



Technical notes on using Analog Devices DSPs, processors and development tools
Visit our Web resources <http://www.analog.com/ee-notes> and <http://www.analog.com/processors> or
e-mail processor.support@analog.com or processor.tools.support@analog.com for technical support.

Power Bypass Decoupling of SHARC® Processors

Contributed by Larry Hurst & Alberto Comaschi

Rev 1 – December 5, 2006

Introduction

This EE-Note describes power-decoupling requirements for boards that contain SHARC® processors. It is intended to help system designers to calculate the total supply decoupling capacitance, choose the capacitor values, and place the capacitors for optimal results. The end of this document includes a summarization of power-decoupling guidelines and a board design check list.

EZ-KIT Lite® evaluation system schematics provide a good starting reference. The design database for SHARC processor EZ-KIT Lite evaluation systems, which is available online, contains all the electronic information required for board design, layout, fabrication, and assembly.

Recently, processor core and peripheral clock frequencies have risen to well over 600 MHz, and switching times have decreased to less than 800 ps. An increased emphasis in high-speed serial communications has resulted in power supply noise frequencies well over 1 GHz. Processor terminal voltages have decreased and currents have increased, representing significantly lower load impedances and making power distribution bypassing even more difficult.

Due to these developments, a more aggressive combination of strategies is required to effectively bypass the reduced power supply load impedances represented by current and future generations of processors. The techniques used must consider not only the values and numbers of bypass (decoupling) capacitors but also their physical location, mounting method, and their interaction with the PC board's power and ground planes. Taking these measures during the design phase will help you to avoid intermittent hardware problems that can be very difficult to diagnose or cure once an application has been released to production (or worse, after thousands of units have been installed in the field).

Why Noise is a Problem

The external power supply pins form a common node for thousands of the transistors that comprise the internal processor logic. These same transistors are used for gates, combinational logic, registers and counters, delay lines, state machines, and other logic functions. Although operation tends to be synchronized by a common core clock or peripheral clock, individual components do not switch at the same time due to differing internal delays. Some gates may still be switching while others have completed the associated state change. An internal gate in the "high" state will pass noise from the power supply through to the input of the following gate. If other inputs to the gate are still transitioning through their "undefined" input region, the noise on the supply pins may be amplified further by the gate and may appear as a noise glitch at the output.

A single positive or negative noise pulse on the power pins may propagate through a gate and toggle the contents of an internal register. Variations in operating temperature and manufacturing processes may cause the problem to be extremely difficult to debug. It is better to design defensively and to eliminate these problems at the start.

A further problem with power supply ‘noise’ is its negative effect on the jitter performance of oscillators and phase locked loops (PLLs) and delay locked loops (DLLs). Applications such as professional digital audio and high-speed communications depend upon low jitter clocks; minimizing power supply noise is essential to maintaining the fidelity of analog-to-digital conversion.

Most processors specify a nominal operating voltage tolerance, typically $\pm 5\%$ (see Figure 1). For example, for a recommended V_{DDEXT} of 3.30 V, the limits are $V_{DDEXT}(\text{MIN}) = 3.135 \text{ V}$ and $V_{DDEXT}(\text{MAX}) = 3.465 \text{ V}$, or simply stated as $3.3 \text{ V} \pm 165 \text{ mV}$. Similarly, the nominal core voltage (V_{DDINT}) is typically $1.2 \text{ V} \pm 60 \text{ mV}$, so that the voltage regulator must maintain the voltage between 1.140 V and 1.260 V. Because these limits include both the DC supply voltage error and the peak amplitude of any noise, you will need to account for deviation in the regulator output voltage by reducing the peak noise voltage. A typical voltage regulator precision is $\pm 2.5\%$; therefore, the peak noise voltage should not exceed 2.5% of the supply voltage.

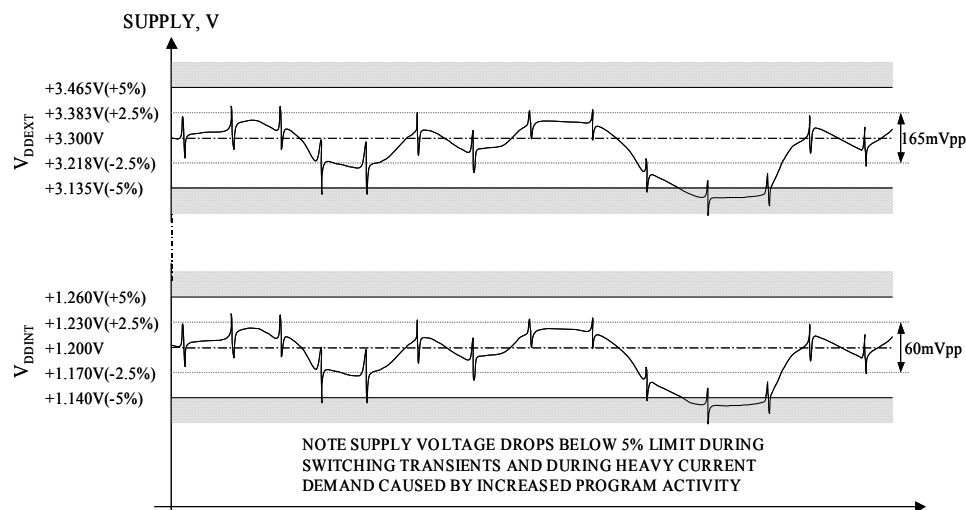


Figure 1. Supply noise vs. permitted tolerance

As described above, the objective of good power supply bypassing is to keep peak noise voltage amplitudes under 82.5 mV for 3.3V supplies (under 30 mV for 1.2V supplies) over the *full* frequency spectrum, from DC to over one Gigahertz.

Minimizing power supply noise is necessary for reliable operation of the processor, and it also minimizes electromagnetic interference (EMI) electromagnetic compatibility (EMC) compliance problems.

Calculating the Required Capacitance

Processor Target Impedance (V_{DDINT})

The load represented by a processor is fairly complex. It depends on the processor's operating mode (for example, idle), clock rates, and code execution. Another load component is the high current and high frequency spectrum transitions typical of CMOS processes that occur at the processor's clock edges. In order to maintain the voltage at the processor's pins within its operating range, the power supply system should have a target impedance proportional to the allowed ripple, load current, and voltage.

