



XAPP489 (v1.0) October 31, 2006

Four- and Six-Layer, High-Speed PCB Design for the Spartan-3E FT256 BGA Package

Summary

This application note addresses low-cost, four- to six-layer, high-volume printed circuit board (PCB) layout for a Spartan™-3E FPGA in the FT256 1 mm BGA package. The impact of high-speed signals and signal integrity (SI) considerations for low layer count PCB layouts is also reviewed. This application note is intended for design engineers, managers, and PCB layout staff, who are already familiar with SI related design issues. It focuses on the Spartan-3E device in the FT256 package, but the information applies to the equivalent FG256 package, and the general guidelines can be used to optimize board layout for other devices and packages.

Introduction

PCB costs are fundamentally driven by two factors: available manufacturing capabilities and volume. Rules for designing a low-cost PCB are dictated by what is manufacturable in PCB production facilities for the minimal price. This reality also dictates the number of PCB layers implemented in a circuit that allow it to remain at a low cost and suitable for high-volume manufacturing. Unfortunately, the market demand for increased package pin count on programmable logic adds to the pressure on PCB layout costs because of smaller form factors. Despite these factors, a board using the Spartan-3E FPGA in the FT256 1 mm ball grid array (BGA) package can be designed at lowest cost with four layers.

Board designs that use exotic design rules, such as 1 mil trace and space, have limited manufacturing options and very high costs. One mil is one thousandth of an inch, also known as a *thou*, and is equal to 0.0254 mm. Some North American vendors might be able to manufacture boards using these rules, and it is not likely that moving the PCB fabrication to mainstream production facilities in Asia will provide major cost reductions. With higher volume production, more vendors are interested in manufacturing the boards and providing cost reductions, but the time needed to achieve acceptable costs might exceed the product life. This application note offers solutions to improve the manufacturing options and lower costs.

PCB Design Rules to Minimize Cost

Table 1 provides an indication of the PCB fabrication cost for different layer counts and fabrication locations in high volumes. The more layers there are, the higher the fabrication cost. Fabrication cost varies according to volume and market conditions.

Table 1: PCB Fabrication Cost

Fabrication Location	Cost in U.S. Dollars/Square Inch ⁽¹⁾		
	2 Layers	4 Layers	6 Layers
Domestic U.S.	0.18	0.22	0.24
Asia	0.13	0.16	0.17

Notes:

1. These costs were assessed in mid-2004 and assume 100% panel utilization, not including unusable areas, such as scribe and breakaway area.

© 2006 Xilinx, Inc. All rights reserved. All Xilinx trademarks, registered trademarks, patents, and further disclaimers are as listed at <http://www.xilinx.com/legal.htm>. All other trademarks and registered trademarks are the property of their respective owners. All specifications are subject to change without notice.

NOTICE OF DISCLAIMER: Xilinx is providing this design, code, or information "as is." By providing the design, code, or information as one possible implementation of this feature, application, or standard, Xilinx makes no representation that this implementation is free from any claims of infringement. You are responsible for obtaining any rights you may require for your implementation. Xilinx expressly disclaims any warranty whatsoever with respect to the adequacy of the implementation, including but not limited to any warranties or representations that this implementation is free from claims of infringement and any implied warranties of merchantability or fitness for a particular purpose.

Table 2 lists widely accepted manufacturing rules to minimize fabrication cost. For example, the tolerance on small traces is difficult to control. Designing a PCB with a trace width of not less than 5 mils, as shown in Table 2, results in minimized fabrication cost.

Table 2: Manufacturing Rules to Minimize Cost

Feature	Imperial (mils)	Metric (mm)
BGA Ball Pad Diameter	15	0.381
Trace Width	5	0.127
Space	5	0.127
Via Pad Diameter	23	0.584
Via Hole Diameter	12	0.305

Ball Pitch, Escape Routes and Fanout

The FT256 BGA package, which has a 1 mm/40 mil BGA ball pitch, follows the manufacturing rules for trace width and space in Table 2. All signals of the FT256 package are located in only two layers (for example, top and bottom), as shown in Figure 1.

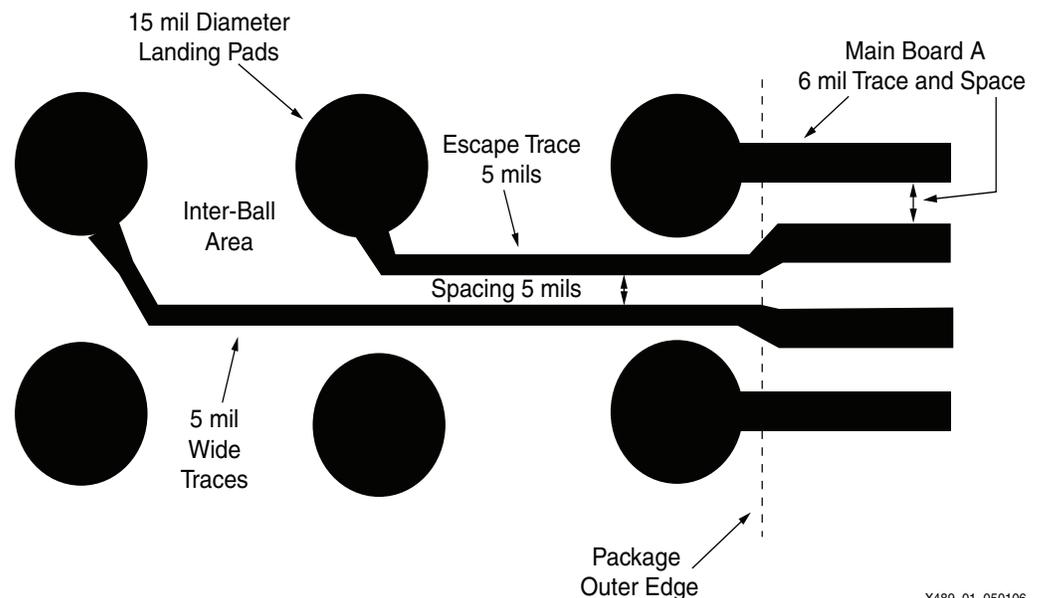
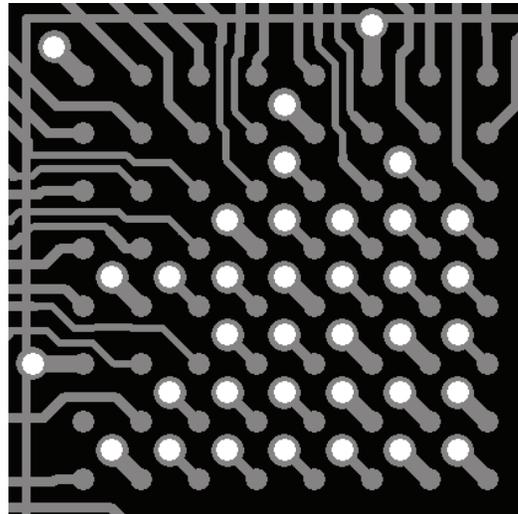


