5

High Performance Power Distribution Systems

Supplying power to high performance integrated circuits has become a challenging task. The system supplying power to an integrated circuit greatly affects the performance, size, and cost characteristics of the overall electronic system. This system is comprised of interconnect networks with decoupling capacitors on a printed circuit board, an integrated circuit package, and a circuit die. The entire system is henceforth referred to as the *power distribution system*. The design of power distribution systems is described in this chapter. The focus of the discussion is the overall structure and interaction among the various parts of the system. The impedance characteristics and design of on-chip power distribution networks, the most complex part of the power distribution system, are discussed in greater detail in the following chapters.

The chapter is organized as follows. A typical power distribution system for a high power, high speed integrated circuit is described in Section 5.1. A circuit model of a power distribution system is presented in Section 5.2. The output impedance characteristics are discussed in Section 5.3. The effect of a shunting capacitance on the impedance of a power distribution system is considered in Section 5.4. The hierarchical placement of capacitors to satisfy target impedance requirements is described in Section 5.5. Design strategies to control the resonant effects in power distribution networks are discussed in Section 5.6. A purely resistive impedance characteristic of power distribution systems can be achieved using an impedance compensation technique, as described in Section 5.7. A case study of a power distribution system is presented in Section 5.8. Design techniques to enhance the impedance characteristics of power distribution networks are discussed in Section 5.9. The

limitations of the analysis presented herein are discussed in Section 5.10. The chapter concludes with a summary.

5.1 Physical structure of a power distribution system

A cross-sectional view of a power distribution system for a high performance integrated circuit is shown in Fig. 5.1. The power supply system

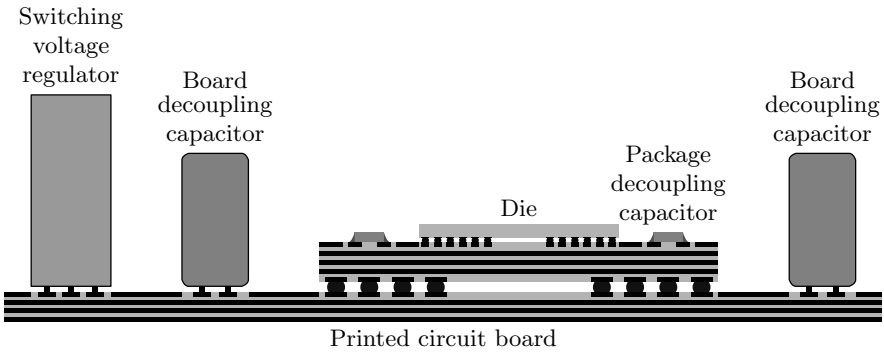

Fig. 5.1. A cross-sectional view of a power distribution system of a high performance integrated circuit.

spans several levels of packaging hierarchy. It consists of a switching voltage regulator module (VRM), the power distribution networks on a printed circuit board (PCB), on an integrated circuit package, and on-chip, plus the decoupling capacitors connected to these networks. The power distribution networks at the board, package, and circuit die levels form a conductive path between the power source and the power load. A switching voltage regulator converts the DC voltage level provided by a system power supply unit to a voltage V_{dd} required for powering an integrated circuit. The regulator serves as a power source, effectively decoupling the power distribution system of an integrated circuit from the system level power supply. The power and ground planes of the printed circuit board connect the switching regulator to the integrated circuit package. The board-level decoupling capacitors are placed across the power and ground planes to provide charge storage for current transients faster than the response time of the regulator. The board power and ground interconnect is connected to the power and ground networks of the package through a ball grid array

(BGA) or pin grid array (PGA) contacts. Similar to a printed circuit board, the package power and ground distribution networks are typically comprised of several metal planes. High frequency (*i.e.*, with low parasitic impedance) decoupling capacitors are mounted on the package, electrically close to the circuit die, to ensure the integrity of the power supply during power current surges drawn by the circuit from the package level power distribution network. The package power and ground distribution networks are connected to the on-chip power distribution grid through a flip-chip array of bump contacts or alternative bonding technologies [103]. On-chip decoupling capacitors are placed across the on-chip power and ground networks to maintain a low impedance at signal frequencies comparable to the switching speed of the on-chip devices.