A typical case: ripple of 2.5%, $V_{DD} = 1.2\text{ V}$, and current = 500 mA:

$$\text{Target Impedance} = \frac{V_{DD}(\text{MIN}) * \text{Ripple}}{I_{DD}(\text{MAX})} = \frac{1.2 * 0.025}{0.500} = 60\text{m}\Omega \quad [1]$$

Because of the broad frequency spectrum of the load, a simple voltage regulator cannot be used to keep the target impedance constant throughout the load frequency spectrum. Typically, the voltage regulator is effective to about 100 KHz, and bulk capacitors are required to keep the target impedance in the frequency range of 10 KHz to 1 MHz. Ceramic capacitors are required for frequencies above 1 MHz and into the GHz range. Another effective high-frequency capacitor is created by the power planes. See Figure 2.




Figure 2. Flat target impedance vs. frequency, and decoupling elements frequency range

Calculating the Total Required Bulk Capacitance (V_{DDINT})

You need to provide enough capacitance (generally described as “bulk”) to maintain the processor terminal voltage during long duration step load changes due to voltage regulator delays. This may require a detailed knowledge of the regulation characteristics; generally, assume that an external linear or switch-mode regulator takes at least 100 μs to respond to a large change in load current. Sufficient storage capacitance must be provided to prevent the supply voltage from ‘sagging’ during this delay. Allow a drop of no more than 2.5% in V_{DD} (for example, 82.5 mV for a 3.3V supply, or 30 mV for a 1.2V supply).

To find the required, *total* supply bypass capacitance, treat the processor as a “noise generator” of a known source impedance. Then determine the capacitive reactance necessary to bypass this source

impedance, keeping the terminal noise voltage within the data sheet limits. This is sometimes referred to as “target impedance”. The difficulty with this approach is the ill-defined nature of the internal core load impedance, represented by the processor; a conservative approximation assumes that the minimum impedance can not be less than the minimum supply voltage divided by the maximum load current stated in the data sheet.

For a typical example for core supply: $V_{DD} = 1.2$ V and current = 500 mA:

$$|Z_{DSP}| \geq \frac{V_{DD}(MIN)}{I_{DD}(MAX)} = \frac{1.2}{500mA} = 2.4\Omega \quad [2]$$

If we know the processor’s supply tolerance (typically $\pm 5\%$) and the voltage regulator accuracy and frequency response (typically $\pm 2.5\%$ and 100 kHz, respectively), we can determine the minimum total bypass capacitance within the limits for the processor, as follows:

$$C_{BPTOT} = \frac{I_{DD}(MAX)}{2f \times ripple \times V_{DD}(MIN)} F \quad [3]$$

For example: processor core voltage (V_{DDINT}) = 1.2 V, I_{DDINT} = 500 mA. Hence, $Z_{DSP} \geq 2.4 \Omega$ and $Z_{TARGET} < 60$ m Ω .)

$$C_{BPTOT} \geq \frac{500(mA)}{2 \times 100(kHz) \times 0.025 \times 1.2(V)} \mu F$$

$$C_{BPTOT} \geq 83.3 \mu F$$

If the maximum change in I_{DD} is known, we can determine the minimum total bypass capacitance. For example, if the change in processor load current from $I_{DD}(MIN)$ to $I_{DD}(MAX)$ is determined to be 250 mA, the minimum capacitance can be found, as follows:

$$C_{BP(MIN)} = \Delta I_{DD} \times \frac{\Delta t}{\Delta V_{DD}} F \quad [4]$$

where:

$$\Delta I_{DD} = I_{DD}(MAX) - I_{DD}(MIN)$$

$$\Delta t = \text{Power Supply Response Time}$$

$$\Delta V_{DD} = \text{Max Permitted Voltage Sag}$$

Example:

$$C_{BP(MIN)} = 250(mA) \times \frac{10(\mu s)}{82.5(mV)} \mu F$$

$$\therefore C_{BP} \geq 30.3 \mu F$$

The minimum *total* bypass/decoupling capacitance should be the greater of [Equation 3](#) or [Equation 4](#).

Another consideration in determining an effective bypassing strategy is the effect of the parasitic inductance and resistance associated with component packages and mounting methods. Fundamental core clock speeds are now well over 500 MHz with rise and fall times of switching transients under 800 ps.

At these frequencies (100 MHz and above), the parasitic impedance of processor packages and bypass capacitors are very significant and must be taken into account.

For example, the inductance of a 10-mm-long, 7-mil-wide copper trace is about 3 nH; this has a reactance greater than $7\ \Omega$ at 400 MHz. The total mounted series self-inductance of a single 100-nF 0603 SMD bypass capacitor is about 1.2 nH, resulting in a self-resonant frequency less than 15 MHz. Above 15 MHz, the capacitor reactance is inductive, and at 400 MHz, the total impedance is typically greater than $3\ \Omega$ and rising. A target bypass impedance of less than $60\ m\Omega$ will not provide effective bypassing.

A partial solution to this problem is to use a larger number of low value capacitors, ranging in value between 470 pF and 100 nF. To maintain as low an impedance as possible (at as high a frequency as possible), the number of capacitors of each value is varied inversely with their value, but the values should not be separated by more than a factor of ten to minimize anti-resonance or parallel resonance effects. Even with a large number of capacitors in parallel, the impedance may still be too high at frequencies above 500 MHz and the distributed capacitance of the power/ground planes must be relied upon to complete the bypassing.

Decoupling the Analog Supply ($A_{VDD} / 1.2VDC$)

The analog supply pin ($A_{VDD} / 1.2\ VDC$) powers the processor's internal clock generator PLL. The load represented by this circuitry is fairly stable; therefore, no calculations are needed for the bypass capacitance. However, in order to produce a stable clock, PCB designs need an external filter circuit for the A_{VDD} pin. Place the filter components as close as possible to the A_{VDD} / A_{VSS} pins. For an example circuit, see Figure 3. (A recommended ferrite chip is the muRata BLM18AG102SN1D). To reduce noise coupling, the PCB should use a parallel pair of power and ground planes for V_{DDINT} and GND . Use wide traces to connect the bypass capacitors to the analog power (A_{VDD}) and ground (A_{VSS}) pins. Note that the A_{VDD} and A_{VSS} pins specified in Figure 3 are inputs to the processor and not the analog ground plane on the board; connect the A_{VSS} pin directly to digital ground (GND) at the chip.

