High-Speed PCB Design Considerations

Technical Note

Disclaimers

Lattice makes no warranty, representation, or guarantee regarding the accuracy of information contained in this document or the suitability of its products for any particular purpose. All information herein is provided AS IS, with all faults and associated risk the responsibility entirely of the Buyer. Buyer shall not rely on any data and performance specifications or parameters provided herein. Products sold by Lattice have been subject to limited testing and it is the Buyer's responsibility to independently determine the suitability of any products and to test and verify the same. No Lattice products should be used in conjunction with mission- or safety-critical or any other application in which the failure of Lattice's product could create a situation where personal injury, death, severe property or environmental damage may occur. The information provided in this document is proprietary to Lattice Semiconductor, and Lattice reserves the right to make any changes to the information in this document or to any products at any time without notice.

Contents

Figures

Acronyms in This Document

A list of acronyms used in this document.

Acronym	Definition
AC	Alternating Current
AMP	AMP Inc.
BGA	Ball Grid Array
CML	Current Mode Logic
DC Drop	Voltage Drop
EMC	Electro Magnetic Compatibility
EMI	Electro Magnetic Interference
ESL	Equivalent Series Inductance
ESR	Effective Series Resistance
GND	Ground
HVLP	High Volume Low Pressure
IC	Integrated Circuit
LVDS	Low-Voltage Differential Signaling
PCB	Printed Circuit Board
SerDes	Serializer/Deserializer
SOSA	Sensor Open Systems Architecture
SPICE	Simulation Program with Integrated Circuit Emphasis
TDR	Time Domain Reflectometer
Tx	Transmitter
Rx	Receiver

1. Introduction

The backplane is the physical interconnection where typically all electrical modules of a system converge. Complex systems rely on the wires, traces, and connectors of the backplane to handle large amounts of data at high speed. The communication between the various backplane modules depends on the inherent electrical characteristics such as impedance, capacitance, and inductance derived from connectors, trace lengths, vias, and termination, to name a few. An extremely important factor in designing a distributed-load high-performance backplane, is a basic understanding of the design practices used to ensure good signal integrity.

This technical note examines some basic differences in interconnection topologies. It describes the various issues that should be considered while designing a backplane and focuses on the critical aspects of point-to-point transmission lines that are run through a backplane. These aspects include PCB line structure, vias, device packaging and backplane connectors. A PCB design checklist is provided to aid the designer. Some frequency specific discussion and guidelines are given. This document also discusses Lattice Semiconductor's FPGA product line and its SerDes high-speed backplane interfaces. These provide high speed serial streams through CML differential buffers.

2. Backplane Topology and Overview

Three different system interconnection topologies are normally used in backplanes today. These are multi-point, multi-drop and point-to-point. Traditionally, systems have used multi-point/ multi-drop connection topologies, which provided efficient interconnection and communication between multiple devices, with a single net (node), as shown in Figure 2.1.
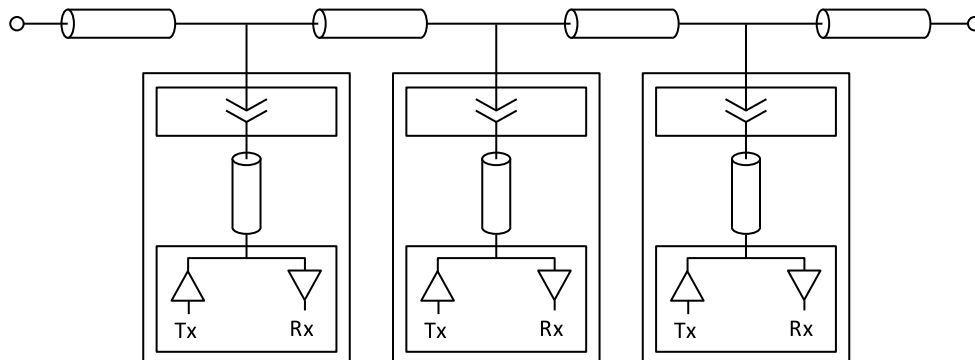

Figure 2.1. Multi-Point Backplane Illustration

Unfortunately, this net structure provides serious data-rate limitations. Each net contains tees or branches at each point where the card connects to the backplane. These tees provide transmission line discontinuities and mismatches along the backplane signal path. As a result, large signal reflections occur at every card to backplane interface. These reflections can propagate back and forth for substantial time periods and severely degrade signal integrity at higher speeds. Acceptable signal communication is normally achieved by waiting for the reflected signals to settle out, for each bit of transmitted data. This imposes significant speed limitations. For this reason, multi-point and multi-drop backplane topologies generally have speed limits below 100 Mbps. This limit can easily drop below 10 Mbps as physical line lengths and the number of card slots increase.

The point-to-point interconnection topology eliminates the signal path branches described above. The resulting signal reflections are eliminated and maximum data rates are increased dramatically. This type of backplane interconnection can be used with data rates to 25 Gbps and above, with careful design methods.

The disadvantage of this approach is an increase in the number of backplane nets and card port interfaces that may be needed. The single net connection between n cards, in a multi-point backplane, must be replaced with n(n-1) unidirectional point-to-point links. Each card must provide n-1 transmit and n-1 receive ports for full system interconnectivity. As an example, a four PCB module system with full interconnectivity is shown in Figure 2.2.
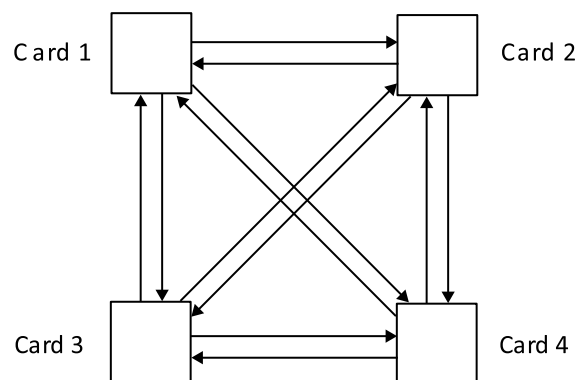
Figure 2.2. Four Card Point-to-Point Interconnection System

Each card must provide three transmitters and three receiver ports. Each arrowed line represents a point-to-point backplane net.

Recent communication equipment designs have shown a rapidly growing need for higher bandwidth interconnection between PCB modules. The fast evolving IC technology, with its multi-gigabit processing and driving capabilities, has made point-to-point backplane the topology of choice for many of today's new hardware systems. Both serial and parallel data structures can be supported with this topology. Lattice Semiconductor has introduced several IC products with multiple ports, each designed with Gbps backplane drive capability. These devices are described later in this document. The remainder of this document focuses on PCB design aspects of the point-to-point backplane interconnection links.