The physical structure of the power distribution system is hierarchical. Each tier of the power distribution system typically corresponds to a tier of packaging hierarchy and consists of a power distribution network and associated decoupling capacitors. The hierarchical structure of the power distribution system permits the desired impedance characteristics to be obtained in a cost effective manner, as described in the following sections.

5.2 Circuit model of a power distribution system

A simplified circuit model of a power supply system[1] is shown in Fig. 5.2. This lumped circuit model is effectively one-dimensional, where each level of the packaging hierarchy is modeled by a pair of power and ground conductors and a decoupling capacitor across these conductors. A one-dimensional model accurately describes the impedance characteristics of a power distribution system over a wide frequency range [104], [105], [106]. The conductors are represented by the parasitic resistive and inductive impedances; the decoupling capacitors are represented by a series RLC circuit reflecting the parasitic impedances of the capacitors. The italicized superscripts "p" and "g" refer to the power and ground conductors, respectively; superscript "C" refers to the parasitic impedance characteristics of the capacitors. The subscripts "r," "b," "p," and "c" refer to the regulator, board, package, and on-chip

[1] The magnetic coupling of the power and ground conductors is omitted in the circuit diagrams of a power distribution system. The inductive elements shown in the diagrams therefore denote a net inductance, as described in Section 2.1.4.

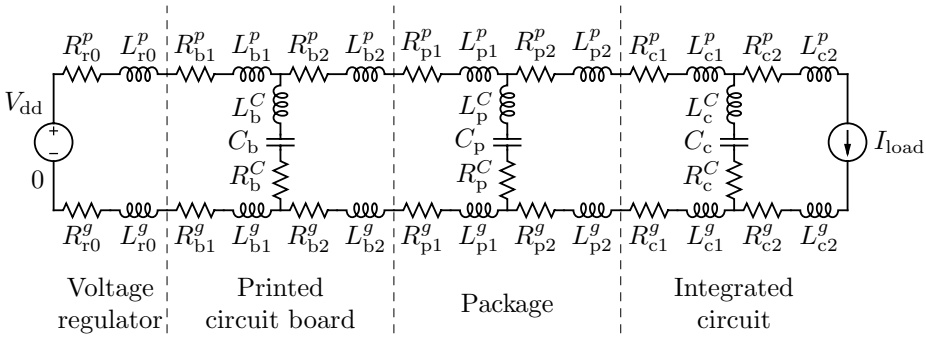

Fig. 5.2. A one-dimensional circuit model of the power supply system shown in Fig. 5.1.

conductors. Subscript 1 refers to the conductors upstream of the respective decoupling capacitor, with respect to the flow of energy from the power source to the load. That is, the conductors denoted with subscript 1 are connected to the appropriate decoupling capacitor at the voltage regulator. Subscript 2 refers to the conductors downstream of the appropriate decoupling capacitor. For example, R^g_{b1} refers to the parasitic resistance of the ground conductors connecting the board capacitors to the voltage regulator, while L^p_{p2} refers to the parasitic inductance of the power conductors connecting the package capacitors to the on-chip I/O contacts. Similarly, L^C_p refers to the parasitic series inductance of the package capacitors.

The model shown in Fig. 5.2 can be reduced by combining the circuit elements connected in series. The reduced circuit model is shown in Fig. 5.3. The circuit characteristics of the circuit shown in Fig. 5.3 are related to the characteristics of the circuit shown in Fig. 5.2 as

$$R^p_r = R^p_{r0} + R^p_{b1} \qquad L^p_r = L^p_{r0} + L^p_{b1} \tag{5.1a}$$
$$R^p_b = R^p_{b2} + R^p_{p1} \qquad L^p_b = L^p_{b2} + L^p_{p1} \tag{5.1b}$$
$$R^p_p = R^p_{p2} + R^p_{c1} \qquad L^p_p = L^p_{p2} + L^p_{c1} \tag{5.1c}$$
$$R^p_c = R^p_{c2} \qquad L^p_c = L^p_{c2} \tag{5.1d}$$

for conductors carrying power current and, analogously, for the ground conductors,