Figure 1: Two Signal Layers of the FT256 Package

The FT256 package has an array size of 16 x 16 balls. The ability to escape the three outer ball rings on a single layer, ignoring the peripherally placed power balls, allows for a maximum escape of 156 signals on the top layer alone. Figure 2 shows one quadrant of a complete three outer-row signal fanout for the top signal layer of an XC3S500E demo board.

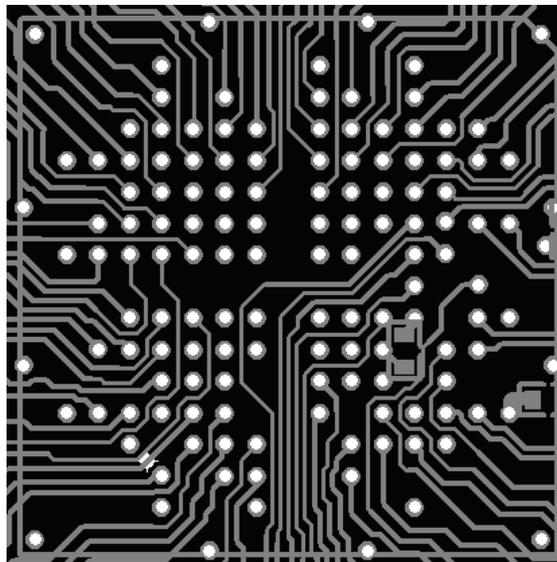


X489_02_033006

Figure 2: One Quadrant of the Top Layer

The fanout pattern is symmetric and applies to all four quadrants. All traces and via patterns fan away from the array center.

By repeating the exercise using vias for the remaining inner ball rows on the bottom layer, a maximum of 84 signal pins can escape. [Figure 3](#) shows the escape routes used on the bottom layer of the demo board. Because the board below the center of the Spartan-3E FT256 package is populated with ground pins, the escape routes of the FT256 package can all be directly connected through adjacent vias to ground. Outside the immediate BGA package area, the escape traces can be fanned out to at least 6 mils for trace plus 6 mils for space.



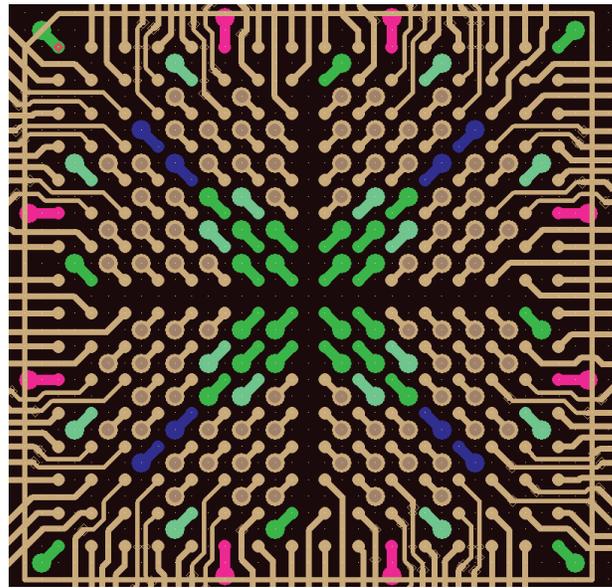
X489_03_032306

Figure 3: Trace Escape Routes on the Bottom Layer

Via Placements

Signal vias should be limited to the center of the ball array whenever possible, with the via toward the outside of the land pad. Vias can cause routing blockages if their placement and size are not controlled for efficient routing of escape routes. Most Spartan-3E FPGA power pins are located at the center of the array, so that the required vias for these pins cause the fewest possible routing blockages.

Figure 4 shows the complete fanout for the XC3S500E FT256 package. Because the symmetric via pattern allows balls on two to three rows to be brought out on the top layer, all I/O signals are brought out using only two signal layers. The vias for the I/O signals are shown in tan. All other vias are for power, as defined in Table 3.



X489_04_081206

Figure 4: Via Placement and Fanout

Table 3: Via Types Key

Color		Voltage (V)	Signal
Green		0	Ground
Pink		2.5	V _{CCAUX}
Dark Blue		1.2	V _{CCINT}
Aqua		Various	V _{CCO}
Tan		Various	I/O

Figure 5 shows only one quadrant of the vias in the FT256 package to display a larger view of the via pattern and symmetry.

