 By Michael C Greim, Mercury Computer Systems Inc

LOWER VOLTAGES, HIGHER CURRENT TRANSIENTS, AND HIGHER CLOCK RATES RENDER RULES OF THUMB USELESS—EVEN HARMFUL—FOR DESIGNING POWER-DISTRIBUTION-SYSTEM DECOUPLING NETWORKS. ONCE YOU UNDERSTAND THE ISSUES, HOWEVER, YOU CAN ADOPT A SIMPLE DESIGN APPROACH THAT WORKS BETTER AND SAVES MONEY AND TIME.

High-end digital systems give a thumbs down to rules of thumb

As edge rates, noise margins, and power-supply voltages decrease and component power requirements and clock rates increase, power-system design becomes increasingly important to products' signal integrity. Although a signal-integrity strategy requires analyzing signals' performance in a system, many of the analysis tools assume that the power planes provide constant voltages from appropriate, frequency-independent sources. At the same time, rules of thumb still govern many approaches to power-system design and decoupling. Because it fails to recognize the ways in which system elements operate and interact, the all-too-prevalent more-is-better approach leads to many signal-integrity failures. The following material describes the basic elements of an effective, contemporary power-distribution system (PDS) and outlines a systematic approach to selecting the system components and deploying them within the system.

Over the last decade, power-supply voltages for advanced-technology ICs have steadily declined—from 5V at the beginning of the decade to a little more than 1V today. At the same time, the power dissipated per square centimeter of pc-board area has risen by a factor of almost 20 (**Table 1**). The result of this evolution is that the maximum acceptable PDS output impedance has dropped by a factor of more than 100. During the same period, rising clock rates have significantly extended the range of frequencies over which PDS output impedances must remain low and controlled. For pc-board designers, these requirements have transformed a task that previously required little or no analysis to one that requires careful signal-integrity investigation. Rules of thumb no longer do the job.

POWER-DISTRIBUTION SYSTEMS AND SIGNAL INTEGRITY

The PDS affects signal integrity in several ways. The power system provides operating current for the drivers and receivers that make up the circuit, provides a reference for the signal conductor, and forms the medium in which return currents flow. An improperly designed PDS can exhibit supply droop, which can degrade driver rise and fall times and re-

Figure 1

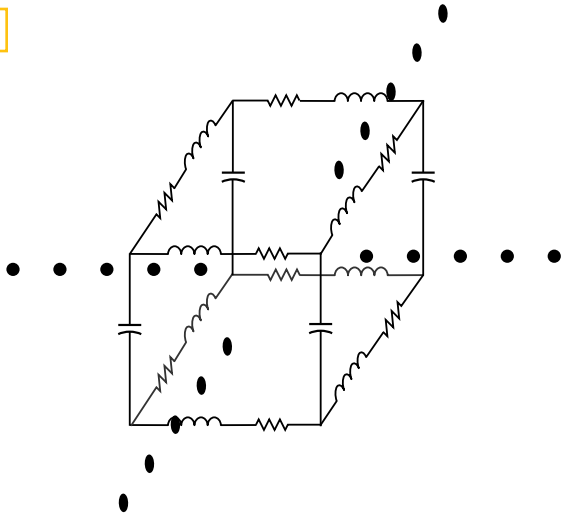

You can regard power planes as distributed networks containing shunt capacitance and series resistance and inductance. Shunt conductance across the capacitance is also present, but its effects are usually negligible.

duce noise margins. Signal-integrity simulations generally fail to catch power-plane noise, which can couple onto signals. Such erroneous simulations occur because many of the analysis tools assume that the power planes are ideal conductors that act as rock-solid references, which can provide unlimited current on demand. Unfortunately, power planes are not ideal. They can supply only limited steady-state and transient power. The planes also exhibit distributed parasitic elements, such as series and shunt resistance, series inductance, and shunt capacitance. These parasitics affect how the planes interact with signal conductors and with each other (**Figure 1**).