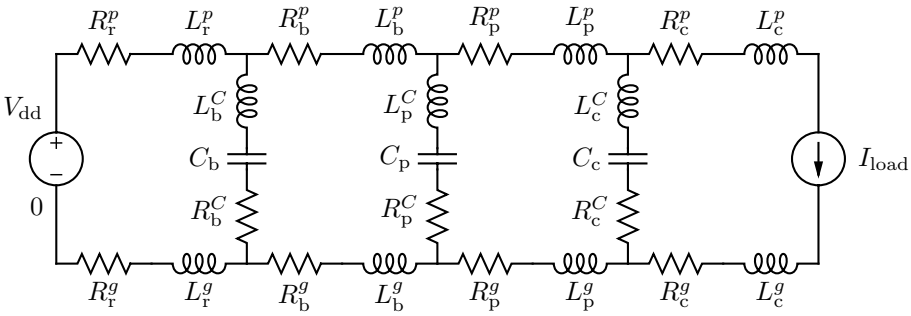

Fig. 5.3. A reduced version of the circuit model characterizing a power supply system.

$$R^g_r = R^g_{r2} + R^g_{b1} \qquad L^g_r = L^g_{r2} + L^g_{b1} \tag{5.1e}$$
$$R^g_b = R^g_{b2} + R^g_{p1} \qquad L^g_b = L^g_{b2} + L^g_{p1} \tag{5.1f}$$
$$R^g_p = R^g_{p2} + R^g_{c1} \qquad L^g_p = L^g_{p2} + L^g_{c1} \tag{5.1g}$$
$$R^g_c = R^g_{c2} \qquad L^g_c = L^g_{c2}. \tag{5.1h}$$

As defined, the circuit elements with subscript "r" represent the output impedance of the voltage regulator and the impedance of the on-board current path from the regulator to the board decoupling capacitors. The elements denoted with subscript "b" represent the impedance of the current path from the board capacitors to the package, the impedance of the socket and package pins (or solder bumps in the case of a ball grid array mounting solution), and the impedance of the package lines and planes. Similarly, the elements with subscript "p" signify the impedance of the current path from the package capacitors to the die mounting site, the impedance of the solder bumps or bonding wires, and, partially, the impedance of the on-chip power distribution network. Finally, the resistors and inductors with subscript "c" represent the impedance of the current path from the on-chip capacitors to the on-chip power load.

The board, package, and on-chip power distribution networks have significantly different electrical characteristics due to the different physical properties of the interconnect at the various tiers of the packaging hierarchy. From the board level to the die level, the cross-sectional dimensions of the interconnect lines decrease while the aspect ratio increases, producing a dramatic increase in interconnect density. The inductance of the board level power distribution network, *i.e.*, L^p_{b1}, L^g_{b1}, L^p_{b2}, and L^g_{b2}, is large as the power and ground planes are typically

separated by tens or hundreds of micrometers. The effective output impedance of the voltage switching regulator over a wide range of frequencies can be described as $R+j\omega L$ [104], [107], [108], hence the presence of R^p_{r2} and R^g_{r2}, and L^p_{r2} and L^g_{r2} in the model shown in Fig. 5.2. The inductance of the on-chip power distribution network, L^p_{c1}, L^g_{c1}, L^p_{c2}, and L^g_{c2}, is comparatively low due to the high interconnect density. The inductance L^p_{p1}, L^g_{p1}, L^p_{p2}, and L^g_{p2} of the package level network is of intermediate magnitude, larger than the inductance of the on-chip network but smaller than the inductance of the board level network. The inductive characteristics of the circuit shown in Fig. 5.3 typically exhibit a hierarchical relationship: $L^p_b > L^p_p > L^p_c$ and $L^g_b > L^g_p > L^g_c$. The resistance of the power distribution networks follows the opposite trend. The board level resistances R^p_{b1}, R^g_{b1}, R^p_{b2}, and R^g_{b2} are small due to the large cross-sectional dimensions of the relatively thick board planes, while the on-chip resistances R^p_{c1}, R^g_{c1}, R^p_{c2}, and R^g_{c2} are large due to the small cross-sectional dimensions of the on-chip interconnect. The resistive characteristics of the power distribution system are therefore reciprocal to the inductive characteristics, $R^p_b < R^p_p < R^p_c$ and $R^g_b < R^g_p < R^g_c$. The physical hierarchy is thus reflected in the electrical hierarchy: the progressively finer physical features of the conductors typically result in a higher resistance and a lower inductance.