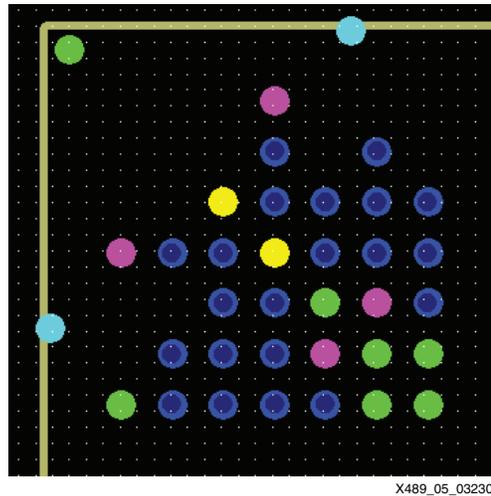


Figure 5: One Quadrant of the Via Placement

Decoupling Capacitor Placement

On any board, the required proximity of the ceramic decoupling capacitors is determined from the ratio of the signal rise time to the time it takes for the same signal to propagate the length of the trace between a capacitor lead to the power pin that it decouples. This topic is covered in more detail in [“Identification and Layout for High-Speed Signals,” page 11](#). The basic rule is that capacitors become less effective the further they are from the package power pin. This effect is proportional to both the frequency and the inductance, which is proportional to the trace length between the capacitor and power pin, as shown in [Equation 1](#).

$$Z = 2\pi FL \quad \text{Equation 1}$$

Z is the impedance in ohms, F is the frequency in Hertz, and L is the inductance in Henries. [Figure 6](#) shows these critical distances for a design with approximately 200-300 ps edge rates. In a design with a slower edge rate, the separation distance can easily be an inch or two without causing noise problems.

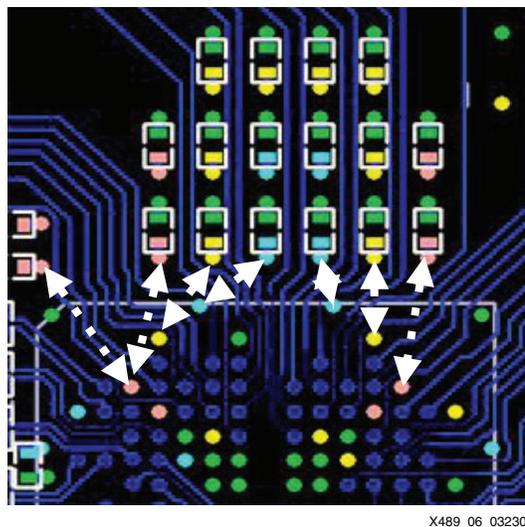
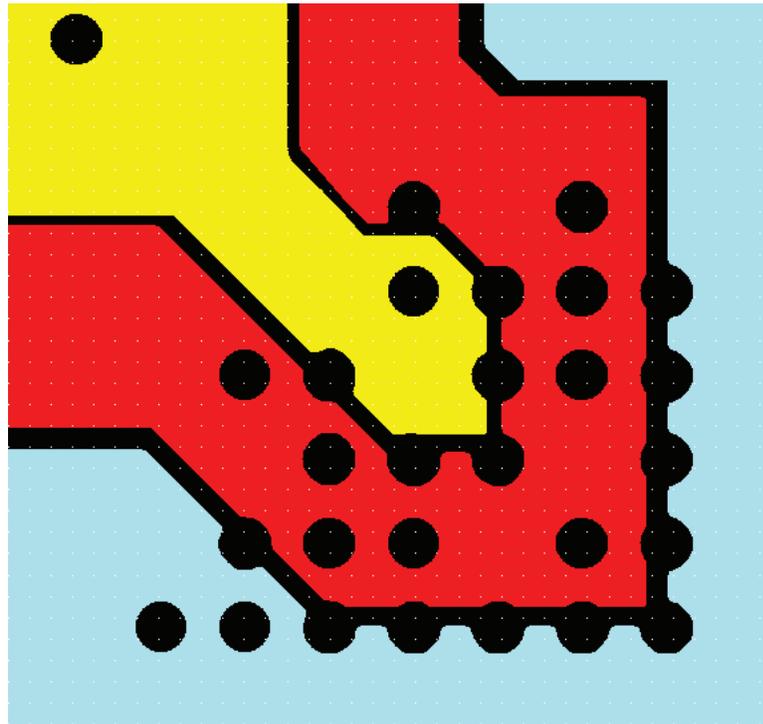


Figure 6: Capacitor Placement

On a two-signal-layer board, routing space is a premium. The FT256 package provides 100% signal fanout on two signal layers. The vias in [Figure 6](#) are 0603-sized, Imperial 100 nF ceramic capacitors. For high-volume, low-cost boards, 0402-sized, Imperial capacitors are acceptable for assembly.

Power Planes

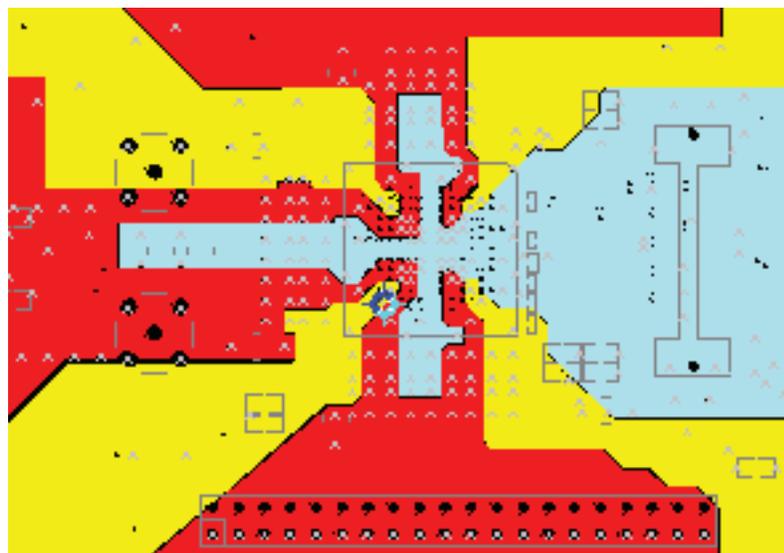
The Spartan-3E high-speed design successfully uses a one-layer, triple split plane. One quadrant of the design is shown in Figure 7. The design has 200-300 ps edge rates and a switching noise of 80 simultaneously switching LVCMOS 3.3V I/Os.