To properly design the PDS, you must first understand the target impedance and frequency-range requirements. You can easily use Ohm's law to calculate the PDS' target impedance. Multiplying the supply voltage by the maximum tolerable amount of ripple and dividing the product by the transient-current requirement yields an effective first-pass approximation of the target impedance.

PDS impedance=Z_{PDS}=((power-supply voltage)*(% ripple/100))/(required current) =1.8V*0.05/50A=1.8 mΩ, for a 1.8V supply with 5% peak-to-peak ripple and 50A transient-peak load current.

To calculate the PDS' upper operating-frequency limit (you can assume that the lower limit is dc), you must understand the system's instantaneous-current requirements. Basic signal theory says that you can calculate a square-wave signal's bandwidth (BW) as BW=0.35*1/T_{RISE}. For a 622-Mbps low-voltage differential signal (LVDS) with a peak toggle rate of 311 MHz and 500-psec rise and fall times, the bandwidth translates to: BW_{LVDS}=0.35*1/500 psec=700 MHz.

A contemporary PDS consists of four basic elements that control the power-bus impedance in different frequency bands. A properly designed PDS meets its target impedance requirements within each frequency band and in the transitions between them. The transitions between bands are where most decoupling schemes encounter problems. **Figure 2** illustrates the basic elements of the PDS and its effective regions.

The simplest and most useful model of any capacitor consists of the capacitor's nominal capacitance in series with the component's effective series resistance (ESR) and effective series inductance (ESL). These parasitic elements, which are functions of the device material and package style, determine the capacitor's performance at various frequencies.

The capacitor's lowest impedance occurs at the resonant frequency, at which the device's capacitive and inductive portions have equal and opposite reactances (**Figure 3**). The equations also underscore the importance of low inductance in high-frequency decoupling networks. For a given capacitor value, higher ESL lowers the device's resonant frequency and increases the decoupling impedance at frequencies above the resonant frequency. This data demonstrates how focusing on ESR and ignoring ESL can cause a power system to fail.

For a given value of C, a small ESR coupled with a high ESL yields an impedance profile that has a deep, narrow notch at the resonant frequency. When you connect in parallel several capacitors whose nominal values differ sufficiently, the resulting impedance peaks can compromise system performance. **Figure 4** illustrates a common decoupling topology: a few tantalum capacitors plus a large number of 0.22- and 0.01-μF ceramic-chip capacitors. Note the large impedance spike between the two chip capacitors' resonant frequencies.

Figure 5 illustrates how focusing on reducing inductance and moving the resonant frequencies closer together in the area of concern permits more effective decoupling with fewer capacitors. Focusing on these design elements reduces the power-system impedance at the problem

Figure 2

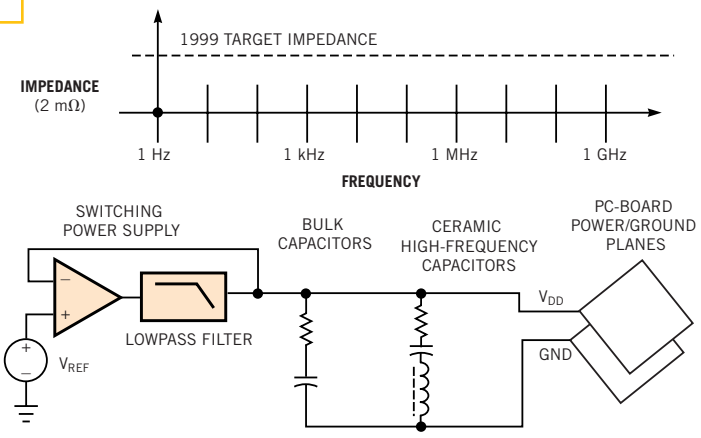

For the purposes of signal-integrity analysis, a power-distribution system comprises the power supply, "bulk" electrolytic capacitors, ceramic-chip capacitors, and the power and ground planes within the pc board.