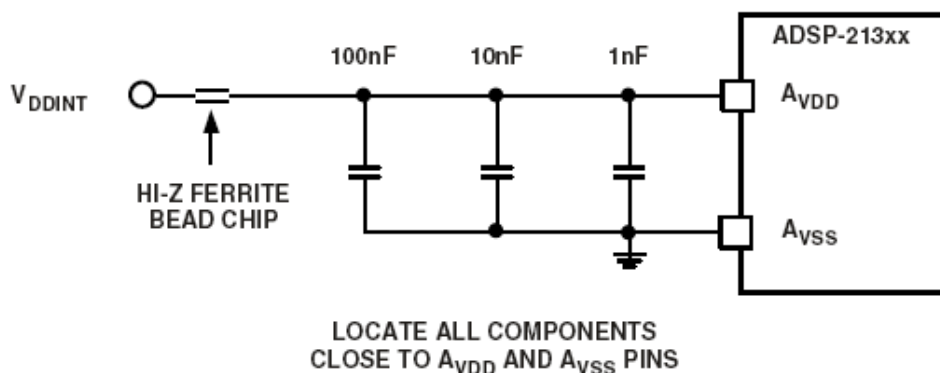


Figure 3. Analog power (A_{VDD}) filter circuit

Bypassing V_{DDEXT} – Allowing for Output Load Capacitance

Consider a worst-case scenario where a 16-bit external data bus is switching all active output lines from zeros to ones (i.e., 0x0000 to 0xFFFF) simultaneously. Assume that all data lines have 30 pF of total load capacitance and that the output voltage on all lines must rise from 0.4 V to $\geq 2.8\ V$ in less than 800 ps

(800 ps is a typical rise/fall time for processor data lines). Also assume that the external power supply (V_{DDEXT}) is 3.3 V and has a maximum permitted droop of 2.5% (i.e., 82.5 mV).

$$\Delta Q_{CH} = n \times C_L \times \Delta V_O = \Delta Q_{DISCH} = C_{BYP} \times \delta V_{DDEXT} = \int_0^{t_{CHRG}} i_{DDEXT}(t) dt \quad [5]$$

$$\therefore C_{BYP} = \frac{n \times C_L \times \Delta V_O}{\delta V_{DDEXT}} = \frac{16 \times 30 \times 10^{-12} \times 2.4(V)}{0.025 \times 3.3(V)} = 13.96 nF$$

Although 13.96 nF does not seem to be a problem, we can be generous and allocate, say, four 1.0 nF capacitors on each side of the processor package, yielding a total of 16 nF. The problem would be solved with a small margin of safety. So, what is the problem? We already expect to provide several hundred microfarads of capacitance to allow for the slow voltage supply regulator response time, so it can't be that 13.96 nF is not enough.

Note that there is no term in [Equation 5](#) for the actual charge/discharge *time*. The equation is merely a ratio of charges based upon total capacitance and voltage change, assuming that the transfer of charge will be complete when required. The equation does not account for the series inductance for component lead connections or the resistance of the driver. This is a limitation of this method of calculation. Although the calculated capacitance seems reasonable and practical for this particular example, only a limited amount of time is available to transfer the required charge. The output drivers see the total distributed load capacitance during the rise time of the output signal. In order to preserve the fast rise times of output signals, considerable thought must be given to **minimizing the series inductance of all components**. Power supply pins on the processor must be connected *directly* to their respective power/ground plane pins with vias located within the mounting pad itself or directly adjacent to the pad. As mentioned earlier, avoid copper traces, which add too much inductance.

Medium value bypass capacitors ranging from 10 nF to 1 μ F can provide the energy to maintain the supply voltage for frequencies of 100 kHz to about 100 MHz. Above 100 MHz, the ESL of capacitors 10 nF and larger limits their effectiveness. A combination of low ESL and low ESR surface-mount capacitors in the range 470 pF to 1 nF, mounted around the periphery of the processor or on the underside of the PC board are the most effective.

Inherently, the PCB power/ground plane pair has very low self-inductance and significant distributed capacitance and can supply the initial sub-nanosecond rise of current during the rapid turn-on of the output signals. It extends the bypassing regime from about 500 MHz to over 1 GHz. More will be said about power/ground planes later.

High Frequency V_{DDINT} Bypassing – Noise Generator and Equivalent Circuit

Consider a processor and its associated core (internal) power supply (see [Figure 4](#)) before adding bypass capacitors. The nominal operating voltage is $V_{DDINT}(NOM)$, and the current is $I_{DDINT}(t)$. It is expected that the internal clock frequencies and associated harmonics will exceed 500 MHz. At these frequencies, the power supply will maintain an average voltage and current, but the impedance of the power distribution network and the slow speed of the voltage regulators will allow considerable transient noise voltages to develop across the processor power pins. These voltages will be proportional to the instantaneous impedance of the processor, $Z_{DSP}(t)$.

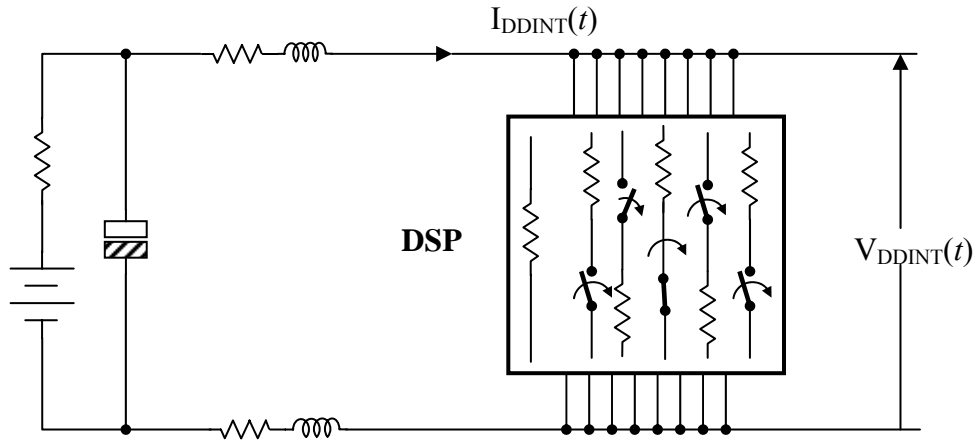


Figure 4. Switched resistor approximation of processor noise model

The operation of the processor can be compared to a very large number of switched resistors distributed throughout the IC, with the switching frequency and duty cycle of the switches constantly changing. This generates a load current that resembles a “pulse-width” and “amplitude modulated” waveform resulting in a broadband noise voltage, $V_{DDINT}(t)$, across the processor terminals. The ‘impedance’ of the processor varies from $Z_{DSP}(MIN) = V_{DDINT}(MIN)/I_{DDINT}(MAX)$ up to $Z_{DSP}(MAX) = V_{DDINT}(MAX)/I_{DDINT}(MIN)$. Making assumptions about the statistical distribution of the impedance with time, $Z_{DSP}(AVG) = V_{DDINT}(MEAN)/I_{DDINT}(MEAN)$. Since it is difficult to determine the absolute instantaneous limits of $V_{DDINT}(t)$ and $I_{DDINT}(t)$, assume a conservative approach and let $Z_{DSP}(MIN) \ll Z_{DSP}(AVG) \ll Z_{DSP}(MAX)$.