X489_07_032606

Figure 7: One Quadrant of the One-Layer, Triple Split Plane

The final plane requirements are determined by the required I/O and V_{CCO} voltages. If the V_{CCO} voltage is 2.5V, the layout is simple because only two planes are required. An example split plane is shown in Figure 8. Plane requirements may also be driven by other SI concerns, such as the need to avoid shapes that resonate at selected frequencies or the need to avoid splits when routing traces on adjacent layers.



X489_08_032606

Figure 8: Example Split Plane

The plane areas should be contiguous, without breaks, if possible. The light blue and red plane areas are clearly contiguous in [Figure 8](#). [Figure 9](#) shows a larger view of the board in which the yellow plane area can also be seen as contiguous.

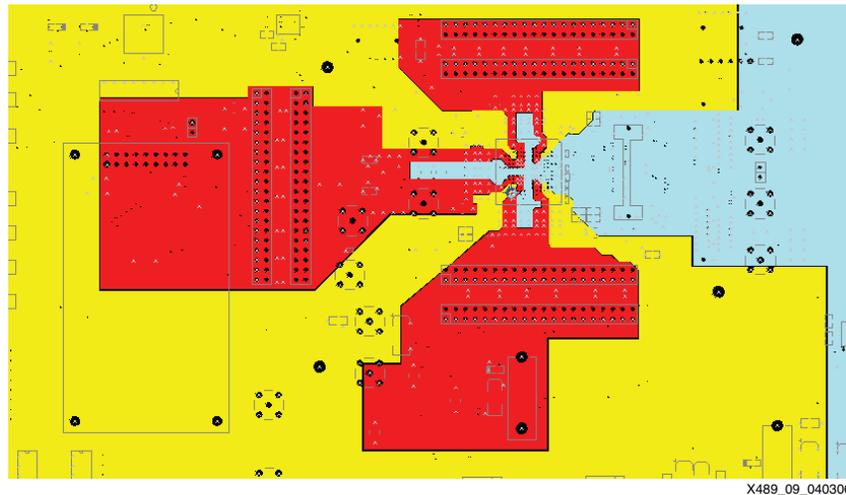


Figure 9: Contiguous Yellow Plane Area

The solid ground plane used on this board is shown in [Figure 10](#).

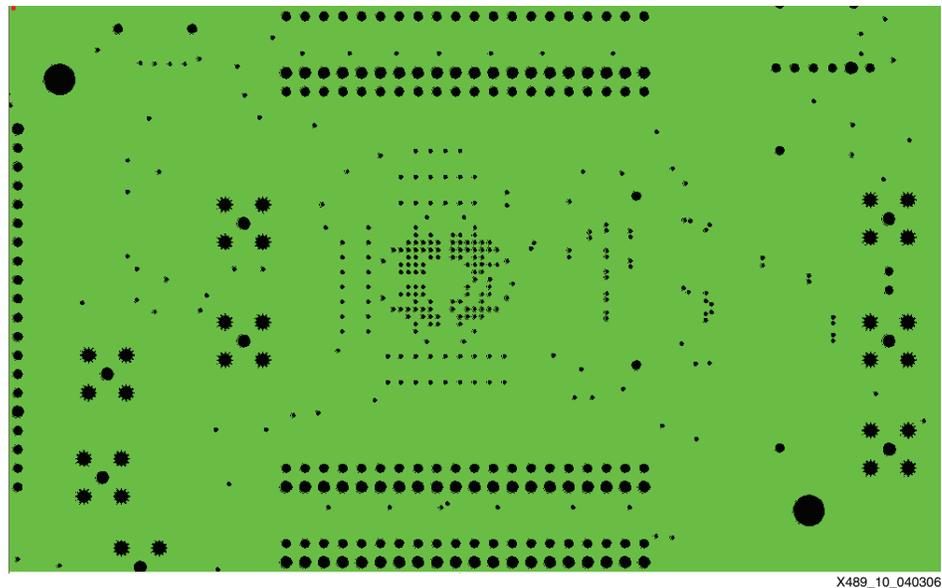


Figure 10: Ground Plane

Four-Layer Board Stack Example

For a low-cost, four-layer board, the planes shown in [Figure 7](#) through [Figure 10](#) can be stacked as described in [Table 4](#).

Table 4: Four-Layer Board Stack

Layer Number	Signal/Power
1	Signal
2	GND
3	$V_{CCINT}/V_{CCAUX}/V_{CCO}$
4	Signal

In this particular case, the board has high-speed LVDS signals running with $V_{CCO} = 2.5V$ and requires a $100\ \Omega$, differential-controlled impedance. This impedance requirement is met using the GND plane coupled to the top-layer LVDS signals and a $2.5V\ V_{CCO}$ plane area on Layer 3, coupling to the LVDS signals on Layer 4.

Six-Layer Board Stack Example

In a six-layer board, various stack combinations can be used for the power planes. [Table 5](#) shows a six-layer stack optimized for decoupling. In this stack, four plane layers provide three decoupling capacitors. A poor stack choice would be to put two ground planes together at layers 3 and 4, because the effective plane decoupling is significantly reduced. Flooding unused and complementary board areas on layer 1 and layer 2, which are connected to plane voltages, can also be used to increase the effectiveness of power plane decoupling. The same technique can be used on layers 5 and 6.