Figure 3

$|Z_C|=\sqrt{(ESR^2+(\omega ESL-1/\omega C)^2}$, WHERE $\omega=2\pi f$.

$f_{RESONANCE}=1/(2\pi*\sqrt{(L*C)})$.

This capacitor-equivalent circuit consists of three elements in series–the device's nominal capacitance, its ESL, and its ESR. Sometimes, you see the circuit drawn with another element–a conductance shunting the capacitor. In some cases, that shunt conductance is frequency-dependent.

TABLE 1–SUPPLY-VOLTAGE DECREASES

Year	Voltage (V)	Power (W)	Current (A)	Z_{TARGET} (Ω)	Frequency (MHz)
1990	5	5	1	250	16
1993	3.3	10	3	54	66
1996	2.5	30	12	10	200
1999	1.8	90	50	1.8	600
2002	1.2	180	150	0.4	1200

Figure 4

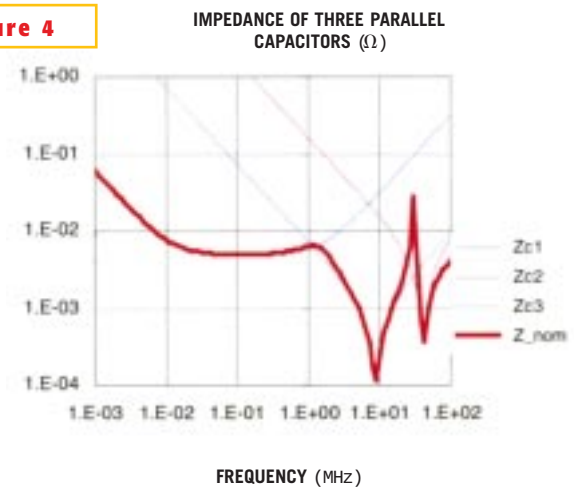

The impedance of this parallel combination of three types of capacitors is well-behaved at frequencies below 1 MHz. At higher frequencies, you see series-resonant valleys and parallel-resonant peaks. These devices have low ESR and relatively high ESL. The three types are 330 μF tantalum (eight in parallel), 0.22 μF ceramic chip (100 in parallel), and 0.01 μF ceramic chip (100 in parallel). The two types of chip capacitors have an ESL of 1.5 nH.

Figure 5

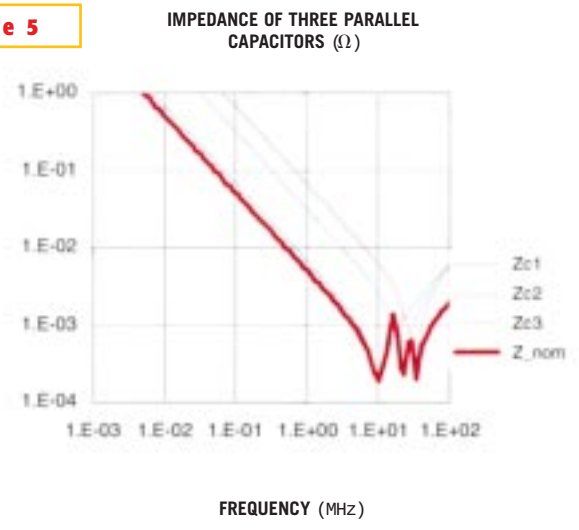

Selecting capacitors for lower ESL than those in Figure 4 allows you to use fewer devices and to achieve a lower impedance over a broader range of frequencies. This setup uses chip capacitors whose ESL is 0.5 nH (50 units of each of three values: 0.47, 0.1, and 0.047 μF).

frequency by more than an order of magnitude. This example shows how adding more of the wrong kind of decoupling capacitors merely takes up board space. A by-product of proper decoupling is that reducing the number of capacitors reduces the number of vias and leaves more room for routing signal traces, thus improving signal integrity.

