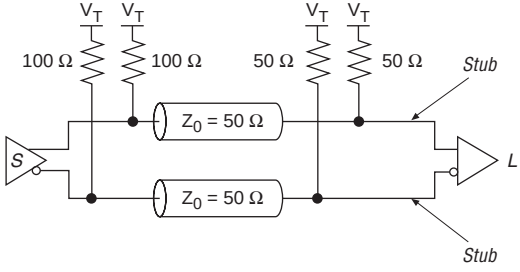

Figure 34 shows the differential pair, fly-by termination scheme for 3.3-V PCML.

Figure 34. Differential Pair (3.3-V PCML) Fly-By Termination

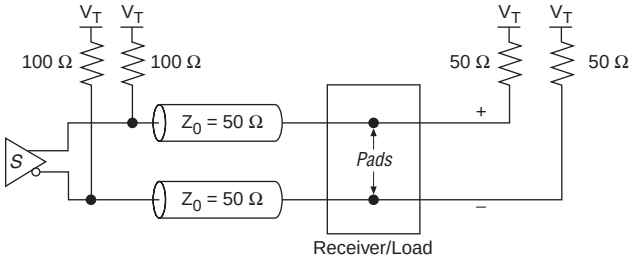

Refer to the *Board Design Guidelines for LVDS Systems White Paper* for more information on terminating differential signals. Refer to *AN 209: Using Terminator Technology in Stratix & Stratix GX Devices* for more information on Stratix and Stratix GX on-chip termination.

On-Chip Termination for Stratix GX Transceivers

Stratix GX devices also have on-chip resistors designed to support termination for several I/O standards. On-chip resistors simplify the task of board design by freeing board space and offering more freedom in signal routing. Also, because the distance from the resistor to the signal pin is smaller, the on-chip resistors reduce stub reflection. Consequently, Stratix GX devices provide better load and/or source termination, which leads to better signal integrity.

All transceivers in Stratix GX devices have programmable, per channel internal termination resistors, which can be programmed to be either 50-, 60-, or 75-Ω single-ended resistors. In differential mode, the resistors generate 100-, 120-, or 150-Ω termination. Figures 35 and 36 show receiver and transmitter on-chip termination schemes for Stratix GX devices.

Because different I/O standards require different termination resistors, programmable internal termination resistors can be helpful. For example, XAUI and Infiniband applications require 100-Ω differential termination, while Gigabit Ethernet and Fibre Channel require 150-Ω differential termination.

Figure 35. Receiver On-Chip Termination Scheme for Stratix GX Devices

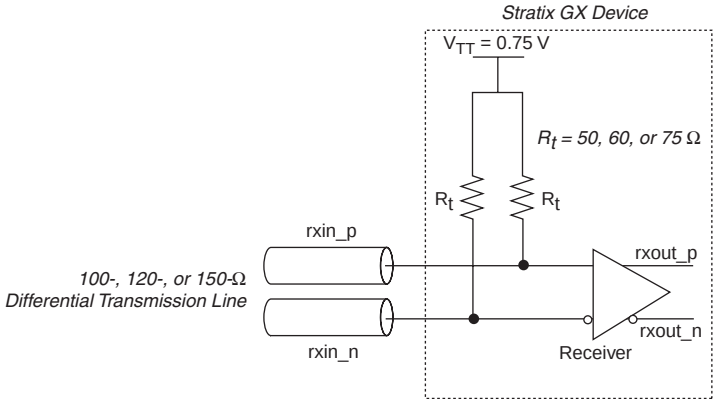

You can bypass the on-chip resistors used in the receiver and use external resistors.

Figure 36. Transmit On-Chip Termination Scheme for Stratix GX Transceivers

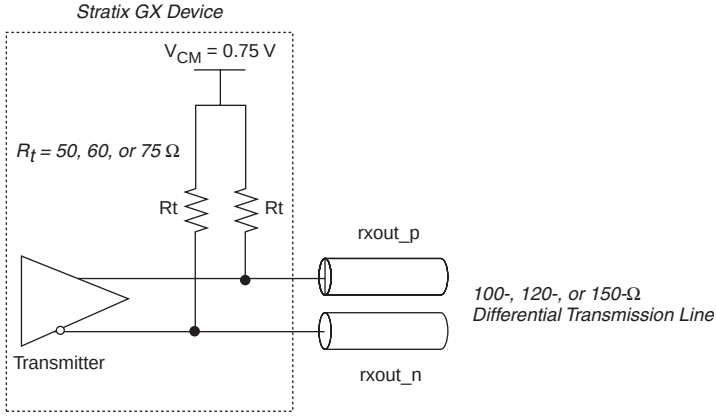

For non-transceiver I/O pins in Stratix GX devices, the termination scheme is the same as with Stratix devices.