5.3 Output impedance of a power distribution system

To ensure a small variation in the power supply voltage under a significant current load, the power distribution system should exhibit a small impedance at the terminals of the load within the frequency range of interest, as shown in Fig. 5.4. The impedance of the power distribution system as seen from the terminals of the load circuits is henceforth referred to as the impedance of the power distribution system. In order to ensure correct and reliable operation of an integrated circuit, the impedance of the power distribution system is specified to be lower than a certain upper bound Z_0 in the frequency range from DC to the maximum frequency f_0 [107], as illustrated in Fig. 5.6. Note that the maximum frequency f_0 is determined by the switching time of the on-chip signal transients, rather than by the clock frequency f_{clk}. The shortest signal switching time is typically smaller than the clock period by at least an order of magnitude; therefore, the maximum frequency of interest f_0 is considerably higher than the clock frequency f_{clk}.

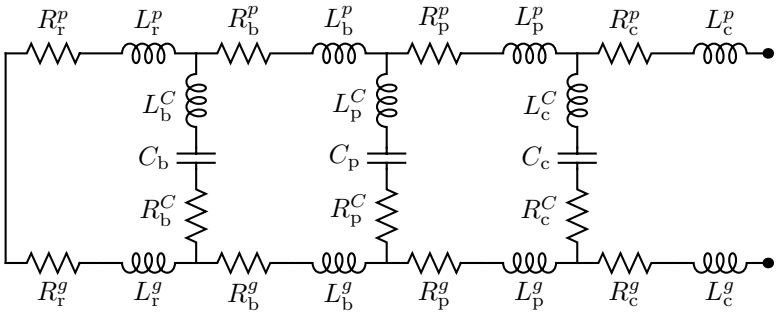

Fig. 5.4. A circuit network representing a power distribution system impedance as seen from the terminals of the power load. The power source at the left side of the network has been replaced with a short circuit. The terminals of the power load are shown at the right side of the network.

The objective of designing a power distribution system is to ensure a target output impedance characteristic. It is therefore important to understand how the output impedance of the circuit shown in Fig. 5.4 depends on the impedance of the comprising circuit elements. The impedance characteristics of a power distribution system are analyzed in the following sections. The impedance characteristics of a power distribution system with no capacitors are considered first. The effect of the decoupling capacitors on the impedance characteristics is described in the following sections.

A power distribution system with no decoupling capacitors is shown in Fig. 5.5. The power source and load are connected by interconnect

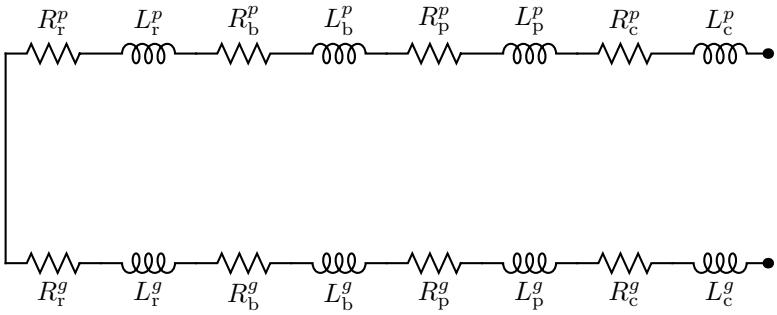

Fig. 5.5. A circuit network representing a power distribution system without decoupling capacitors.