At very high frequencies, the power supply network behaves like a constant current supply; the transient noise voltages developed across the processor terminals will have a statistical amplitude distribution similar to $Z_{DSP}(t)$. Taking the minimum processor impedance, $Z_{DSP}(MIN) = V_{DDINT}(MIN)/I_{DDINT}(MAX)$ as a reference, a parallel shunt capacitor will reduce the noise voltage in proportion to the ratio of capacitive reactance to $Z_{DSP}(MIN)$, hence:

$$\left(\frac{X_C}{Z_{DSP}(MIN)} \right) \leq \left(\frac{\delta V_{DD}}{V_{DDINT}} \right) \quad [6]$$

$$\frac{1}{2\pi f_{MIN} C_{BP} \times Z_{DSP}(MIN)} \leq \left(\frac{\delta V_{DD}}{V_{DDINT}} \right)$$

$$\frac{I_{DDINT}(MAX)}{2\pi f_{MIN} C_{BP} \times Z_{DSP}(MIN)} \leq \left(\frac{\delta V_{DD}}{V_{DDINT}} \right)$$

$$C_{BP} \geq \frac{I_{DDINT}(MAX) \times V_{DDINT}}{(2\pi f_{MIN} \times V_{DDINT} \times \delta V_{DD})}$$

$$C_{BP} \geq \frac{I_{DDINT}(MAX)}{(2\pi f_{MIN} \times \delta V_{DD})} F$$

As an example, assume an ADSP-21365 SHARC processor with typical operating conditions so that $V_{DDINT} = 1.2$ V and $I_{DD}(MAX) = 1.1$ A, at a minimum core clock frequency of 100 MHz. This device is

available in a 136-ball BGA package with 13 balls allocated to V_{DDINT} . Let the required ripple factor (RF) be 2.5%; therefore, $\delta V_{DD} = 30$ mV.

$$C_{BP} = \frac{1.1}{\left(2 \times \pi \times 100 \times 10^6 \times 30 \times 10^{-3}\right)}$$

$$\therefore C_{BP} = 58.4 \text{ nF}$$

A practical choice would be to choose 16 capacitors, allocating four to each side of the processor package. Note that we are not attempting to allocate a certain number of capacitors to each ‘ball’ or ‘pin’. Instead, we are seeking to distribute capacitance around the periphery of the IC package uniformly.

$$C_{BP}' = \frac{C_{BP}(\text{nF})}{n} \text{ nF}$$

$$C_{BP}' = \frac{58.4}{16}$$

$$\therefore C_{BP}' = 3.65 \text{ nF}$$

Hence, use sixteen 3.9 nF capacitors.

Unfortunately, it’s not this easy! If we plot the true impedance of a 0603 SMD 3.9 nF capacitor over the range of frequencies of interest, taking into account the equivalent series resistance (ESR) and inductance (ESL), the true impedance above 300 MHz, is typically greater than 1200 m Ω and worsens as the frequency increases. Sixteen 3.9 nF capacitors gives an impedance of approximately 75 m Ω , which is considerably more than the required 27 m Ω ($k \times V_{DD(INT)}/I_{DD(MAX)}$), but is not low enough to guarantee the supply noise to be less than 2.5%. Using twenty larger value capacitors (e.g., 10 nF or 100 nF) is worse because of their higher ESL.

Fortunately, several factors come to our aid. First, instantaneous switching of all outputs does not immediately connect the power supply pins to the output pins due to current limiting in the drivers and internal self-inductance of the power supply and output signal bond-out and pins. Second, as frequency increases, the contribution of high-frequency noise decreases as a fraction of the total noise on the power supply. A third consideration is that the internal logic becomes progressively less responsive to noise as the frequency increases above the transition frequency of the active devices. A final contributor (that takes over) is the extremely low characteristic impedance of the power/ground plane pair. For placement of these capacitors refer to [PCB Layout and Capacitor Mounting](#) and [Capacitor Effective Bypass Radius](#) later in this document.

How Much Bypassing is Too Little?

The discussion above presents the processor as a broadband noise generator with defined source impedance and suggests a simple method of determining the number and size of bypass capacitors to provide a low-impedance bypass impedance that will reduce the noise and ripple across the terminals. The important consideration now is to determine a network of capacitors that will maintain this impedance across all frequencies of interest.

In the example above, the required total bypass capacitor impedance was calculated to be ≤ 27 m Ω . Not only is this a very low impedance, but it is also difficult to maintain over a wide band of frequencies. With processor core clock frequencies in excess of 300 MHz, the bypass network should attempt to

maintain this impedance from well below 100 kHz up to at least 500 MHz. (Typically, below 100 kHz, it is expected that the voltage regulator will be effective.)

Capacitive reactance (X_C) is calculated from the equation, $1/\omega C$; hence, at 300 MHz, for a reactance of 27.3 mΩ, it would appear that single capacitor of about 20 nF is all that is needed. Unfortunately, the inherent resistance and inductive reactance of capacitors require the use of multiple capacitors in parallel to achieve this low impedance.

The impedance of a single capacitor, where ESR is the equivalent series resistance, is given by the following equation.

$$|Z_C(f)| = \sqrt{(ESR)^2 + \left(\omega(ESL) - \frac{1}{\omega C}\right)^2} \quad [7]$$

Figure 5 shows a plot of the net impedance $|Z_{NET}(f)|$ of multiple surface-mounted capacitors listed in Table 1, assuming typical values of ESR and ESL, and including PCB via inductance. While not quite meeting our specification for 27.3 mΩ, the mean value is about 40 mΩ, up to 300 MHz, and the additional benefit of a power/ground plane pair (not shown) would further reduce the impedance over this frequency. Note in Table 1 and Figure 5 that none of the individual groups of capacitors would have an impedance close to 27.3 mΩ, especially above 300 MHz; hence, the necessity for multiple capacitors in parallel. Also note that above 100 MHz, 680 pF is more effective than 2.7 nF or 100 nF.