BULK CAPACITANCE

Bulk capacitance supplies current and maintains the output voltage at frequencies above those at which the power supply can respond and below those at which chip capacitors are effective. Generally, this range extends from tens of kilohertz to almost 10 MHz. **Figure 6** shows the impedance versus frequency plot of nine T510-series tantalum capacitors in parallel.

As a PDS designer, you must determine how much decoupling capacitance and how many devices the pc board requires. You can determine the required capacitance from the transient-current requirements, the supply's response time, and the supply's allowable output-voltage ripple. This calculation translates to another basic circuit-theory equation: C=I*dt/dV; that is, capacitance=transient current*(supply-response time/supply ripple).

As an example, suppose that a PDS' bulk decoupling networks must respond in 90 μsec to a 4A current transient. The allowable ripple should not exceed 5% of the 1.8V rail voltage. These requirements translate to: C_{BULK}=4A*(90 μsec/0.09V) =4000 μF.

Once you calculate the capacitance, you must decide how many devices to use to achieve the target capacitance value. To determine the appropriate number of capacitors, compare the bulk capacitors' ESR to the system's target impedance. The decoupling impedance decreases by half each time you double the number of capacitors. By comparing the parallel ESR (based on the capacitors' nominal value) to the system requirements, you can determine how many devices you need to achieve your design goals.

CERAMIC-CHIP CAPACITORS

Several factors influence the selection and distribution of a PDS' chip capacitors. The main selection criterion remains satisfying the PDS impedance target across the frequency range. Ceramic-chip capacitors are effective from less than 10 MHz to several hundred megahertz (**Figure 6**). The **figure** shows that, with an array of 47-nF reverse-form-factor (0508-style) capacitors, the PDS impedance reaches unacceptable levels just below 100 MHz. At frequencies above this point, the network becomes ineffective because of the packages' inherent inductance, the pc-board attachment method, and the low capacitance. Reducing the value of the capacitors that constitute the array does not help. If you reduce a capacitor's value but require the component to deliver the same current, you increase dV/dt and, hence, ripple. As the following sections reveal, other decoupling methods prove more effective at higher frequencies.

Capacitor material is another important selection consideration. Designers most frequently encounter NP0, X7R, and Y5V ceramic-dielectric types. For most decoupling applications, X7R is the best choice. X7R has higher ESR and poorer temperature and voltage coefficients of capacitance than does NP0. Nevertheless, X7R combines adequate temperature and voltage performance with a wider capacitance range and higher capacitance per unit volume than NP0. Although it produces still higher capacitance per unit volume, the Y5V dielectric's temperature and voltage performance are frequently unacceptable.

Often, the most important choice in selecting a capacitor is package style, which strongly influences the ESL for a given capacitance. Low ESL is critical at high frequencies because it determines the capacitor's resonant frequency and impedance roll-off characteristics above

the resonant frequency. **Table 2** (pg 110) indicates the ESL of several common chip-capacitor package types. For a given capacitance, a range of resonant frequencies and impedance roll-off characteristics is possible, depending on the package style and mounting technique. Availability, leadtime, and cost often lead to the selection of higher-impedance devices, even though they do not perform as well as their lower inductance counterparts.

Figure 6

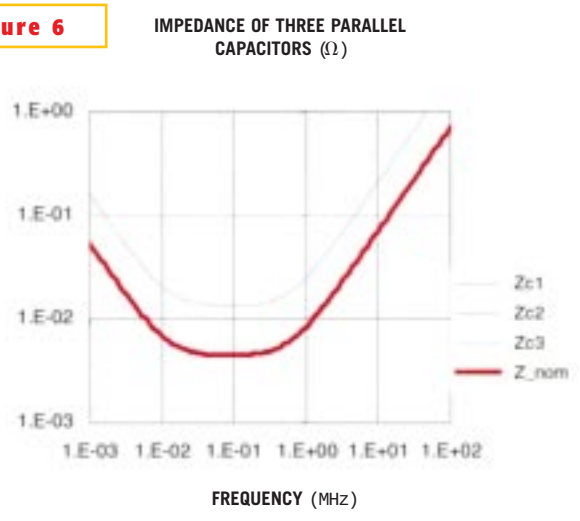

Tantalum capacitors exhibit three distinct impedance regions. At low frequencies, they appear capacitive. At high frequencies, they appear inductive. In a middle range, they appear resistive. This plot shows the impedance versus frequency of three T510 devices–individually (upper curve) and in parallel (lower curve).