with resistive and inductive parasitic impedances. The magnitude of the impedance of this network is

$$|Z_{\text{tot}}(\omega)| = |R_{\text{tot}} + j\omega L_{\text{tot}}|, \tag{5.2}$$

where R_{tot} and L_{tot} are the total series resistance and inductance of the power distribution system, respectively,

$$R_{\text{tot}} = R^p_{\text{tot}} + R^g_{\text{tot}}, \tag{5.3}$$
$$R^p_{\text{tot}} = R^p_{\text{r}} + R^p_{\text{b}} + R^p_{\text{p}} + R^p_{\text{c}}, \tag{5.4}$$
$$R^g_{\text{tot}} = R^g_{\text{r}} + R^g_{\text{b}} + R^g_{\text{p}} + R^g_{\text{c}}, \tag{5.5}$$
$$L_{\text{tot}} = L^p_{\text{tot}} + L^g_{\text{tot}}, \tag{5.6}$$
$$L^p_{\text{tot}} = L^p_{\text{r}} + L^p_{\text{b}} + L^p_{\text{p}} + L^p_{\text{c}}, \tag{5.7}$$
$$L^g_{\text{tot}} = L^g_{\text{r}} + L^g_{\text{b}} + L^g_{\text{p}} + L^g_{\text{c}}. \tag{5.8}$$

The variation of the impedance with frequency is illustrated in Fig. 5.6. To satisfy the specification at low frequency, the resistance of the power distribution system should be sufficiently low, $R_{\text{tot}} < Z_0$. Above the frequency $f_{L_{\text{tot}}} = \frac{1}{2\pi}\frac{R_{\text{tot}}}{L_{\text{tot}}}$, however, the impedance of the power distribution system is dominated by the inductive reactance $j\omega L_{\text{tot}}$ and increases linearly with frequency, exceeding the target specification at the frequency $f_{\max} = \frac{1}{2\pi}\frac{Z_0}{L_{\text{tot}}}$.

The high frequency impedance should be reduced to satisfy the target specifications. Opportunities for reducing the inductance of the interconnect structures comprising a power system are limited. The feature size of the board and package level interconnect, which largely determines the inductance, depends on the manufacturing technology. Furthermore, the output impedance of the voltage regulator is highly inductive and difficult to reduce.

The high frequency impedance is effectively reduced by placing capacitors across the power and ground conductors. These shunting capacitors terminate the high frequency current loop, permitting the current to bypass the inductive interconnect, such as the board and package power distribution networks. The high frequency impedance of the system as seen from the load terminals is thereby reduced. The capacitors effectively "decouple" the high impedance parts of the power distribution system from the load at high frequencies. These capacitors are therefore commonly referred to as decoupling capacitors. Several stages of decoupling capacitors are typically used to confine the output impedance within the target specifications.

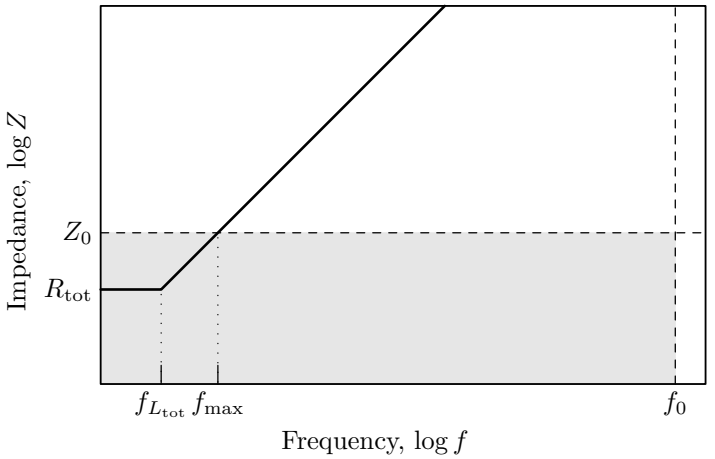

Fig. 5.6. Impedance of the power distribution system with no decoupling capacitors. The shaded area denotes the target impedance specifications of the power distribution system.

In the following sections, the impedance characteristics of a power distribution system with several stages of decoupling capacitors are described in several steps. The effect of a single decoupling capacitor on the impedance of a power distribution system is considered in Section 5.4. The hierarchical placement of the decoupling capacitors is described in Section 5.5. The impedance characteristics near the resonant frequencies are examined in Section 5.6.

5.4 A power distribution system with a decoupling capacitor

The effects of a decoupling capacitor on the output impedance of a power distribution system are the subject of this section. The impedance characteristics of a power distribution system with a decoupling capacitor is described in Section 5.4.1. The limitations of using a single stage decoupling scheme are discussed in Section 5.4.2.

