Printed Circuit Board Considerations

Layout Considerations

The PC board is no longer just a means to hold ICs in place. At today's high clock rates and fast signal transitions, the PC board performs a vital function in feeding stable supply voltages to the IC and in maintaining signal integrity between devices.

VCC and Ground Planes

Since CMOS power consumption is dynamic, it is a non-trivial task to assure stable supply voltages at the device pins and to minimize ground differentials. A multi-layer PC board is a must, with four layers for the simplest circuits, 6 to 12 layers for typical boards. Ground and V_{CC} must each be distributed in complete layers with few holes. Slots in these layers would cause an unacceptable inductive voltage drop, when the supply current changes at a rate of 1 A/ns, or even faster. Besides an uninterrupted ground plane, Virtex-II devices require one plane for V_{CCINT} (1.5 V) plus one plane for V_{CCAUX} (3.3 V). V_{CCO} can be distributed on wide signal traces with sufficient bypass capacitors.

Beyond low resistance and inductance, ground and V_{CC} planes combined can also provide a small degree of V_{CC} decoupling. The capacitance between two planes is ~100 pF/inch^2 or ~15 pF/cm^2, assuming 10 mil (0.25 mm) spacing with FR4 epoxy.

V_{CC} Decoupling

Fast changing Icc transitions must be supplied by local decoupling capacitors, placed very closely to the V_{CC} device pins or balls. These capacitors must have sufficient capacitance to supply Icc for a few ns and must have low intrinsic resistance and inductance. X7R or NPO ceramic surface-mounted capacitors of 0.01 to 0.1 µF, one per V_{CC} device pin, are appropriate. 0.1 µF can supply 1A for 2ns with a 20 mV voltage droop.

1A • 2ns = 2 nanocoulomb = 100 nF • 0.02 V

Low impedance at >100 MHz is important, but capacitance variation with temperature is acceptable. These small capacitors are the first-line source for Icc, and they must be placed very close to the V_{CC} pins. A half-inch or 10 mm trace represents an inductance of several nanohenries, defeating the purpose of the decoupling capacitor. Backing up this local decoupling is one tantalum capacitor of 10 to 100 µF, able to supply multiple amperes for about 100 ns.

Finally, each board needs a power-supply decoupling electrolytic capacitor of 1000 to 10,000 µF able to supply even more current for a portion of the supply switching period. As described below, larger capacitors inevitably have higher series resistance and inductance, which is the reason for the above-mentioned hierarchy of supply decoupling. As a general rule, multiple capacitors in parallel always offer lower resistance and inductance than any single capacitor.

Decoupling Capacitors

The ideal decoupling capacitor would present a short circuit to ground for all ac signals. A real capacitor combines a given amount of capacitance with unavoidable parasitics, a small series resistance and inductance. At low frequencies, the composite impedance is capacitive, i.e., it decreases with increasing frequency. At high frequencies, it is inductive and increases with frequency, making the decoupling ineffective. In-between, there is the LC resonant frequency, where the capacitor looks like a small resistor.

Different technologies provide different trade-offs between desirable features like small size and high capacitance, and undesirable features like series resistance and inductance. Electrolytic and tantalum capacitors offer the largest capacitance in a given physical size, but also have the highest inductance. This makes them useful for decoupling low frequencies and storing large amounts of charge, but useless for high frequency decoupling. Surface-mount ceramic capacitors, on the other hand, offer the lowest

inductance and the best high-frequency performance, but offer only a small amount of capacitance, less than a microfarad.

Figure 4-41 shows the frequency-dependent impedance and resistance of a typical electrolytic capacitor of 1500 µF, while Figure 4-42 and Figure 4-43 show the equivalent data for ceramic bypass capacitors of 33,000 and 3,300 pF, respectively. Note that the resonant frequency for the small ceramic bypass capacitor at 100 MHz is 10,000 times higher than the resonance frequency of the large electrolytic capacitor at 10 KHz. For more technical information on decoupling capacitors, see the manufacturers' websites.