For more information on designing termination schemes with Stratix and Stratix GX devices, refer to *AN 209: Using Terminator Technology in Stratix & Stratix GX Devices*.

Simultaneous Switching Noise (SSN)

As digital devices become faster, their output switching times decrease, causing higher transient currents in outputs as the devices discharge load capacitances. These higher transient currents result in a board-level phenomenon known as ground bounce.

Because many factors contribute to ground bounce, you cannot use a standard test method to predict its magnitude for all possible PCB environments. You can only test the device under a given set of conditions to determine the relative contributions of each condition and of the device itself. Load capacitance, socket inductance, and the number of switching outputs are the predominant factors that influence the magnitude of ground bounce in FPGAs.

Altera requires 0.01- to 0.1-µF surface-mount capacitors in parallel to reduce ground bounce. Add an additional 0.001-µF capacitor in parallel to these capacitors to filter high-frequency noise (>100 MHz).

Altera recommends that you take the following steps to reduce ground bounce and V_{CC} sag:

- Configure unused I/O pins as output pins, and drive the output low to reduce ground bounce. This configuration will act as a virtual ground.
- Configure the unused I/O pins as output, and drive high to prevent V_{CC} sag.
- Create a programmable ground or V_{CC} next to switching pins.
- Reduce the number of outputs that can switch simultaneously and distribute them evenly throughout the device.
- Manually assign ground pins in between I/O pins. (Separating I/O pins with ground pins prevents ground bounce.)
- Turn on the slow slew rate logic option when speed is not critical.
- Eliminate sockets whenever possible.
- Depending on the problem, move switching outputs close to either a package ground or `VCC` pin. Eliminate pull-up resistors, or use pull-down resistors.
- Use multi-layer PCBs that provide separate V_{CC} and ground planes to utilize the intrinsic capacitance of GND-V_{CC} plane.
- Create synchronous designs that are not affected by momentarily switching pins.
- Add the recommended decoupling capacitors to V_{CC}/GND pairs.
- Place the decoupling capacitors as close as possible to the power and ground pins of the device.
- Connect the capacitor pad to the power and ground plane with larger vias to minimize the inductance in decoupling capacitors and allow for maximum current flow.
- Use wide, short traces between the vias and capacitor pads, or place the via adjacent to the capacitor pad (see Figure 37).

Figure 37. Suggested Via Location that Connects to Capacitor Pad

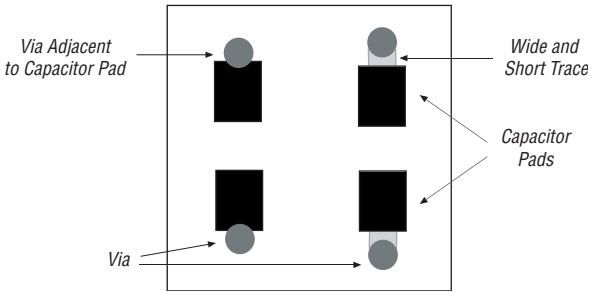

- Traces stretching from power pins to a power plane (or island, or a decoupling capacitor) should be as wide and as short as possible. This reduces series inductance, and therefore, reduces transient voltage drops from the power plane to the power pin. Thus, reducing the possibility of ground bounce.
- Use surface-mount low effective series resistance (ESR) capacitors to minimize the lead inductance. The capacitors should have an ESR value as small as possible.
- Connect each ground pin or via to the ground plane individually. A daisy chain connection to the ground pins shares the ground path, which increases the return current loop and thus inductance.

Refer to the *Minimizing Ground Bounce & V_{CC} Sag White Paper* for more information about ground bounce and V_{CC} sag.

Power Filtering & Distribution

You can reduce system noise by providing clean, evenly distributed power to V_{CC} on all boards and devices. This section describes techniques for distributing and filtering power.

Filtering Noise

To decrease the low-frequency (< 1 kHz) noise caused by the power supply, filter the noise on power lines at the point where the power connects to the PCB and to each device. Place a 100-µF electrolytic capacitor where the power supply lines enter the PCB. If you use a voltage regulator, place the capacitor immediately after the pin that provides the V_{CC} signal to the device(s). (Capacitors not only filter low-frequency noise from the power supply, but also supply extra current when many outputs switch simultaneously in a circuit.)