BOARD-LAYOUT CONSIDERATIONS

An important and frequently overlooked effect in PDS design is the device-to-pc-board attachment method, which can profoundly affect capacitors' effective inductance. Thanks to recent reductions in capacitor inductance, you often find low-inductance capacitors attached to the circuit board through a fan-out etch whose inductance dwarfs the device inductance. As expected, this effect significantly affects the PDS' high-frequency output impedance. In the impedance-corrected circuit of **Figure 5**, you can see the effects of incorrectly mounted capacitors that have the proper capacitance and ESL. Long etch fans outward from the end of the pads. Mounting the capacitors in this way raises the inductance by a factor of 10 and adversely affects the frequency response.

A little-known fact is that the placement of the fan-out etch and vias that connect a decoupling capacitor to the power plane and ground can substantially affect the device's series inductance. To understand this phenomenon, you must understand the importance of minimizing the loop inductance. **Figure 7** illustrates the parasitic loop-inductance path for a mounted SMD capacitor. From the circuit's mounting description (**Figure 8**), you can see that routing the fan-out etch outward from the end of the capacitor maximizes the loop area for the capacitor size and fan-out etch length. Moving this escape etch to the inner side of the pad edge substantially reduces the loop area and the corresponding inductance. To minimize the inductance, you'd eliminate the fan-out etch and place the via within the pad at its inner edge. Unfortunately, this mounting method can cause problems in manufacturing. Making the via tangent to the inner edge of the pad is a good compromise. This scheme easily achieves an attachment inductance of less than 1 nH (**Figure 9**).

PLANE CAPACITANCE 101

The role of plane capacitance in power distribution has recently received considerable attention. As thinner laminates have become less expensive and more common, designers have recognized that the capacitance between power and ground planes can provide appreciable decoupling. More important, this distributed capacitance remains effective at frequencies so high that chip capacitors become ineffective. Making the plane areas larger than the plane-to-plane spacing yields significant capacitance. Moreover, these "embedded capacitors" exhibit extremely low ESR and ESL. The following equation defines the plane capacitance per unit area: $C_{PLANE}=(0.225*\epsilon_R*A)/d$, where ϵ_R=relative electric permittivity of the material, A=area of the shared power planes in square inches, D=separation between the planes, in inches, and C_{PLANE}=capacitance per square inch in pF.

Two power planes, separated by 2 mils of standard FR-4 material ($\epsilon_R=4.5$) provide more than 500 pF of capacitance per square inch. With the low

Figure 7

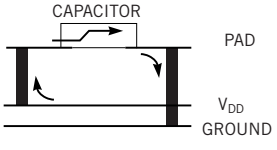

As this cross-sectional view of a mounted chip capacitor shows, the path that the current follows from V_{CC} or V_{DD} to ground can be substantial, thanks simply to the placement of the power and ground planes near the board surface farther from the capacitor. The ESL of the mounted device depends strongly on the area of this loop.