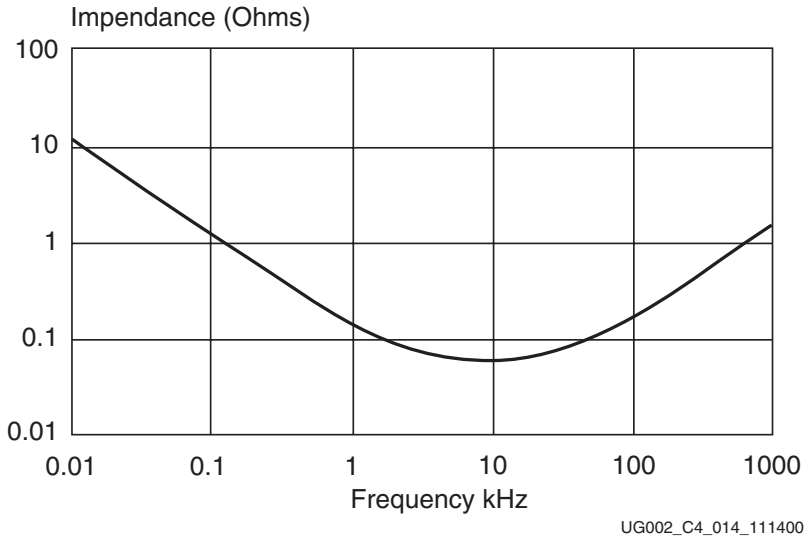

Figure 4-41: **1500 µF Electrolytic Capacitor Frequency Response Curve**

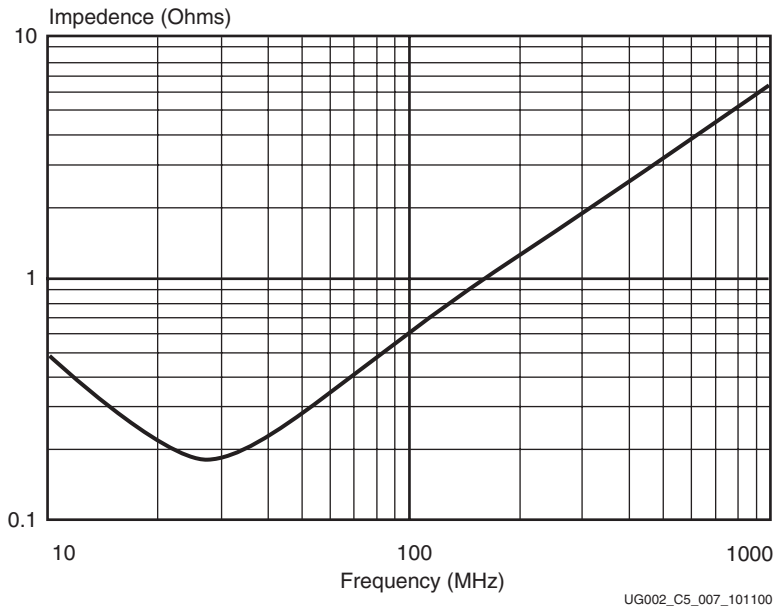

Figure 4-42: **33000 pF X7R Component Frequency Response Curve**

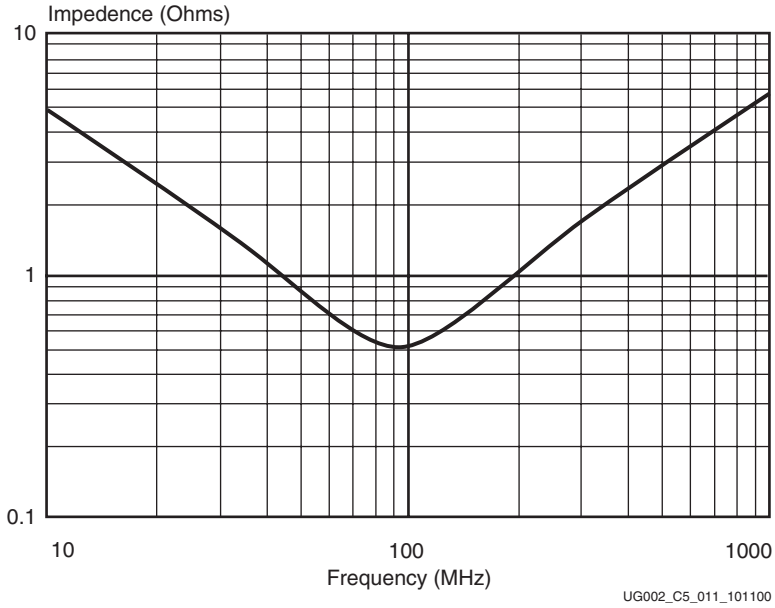

Figure 4-43: **3300 pF X7R Component Frequency Response Curve**

Transmission Line Reflections and Terminations

A PC board trace must be analyzed as a transmission line. Its series resistance and parallel conductance can generally be ignored, but series inductance and parallel capacitance per unit length are important parameters. Any signal transition (rising or falling edge) travels along the trace at a speed determined by the incremental inductance and capacitance.

For an outer-layer trace (air on one side) the propagation delay is 140 ps/inch, or 55 ps/cm. For an inner-layer trace (FR4 with ε=4.5 on both sides), the propagation delay is 180 ps/inch, or 70 ps/cm.

The voltage-to-current ratio at any point along the transmission line is called the characteristic impedance Z_0. It is determined by w/d, the ratio of trace width w to the distance d above the ground or V_{CC} plane.