3. Point-to-Point Backplane Signal Path Structure

The typical point-to-point topology utilizes a simple, single-path interconnection structure that extends from a transmitting device on one card, across a backplane to a receiving device on another card. The physical path for such an interconnection is illustrated in Figure 3.1.
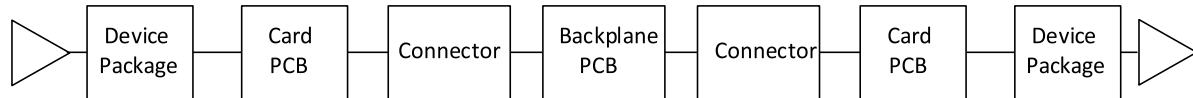
Figure 3.1. Interconnection Link Physical Structure

The point-to-point interconnection elements are all serially connected and provide a single signal path. Each element may be thought of as a transmission line segment. Ideally, by controlling and matching the characteristic impedance of each line segment, a uniform electrical signal path is created. Signals can then propagate the entire path length with no reflections occurring. Adding a resistive termination at the receive device input, with value equal to the characteristic impedance of the line elements, would provide a distortionless data link with maximum bandwidth, between the transmit and receive devices.

Each of the elements in Figure 3.1 can be broken down into of a number of sub-elements. The PCB elements for example, which provide the most significant transmission line segments in the data path, consists of metal traces, dielectric layers, ground plane layers, and (inter-layer) vias. Each of these sub-elements is a critical part of the signal path and can cause electrical line discontinuities and signal reflections, if not properly designed. The design aspects of the elements and sub-elements contained in Figure 3.1 are discussed in the following sections.

4. Advantages of Differential Signaling

The inherent system advantages of differential signal interconnection schemes are well known in all fields of electronic design. These advantages are especially important in high bandwidth, high-density hardware systems where very low error-rate data links are required. Differential signaling provides critically needed immunity to common-mode electrical noise that is present at significant levels in most application systems. For example, using differential signaling avoids the classic "ground bounce" noise problem that is experienced with many high density ICs that use single-ended interfaces. It also provides higher noise margins, which lead to lower bit-error rates in digital data links. As the edge rates of signals have increased due to the need to support higher bandwidths, another need for PCB design is to provide return paths for the inductive coupled current of the PCB. Differential signaling helps reduce the "bounce" seen from this inductive current because the current remains localized. This is due to the fact that as one of the leads of the differential pair is sinking, the other is sourcing current, essentially canceling the inductive influences. Differential signal interconnection methods should be used for all critical, high-speed interconnections.

5. Board Design Practices

5.1. Differential Trace Design

Differential signal trace-pairs with controlled impedance can be arranged in a number of different configurations. Most common are the following four figures.
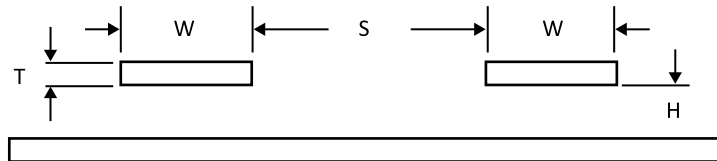
Figure 5.1. Edge Coupled Microstrip (Surface Routing)
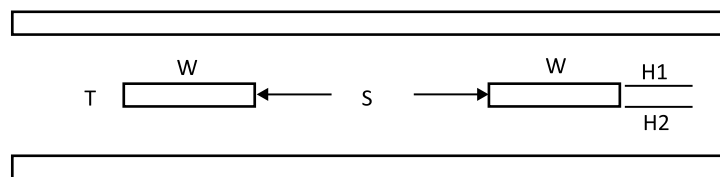
Figure 5.2. Edge Coupled Stripline (Sandwiched Between Two Reference Planes)
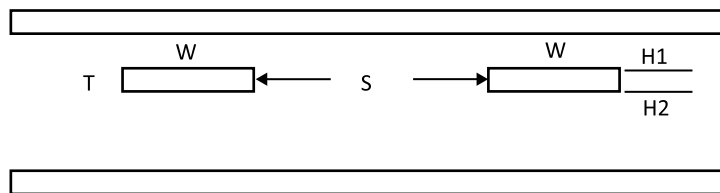
Figure 5.3. Offset Edge Coupled Stripline (Same as Figure 5.2., But Not Centered Between Reference Planes)
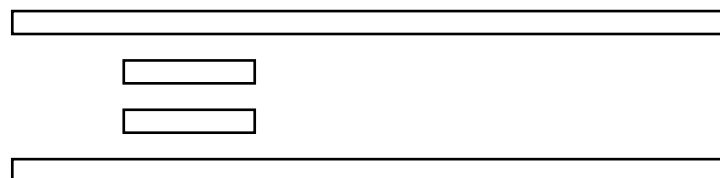
Figure 5.4. Broadside Coupled Stripline (Also Referred to as Dual Stripline)

100 Ω characteristic impedance has become an industry standard value for differential lines used for interconnection. This impedance level lends itself well to PCB structures and other components designs where controlled transmission line impedance must be provided.

A 100 Ω differential line can be constructed with two 50 Ω single-ended lines of equal length.

As the two line traces are brought near each other (as shown in Figure 5.1 through Figure 5.4), the field coupling between the traces reduces the differential mode impedance of the line. To maintain 100 Ω differential impedance, the trace width must be reduced slightly. As a result, the common mode impedance of each trace in a 100 Ω couple-trace differential pair is slightly more than 50 Ω.

Achieving a 100 Ω differential impedance with a coupled pair of traces implies that the single ended impedance of Z0 ranges from 53 to 60 Ω, with a coupling coefficient typically ranging from 1-15%. The relationship between the common mode impedance Z0 and the differential impedance Zdiff is given by the expression Zdiff = 2 Z0 (1 - k)/(1 + k), where k is the trace coupling coefficient.

50 Ω resistors, tied to ground, are normally used to terminate the 100 Ω differential lines. This provides the ideal differential line termination, which is most important since the data links use differential signals. The slight impedance mismatch that occurs in the common mode is usually of little consequence. Normally only noise and crosstalk signals occur in common mode.

5.2. Common Mode Noise Tolerance

In order to prevent common mode noise from converting to differential mode noise, it is important to maintain the symmetry of the differential pair. Reflections and impedance mismatch in the common mode does not affect the differential mode performance as long as the two modes can be kept relatively orthogonal.