Table 5: Six-Layer Board Stack

Layer Number	Signal/Power
1	Signal
2	GND/Signal
3	V_{CCINT}/V_{CCAUX}
4	GND
5	V_{CCO} /Signal
6	Signal

The Importance of Power Planes

Power supply design starts with an understanding of an acceptable level of power supply ripple and noise across the operational frequency range of the target circuit. The top end of this power spectrum is defined by the rise time of the fastest I/O signal, sometimes referred to as the *knee frequency*. It is called the knee frequency because a plot of the spectral power density for random digital signal patterns shows a critical fall-off at this frequency that looks graphically like the shape of a bent human knee. This frequency is defined for all signals, where the rise time is effectively half the period, as shown in [Equation 2](#).

$$F_{knee} = \frac{1}{2}(\text{Rise Time}_{fastest}) \quad \text{Equation 2}$$

For example, a system with a 1 ns signal edge rate, requires that the power supply must support a bandwidth of 500 MHz. This frequency is the upper limit at which impedance calculations must be done to ensure the effective decoupling. If the supply has a lower bandwidth than the knee, the edge rates are also lowered. EMI generated from high-speed digital circuits extends two to three times beyond the knee frequency.

The designer must choose decoupling capacitor values and types along with power planes to cover the power spectrum of the design. For a high-speed board design, the following four decoupling regions can affect the board's noise performance:

1. Electrolytic (or Bulk) Capacitors

[Table 6](#) lists the characteristics of key parameters for electrolytic (or bulk) capacitors, such as aluminum and tantalum.

Table 6: Characteristics of Electrolytic (or Bulk) Capacitors

Parameter	Typical Value	Comment
Value Range	0.1 μ F - Farads	Good
Inductance	10-60 nH	Very High - Good for Lower Frequencies
Resistance	700 m Ω	High - Good for Lower Frequencies
Cost	\$0.05 - \$1 or more	-

2. Ceramic Capacitors

[Table 7](#) lists the characteristics of key parameters for ceramic capacitors.

Table 7: Characteristics of Ceramic Capacitors

Parameter	Typical Value	Comment
Value Range	Picofarads - Microfarads	Good
Inductance	1 nH	Low - Good
Resistance	100 m Ω	Low - Good
Cost	\$0.01 - \$2	Good

3. Power Planes

Suitably designed power planes can provide decoupling at the highest end of the required frequency range. Power planes are simply large-sized capacitors, and their capacitive values follow the same formula for the capacitance of two plates, as shown in [Equation 3](#).

The spacing between the power planes form the plates of a high-quality (low-impedance) decoupling capacitor. Because many other factors, such as the board size, can already be fixed by other design constraints, the critical factors to consider at this stage are the plane separation and the dielectric material used to separate the planes.

$$C_{PowerPlane} = \epsilon_o \times \epsilon_r \times \frac{Area}{PlaneSeparation} \quad \text{Equation 3}$$

where:

- ◆ $C_{PowerPlane}$ is the capacitance of the power planes.
- ◆ ϵ_o is the permittivity for vacuum.
- ◆ ϵ_r is the relative permittivity or separating insulator.
- ◆ Area is the common plane area in square inches.
- ◆ PlaneSeparation is the spacing between the planes in inches.

As an example:

- ◆ $\epsilon_o = 2.25 \times 10^{-13}$
- ◆ $\epsilon_r = 4.7$ (for an FR4 PCB)
- ◆ Area = 9 inches x 10 inches = 90 in² (must be common plane area)
- ◆ PlaneSeparation = 0.003 inches (3 mils)

Thus $C_{PowerPlane} = 32$ nF, or 353 pF per square inch. The capacitance of the printed circuit board power planes, although small in value, adds a high-quality decoupling component, because of its relatively low impedance.

4. Power/Ground Pin Proximity

Close proximity of signals to power pins helps to minimize the ground loop and noise effect. For example, either flood surrounding areas that are free from traces with ground copper or consider assigning otherwise unused connector pins to ground.

Basic Design Rules for Power Planes

A power plane should:

- Have the largest available area
- Have wide connections
- Be coupled closely to ground planes
- Be symmetric
- Have minimal separation from the ground planes
- **Not** have breaks (holes or cut-outs)

Because the principal limitation on the high frequency response of plane capacitance is inductance, narrow power plane areas should be avoided. The goal is to have the largest available area that can be used as a power plane coupled with a complementary ground plane. Device power pins should have vias that go directly into the plane, as soon as design rules allow. Trace lengths between the power pin and the plane via should be minimized. For noise reduction, the best location for power plane vias in BGA-packaged devices are in the pad. Unfortunately, unless vias are already used in pads elsewhere on the board, adding vias in pads that have to be filled (to prevent solder ball wicking into the via hole) can increase the cost of the PCB by approximately 10% to 20%.

Split Plane Area Requirements

Each power rail has its own power decoupling requirements. These should be factored into the proportions of the planes on a split plane layer.

The quickest and easiest plane layout is to dedicate one layer per plane. In cases where time to market is an overriding consideration, these planes present the fastest layout. The layer count can be reduced later during the cost reduction phase.

[Table 8](#) shows approximate power plane decoupling requirements. Power analysis and estimation can be performed using the [Xilinx power tools](#).

Table 8: Power Plane Decoupling Requirements

Rail	Voltage (V)	Normalized Power Ratio ⁽¹⁾
V _{CCINT}	1.2	400%
V _{CCAUX}	2.5	200%
V _{CCO} ⁽²⁾	3.3	100%

Notes:

1. The simulation conditions were:
 - ◆ There was 90% utilization of the CLBs, IOBs, and block RAM.
 - Of the 90%, 60% of the CLB logic toggled at 70 MHz for 25% of the time, and 40% of the CLBs toggled at 150 MHz for 25% of the time.
 - Of the 90%, 60% of the block RAM with 18-bit read/write ports toggled at 70 MHz for 25% of the time, and 40% of the block RAM with 18-bit read/write ports toggled at 150 MHz for 25% of the time.
 - ◆ Half of the multipliers ran at 70 MHz, and the other half ran at 150 MHz. Medium toggle rate.
 - ◆ Half of the DCMs ran at 70 MHz, and the other half ran at 150 MHz. Low-frequency mode.
 - ◆ Of the 90%, 60% of the IOBs, divided equally as inputs and outputs, toggled at 70 MHz for 25% of the time, and 40% of the IOBs, divided equally as inputs and outputs, toggled at 150 MHz for 25% of the time. Outputs were enabled 50% of the time.
2. If multiple I/O levels are required in the design, consider each bank on its own merits.