Figure 8

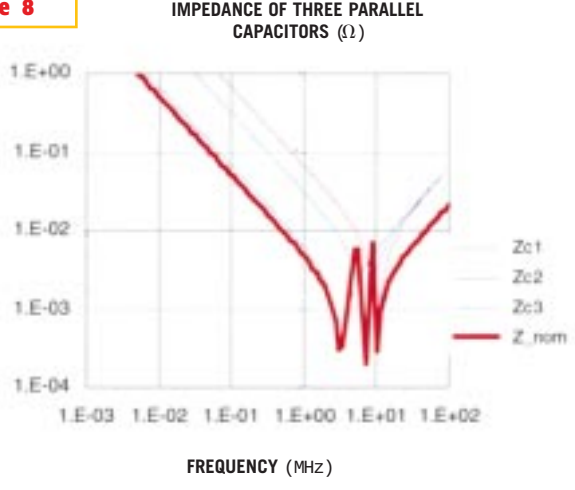

Selecting low-ESL capacitors does no good unless you mount them properly. The inductance of the connections to the power and ground planes ruins the behavior of these devices.

inductance of this structure, the resonant frequency can approach 1 GHz and, compared with mounted chip capacitors, can provide a low rate of impedance roll-off. With standard pc-board core and prepreg-material thickness continuing to decrease, planes separated by a pair of signal layers can have capacitances per square inch of almost 100 pF. Whereas standard pc-board materials can achieve this effect, a number of materials, such as EmCap from Hadco (www.hadco.com/prod03.htm), have material properties targeting embedded-capacitance applications.

Signal-layer-plane fills are the logical extension of plane capacitance. Routed boards often contain large areas of signal layers that are not occupied with traces. In such cases, you can add copper shapes to form plates on adjacent layers. For a dual-stripline pc-board stackup, this practice yields three times the capacitance per unit area, because the shapes couple to each other and to the adjacent planes. In assigning power to these shapes, take care to preserve the alternating-power-and-ground-plane structure (**Figure 10**). Another consideration in implementing this scheme is the effect on trace impedance of the spacing between the plane fill shape and the signal traces on the layer.

SUMMING UP

Because of increases in component power consumption and operating frequency, and decreases in supply voltages and edge rates, rule-of-thumb design no longer yields power-distribution systems that provide adequate signal integrity. For signal-integrity simulations to accurately predict the behavior of the real boards, the PDS must provide a stable reference that presents the appropriate impedance over the logic circuits' frequency range. By paying attention to the following areas, you can achieve a properly designed PDS with little or no additional effort, and often with fewer components.

- Identify the system's transient-current needs, required frequency range, and acceptable voltage ripple.
- Calculate the PDS' target output impedance.
- Understand the role of each of the PDS' basic elements and the frequency range at which each is effective.
- Take care in the transition bands between the four regions to ensure the PDS' output impedance remains below the target values.
- To avoid unanticipated impedance spikes, understand the interaction of ESR and ESL and their impact on the resonant frequency and impedance roll-off.
- Pay attention to the method of mounting the chip capacitors to the pc board. For good, high-frequency performance, you must minimize loop inductance. Short escape etch is not the only factor. The depth to the plane and the distance between the points at which you attach the capacitor to the board are also significant.
- The circuit board plays an essential role in the decoupling strategy. Take advantage of interplane capacitance and plane fills to increase the frequency response of the PDS.
- To validate the scheme, simulate the system behavior over the frequency range of interest. □

REFERENCES

1. Novak, Istvan, "Powering digital boards: distribution and performance," Meeting of the Greater Boston Chapter, Institute for Printed Circuits, Designer's Council, Feb 9, 1999.
2. Matthews, Ron, and M Carpenter,

(continued on pg 110)

Figure 9

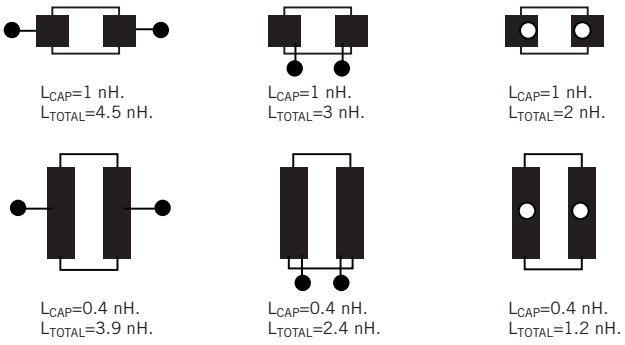

The way you route the power and ground connections to a surface-mounted chip capacitor can cause the mounted device's wiring inductance to vary from 0.8 to 3.5 nH, a range of 4.375 to 1.