For an outer layer trace (microstrip),

Z_0=50 Ω when w = 2d (e.g., w = 12 mil, d = 6 mil),

Z_0=75 Ω when w = d (e.g., both 6 mil = 0.15 mm).

For an inner layer trace between two ground or V_{CC} planes (stripline),

Z_0=50 Ω when w = 0.6•d (e.g., w = 5 mil, d = 8 mil),

Z_0=75 Ω when w = 0.25•d (impractical).

Most signal traces fall into the range of 40 to 80 Ω.

A slow transition treats a short narrow trace as a lumped capacitance of about 2 pF per inch (0.8 pF per cm). However, if the trace is so long, or the signal transition is so fast that the potential echo from the far end arrives after the end of the transition, then the trace must be analyzed as a transmission line.

In this case, the driver sees the trace not as a lumped capacitance, but rather as a pure resistance of Z_0. The signal transition then travels along the trace at the speed mentioned above. At any trace-impedance discontinuity all or part of the signal is reflected back to the origin. If the far end is resistively terminated with R=Z_0, then there is no reflection. If, however, the end is open, or loaded with only a CMOS input, then the transition doubles in amplitude, and this new wave travels back to the driver, where it may be reflected again, resulting in the familiar ringing. Such ringing has a serious impact on signal integrity, reduces noise margins, and can lead to malfunction, especially if an asynchronous signal or

a clock signal crosses the input threshold voltage unpredictably. Two alternate ways to avoid reflections and ensure signal integrity are parallel termination and series termination.

Parallel Termination

Reflections from the far end of the transmission line are avoided if the far end is loaded with a resistor equal to Z_0. A popular variation uses two resistors, one to V_{CC}, one to ground, as the Thevenin equivalent of Z_0. This reduces the load current for one signal level, while increasing it for the other. Parallel termination causes dc power consumption which can be eliminated by inserting a capacitor between the terminating resistor and ground. The value of this capacitor is determined as follows:

Signal transition time << RC << signal level duration

For example, 50 Ω • 120 pF for a 2 ns transition every 20 ns. See Figure 4-44.

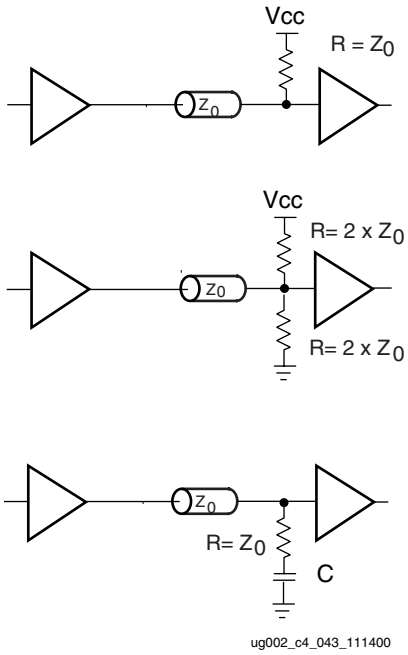

Figure 4-44: **Parallel Termination**

Series Termination

While parallel termination eliminates reflections, series termination relies on the reflection from the far end to achieve a full-amplitude signal. For series termination, the driver impedance is adjusted to equal Z_0, thus driving a half-amplitude signal onto the transmission line. At the unterminated far end, the reflection creates a full-amplitude signal, which then travels back to the driver where it gets absorbed, since the output impedance equals Z_0. See Figure 4-45.

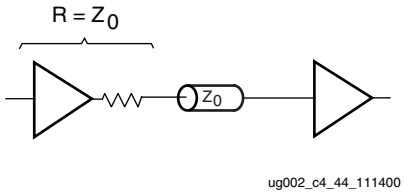

Figure 4-45: **Series Termination**

Series termination dissipates no dc power, but the half-amplitude round-trip delay signal means that there must be no additional loads along the line. Series termination is ideal (and only meaningful) for single-source-single-destination interconnects.

Virtex-II devices offer digitally controlled output impedance drivers and digitally-controlled input termination, thus eliminating the need for any external termination resistors. This feature is extremely valuable with high pin-count, high density packages.

These PC board considerations apply to all modern systems with fast current and voltage transitions, irrespective of the actual clock frequency. The designer of relatively slow systems is more likely caught off-guard by the inherent speed of modern CMOS ICs, where di/dt is measured in A/ns, dV/dt is measured in V/ns, and input flip-flops can react to 1 ns pulses, that are invisible on mid-range oscilloscopes. Powerful tools like HyperLynx can analyze signal integrity on the PC board and can often be amortized by one eliminated board-respin.