The loop area is defined as the area between the signal path and its return path. On differential traces, the signal is on one trace and the return is on another trace. So the loop area is a function of how close the traces are routed together. If we are concerned about EMI emission and susceptibility, which is generally understood as a loop area concern, we must route the traces close together. The more closely we route them to each other, the smaller the loop area is and the less EMI is generated.

One of the primary advantages of differential signals is the signal-to-noise ratio improvement that is obtained. Since the signal is one polarity on one trace and the other polarity on the other trace, the resulting signal at the receiving device is twice what the single-ended signal would be. Ideally common mode rejection at the receiving device is such that the receiving device only responds to the difference in signal level between the two traces. Since noise is typically in the common mode, it is rejected at the receiver and maintains a high differential signal-to-noise ratio.

In order to have good common mode noise rejection, it is important that any noise that is present affects the signals on both traces equally. That is, if noise is coupled into one trace, an equal amount of noise must be coupled into the other trace. Then the common mode rejection capability of the receiving circuit rejects the noise. But if noise is coupled into one trace more strongly than into the other trace, the noise appears as a differential mode signal to the receiver and it can be amplified. The way to ensure that any noise is coupled equally into both traces is to route the two traces very close together. Then, they are both in the same noise environment.

5.3. PCB Trace Impedance Calculation

In the past, calculation of the characteristic impedance for a printed circuit board trace was a complex, error prone process involving complicated calculations and approximations. Nomographs and simplified formulas have been generated to simplify the design process, but are often inaccurate. The most accurate method available is a field solver program (usually 2D, sometimes 3D), which solves Maxwell's equations directly over the volume of PCB under consideration using finite elements. These simulations can be verified in hardware with a Time Domain Reflectometer (TDR) measurement device.

One example of a 2D field solver program is the Si9000e program from Polar Instruments. This is available as a licensed version, and can be downloaded for testing on a laptop. There are also many free trace impedance calculators on the web for many types of micro-strip, stripline, and coplanar waveguides. There is a good PCB Toolkit of Saturn PCB Design INC., which you can download for free.

Even when using a field solver, there are still uncertainties in the impedance calculation arising from variations in the effective dielectric constant in the glass fiber, prepreg, and epoxy used in typical FR4 manufacturing. The average dielectric constant of FR4 varies from 4.2 to 4.5, depending on the material, exact location, and construction method used.

Verification of the impedance on a PCB depends on actual measurements of typical copper traces. Some manufacturers use an auxiliary test section of a PCB called a coupon, which is a long rectangular test section of PCB with pads designed to accept probes from a TDR measuring instrument. It is possible for manufacturing errors to result in over-etching of the copper traces, with a resultant impedance error. Monitoring production quality with coupons can prevent this type of problem.

5.4. Examples of PCB Trace Impedance Calculation

As a calculation example, the differential impedance for a pair of 10 mil traces in 1/2 ounce copper with 10 mil spacing that is on an FR4 substrate with a spacing of 15 mils above the ground plane (microstrip). The copper thickness (T) is 1.7 mils. Figure 5.5 shows the parameters. Note that this example uses the Saturn PCB Toolkit transmission line calculator product mentioned previously.
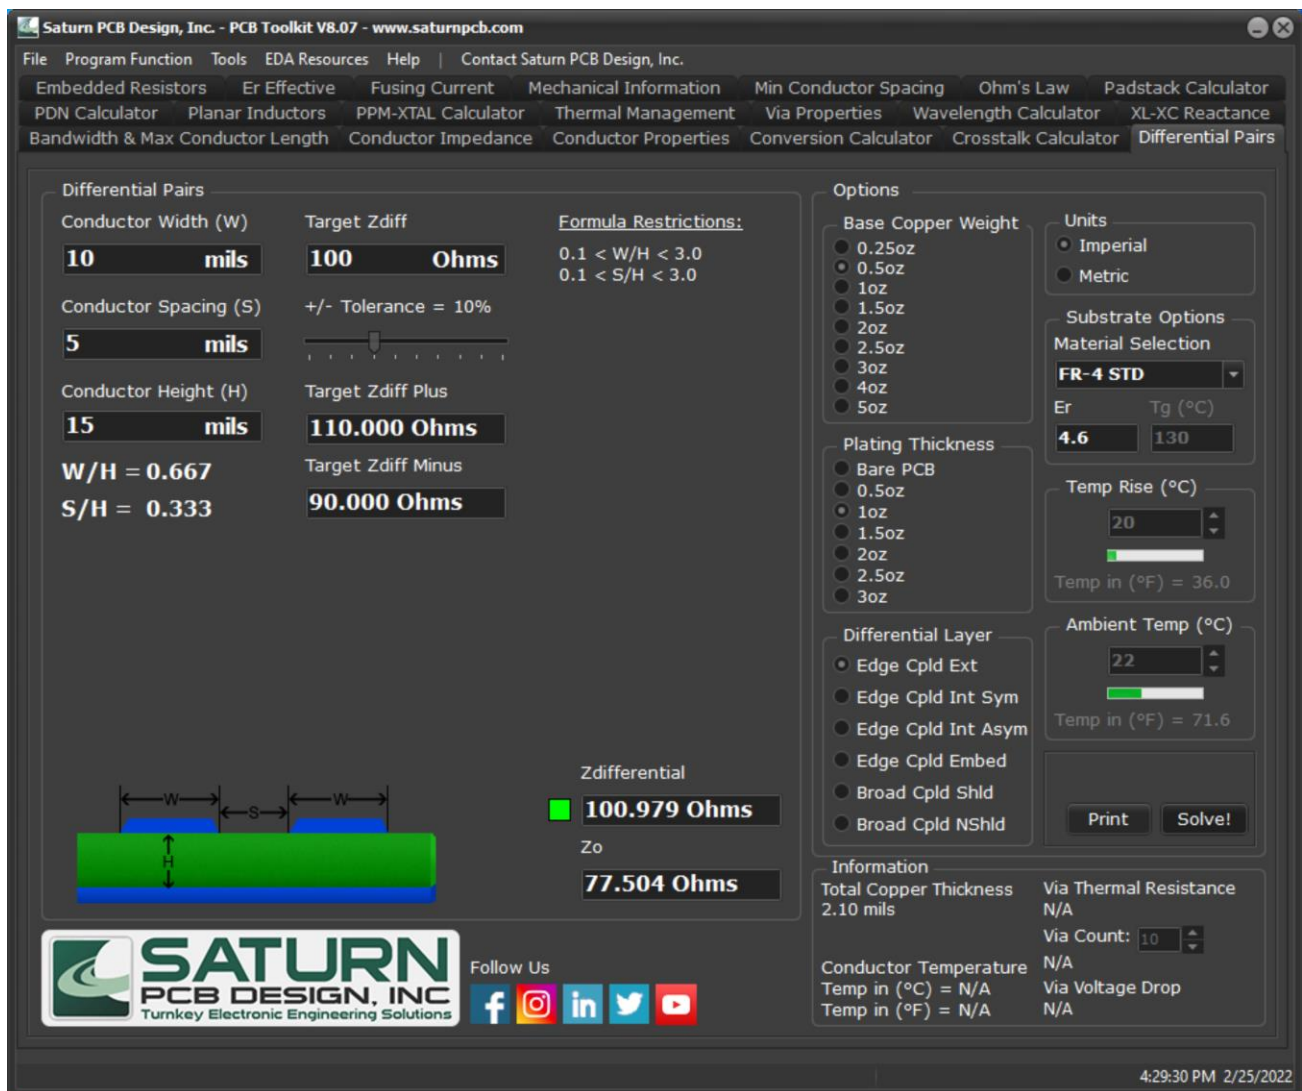