To filter power supply noise, use a non-resonant, surface-mount ferrite bead large enough to handle the current in series with the power supply. Place a 10- to 100-µF bypass capacitor next to the ferrite bead (see Figure 38). (If proper termination, layout, and filtering eliminate enough noise, you do not need to use a ferrite bead.) The ferrite bead acts as a short for high frequency noise coming from the V_{CC} source. Any low frequency noise is filtered by a large 10-µF capacitor after the ferrite bead.

Figure 38. Filtering Noise with a Ferrite Bead

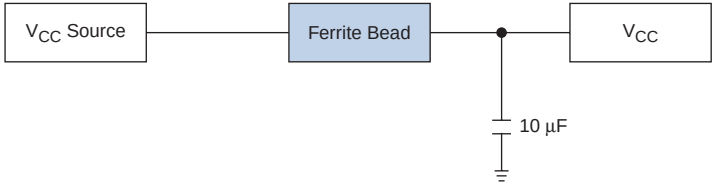

Usually, elements on the PCB add high-frequency noise to the power plane. To filter the high-frequency noise at the device, place decoupling capacitors as close as possible to each V_{CC} and GND pair.

Refer to the *Operating Requirements for Altera Devices Data Sheet* for more information on bypass capacitors.

Power Distribution

A system can distribute power throughout the PCB with either power planes or a power bus network.

You can use power planes on multi-layer PCBs that consist of two or more metal layers that carry V_{CC} and GND to the devices. Because the power plane covers the full area of the PCB, its DC resistance is very low. The power plane maintains V_{CC} and distributes it equally to all devices while providing very high current-sink capability, noise protection, and shielding for the logic signals on the PCB. Altera recommends using power planes to distribute power.

The power bus network—which consists of two or more wide metal traces that carry V_{CC} and GND to devices—is often used on two-layer PCBs and is less expensive than power planes. When designing with power bus networks, be sure to keep the trace widths as wide as possible. The main drawback to using power bus networks is significant DC resistance.

Altera recommends using separate analog and digital power planes. For fully digital systems that do not already have a separate analog power plane, it can be expensive to add new power planes. However, you can create partitioned islands (split planes). Figure 39 shows an example board layout with phase-locked loop (PLL) ground islands.

Figure 39. Board Layout for General-Purpose PLL Ground Islands

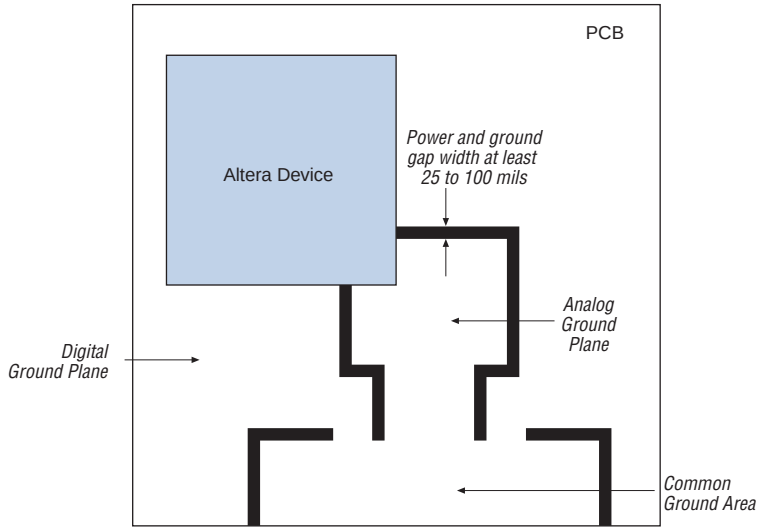

If your system shares the same plane between analog and digital power supplies, there may be unwanted interaction between the two circuit types. The following suggestions will help to reduce noise:

- For equal power distribution, use separate power planes for the analog (PLL) power supply. Avoid using trace or multiple signal layers to route the PLL power supply.
- Use a ground plane next to the PLL power supply plane to reduce power-generated noise.
- Place analog and digital components only over their respective ground planes.
- Use ferrite beads to isolate the PLL power supply from digital power supply.

EMI

Electromagnetic interference (EMI) is directly proportional to the change in current or voltage with respect to time. EMI is also directly proportional to the series inductance of the circuit. Every PCB generates EMI. Precautions such as minimizing crosstalk, proper grounding, and proper layer stack-up can significantly reduce EMI problems.