5.4.1 Impedance characteristics

A power distribution system with a shunting capacitance is shown in Fig. 5.7. The shunting capacitance can be physically realized with

a single capacitor or, alternatively, with a bank of several identical capacitors connected in parallel. Similar to a single capacitor, the

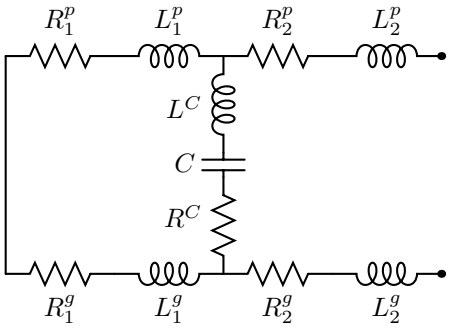

Fig. 5.7. A circuit model of a power distribution network with a decoupling capacitor C.

impedance of a bank of identical capacitors is accurately described by a series RLC circuit. The location of the capacitors partitions the power network into upstream and downstream sections, with respect to the flow of energy (or power current) from the power source to the load. The upstream section is referred to as stage 1 in Fig. 5.7; the downstream section is referred to as stage 2. Also note that the overall series inductance and resistance of the power distribution network remains the same as determined by (5.3) to (5.8): $L_1^p + L_2^p = L_{\text{tot}}^p$, $L_1^g + L_2^g = L_{\text{tot}}^g$, $R_1^p + R_2^p = R_{\text{tot}}^p$, and $R_1^g + R_2^g = R_{\text{tot}}^g$. The impedance of the power distribution network shown in Fig. 5.7 is

$$Z(\omega) = R_2 + j\omega L_2 + (R_1 + j\omega L_1) \,\|\, \left(R^C + j\left(\omega L^C - \frac{1}{\omega C}\right)\right), \quad (5.9)$$

where $R_1 = R_1^p + R_1^g$ and $L_1 = L_1^p + L_1^g$; analogously, $R_2 = R_2^p + R_2^g$ and $L_2 = L_2^p + L_2^g$.

The impedance characteristics of a power distribution network with a shunting capacitor are illustrated in Fig. 5.8. At low frequencies, the impedance of the capacitor is greater than the impedance of the upstream section. The power current loop extends to the power source, as illustrated by the equivalent circuit shown in Fig. 5.9(a). Assuming that the parasitic inductance of the decoupling capacitor is significantly smaller than the upstream inductance, *i.e.*, $L^C < L_1$, the capacitor has a lower impedance than the impedance of the upstream section

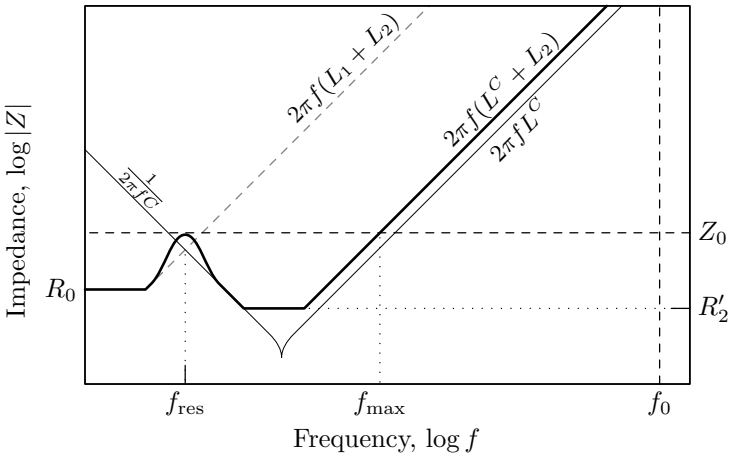

Fig. 5.8. Impedance of the power distribution system with a decoupling capacitor as shown in Fig. 5.7. The impedance of the power distribution system with a decoupling capacitor (the black line) exhibits an extended region of low impedance as compared to the impedance of the power distribution system with no decoupling capacitors, shown for comparison with the dashed gray line. The impedance of the decoupling capacitor is shown with a thin solid line.