Figure 5.5. Differential Microstrip Example with Saturn PCB Toolkit Impedance Calculator Tool

In a typical FR4 PCB, there are three types of differential pair routing encountered. For connections to surface mounted components, edge coupled microstrips may be needed, while connections between through-hole components or via pairs can use stripline and offset stripline. Dual stripline with broadside coupling should be avoided, since this configuration is subject to differential noise coupling from the reference planes. Another problem with broadside coupling is that any asymmetry in the PCB manufacturing can result in an asymmetric trace impedance, causing mismatch in the effective electrical length, even if the physical lengths match exactly. Using edge coupled differential pairs makes it is easier to maintain symmetry.

Vias, connectors, and component pads all introduce impedance discontinuities into the signal path. This can be measured with TDR techniques.

In order to avoid crosstalk, when laying out differential pair trace with spacing of value S, it is recommended to place other pairs no closer than a distance of 3 S, and preferably a distance of 4 S if possible. This rule can be relaxed if a differential pair is only briefly in proximity with another pair, such as at a connector or via layer switch.

5.5. PCB Design Checklist

1. Use 100 Ω differential impedance pairs on PCB. Controlled impedance lines should be specified in the PCB layout mechanical drawing.

2. Match trace lengths in a pair with tolerance of 20% of the signal rise/fall time.

3. Use connectors that are designed and characterized at the highest data frequency. (Vendors should provide characterization and model data.)

4. Use stripline construction with ground/power planes above and below the differential pairs. The ground and power planes also provide return paths for signal currents.

5. Use edge-coupled pairs in PCBs; try to avoid broadside coupled pairs.

6. Use 3 S or 4 S separation rules between pairs to avoid crosstalk and excess coupling. Use offset stripline routing to get higher density of differential pairs with each routing layer running orthogonal to each other.

6. PCB Layer Design (Board Stack-up)

Multi-layer boards are a must in both daughter board and backplane design. The multiple metal layers facilitate high connection density, minimum crosstalk, and good ElectroMagnetic Compatibility (EMC). These factors are key to achieving good signal integrity for all the signal interconnections. Ideally, all signal layers should be separated from each other by ground or power planes (metal layers). This minimizes crosstalk and provides homogeneous transmission lines, with properly controlled characteristic impedance, between devices and other board components. Best performance is obtained when using dedicated ground and power plane layers that are continuous across the entire board area. When it is not feasible to provide ground or power planes between signal layers, great care must be taken to ensure signal line coupling is minimized. Orthogonal routing on adjacent signal layers minimizes coupling and should be used. CAD tools, which predict line coupling and signal crosstalk, can be very helpful in this type of design.

6.1. Vias

Vias generally provide two purposes. One is used for mounting a through-hole component to a board. The second is to interconnect traces on different metal layers. Electrically, vias are often modeled as having an inductive and capacitive parasitic value. Smaller vias have lower capacitance. Short length, larger diameter vias have lower inductance. Both parasitic elements can have detrimental affects, but it is often the inductance parasitic element that provides an unexpected series impedance that creates problems.

Upper layers should be used for high priority supplies. By placing high transient current supplies vertically closer to the device, this decreases the distance the currents need to travel through vias. Ground planes should also be adjacent to high transient current power planes to reduce inductance and couple the high-frequency noise.

6.2. Return Paths

Often designers misunderstand the use of ground planes as the best return for signal currents. Ground planes alone do not ensure a high-quality ground reference for high-speed AC circuits. PCB stack up should consider a signal reference to both ground and the power supply for the circuit. Layers above and below the signal layer should reference these power nodes accordingly. Slots should not interrupt the planes or they can possibly force current to find an alternate return path. This undesired return path could cause a localized "bounce" on the power or ground plane that can possibly be capacitively coupled to all signals adjacent to the planes.

7. Decoupling and Bypassing

Traditional methods for providing local power supply decoupling involve placing capacitors near the device in locations that are convenient based on the routing of the board, and applying some predetermined ratio of caps to power supply pins. Rule of Thumb provides many cap values in different decades like 0.1 µF, 0.01 µF, and 10 µF for each power rail with the smaller two values at each power pin. Unfortunately, the higher switching speeds of complex FPGA designs may render such typical ratios less than useful. Today's high-speed designs produce fast edge rates and large output loading leaving the decoupling Rule of Thumb guidelines less than optimal. Careful planning and analysis should be performed to ensure that sufficient decoupling is provided. Simulation with a power integrity solver like HyerLynx DC Drop is always a good idea to catch power plane issues and fix them in the PCB layout before you fabricate the PCB.

Among the FPGA device power pins are supplies that source power to the FPGA core, configuration logic, I/O buffers, phase-locked loops, and specialized SERDES power supplies. Dependent on the design intent of the FPGA device, a designer must pay strict attention to the PCB power distribution. Understanding that any unintentional coupling between supplies from high-speed switching currents can cause very undesirable performance problems. The FPGA also provides many high-speed ASIC-like I/O buffers. The interfaces that are used with these buffers are used in a variety of communication protocol bridges and memory interconnections. Some of the interfaces use terminated transmission lines. These terminations pose many concerns that need to be addressed in the power distribution scheme. They include low-impedance output termination voltages and quiet input voltage references. These also need correct decoupling to meet the performance expectations.