Place each signal layer in between the ground plane and power plane. Inductance is directly proportional to the distance an electric charge has to cover from the source of an electric charge to ground. As the distance gets shorter, the inductance becomes smaller as well. Therefore, placing ground planes close to a signal source reduces inductance and helps contain EMI. Figure 40 shows an example of an eight-layer stack-up. In the stack-up, the stripline signal layers are the quietest because they are centered by power and GND planes. A solid ground plane next to the power plane creates a set of low ESR capacitors. With IC edge rates becoming faster and faster, these techniques help to contain EMI.

Figure 40. Example Eight-Layer Stack-Up

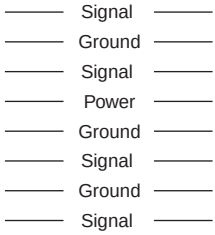

Component selection and proper placement on the board is very important to controlling EMI.

The following guidelines can help reduce EMI:

- Select low-inductance components, such as surface mount capacitors with low ESR, and effective series inductance (ESL).
- Use proper grounding for the shortest current return path.
- Use solid ground planes next to power planes.
- In unavoidable circumstances, use respective ground planes next to each segmented power plane for analog and digital circuits.

Additional FPGA-Specific Information

This section provides additional FPGA-specific configuration, Joint Test Action Group (JTAG), and test point information recommended by Altera for board design and signal integrity.

Configuration

The DCLK signal is used in configuration devices and passive serial (PS) and passive parallel synchronous (PPS) configuration schemes. This signal drives edge-triggered pins in Altera devices. Therefore, any overshoot, undershoot, ringing, crosstalk or other noise can affect configuration. Use the same guidelines for designing clock signals to route the DCLK trace (see "Signal Trace Routing" on page 13). If your design uses more than five configuration devices, Altera recommends using buffers to split the fan-out on the DCLK signal.

JTAG

As PCBs become more complex, testing becomes increasingly important. Advances in surface mount packaging and PCB manufacturing have resulted in smaller boards, making traditional test methods such as external test probes and "bed-of-nails" test fixtures harder to implement. As a result, cost savings from PCB space reductions can be offset by cost increases in traditional testing methods.

In addition to boundary scan testing (BST), you can use the IEEE Std. 1149.1 controller for in-system programming. JTAG consists of four required pins, test data input (TDI), test data output (TDO), test mode select (TMS), and test clock input (TCK) as well as an optional test reset input (TRST) pin.

Use the same guidelines for laying out clock signals to route TCK traces. Use multiple devices for long JTAG scan chains. Minimize the JTAG scan chain trace length that connects one device's TDO pins to another device's TDI pins to reduce delay.

Refer to *AN 39: IEEE 1149.1 (JTAG) Boundary-Scan Testing in Altera Devices* for additional details on BST.

Test Point

As device package pin density increases, it becomes more difficult to attach an oscilloscope or a logic analyzer probe on the device pin. Using a physical probe directly on to the device pin can damage the device. If the ball grid array (BGA) or FineLine BGA® package is mounted on top of the board, it is difficult to probe the other side of the board. Therefore, the PCB must have a permanent test point to probe. The test point can be a via that connects to the signal under test with a very short stub. However, placing a via on a trace for a signal under test can cause reflection and poor signal integrity.

Refer to *AN 175: SignalTap Analysis in the Quartus II Software* for more information on the SignalTap® embedded logic analyzer.

Summary

You must carefully plan out a successful high-speed PCB. Factors such as noise generation, signal reflection, crosstalk, and ground bounce can interfere with a signal, especially with the high speeds that Altera devices transmit and receive. The signal routing, termination schemes, and power distribution techniques discussed in this application note can help you design a more effective PCB using high-speed Altera devices.

References

Johnson, H. W., and & Graham, M., "*High-Speed Digital Design.*" Prentice Hall, 1993.

Hall, S. H., Hall, G. W., and McCall J. A., "High-Speed Digital System Design." John Wiley & Sons, Inc. 2000.

Revision History

The information contained in *AN 224: High-Speed Board Layout Guidelines* version 1.1 supersedes information published in previous versions.

AN 224: High-Speed Board Layout Guidelines version 1.1 contains the following change: updated the legend in Figure 4.

101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
http://www.altera.com
Applications Hotline:
(800) 800-EPLD
Literature Services:
lit_req@altera.com

Copyright © 2002 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, mask work rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

I.S. EN ISO 9001