Signal Integrity

Today 90 nm FPGAs have extremely high-performance I/Os. The Spartan-3E FPGA has I/Os that can switch LVDS at data rates in excess of 800 MHz. Being programmable, the same I/O cells can obviously be switched at much lower frequencies. Unfortunately lower frequency does not always equal low speed. While signal integrity is indeed an intrinsic high clock rate issue, it is not the clock frequency that determines whether I/Os generate noise needing to be terminated. Instead the output driver edge rates or rise times are the factors that determine whether traces need termination. A design clocked at 1 MHz can have serious signal integrity issues.

As an analogy, consider being hit on the head with a heavy object. It is not how often the collision happens that causes pain. Pain is caused by the speed at which the object strikes the head. A light tap on the head ten times a second is no more painful than being tapped lightly on the head once a second. However, even a single heavy, fast blow is painful.

Frequency is not the key factor here; the critical issues are the power or force and the rate of delivery.

Exactly the same is true with I/O circuits. An output transition can be slow, and the input circuit likely detects and receives it without problem. But speed up the transition too much, and not only does the input circuit start to see noise (pain), but an undamped, fast transition can potentially overstress or damage circuitry! The danger to I/O circuits with excessively fast signal transitions is *overshoot* and *undershoot* noise. Severe and prolonged overshoot or undershoot can potentially damage I/O structures over time, because the voltages force high currents into or out of the device.

Readers who are unfamiliar with SI issues are referred to *High-Speed Digital Design* by Drs. Howard Johnson and Martin Graham [Ref 4] and also XAPP623 [Ref 5], an application note on power supply design and layout.

Identification and Layout for High-Speed Signals

A *high-speed signal* is fundamentally defined by the signal's rising edge rate, and its relationship to two additional parameters:

- Signal propagation speed along its PCB trace
- The full length of the signal trace

The above qualifications do NOT include a minimum clock frequency. Instead for signals with fast edge rates, the physical properties of the board are used to determine whether the *fast* rates are possible.

Even signals with fast edge rates, which are toggled at very low frequencies, can cause serious logic problems. With high-speed signal operation, the time it takes for a signal's rising edge to transition between 10% and 90% V_{OH} is close to the time the same signal takes to propagate the length of its trace on a PCB. For this case, designers have to either use signal termination techniques where the traces are described as *transmission lines*, or determine a maximum trace length where traces can be treated as *low speed* or *lumped* circuits.

What Qualifies as High-Speed?

This section explains how to recognize a high-speed signal and whether signal integrity issues must be considered during PCB layout.

When designing a PCB, the designer needs to address the following issues:

- When is the signal rise time, T_r , treated as high speed?
- What can be done to the circuit board layout to avoid treating the signal as high speed?
- Can SI simulation and transmission line models be avoided?

A simple example shows how to decide whether or not high-speed signals and signal integrity are of concern for the PCB layout. In this example, the circuit board design has the following characteristics:

- I/O Type: Spartan-3E IBIS Model, LVCMOSS33, Fast, 12 mA (strong-fast model)
- Signal Rise Time or *Edge Rate*: $T_r = 0.225$ ns
- Signal Propagation Speed: $T_{PD} = 1.8$ ns/foot
- PCB Trace Type: Microstrip
- PCB Dielectric Material: FR4

The designer uses [Equation 4](#) to calculate the edge rate length (L_{Tr}).

$$\begin{aligned}
 L_{Tr} &= \frac{T_r}{T_{PD}} \\
 &= \frac{0.225 \text{ ns}}{1.8 \text{ ns/foot}} \\
 &= 1.5 \text{ inches (38 mm)}
 \end{aligned}
 \tag{Equation 4}$$

Circuits smaller than $L_{Tr}/6$ can be considered lumped and NOT transmission lines, so

$$L_{Tr}/6 = 0.25 \text{ inches (6.35 mm)} \tag{Equation 5}$$

This *critical* length is directly proportional to the edge rate, T_r , so a circuit board with a signal with a slower edge rate, such as 5 ns, can have a trace length that is 10 times longer before signal integrity becomes a concern.

In a second example, the board design characteristics are:

- I/O Type: Spartan-3E IBIS Model, LVTTTL, Slow, 4 mA (strong-fast model)
- Signal Rise Time or *Edge Rate*: $T_r = 2.3$ ns
- Signal Propagation Speed: $T_{PD} = 1.8$ ns/foot
- PCB Trace Type: Microstrip
- PCB Dielectric Material: FR4

Substituting the new example values into [Equation 4](#) and [Equation 5](#) yields:

$$\begin{aligned}
 L_{Tr} &= \frac{T_r}{T_{PD}} \\
 &= \frac{2.3 \text{ ns}}{1.8 \text{ ns/foot}} \\
 &= 15.3 \text{ inches (389 mm)}
 \end{aligned}
 \tag{Equation 6}$$

$$L_{Tr}/6 = 2.55 \text{ inches (65 mm)} \tag{Equation 7}$$

In conclusion, signals must be treated as high speed and traces terminated, when the time taken for an output to rise from 10% to 90% of its final value is less than around 1/6 the time that this level change takes to travel along the full length of its PCB trace and return.

Circuit board layout can be simplified and area kept to a minimum by keeping trace lengths within this limit. In these circumstances, extensive use of SI simulation and transmission line models can be avoided. However, if the fastest output rise times cannot be determined from data sheets or IBIS models, some modeling might be required to accurately determine a maximum permissible trace length for any given signal.