Figure 10

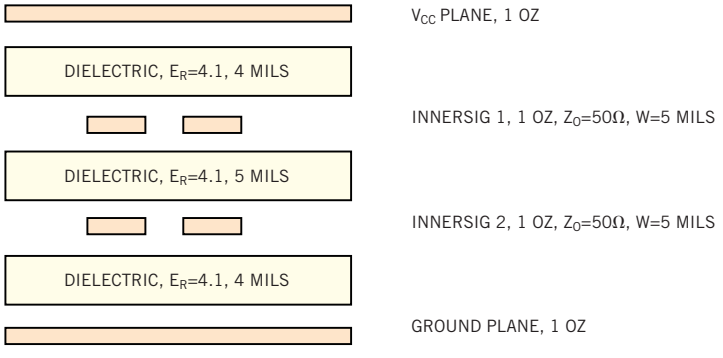

In this four-layer pc board, the power plane is on top, the ground plane is on the bottom, and the signals are on the two inner layers. You see an end view of two side-by-side striplines, both of which lie between the plates of the capacitor that consists of the power and ground planes and the dielectric material between them.

"Spectraquest power-plane designer: the future of power-plane design," Cadence Design Systems Inc.

3. Smith, Larry, et al, "Power-distribution-system design methodology and capacitor selection for modern CMOS technology," Sun Microsystems Inc, August 1999.

4. Hadco Corp, "The buried capacitance design guide," Santa Clara, CA, June 1996.

5. Roy, Tanmoy, and L Smith, "ESR and ESL of ceramic capacitor applied to decoupling applications," Sun Microsystems, October 1998.

6. Prymak, John D, "Advanced decoupling using ceramic MLC capacitors," AVX Corp.

7. Brooks, Douglas G, "ESR and bypass-capacitor self-resonant behavior: how to select bypass caps," UltraCAD Design Inc, Feb 21, 2000.

8. Buchanan, James, *Signal and Power Integrity in Digital Systems: TTL, CMOS, and BiCMOS*, McGraw Hill, New York, 1995, ISBN 0-07-008734-2 .

9. Johnson, Howard W, and M Graham, *High-Speed Digital Design, a Handbook of Black Magic*, PTR Prentice Hall, Englewood Cliffs, NJ, 1993, ISBN 0-13-395724-1.

10. www.qsl.net/wb6tpu/si_documents/docs.html.

11. http://home.att.net/~istvan.novak/.

12.www.avxcorp.com/TechInfo_catlisting.asp.

13. www.ultracad.com/.

14. ftp://download.intel.com/design/PentiumIII/xeon/applnots/24509501.pdf.

15.www.chips.ibm.com/products powerpc/library/an022.pdf.

TABLE 2–TYPICAL INDUCTANCE VALUES FOR CHIP-CAPACITOR PACKAGES

Package type	Associated Inductance (pH)
Axial Lead MLC	2000
1206 SMD	1250
0805 SMD	1050
1210 SMD	980
0603 SMD	870
0612 SMD	610
0508 SMD	600
AVX DCAP	50

Notes:
MLC=multilayer capacitor.
SMD=surface-mounted device.
DCAP=A trademark of AVX Corp.

AUTHOR'S BIOGRAPHY

Michael C Greim is a consulting hardware engineer for Mercury Computer Systems Inc (Chelmsford, MA). He has a BSEE from Boston University (Boston) and an MS in engineering management from the Gordon Institute of Tufts University (Medford, MA). His background includes high-speed logic design, signal-integrity analysis, and multiprocessor and memory-system design. You can reach him at mgreim@mc.com.