$R_1 + j\omega L_1$ above the frequency $f_{\text{res}} \approx \frac{1}{2\pi}\frac{1}{\sqrt{L_1 C}}$. Above the frequency f_{res}, the bulk of the power current bypasses the impedance $R_1 + j\omega L_1$ through the capacitor, shrinking the size of the power current loop, as shown in Fig. 5.9(b). Note that the decoupling capacitor and the upstream stage form an underdamped LC tank circuit, resulting in a resonant[2] peak in the impedance at frequency f_{res}. The impedance at the resonant frequency f_{res} is *increased* by Q_{tank}, the quality factor of the tank circuit, as compared to the impedance of the power network at the same frequency without the capacitor. Between f_{res} and $f_R = \frac{1}{2\pi}\frac{1}{(R_2+R^C)C}$, the impedance of the power network is dominated by the capacitive reactance $\frac{1}{\omega C}$ and decreases with frequency, reaching $R_2' = R_2+R^C$. The parasitic resistance of the decoupling capacitor R^C should therefore also be sufficiently low, such that $R_2 + R^C < Z_0$, in order to satisfy the target specification. At frequencies greater than f_R, the

[2] The circuit impedance drastically increases near the resonant frequency in parallel (tank) resonant circuits. The parallel resonance is therefore often referred to as the *antiresonance* to distinguish this phenomenon from the series resonance, where the circuit impedance *decreases*. Correspondingly, peaks in the impedance are often referred to as antiresonant.

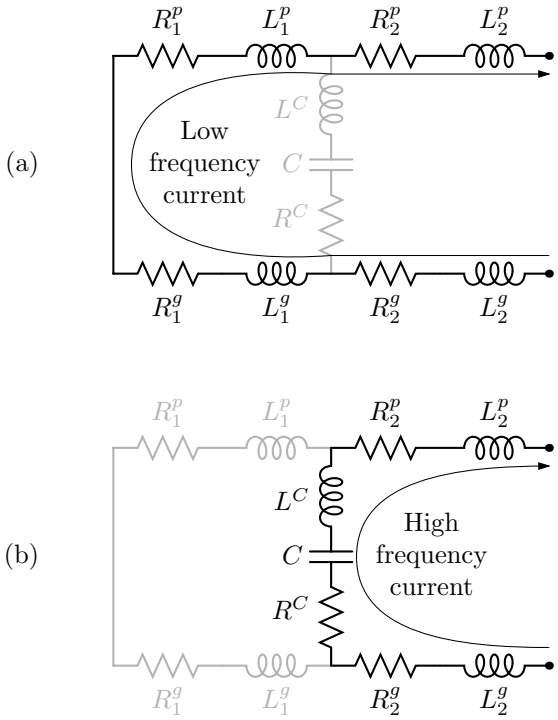

Fig. 5.9. The path of current flow in a power distribution system with a decoupling capacitor. (a) At frequencies below $f_{\text{res}} = \frac{1}{2\pi}\frac{1}{\sqrt{L_1 C}}$, the capacitive impedance is relatively high and the current loop extends throughout an entire network. (b) At frequencies above f_{res}, the capacitive impedance is lower than the impedance of the upstream section and the bulk of the current bypasses the upstream inductance $L_1^p + L_1^g$ through the decoupling capacitor C.

network impedance increases as

$$Z(\omega) = R_2 + R^C + j\omega(L_2 + L^C). \tag{5.10}$$

A comparison of (5.10) with (5.2) indicates that placing a decoupling capacitor reduces the high frequency inductance of the power distribution network, as seen by the load, from $L_{\text{tot}} = L_1 + L_2$ to $L_2' = L_2 + L^C$.

The output resistance R_2' and inductance L_2' as seen from the load are henceforth referred to as the *effective* resistance and inductance of the power distribution system to distinguish these quantities from the *overall series* resistance R_{tot} and inductance L_{tot} of the system (which are the effective resistance and inductance at DC). The shunting capacitor "decouples" the upstream portion of the power distribution

system from the power current loop at high frequencies, decreasing the maximum impedance of the system. The frequency range where the impedance specification is satisfied is correspondingly increased to $f_{\max} = \frac{1}{2\pi}\frac{Z_0}{L_2'}$.