Modeling the Trace

In circuits where the trace length is less than $L/6$, the trace can be treated or modeled as a simple RC or *lumped* circuit, where the Rs and Cs add and combine as in low-speed RC circuits. For traces where the length L is greater than $L/6$, the trace must be modeled as a transmission line. $L/6$ is used to determine the threshold between a trace being considered as a lumped circuit and where a *transmission line* model is more appropriate. $L/6$ is a conservative guideline, and $L/4$ is the commonly accepted upper limit.

However, considering the trace as a transmission line makes the math more complex, and signal reflections need to be considered. A *reflection* is simply an echo of the transition, just like an audible echo that returns from the end of the trace before the transition edge has completed or finished. In a low speed or lumped circuit, the distances are so small and the echo happens so quickly that the transition and the echo are indistinguishable. Like audible echoes, when the echo has the combined properties of high amplitude and significant delay, the result is a confused or corrupted message. In the lumped case, the components may be treated as responding together in time, and hence their response is as a single lump. A high-speed system is called *distributed* because the response of each component has to be analyzed in a more independent manner. The reason is that each component at such high signal speeds can see a different voltage from other portions of the circuit depending upon the distance of it from the other components, hence the term *distributed*.

The problems with reflections or echoes in high-speed circuits are analogous and result in the same message corruption. In the high-speed signal case, the message is of logical ones and zeros. Like audible echoes, the problem may be solved simply by stopping the reflection from occurring at the far end so that it does not return. In the case of a sound echo, an absorbent material may be used to dampen or kill the reflected sound energy. In the case of high-speed electrical signals, the same effect is achieved with a suitable resistive value applied to absorb the signal energy at the far end. There are several different techniques for signal termination, but they all aim to absorb the signal energy so that the reflection does not cause signal confusion. Without suitable termination, signal reflections can cause anything from a continuous and complete loss of the signal, literally a zero volt signal amplitude to a double voltage signal amplitude. What a receiving logic gate sees is entirely dependent upon its relative position on the trace. In this worst case scenario, the effect on the signal is like a resonant sound tube, such as an organ pipe or tuning fork. Instead of the wave or sound dying slowly away, the signal transition persists as a standing wave or resonates.

Inductors as Resistors

At high frequencies, other characteristics, such as the inductance in a wire, trace, package lead, or pin, can result in effects that would otherwise be insignificant at lower frequencies. For example, consider an oscilloscope probe with two inches of ground wire from the probe to the package pins on a device being examined. If the signal transitions being observed are at 100 MHz and the wire approximates to an inductance $L = 200$ nH, then that simple ground wire acts like a resistor to the signal edges, as shown in [Equation 8](#).

$$Z = 2\pi FL \quad \text{Equation 8}$$

In this example, Z is on the order of $2 * 3.14 * 100^6 * 200^{-9} = 120 \Omega$. If the output driver in this example is CMOS and switches between 0 and 3.3V with an output impedance of approximately 30Ω , the observed output swing is noticeably different from what was expected. In high-speed designs, care must be taken to ensure that the power supply can respond to the signal's highest bandwidth requirements. For a fast signal transition to propagate along a PCB trace, the board's power supply must be able to respond at least as fast as the required transition frequency. Less obvious is that this response must be supported along the full length of the high-speed trace. Decoupling and impedance must therefore be considered at the trace and component level and not just for the board as a whole.

Characterizing Board Noise Quality

In a synchronous design, noise that occurs close to the clock edge is the main problem that is targeted by a good power supply design. In practice, unless a design has multiple asynchronous clocks, at least for the digital portion of a design, the clock edge is the dominant window of time when switching currents generate noise. The following subsections describe methods for determining the effect of board noise.

Acceptable Noise

Even apparently severe noise levels seen between clock edges can often be ignored, as long as the noise settles before the all important clock edge(s). Of course, there can be other non-digital circuits where such noise might be a problem, and serious switching noise can occur at any part in the clock cycle for circuits that have multiple asynchronous clocks.

The Problem with Asynchronous Clocks

Two or more asynchronous clocks are completely independent of each other. Because all clocks drift, these clocks drift randomly in their phase relationship. The unavoidable problem with asynchronous clocks is that at some point in time the two clocks drift to where one switches and generates noise at the same time as the other clock is in its critical transition region and all the circuits switched by it at this point are most vulnerable to noise. This problem is similar to the issue of metastability and the need to synchronize signals transmitted between two circuits with asynchronous clocks.

Simulation

Simulation can be extremely valuable in modeling noise. In particular, IBIS and a simulator such as HyperLynx (Mentor Graphics) can help identify noise problems due to poor board layout. SPICE is a powerful modeling tool to use for simulations but is difficult to learn. IBIS is a dominant standard for I/O noise simulation, and most IC vendors provide IBIS for the following reasons:

- Ease of use and support
- Widespread tool support
- Performance

IBIS simulation is faster than SPICE. IBIS represents I/O circuit behavior using a text-based database. I/O simulation requires a database read/look-up, whereas SPICE can require heavy computation.

- Protection of confidential circuit design information

SPICE netlists are circuit netlists, representative of the actual circuits used by the IC vendors. In contrast, IBIS uses a behavioral representation, implemented as an IV data look-up table, rather than a direct transistor-by-transistor description.

IBIS is useful even without a suitable IBIS simulation tool. IBIS waveforms can be viewed using a freely available IBIS viewer. The IBIS database for an I/O circuit contains information on edge switching rates that is the key to understanding signal integrity. The IBIS viewer can graphically display these edge rates and the signal transition waveforms for a given load of 50 Ω . The importance of signal edge rates is explained in more detail in [“Signal Integrity,” page 11](#).

Network Analyzers

Network analyzers used to be a tool confined to RF labs and were considered as somewhat esoteric because of the relatively limited market. However, with deep sub-micron, 90 nm and below, they are extremely useful in characterizing a board's power supply and should be considered as a tool to provide a sanity check for power supply simulations and for debugging noise problems. The correct application of this tool is beyond the scope of this note, but the principle of operation and board layout requirements to support their use is introduced here. A

board's power supply is essentially a big, distributed filter circuit with the ideal output being perfectly flat up to the highest operational and critical frequency on the board. This frequency is the knee frequency, described in "The Importance of Power Planes," page 8.