C ₀	20 × 680 pF, NP0 ceramic, assume ESL ≅ 1.2 nH, ESR ≅ 200 mΩ
C ₁	10 × 2.7 nF, NP0 ceramic, assume ESL ≅ 1.2 nH, ESR ≅ 60 mΩ
C ₂	12 × 10 nF, X7R ceramic, assume ESL ≅ 2 nH, ESR ≅ 150 mΩ
C ₃	8 × 100 nF, X7R ceramic, assume ESL ≅ 2.5 nH, ESR ≅ 50 mΩ
C ₄	4 × 2.2 μF, tantalum, assume ESL ≅ 10 nH, ESR ≅ 150 mΩ

Table 1. Net impedance of five groups of capacitors

Figure 5 shows the series resonant frequencies of capacitors C_{3..0} as indicated by the impedance dips at 10 MHz, 35 MHz, 89 MHz, and 170 MHz, respectively. The series resonant frequency of the 2.2 μF tantalum electrolytic capacitor (C₄), which is less than 1 MHz, is not shown. Above 200 MHz, the impedance continues to rise and depends upon the low inductance and distributed capacitance of the power and ground plane pair to complete the broadband low-impedance bypassing. The peaks in the impedance plot, which are due to parallel resonance between adjacent capacitor values, is the main reason why capacitor values should differ by more than an order of magnitude.




Figure 5. Example: capacitor net impedance magnitude vs. frequency

Although it would seem that it is only necessary to use a sufficient number of capacitors, of each value, there are still a few other considerations to meet the 27.3 mΩ requirement.

One consideration is the parallel or anti-resonant impedance between capacitors of different values. In this case, the impedance can be higher at the parallel resonant frequency than the impedance of each individual capacitor and result in high-frequency ringing on the power supply. To eliminate this possibility, capacitor values should *not* be separated by more than a factor of ten (e.g., 100 pF, ~1 nF, ~10 nF, ~100 nF). Also, several electrolytic capacitors should be distributed around the PC board. The effective “parallel ESR” of several smaller electrolytic capacitors and the load impedance of the processor will assist in damping oscillatory behavior due to parallel resonance.

Other considerations including cost and available PCB area are also important drawbacks that limit the number of capacitors that can be added. Adding too many capacitors may make efficient routing of signal traces impossible.

Bypass Capacitor Notes

The choice for effective bypassing between 10 MHz and 500 MHz are multi-layer ceramic chip (MLCC) SMD capacitors with X7R or NP0 (COG) dielectric. MLCC capacitors are readily available in a range of standard packages from several manufacturers. Other dielectric types are available, but they offer little advantage or may have excessive ESR or high cost.

Table 2 lists the size and ESL of the SMD capacitors most likely to be used for high-frequency bypassing. The ESL is for the capacitor only; vias typically add about 1 nH. Mount the capacitors as close to the IC

package as practical, although distances of up to 10 to 15 mm are not so critical as long as full dedicated power and ground planes are used.

EIA Size Code	Length	Width	ESL
	mm	mm	nH (typ)
0402	1.00±0.05	0.50±0.05	< 0.10
0508	1.25±0.15	2.00±0.15	< 0.60
0603	1.60±0.10	0.80±0.10	0.87
0612	1.60±0.20	3.20±0.20	0.61
0805	2.00±0.10	1.25±0.10	1.05
1206	3.20±0.15	1.60±0.15	1.25
1210	3.20±0.30	2.50±0.20	0.98

Table 2. SMD capacitors for high-frequency bypassing

Of particular interest are the 0508, 0612, and 1210 SMD packages (shown in bold type), which have a geometry designed to minimize the inherent ESL and ESR of the capacitor. Sizes smaller than 0402 are available, but not all PCB assembly houses can handle them easily.

While choosing a decoupling capacitor, it is also important to consider the loss of capacitance due to output voltage dc bias. Figure 6 shows the loss of capacitance due to output voltage dc bias for selected XR5 MLCC capacitors.




Figure 6. Capacitance drop vs. bias for ceramic capacitors (source Murata Corp.)

NP0 (or COG) dielectric, which has an ESR of about one-fourth that of an equivalent X7R capacitor, results in packages that are physically larger for a given capacitance. Consider the following:

- Manufacturers often quote 'Q' (or 'DF') ($\tan\delta$) at 1 kHz or 1 MHz, but the data is meaningless at frequencies over 100 MHz. NP0 dielectric has the lowest losses (hence, lowest ESR), but may actually promote very high frequency ringing on the power supply plane due to its high Q-factor.
- Generally, higher voltage and larger packages have lower ESR at resonance.
- Generally, larger capacitors have lower ESR, but also have a lower self-resonant frequency due to higher ESL.
- Generally, smaller packages and shorter, wider, packages have lower ESL, but sizes smaller than 0603 become increasingly difficult to handle during assembly.
- The 1210 package has a lower ESL and ESR than a 1206, 0805, or 1812 package because of its square stubby shape.
- Note that capacitors have particularly low impedance at their series resonant frequency (\cong ESR). Take advantage of this fact by selecting capacitors that are 'tuned' to particularly troublesome EMI noise or clock frequencies

PCB Layout and Capacitor Mounting

Using surface-mount capacitors that have package configurations with the lowest possible ESL is only part of the solution to effective bypassing, it is also necessary to ensure that the method of mounting the capacitor does not add additional self-inductance.

To minimize the total self-inductance of the capacitor, care must be taken in the location of the connecting vias. PCB traces have too much self-inductance and should be avoided. Vias can be located at the opposite ends of the mounting pads, within the mounting pads, or on the sides of the mounting pads (see Figure 7). The goal is to minimize the self-inductance of the mounting by reducing the loop area enclosed by the two vias. Placing the vias at the ends of the package is least desirable. An alternative is to place the vias on one side of the pads or, if room permits, use two vias on opposite sides of each pad.




Figure 7. Optimal placement of vias in PCB pads

Placing the vias within the mounting pad offers the lowest mounted ESL, but may contribute to assembly problems, especially with hollow vias, due to entrapment of flux and air within the via. This can cause the capacitor to 'stand up' during wave soldering operations. One solution to this problem has been the introduction of solid (i.e., filled) vias; this is recommended as the ultimate in low self-inductance connections between bypass capacitors and the power and ground planes.




Figure 8. Comparing bypass capacitor mounting locations

Compare mounting a capacitor at location ‘A’ (see Figure 8) versus mounting it beneath the package at position ‘B’. In location ‘A’, the total inductance seen at point ‘X’, where the processor connects to the power and ground planes, consists of the horizontal distance between point ‘X’ and the position of the vertical vias, plus the inductance of the vias and terminating pads. In position ‘B’, the inductance is due almost entirely to the height of the vias between point ‘X’ and the terminating pads on the bottom side of the board.