7.1. Capacitor Selection

Decoupling capacitors are generally chosen based on the individual capacitive property. However when capacitors are being selected for high-speed designs, the designer should carefully choose capacitors based on other parasitic characteristics such as inductance and resistance. Local decoupling capacitors should have low-effective series resistance (ESR) and low-equivalent series inductance (ESL) while having a large enough capacitance value to supply current to the IC during switching.

Every capacitor has a narrow frequency band where it is most effective as a decoupling capacitor. The frequency bands of some capacitors are wider than others. The effective frequency bandwidth of a capacitor is determined by the ESR and (Q) quality factor. Tantalum capacitors generally have a very wide effective band, while the lower ESR of ceramic X7R and X5R chip capacitors typically have a very narrow effective band. The dielectric material and geometry of the capacitor also determines how well the capacitor can suppress switching noise. Mixing several types of capacitors contribute to the total decoupling effectiveness.

In a typical FPGA board design, the capacitor closest to the power supply supplies the lowest frequency components of the load's varying current. The low frequency energy is decoupled by large electrolytic capacitors and traditionally governed by regulated voltage sources. These larger capacitors are employed as a method of low frequency filtering and to prevent supply droop. The droop is typically due to sections of a design becoming active and covers the lag until the regulator can respond. These big capacitors are usually electrolytic and have a low-frequency response of DC to a couple hundred KHz. Therefore, the close proximity of the capacitor to the FPGA is not critical.

The middle capacitors supply mid-frequency energy with large ceramic or tantalum capacitors. The use of these capacitors with generally a very wide effective band should be in close proximity of the FPGA. These capacitors typically have a response time adequate enough to counteract localized power droop caused by portions of the FPGA becoming active and changing the demand on the supply. Primarily, because of their low ESR, ceramic capacitors are often regarded as superior at high frequencies to tantalum capacitors. However, for decoupling, you can use the tantalum capacitor's ESR to dampen the resonances that result from interaction between the capacitors' ESLs and the pc board's various capacitances. The high ESR acts as a built-in damping resistor and makes the tantalum capacitor a good choice for decoupling.

As the number of decoupling paths increase, so do the number of voltage drops across them and this can result in power bus transients along with the associated common mode emissions. This problem can be minimized with proper power plane design in the area of the ICs. Use of adjacent power and ground planes in adjacent layers of the PCB stack up are capacitively coupled. This power and ground plane acts as an effective high frequency capacitor, and consequently, as an additional energy source which compensates for transient currents

The amount of current transients switching across the power bus increases in complex FPGA designs. These instantaneous current issues are typically associated with simultaneously switching outputs or SSO. Capacitors with very little intervening inductance supplies the localized high frequency energy and are needed to decouple noise from the switching currents of the power bus. The decoupling needed to prevent the instantaneous currents from attacking the device supplies are required to be placed directly by the FPGA. Many smaller caps used in parallel function as local energy storage for the device. Only a small amount of energy is stored in them and cannot provide DC power. But the function of this localized energy storage is to respond very quickly to changing current demands rather than large capacitors having more storage but poor response.

7.2. Localized Decoupling Considerations

Localized passive filtering is recommended to provide necessary isolation from high frequency power supply noise. The filter networks are recommended on the analog and high-speed transceiver (SerDes) supplies. These filter networks should include a series ferrite beads like a Murata BLM41P or BLM18A EMIFIL inductors. The impedance seen by the load is important because large output impedance causes the load noise current to be translated into a large noise voltage. Typically, the ferrite bead provides good isolation. This limits the energy between the source voltage and device supply. However, this requires a large bypass capacitor in order to keep the output impedance at a reasonable level. The smallest inductance that gives the required isolation should be used. The filter network should be AC coupled to GND through an appropriate 10-22-uF capacitor.
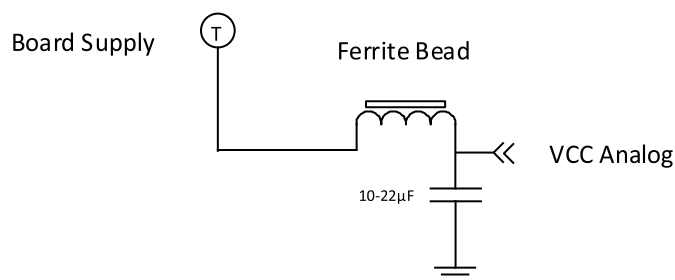

Figure 7.1. Passive Filter Network

7.3. Proper Decoupling Capacitor Placement

Increased spacing between the device and the decoupling capacitor increases the current path distance to the power and ground planes, consequently affecting the inductance of the current path between the device and the capacitor. To optimize the effectiveness of decoupling capacitors, surface-mounted capacitors mounted on the bottom-side of the PCB keeps the parasitic effects to a minimum. Placing capacitors directly underneath the BGA package improves the high frequency response of very small value capacitors.

Using surface mounted capacitors, the layout should not be allowed to reduce its effectiveness by connecting it through long, skinny traces leading to the power and ground. Use large or multiple smaller vias, and use short and fat traces to capacitors where possible.

Underside routing should be carefully done to place capacitors directly on package ball vias to device power pins. This technique reduces the distance of the current path traveled. The exposed metal on the surface of a PCB where surface-mount devices are soldered or land should have the shortest possible distance to connect to the device. The best practice is to eliminate any trace connections to capacitors. However, due to PCB assembly limits this is not always possible. The round-trip delay to the capacitor should be very small. For a particular frequency, the distance to the capacitor should not be greater than one-quarter of a wavelength. If the capacitors placement is greater than one-quarter wavelength then the energy transferred to the FPGA is negligible.

Use a layout method that hides signal ball vias and creates bottom layer islands that allows ample space to build large low-inductance areas that can accommodate several surface mount capacitors. The following Figure 7.2 depicts a BGA layout with underside decoupling capacitors.
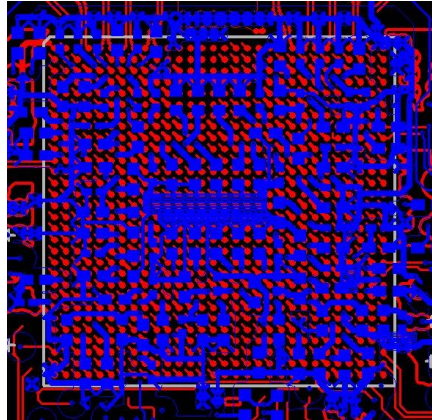
Figure 7.2. Decoupling Capacitor Placement Example

Use the following general checklist for making decoupling considerations.

1. Identify all high current sources and sinks and identify their return paths.

2. During layout, maximize the trace width to minimize inductance by mutual coupling. If possible, lay out power supply busses in a grid or a plane. Avoid long serial supply traces.