JTAG Configuration and Test Signals

Poor signal integrity and limitations of devices in a JTAG scan chain can reduce the maximum JTAG test clock (TCK) rate and reliability of JTAG-based configuration and test procedures. The JTAG TCK and test mode (TMS) signals must be buffered, distributed, and routed with the same care as any clock signal especially for long JTAG scan chains. The devices in a JTAG scan chain should be ordered such that the connections from the TDO of one device to the TDI of the next device are minimized. When high-speed JTAG-based configuration for the Virtex-II devices is required, devices with lower-specified maximum TCK rates can be placed in a separate JTAG scan chain.

Crosstalk

Crosstalk can happen when two signals are routed closely together. Current through one of the traces creates a magnetic field that induces current on the neighboring trace, or the voltage on the trace couples capacitively to its neighbor. Crosstalk can be accurately modeled with signal integrity software, but two easy to remember rules of thumb are:

- Crosstalk falls off with the square of increasing distance between the traces.
- Crosstalk also falls off with the square of decreasing distance to a ground plane.

$$\text{Peak Crosstalk Voltage} = \frac{DV}{1 + (D/H)^2}$$

where

DV is the voltage swing

D is the distance between traces (center to center)

H is the spacing above the ground plane

Example:

3.3V swing, and two stripline traces 50 mils apart and 50 mils above the ground plane.

Peak Crosstalk Voltage = $(3.3\ V)/(1 + (0.05/0.05)^2) = 1.65\ V$

This can cause a false transition on the neighboring trace. Separating the trace by an additional 50 mils is significantly better:

Peak Crosstalk Voltage = $(3.3\ V)/(1 + (0.1/0.05)^2) = 0.66\ V$

Signal Routing to and from Package Pins

Signal escaping (traces leaving the pin/ball area) can be quite difficult for the large FG and flip-chip packages. The number of signal layers required to escape all the pins depends on the PCB design rules. The thinner the traces, the more signals per layer can be routed, and the fewer layers are needed. The thinner traces have higher characteristic impedance, so choose an impedance plan that makes sense, and then be consistent. Traces from 40 to 80 ohms are common.

If only one signal can be escaped between two pads, only two rows of pins can be escaped per layer. For FG packages (1.0mm pitch) one signal of width 5 mils (0.13mm) can be

escaped between two pads, assuming a space constraint equal to the trace width. For a discussion of signal routing specific to Virtex-II devices, see www.xilinx.com for currently available application notes.

As packages are able to handle more I/Os with a minimum increase in size, the signal integrity of those signals must be considered, regardless of clock frequency. Especially with the largest packages, precise PCB layer stackup is required. Parameters such as board material, trace width, pad type, and stackup must be defined based on simulation, and the fabrication drawings must be marked with "precise layer stackup" and the stackup specified. A number of board-level signal integrity simulators exist, and careful attention to PCB design rules creates a robust design with low EMI and high signal reliability.

Board Routability Guidelines

Board-Level BGA Routing Challenges

Xilinx ball grid array (BGA) wire-bond and flip-chip packages contain a matrix of solder balls (see Figure 4-46). These packages are made of multilayer BT substrates. Signal balls are in a perimeter format. Power and ground pins are grouped together appropriately.

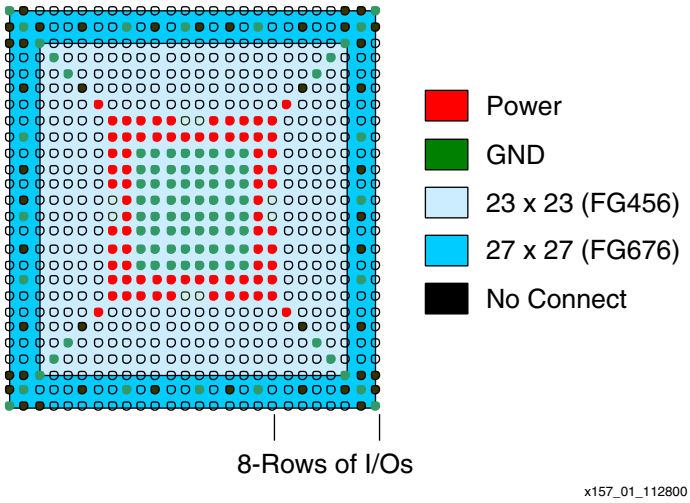

Figure 4-46: **Fine-Pitch BGA Pin Assignments**

The number of layers required for effective routing of these packages is dictated by the layout of pins in each package. If several other technologies and components are already present on the board, the system cost is factored with every added board layer. The intent of a board designer is to optimize the number of layers required to route these packages, considering both cost and performance. This section provides guidelines for minimizing required board layers for routing BGA products using standard PCB technologies (5 mils-wide lines and spaces or 6 mils-wide lines and spaces).

For high performance and other system needs, designers can use premium technologies with finer lines/spaces on the board. The pin assignment and pin grouping scheme in BGA packages enables efficient routing of the board with an optimum number of required board layers.