In both cases, the inductance of the vias connecting the *processor* to the power/ground planes is the same and can be ignored. Moving the planes closer to the surface on the processor side is advantageous. The capacitor is a standard 1 nF X7R ceramic 0603 SMD with an assumed total mounted inductance (L_{MOUNT}) of 1.0 nH.

Analysis of Total Inductance at Point ‘X’ for Capacitor at Position ‘A’

The total inductance (L_{TOT}) is the sum of L_{MOUNT} and the inductance of the vertical vias (L_{VVA}) and the inductance of the horizontal planes (L_{HP}). Thus, from dimensions in Figure 8, the inductance of a pair of closely spaced vias is given by:

$$\begin{aligned}
 L_{VVA} &= 5.08 \times h_A^2 \times \left(\frac{2}{d} - \frac{1}{C_L} \right) \text{ nH} \\
 &= 5.08 \times 0.012^2 \times \left(\frac{2}{0.012} - \frac{1}{0.063} \right) \\
 &= 0.11 \text{ nH}
 \end{aligned}$$

The inductance contribution of the horizontal distance enclosed by the planes is given by:

$$\begin{aligned}
 L_{HP} &= 10.16 \times s \times \ln \left(\frac{2L}{d} \right) \text{ nH} \\
 &= 10.16 \times 0.004 \times \ln \left(\frac{2 \times 0.394}{0.012} \right) \\
 &= 0.17 \text{ nH}
 \end{aligned}$$

Hence, total inductance for capacitor ‘A’ at point ‘X’ is:

$$\begin{aligned} L_{ATOT} &= L_{MOUNT} + L_{VVA} + L_{HP} \quad nH \\ &= 1.0 + 0.11 + 0.17 \\ &= 1.28 \quad nH \end{aligned}$$

Analysis of Total Inductance at Point ‘X’ for Capacitor at Position ‘B’

The total inductance (L_{BTOT}) is the sum of L_{MOUNT} and the inductance of the vertical vias (L_{VVB}):

$$\begin{aligned} L_{VVB} &= 5.08 \times h_B^2 \times \left(\frac{2}{d} - \frac{1}{C_L} \right) \quad nH \\ &= 5.08 \times 0.046^2 \times \left(\frac{2}{0.012} - \frac{1}{0.063} \right) \\ &= 1.62 \quad nH \end{aligned}$$

Hence, total inductance for capacitor ‘B’ at point ‘X’ is:

$$\begin{aligned} L_{BTOT} &= L_{MOUNT} + L_{VVB} \quad nH \\ &= 1.0 + 1.62 \\ &= 2.62 \quad nH \end{aligned}$$

Note that the inductance of the vias is influenced critically by the square of their length. Although the top side capacitor at position ‘A’ is mounted farther away, it has less than half the series inductance of the capacitor mounted on the bottom side of the board at position ‘B’. This means capacitors can be mounted safely at a reasonable distance around the periphery of the processor package and need not to be mounted directly underneath. This calculation is only an approximation because the effects of EM field fringing at discontinuities has been ignored, but it does indicate the advantage of using closely spaced power and ground planes and the effect of via inductance. Accurate results can be obtained by using 2-D or 3-D EM field modeling software.

Figure 9 shows a suggested layout with sixteen 0603 SMD capacitors, for each of the V_{DDINT} and V_{DDEXT} power planes, distributed around a 12-mm x 12-mm, 136-ball BGA package. The capacitors are spaced approximately 1 mm from the edge of the IC package and allow sufficient room for two 5-mil signal traces to pass between them. Note that using 0402 sized capacitors with more aggressive PCB layout routing enable all signal balls to be brought-out to the top layer.

Multiple source termination resistors can be placed inside the bypass capacitor perimeter without significant effect on the bypassing effectiveness, provided that power and ground planes are used.




Figure 9. Example bypass capacitor layout for SMD 12x12 BGA package (not to scale)

For leaded packages such as the LQFP, locate power and ground vias at the “heel” end of the PCB lead pads rather than the “toe” end, as this will help to reduce the individual lead inductance by approximately 500 pH. This can also reduce the amplitude of ground bounce and supply dip problems.

Figure 10 shows a suggested PCB layout for an LQFP package. Note the compact layout around the crystal/oscillator components and the location of the PLL A_{VDD} supply bypassing components.

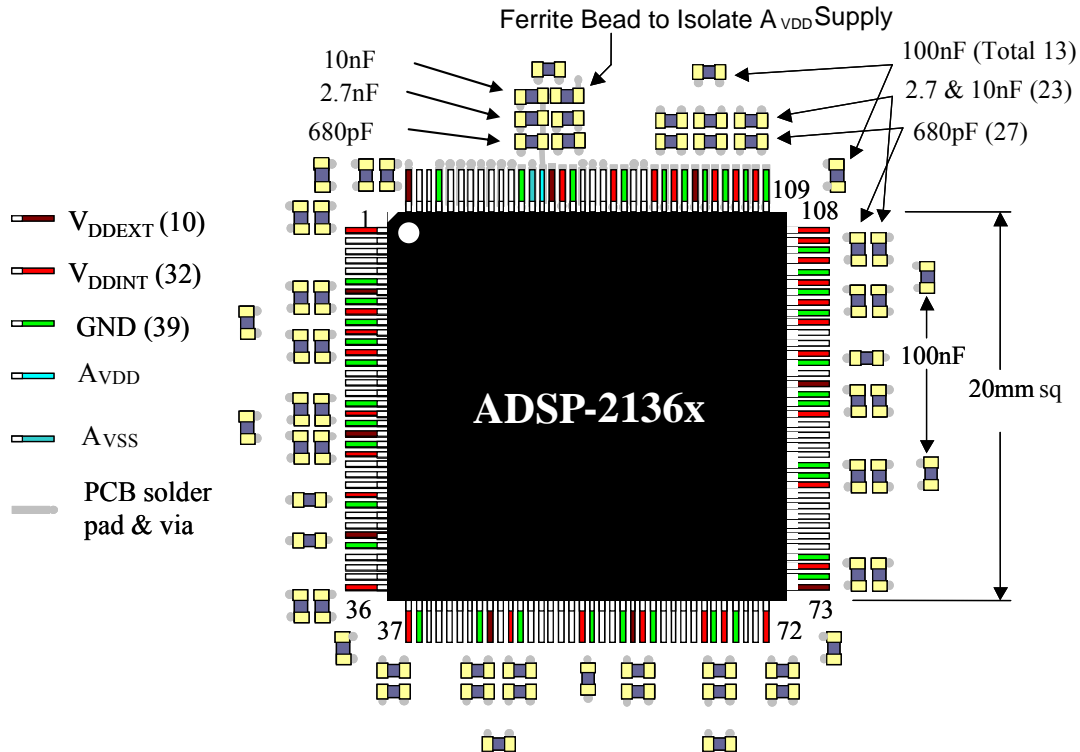


Figure 10. Example bypass capacitor layout for LQFP package (not to scale)