3. Utilize passive filter networks on analog supplies using series ferrite bead inductors and proper AC decoupling capacitors.

4. Bypass all high current sources and sinks with capacitors that work well at the frequencies of interest. Ceramic capacitors are good for this application for those capacitors placed very close to the FPGA because they are inexpensive, small, and work well at high frequencies. Use distributed equivalent capacitors to reduce ESR and ESL parasitics. Use tantalum capacitors for counteracting localized power droop. These capacitors work well when placed nearby the FPGA to provide energy storage.

5. Choose the decoupling elements based on required isolation and frequency response requirements. If using a simple series inductor, its value should be as small as possible. Avoid high-Q inductors. In this application, low-Q is desirable.

8. Special Design Considerations at >622 Mbps

8.1. Line Loss and Impedance Discontinuities

At data rates of 622 Mbits and higher, the skin effect is extremely important for signal conduction. Small traces on a PCB (like four or five mil widths) exhibits significant signal attenuation over long distances. Over-etching of a PCB can produce narrow traces that can reduce the signal amplitude available at the receiver. The end result is that, to the designer, the interconnection between devices resembles a badly designed low pass filter, with attenuation which increases with frequency. For this reason, the longer the backplane, the wider the signal traces should be. Long backplane traces (more than 20 in.) should have trace widths of 10 or 12 mils.

Connectors and vias in the signal path introduce discontinuities that resemble lumped elements in an electrical model. One way to take this into account is to perform SPICE or HyperLynx simulation of the backplane system using lossy transmission line models, and manufacturer supplied models of the connectors, signal drivers, and signal receivers.

8.2. High-Speed Connectors

Many connectors have been tried and discarded in high-speed applications. Surprisingly, some outdated connector designs have been found to be usable at gigabit data rates. An example of this is the DB-9 connector, which is sometimes found in Fibre Channel products. A more modern approach is to use controlled impedance connectors specifically designed for high-speed data, where abundant ground connections and shielding features reduce the noise and impedance discontinuities seen in older connectors.

Examples of these are the AMP Mictor connectors, and the 2 mm standard backplane connector families that are available from various suppliers for the 2 mm hard metric backplane standard (such as the AMP HS3 connector or the Tyco RT3). These are available in both vertical and horizontal configurations. Several evaluation cards use an unshielded 2 mm connector (AMP 636120-1) which provides good performance for driving twinax cables up to a distance of 65 ft (at 622 Mbps), error free. The Tyco RT3 connector is rated at 25 Gbps or more.

8.3. Device Packaging

Transmitter and receiver device packaging parasitic reactances are important to signal integrity. Wire-bond and package substrate inductance and capacitance should be included in device SPICE models. Simulation of package parasitics has shown that impedance transformation and signal reflections can result at higher frequencies. Pin location in larger packages can have a significant impact on the parasitic values of the model. Receivers with internal device terminations, such as those provided in Lattice LVDS and CML buffers, were found to have superior performance when compared to receivers which require external resistor termination components.

8.4. High-Speed Copper Cables

High-performance cables generally far outperform PCB interconnections in terms of bandwidth and signal attenuation. This is because a high-performance cable uses expanded Teflon dielectric (PTFE), silver-plated conductors, and low-loss shielding material. These cables are also engineered with a conductor geometry that is usually extremely close to the optimal position for the desired bandwidth and characteristic impedance.

One cable that performs extremely well is the W. L. Gore DXSN2112 Eye-opener Plus cable. This high performance cable is engineered specifically for data transmission at 622 Mbps. Unfortunately, this cable is not easily assembled with connectors using simple hand tools. Complete cables with connectors may be ordered directly from Gore. The cable assembly that matches the 2 mm backplane connector (AMP 636120- 1) on many Lattice evaluation cards is Gore part number 2MMA3106.

The latest high-speed cable assemblies from TE Connectivity are the QSFP28/56 and SFP28/56. These cable assemblies support aggregate data rates of 25, 50, 100, and 200 Gbps. These cable assemblies are at the forefront of the next generation connectivity, meeting 100G Ethernet and Infiniband Enhanced Data Rate (EDR) requirements.

Other Flyover Twinax cable assemblies are available from Samtec in FQSFP and FQSFP-DD formats for up to 800 Gbps 112Gbps PAM4 aggregate for eight channels, which have 16 pairs high-speed, 20 low-speed.
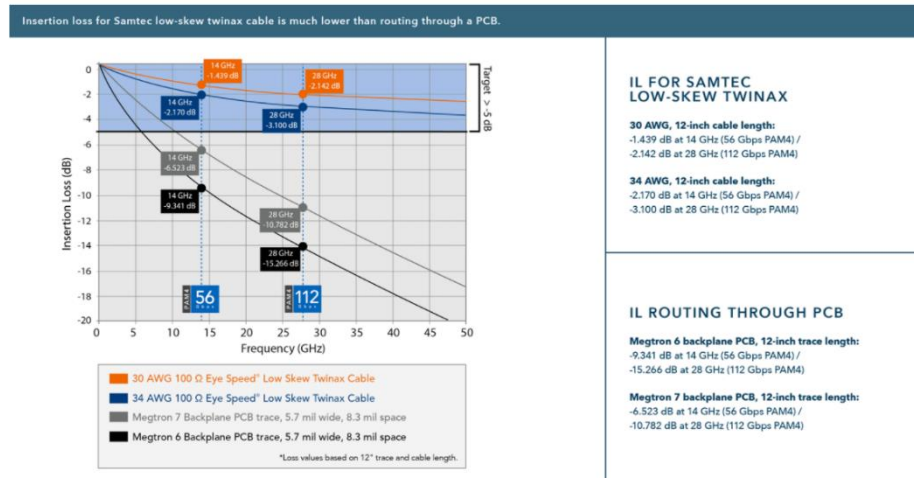
Figure 8.1. Insertion Loss Comparison for Flyover Twinax Cables versus Megtron 6 and Megtron 7 with 12 in Traces

9. Special Design Considerations at >2.5 Gbps

At 2.5 Gbps and greater, the design problem becomes substantially more difficult. The higher copper and dielectric losses occurring at these frequencies, generally limit PCB interconnection lengths to about 40 inches. The greatest care in all aspects of PCB layer and layout design is required at these frequencies.