As shown in Figure 11, a network analyzer is two instruments in one: a signal generator that sweeps the frequency spectrum and a spectrum analyzer that measures the filtered output. This measurement can be made with a single access point, typically an SMA, to the plane being measured, but many designers prefer to use two SMA test points: one to inject the sweep signal and the other to receive it. The best place for the input is near to the power supply plane entry on the board. The best place for the output SMA for the spectrum analysis is on the side of the FPGA furthest away from the power supply. The SMA connectors are only required during board characterization and can be eliminated on the production design to save costs.

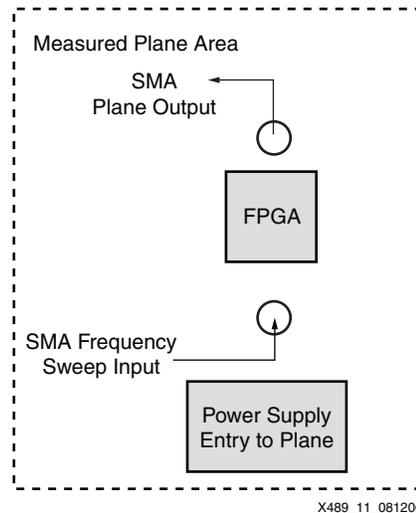


Figure 11: Network Analyzer

Board Design Strategies

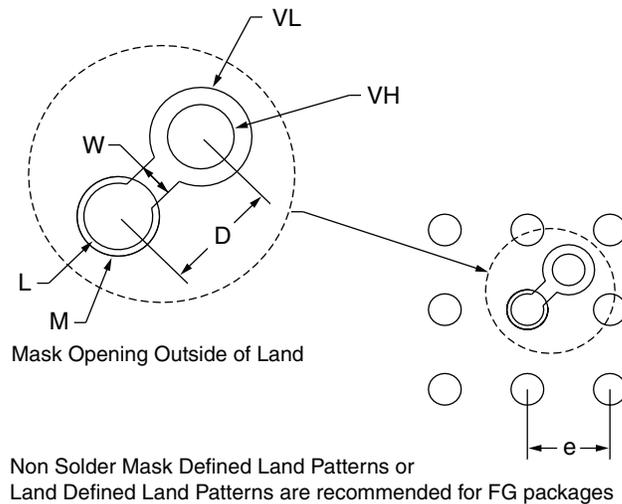
A simulation is only as good as the model's accuracy. Because such models have considerable unknowns or can miss key elements, the following strategies are recommended:

1. Simulation
2. Validation
3. Characterization
4. Cost reduction

After validating and characterizing the board design, the model is fixed as required. Be aware of functional-only test patterns, which do not simultaneously exercise the board's full power and speed requirements.

Board Routing Strategy

Xilinx provides the diameter of a land pad on the component side. This information is required prior to the start of the board layout to design the board pads to match the component side land geometry. The typical values of these land pads are shown in Figure 12 and summarized in Table 9.



X489_12_041806

Figure 12: Suggested Board Layout of Soldered Pads for 1.0 mm Pitch BGA

Table 9: Recommended PCB Design Rules

	FT256/FG256	
	Dimensions in mm	Dimensions in mils
Component Land Pad Diameter (SMD) ⁽¹⁾	0.40	15.7
Solder Land (L) Diameter	0.40	15.7
Opening in Solder Mask (M) Diameter	0.50	19.7
Solder (Ball) Land Pitch (e)	1.00	39.4
Line Width Between Via and Land (w)	0.13	5.1
Distance Between Via and Land (D)	0.70	27.6
Via Land (VL) Diameter	0.61	24.0
Through Hole (VH) Diameter	0.300	11.8
Pad Array	Full	
Matrix or External Row	16x16	
Periphery Rows	-	

Notes:

1. Component land pad diameter refers to the pad opening on the component side (solder mask defined).

Conclusion

This application note is an overview of how low cost but very high speed, four- and six-layer PCBs can be successfully designed using Spartan-3E devices in the FT256 1 mm BGA package. Also reviewed are areas for special care and attention during PCB layout to avoid problems with signal integrity.

Related Documents

The following links provide additional information useful to this application note:

1. Pads PCB format (Mentor Graphics) layout files with the XC3S500E FT256, as shown in this note, with the essential fanout and plane definitions available at www.xilinx.com/bvdocs/appnotes/xapp489.zip.
2. HyperLynx IBIS simulation tool:
http://www.mentor.com/products/pcb/analysis_verification/hyperlynx/hyperlynx_software_eval.cfm
3. Signal integrity definitions:
<http://www.sdsmt.edu/syseng/ee/courses/ee690/690review2.pdf>
4. Johnson, Howard and Martin Graham. 1993. *High-Speed Digital Design*. Prentice Hall PTR.
5. Power supply simulations:
[XAPP623](#), *Power Distribution System (PDS) Design: Using Bypass/Decoupling Capacitors*
6. SPICE modeling example of a power supply simulation netlist as described in XAPP623:
<http://www.xilinx.com/bvdocs/appnotes/xapp623.zip>
7. [UG112](#), *Device Package User Guide*
8. [DS312](#), *Spartan-3E FPGA Family Data Sheet*
9. Graphical and text pinout files:
http://www.xilinx.com/bvdocs/publications/s3e_pin.zip
10. Xilinx power tools for power analysis and estimation:
<http://www.xilinx.com/power>

Acknowledgments

Xilinx wishes to thank Gary Lawman of Tentmaker Systems (www.tentmakersystems.com) and the following PCB industry advisors: Greg Halvorson, President, Streamline Circuits (www.streamlinecircuits.com) and Don Philips of Merix (www.merix.com).

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
10/31/06	1.0	Initial Xilinx release.