Capacitor Effective Bypass Radius

When a bypass capacitor is connected by short vias directly between the power and ground planes, it is useful to determine the maximum effective bypass radius from the processor power pins. One approach is to consider the propagation velocity (v_P) of a wavefront traveling between the power and ground planes and the corresponding wavelength at the series resonant frequency of the capacitor. To be effective up to the series resonant frequency, mount the capacitor within 2% of a wavelength. The propagation velocity is given by the following equation.

$$v_P = \frac{c}{\sqrt{\epsilon_r}} \text{ m/s}$$

For a typical FR-4 PC board, ϵ_r is ~ 4.2 , so v_P is about 146×10^6 m/s. The wavelength within the power/ground planes is given by:

$$\lambda = \frac{v_P}{f_{SR}} = \frac{146 \times 10^3}{f_{SR} \text{ (MHz)}} \text{ mm}$$

$$\therefore r \approx 0.02 \times \lambda = \frac{2.92 \times 10^3}{f_{SR} \text{ (MHz)}} \text{ mm}$$

For a typical 0603 SMD 470-pF capacitor, the mounted resonant frequency (f_{SR}) is about 300 MHz; hence, $\lambda = 487$ mm, and the capacitor should be located within ~ 9.7 mm of the processor power pins to be effective. Equal value or larger capacitors are no more effective when mounted closer than this distance.

For example, a 10-nF 0603 SMD capacitor has an f_{SR} of about 30 MHz; hence, it is recommended that the maximum mounting distance be about 97 mm (~ 3.8 inches) from the power pins.

This means that *not all capacitors* have to be mounted alongside the processor power pins and provides you latitude in placement to allow for high track densities near the processor. This reasoning applies only when the capacitor is connected directly to the power and ground planes by short vias.

What Effect a Plane or Two?

As an approximate guide, PCB traces and vias add inductance at a rate of about 1 nH/mm. The loop inductance between two points on a parallel pair of copper planes, due only to the planes themselves, is directly proportional to the distance between the planes and is as low as it is possible to make it.

The power and ground pins of the processor package and the bypass capacitors should be connected to their respective planes by the most direct vias possible, in order to take advantage of the low loop inductance of the planes. Generally, this means locating the via within the component's mounting pad. For BGA packages, this is not recommended usually because of assembly difficulties, but the via should still be located within half the diagonal pitch of the ball array.

Though the main advantage of a power/ground plane pair is the reduction in inductance, the two parallel planes form a very effective high-frequency capacitor with very low ESL. The capacitance can be estimated from the following equation.

$$C_P = \frac{\epsilon_0 \epsilon_r A}{d} \text{ F}$$

where:

$$\epsilon_0 = 8.854 \text{ pF/m} \quad (\text{permittivity of free space})$$

$$\epsilon_r = 4.2 \quad (\text{relative permittivity for FR-4 dielectric})$$

$$A = \text{area, m}^2$$

$$d = \text{separation, m}$$

For a typical multi-layer FR-4 board with a separation distance of 4 mils between the inner power and ground planes, the distributed capacitance is ~360 nF/m² (or 240 pF/in²), as shown below.

$$C_P = \frac{8.854 \times 4.2}{0.004 \times 0.0254} \text{ pF/m}^2$$

$$\therefore C_P = \frac{1464.1}{4(\text{mils})} \cong 366 \text{ nF/m}^2$$

For a small 3" x 4" PC board, this represents about 2.9 nF of extremely effective, low-inductance bypass capacity that is effective up to at least 1 GHz.

At the same time, the inductance of a pair of parallel planes is already extremely low, and is reduced even further as spacing is reduced. In the absence of high-permeability material, inductance is given by the following equation:

$$L_P = \mu_0 \times d \text{ H/sq}$$

where:

$$\mu_0 = 4\pi 10^{-7} \text{ H/m} \text{ (permeability of free space)}$$

$$d = \text{separation, m}$$

Using the same PC board with 4-mil spacing, the inductance of the planes is:

$$L_P = 10.16\pi \times 4(\text{mils}) \text{ pH/sq}$$

$$L_P = 127.7 \text{ pH/sq}$$

Note that the contribution of the plane inductance is much less than for discrete copper traces or vias connecting surface mounted capacitors, which is typically > 1 nH.

At frequencies above approximately 500 MHz, the power and ground planes become more effective than multiple discrete capacitors. Some PC board fabricators offer power/ground spacing as close as 2 mils with FR-4 material and special high-permittivity dielectrics such as 3M's C-Ply that allows power/ground spacing as close as 8 μm , to increase the capacitance and reduce inductance even further. The effectiveness of a PC board's power/ground planes at distributing capacitance and inductance should not be compromised by adding 'slots' or removing areas of the copper unless these modifications have been shown to be effective by prototype measurements or by using a 3-D EM field solver.

Another feature of power/ground plane pairs is their inherent low-characteristic impedance above their cutoff frequency. For a typical pair of planes, the impedance is given by:

$$Z_0 = \frac{120\pi \times d(m)}{\sqrt{\epsilon_R}} = \sqrt{\frac{L_P}{C_P}} \Omega$$

For the example above,

$$Z_0 = \sqrt{\frac{127.7 \times 10^{-12}}{366.01 \times 10^{-9}}} \Omega$$

$$Z_0 = 18.7 \text{ m}\Omega$$

Above the cut-off frequency for a parallel power/ground plane pair, this is probably the lowest realizable impedance that can be achieved. This assumes that the board extends infinitely in all directions. In practice the board has finite dimensions that lead to modal resonances because of reflections from the 'open' edges. This will cause localized areas of greater or lesser impedance than that indicated by the above equation.

Guidelines and Checklist

The following guidelines are derived from the preceding discussion.