9.1. Board Thickness and Vias

Backplane thickness and via design can have significant effects on signal integrity. A backplane thickness of less than 0.200 inches generally gives the best results. Vias used to interconnect between layers create transmission line discontinuities. These vias need to be top and back drilled to reduce the stub effects. PCB designs with high-speed signal traces should be routed on as few layers as possible, thus limiting the number of vias. Thicker boards normally have longer length vias that can cause larger discontinuities, and degrade the signals. Longer vias connecting signal layers that are close together appears as transmission line stubs, attached to the signal path. Stubs have been shown to have a very detrimental affect on signal integrity. Buried vias can be used to reduce this problem in thicker boards, but manufacturing costs for this technology can be prohibitive. Ideally, each signal path through the backplane should be kept on the same layer.

9.2. Board Material

FR4 dielectric loss becomes a significant design factor above 2 Gbps. Another design option is to use low-loss dielectric PCB material, such as Rogers 4350, GETEK, or ARLON. This is approximately double the cost of FR4 PCB material, but can provide increased eye-opening performance when longer trace interconnections are required. As shown from data collected by AMP Inc., Figure 9.1 gauges the improvement in signal eye opening at 2.4 Gbps, as lower loss materials are used. From the figure, FR-4 material may deliver a satisfactory eye opening. It might then be the preferred low cost solution for a particular application.

New materials like MEGTRON 6G & MEGTRON 7G are enabling board speeds upwards of 112 Gbps PAM4. There is also Tacheon 100G materials that are enabling speeds of 100 Gbps+. The Tacheon material from ISOLA is optimized with the use of spread glass to mitigate skew, improve rise times, reduce jitter, and increase eye width/height and that use ultra smooth HVLP (VLP2) 2 um Rz copper that significantly reduces conductor losses.
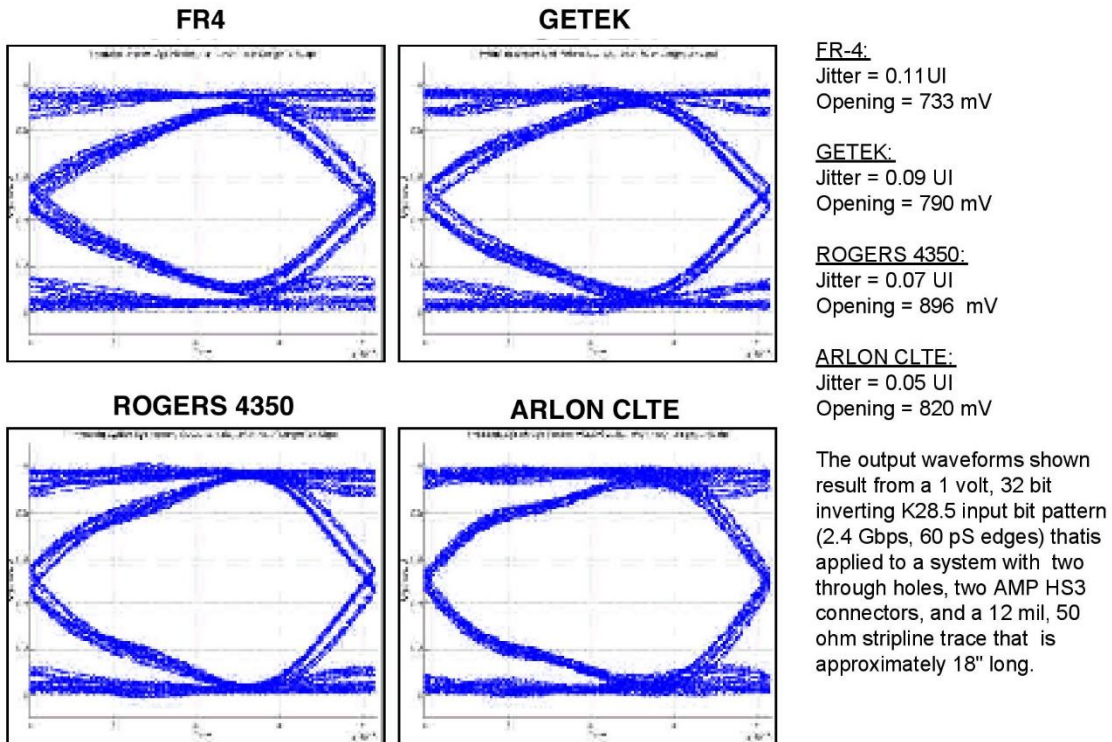
Figure 9.1. System Eye Patterns (2.4Gbs) vs. PCB Dielectric Material

FR-4:
Jitter = 0.11 UI
Opening = 733 mV

GETEK:
Jitter = 0.09 UI
Opening = 790 mV

ROGERS 4350:
Jitter = 0.07 UI
Opening = 896 mV

ARLON CLTE:
Jitter = 0.05 UI
Opening = 820 mV

The output waveforms shown result from a 1 volt, 32 bit inverting K28.5 input bit pattern (2.4 Gbps, 60 pS edges) that is applied to a system with two through holes, two AMP HS3 connectors, and a 12 mil, 50 ohm stripline trace that is approximately 18" long.

10. Special Layout Considerations for Lattice SERDES Devices

10.1. PCB Routing and Board Stack-up

Routing on inner and outer layers have impedance changes inversely proportional to line width and directly proportional to height. The rate of change in impedance with trace height above GND is much slower in a stripline signal compared with a microstrip signal. Stripline has considerably higher (typically one and a half times) propagation times than that of microstrip. To reduce microstrip losses you can also use a co-planar waveguide (CPW) at rates above 50 MHz to 30 GHz to increase edge rates and reduce losses. This uses a microstrip trace with two ground strips on either side of the trace. A coplanar waveguide with a ground reference (CPWG) covering the second layer gives a third return conductor and is a common variant.

A stripline route has a signal sandwiched by FR-4 material and a microstrip has one conductor open to air. A microstrip route is coupled to the ground plane below, which reduces EMI by absorbing some of the electromagnetic field lines. In stripline routing, all of the electromagnetic field lines are coupled to the above and below reference planes thereby reducing EMI significantly. To achieve the same line impedance, the dielectric distance must be greater in stripline layouts compared with microstrip layouts. A higher effective dielectric constant of stripline is due to the sandwich effect as compared to microstrip. Controlled impedance lines stripline traces are narrower than microstrip and it is also difficult to achieve accurate 100 Ω differential impedance on inner layer routing.

The before mentioned trade-offs between the microstrip and stripline routing play a vital role in determining total system jitter and signal strength characteristics. Experimentation has shown that the SerDes receiver input signals can be adequately routed using stripline without any critical system penalties. However, transmitter outputs routed on 1 ounce copper weight, 10-mil wide microstrip, provided the most optimal performance and signal characteristics on a 14 -layer FR-4 board versus experiments with stripline routing. Experimental results have shown maximum eye openings and low jitter exhibited with microstrip. In addition, routing the SerDes signals on the device-side outer layer offers additional benefits as it also reduces discontinuities from additional vias before any potential discontinuities at the on-board connectors.

11. High-Speed Connectors and IC Packaging

Above 1 Gbps, connectors specifically designed for higher frequencies are recommended. Several new controlled impedance backplane connectors have become available, with data-rate capabilities in excess of 25 Gbps. A popular example is the Tyco MULTIGIG TR3 2-mm pitch family, which has been carefully characterized and modeled at frequencies up to 25 GHz. This class of connector provides some additional shielding benefits, which can aid the designer in controlling system noise and crosstalk. The VITA 46 standard defines the connector requirements and OpenVPX standard slot configurations are defined in VITA 65.0 and 65.1. VITA 65.0 is the Sensor Open Systems Architecture or (SOSA) compatibility with VITA 46.

The IC packaging comments of the previous section apply here as well. SPICE modeling of the package parasitics is the best way to evaluate effects on system performance, since measurement probes typically have parasitics equal to or greater than those of the packages used today. Care should be taken to insure that the vendor provided package models are valid through the intended frequency of operation.

12. Pre-Emphasis

Signal pre-emphasis is a means of compensating for the increased PCB loss that occurs at higher frequencies. A simple algorithm can be employed in the line driver to increase transmitted signal amplitude, whenever the data patterns have transitions (and therefore higher frequency content). This function is provided by the CertusPro™-NX, Certus™-NX, LatticeECP5™, LatticeECP5-5G, and CrossLink™-NX SerDes CML drivers.

For longer PCB interconnection trace lengths, a significant increase in eye opening often results. Use of the pre-emphasis can extend the maximum usable interconnection length, or allow the use of lower cost (greater loss) material and components, in system design. A more detailed description of the pre-emphasis feature can be found in Lattice SerDes/PCS technical notes.

13. Receiver Equalization

Receiver equalization has certain advantages over pre-emphasis. Electromagnetic interference is less significant with equalization versus pre-emphasis because the system does not boost a high-frequency signal at transmission. The problem with transmit pre-emphasis is that the high frequency energy added to the signal contributes to much higher levels of crosstalk especially in connectors. Receiver equalization allows adaptability to varying trace lengths, and to various kinds of boards and signal swings.

An optimal solution is one that does not place excessive restrictions on board design layout and thus is one that offers the most robust performance. Especially when driving longer trace lengths, or to just provide a higher level of margin, both Tx pre-emphasis and Rx equalization are required. These are programmable adaptive settings on the Lattice SerDes.

14. Conclusion

PCB backplane interconnections with serial data rates at 25 Gbps and above are possible with today's technology. Lattice devices allow easy system design at rates to 1066 Mbps parallel interfaces. Increased performance with rates at and above 25 Gbps is achievable with Lattice SerDes based products, but greater care is needed in the PCB design. Systems running at these higher data rates may benefit from the use of transmitter pre-emphasis, controlled impedance connectors, and low loss PCB dielectric materials.

References

- CertusPro-NX Family Data Sheet (FPGA-DS-02086)
- ECP5 and ECP5-5G Family Data Sheet (FPGA-DS-02102)
- Mach-NX Family Data Sheet (FPGA-DS-02084)
- CrossLink-NX Family Data Sheet (FPGA-DS-02049)
- iCE40 UltraPlus Family Data Sheet (FPGA-DS-02008)
- ORTx2G5, ORSOx2G5 and ORSPI4 High-Speed Backplane Measurements (TN-1027)
- Transmission of High-Speed Serial Signals Over Common Cable Media (FPGA-TN-1066)
- Power Decoupling and Bypass Filtering for Programmable Devices (FPGA-TN-02115)
- PCB Layout Recommendations for BGA Packages (FPGA-TN-02024)
- Si9000 field solving impedance calculator: www.polarinstruments.com
- W.L. Gore high-performance interconnect products: www.goreelectronics.com
- Samtec Flyover QSFP28 Cable: www.Samtec.com/Flyover
- www.Saturnpcb.com/saturn-pcb-toolkit/

Technical Support Assistance

Submit a technical support case through www.latticesemi.com/techsupport.

Revision History

Revision 6.3. April 2022

Section	Change Summary
Acronyms	Newly added section.
Backplane Topology and Overview	Changed to *This type of backplane interconnection can be used with data rates to 25 Gbps and above*.

Board Design Practices	• In the PCB Trace Impedance Calculation section:
	  • updated the 2D field solver program example to the *Si96000eb program from Polar Instruments* and updated relevant description.
	• In the Examples of PCB Trace Impedance Calculation section:
	  • updated to the latest Product and the latest speed;
	  • updated Figure 5.5. Differential Microstrip Example with Saturn PCB Toolkit Impedance Calculator Tool.
	• In the PCB Design Checklist section:
	  • updated to Use 3 S or 4 S separation rules.

Decoupling and Bypassing	Newly added description about *Rule of Thumb* and *Simulation with a power Integrity solver*.

Special Design Considerations at >622 Mbps	• Newly added description about HyperLynx and Tyco RT3 connector.
	• In the High-Speed Copper Cables section:
	  • newly added description of the latest high speed cable assemblies from TE Connectivity.
	  • newly added Figure 8.1. Insertion Loss Comparison for Flyover Twinax Cables versus Megtron 6 and Megtron 7 with 12 in Traces.

Special Design Considerations at >2.5 Gbps	• In the Board Thickness and Vias section:
	  • newly added These vias need to be top and back drilled to reduce the stub effects.
	• In the Board Material section:
	  • added description regarding new materials increasing the board speed.

Special Layout Considerations for Lattice SERDES Devices	• In the PCB Routing and Board Stack-up section: • added description on how to reduce microstrip losses. • newly added the last paragraph.
High-Speed Connectors and IC Packaging	General update to the first paragraph reflecting the most recent FPGA product information.
Pre-Emphasis	General update to the first paragraph reflecting the most recent FPGA product information.
Conclusion	Updated data rates to 25 Gbps and *allow easy system design at rates to 1066 Mbps parallel interfaces*.
References	Globally updated the list.

Revision 6.2, December 2019

Section	Change Summary
All	• Changed document number from TN1033 to FPGA-TN-02178. • Updated document template.
Disclaimers	Added this section.

Revision 6.1, April 2011

Section	Change Summary
All	Updated for LatticeECP3 FPGA family.

Previous Lattice releases