1. Good power distribution bypassing ensures not only more reliable operation of the processor but also helps to minimize ground bounce, signal integrity, and EMI problems.
2. Use sufficient capacity to allow for low-frequency (i.e., 1 kHz to 100 kHz), program-dependent, current demands. Several 47 μF electrolytic capacitors are suggested, but it will depend upon the application.
3. Place several medium-value electrolytic capacitors around the board to provide damping and medium-frequency (i.e., 100 kHz to 1 MHz) bypassing. Suggested values are 1 μF to 10 μF .
4. Use a range of small-SMD, low-ESL, low-ESR capacitor values with proportionately smaller values than larger for high-frequency bypassing. Capacitor values should not be more than a factor of ten between each value. Suggested values include 470 pF, 2.7 nF, 10 nF, and 100 nF.
5. Use full power/ground plane pairs and distribute 2.7 nF to 10 nF bypass capacitors evenly across the planes at a distance of about one capacitor per approx 400 mm^2 to provide high-frequency (i.e., > 100 MHz) current demands.
6. Use minimal spacing between the power and ground planes to minimize their inductance. As an added benefit, this also maximizes capacitance.
7. Locate the power/ground planes close to the component surface. (Note: because of physical problems with temperature warping, copper plane pairs are normally fabricated in dual pairs that are located symmetrically about the center of the PC board, stack up.)
8. For critical applications, use ‘optimized’ SMD capacitor packages such as size 0508 and 0612.
9. Minimize the contribution of via inductance by locating them within the capacitor PCB mounting pads. If this is not possible (because of assembly limitations), locate them on one side of the pad.
10. For ‘busy areas’ such as processor packages, use multiple capacitors spaced evenly around the periphery of the IC package.
11. Locate bypass capacitors on the same side as the processor, assuming the power/ground plane pair is located closer to the component side of the PC board where the processor is mounted.
12. Beware of critical PC board dimensions that might stimulate modal resonances, especially for board dimensions longer than ~20 cm and frequencies above 200 MHz.

Use the following checklist when designing your board.

Design Checklist	
<input type="checkbox"/>	Determine the number or required layers and provide the proper layer stacking
<input type="checkbox"/>	Determine the voltage regulator frequency response
<input type="checkbox"/>	Determine the Vdd internal bulk capacitance
<input type="checkbox"/>	Determine the Vdd internal high-frequency decoupling capacitance. Select the capacitors. Ensure proper placement and layout
<input type="checkbox"/>	Determine the external Vdd required decoupling capacitance. Select the capacitors. Ensure proper placement and layout
<input type="checkbox"/>	Implement the recommended circuitry for A_{VDD} decoupling and provide recommended layout
<input type="checkbox"/>	Avoid long traces and vias to minimize inductance
<input type="checkbox"/>	Refer to the EZ-KIT Lite evaluation system schematics: ftp://ftp.analog.com/pub/tools/Hardware/Reference_Designs The design database for the SHARC processors' EZ-KIT Lite, contains all the electronic information required for the design, layout, fabrication, and assembly:

References

- [1] “Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology” IEEE Transaction on Advanced Packaging, Larry D. Smith, Douglas W. Forehand et al
- [2] “Decoupling Capacitor Calculations for CMOS Circuits” IEEE Transaction, Larry D. Smith
- [3] “High-Speed Digital Design - A Handbook of Black Magic”, Howard W. Johnson, Martin Graham, Prentice-Hall, 1993, ISBN 0-13-395724-1
- [4] “Signal Integrity Simplified”, Eric Bogatin, Prentice-Hall, 2003, ISBN 0-13-066946-6
- [5] “High-Speed Digital System Design”, Stephen H. Hall, Garrett W. Hall, James A. McCall, John Wiley & Sons, 2000, ISBN 0-471-36090-2
- [6] “High-Speed Signal Propagation, Advanced Black Magic”, Howard W. Johnson, Martin Graham, Prentice-Hall, 2003, ISBN 0-13-084408-X
- [7] “Tolerance Calculations in Power Distribution Networks”, Istvan Novak, Sun Microsystems, Note published in Xcell Journal, Summer 2004. Available at www.xilinx.com
- [8] “Power Distribution System (PDS) Design; Using Bypass/Decoupling Capacitors”, Mark Alexander, Xilinx Application Note, XAPP623 (V2.0) April 2004.
- [9] “High Speed PCB Design & System Design”, Course notes, Instructor Lee Ritchie, UC Berkeley Engineering Extension, May 1998
- [10] “Signal Integrity Issues & Printed Circuit Board Design”, Douglas Brooks, Prentice-Hall, 2003, ISBN-0-13-141884-X
- [11] “Effective Decoupling Radius of Capacitor”, H. Chen, J. Fang, W. Shi, paper presented at 2001 EPEP conference
- [12] “ESR & Bypass Capacitor Self Resonant Behavior – How to Select Bypass Caps”, Douglas Brooks, UltraCAD Design, Inc., 2000
- [13] “Parasitic Inductance of a Bypass Capacitor”, Dr. H. Johnson, page 32, EDN Magazine, July 2000
- [14] “Analysis of Multi-Layered Irregular Power Distribution Planes with Vias Using Transmission Matrix Method”, J. Kim and M. Swaminathan, Dept of Elec & Cmptr Eng’g, Georgia Inst of Tech

- [15] “A Transmission-Line Mode for Ceramic Capacitors for CAD Tools Based on Measured Parameters”, Larry Smith, et. al., Sun Microsystems Inc, Published in ECTC, San Diego, CA, pp331-336, May 2002
- [16] “Right the First Time – A Practical Handbook on High Speed PCB & System Design”, Lee W. Ritchey, Speeding Edge, May 2003, ISBN 0-9741936-0-7

Additional Resources

Most manufacturers of high-speed semiconductor ICs and PC boards have investigated the issues of effective power supply bypassing. The following companies have various application notes and support tools on their Web sites.

1. Analog Devices, Inc: www.analog.com
2. Intel: <http://www.intel.com/hardware/design/solutions/>
3. AMD: www.amd.com
4. Xilinx: www.xilinx.com
5. Sun Microsystems: www.sun.com
6. AVX Corp: www.avxcorp.com
7. Actel: <http://www.actel.com/documents/PowerSupplyAN.pdf>
8. Altera: <http://www.altera.com/support/spt-index.html>
9. Kemet Electronic Corp: www.kemet.com
10. SPECCTRAQuest PCB 2D/3D EM Modeling software: www.specctraquest.com/
11. Allegro High-Speed PCB Design Tools: www.cadencepcb.com
12. HyperLynx High-Speed Signal Integrity & EMI Analysis: www.mentor.com
13. AnSoft: 2-D, 3-D EM Modelling: www.ansoft.com
14. Signal Consulting, Inc.: www.signalintegrity.com

Document History

Revision	Description
<i>Rev 1 – December 5, 2006 by Alberto Comaschi</i>	Initial Release.