5.4.2 Limitations of a single-tier decoupling scheme

It follows from the preceding discussion that the impedance of a power distribution system can be significantly reduced by a single capacitor (or a group of capacitors placed at the same location) over a wide range of frequencies. This solution is henceforth referred to as *single-tier decoupling*. A single-tier decoupling scheme, however, is difficult to realize in practice as this solution imposes stringent requirements on the performance characteristics of the decoupling capacitor, as discussed below.

To confine the network impedance within the specification boundaries at the highest required frequencies, *i.e.*, the $2\pi f_0(L_2 + L^C) < Z_0$, a low inductance path is required between the decoupling capacitance and the load,

$$L_2 + L^C < \frac{1}{2\pi}\frac{Z_0}{f_0}. \tag{5.11}$$

Simultaneously, the capacitance should be sufficiently high to bypass the power current at sufficiently low frequencies, thereby preventing the violation of target specifications above $f_{\max}$ due to a high inductive impedance $j\omega(L_1 + L_2)$,

$$\frac{1}{2\pi f_{\max} C} < 2\pi f_{\max}(L_1 + L_2). \tag{5.12}$$

Condition (5.12) is, however, insufficient. As mentioned above, the decoupling capacitor C and the inductance of the upstream section L_1 form a resonant tank circuit. The impedance reaches the maximum at the resonant frequency $\omega_{\text{res}} \approx \frac{1}{\sqrt{L_1 C}}$, where the inductive impedance of the tank circuit complements the capacitive impedance, $j\omega_{\text{res}} L_1 = -\frac{1}{j\omega_{\max} C}$. Where the quality factor of the tank circuit Q_{tank} is sufficiently larger than unity (*i.e.*, $Q_{\text{tank}} \gtrsim 3$), the impedance of the tank circuit at the resonant frequency ω_{res} is purely resistive and is larger than the characteristic impedance $\sqrt{\frac{L_1}{C}} = \omega_{\text{res}} L_1$ by the quality factor Q_{tank}, $Z_{\text{tank}}(\omega_{\text{res}}) = Q_{\text{tank}}\sqrt{\frac{L_1}{C}}$. The quality factor of the tank circuit is

$$Q_{\text{tank}} = \frac{1}{R_1 + R^C} \sqrt{\frac{L_1}{C}}\,, \tag{5.13}$$

yielding the magnitude of the peak impedance,

$$Z_{\text{peak}} = \frac{1}{R_1 + R^C} \frac{L_1}{C}\,. \tag{5.14}$$

To limit the impedance magnitude below the target impedance Z_0, the decoupling capacitance should satisfy

$$C > \frac{L_1}{Z_0(R_1 + R^C)}\,. \tag{5.15}$$

A larger upstream inductance L_1 (*i.e.*, the inductance that is decoupled by the capacitor) therefore requires a larger decoupling capacitance C to maintain the target network impedance Z_0. The accuracy of the simplifications used in the derivation of (5.14) decreases as the factor Q_{tank} approaches unity. A detailed treatment of the case where $Q \approx 1$ is presented in Section 5.6.

These impedance characteristics have an intuitive physical interpretation. From a physical perspective, the decoupling capacitor serves as an intermediate storage of charge and energy. To be effective, such an energy storage device should possess two qualities. First, the device should have a high capacity to store a sufficient amount of energy. This requirement is expressed in terms of the impedance characteristics as the minimum capacitance constraint (5.15). Second, to supply sufficient power at high frequencies, the device should be able to release and accumulate energy at a sufficient rate. This quality is expressed as the maximum inductance constraint (5.11).

Constructing a device with both high energy capacity and high power capability is, however, challenging. The conditions of low inductance (5.11) and high capacitance (5.15) cannot be simultaneously satisfied in a cost effective manner. In practice, these conditions are contradictory. The physical realization of a large decoupling capacitance as determined by (5.15) requires the use of discrete capacitors with a large nominal capacity, which, consequently, have a large form factor. The large physical dimensions of the capacitors have two implications. The parasitic series inductance of a physically large capacitor L^C is relatively high due to the increased area of the current loop within the capacitors, contradicting requirement (5.11). Furthermore, the large physical size of the capacitors prevents placing the capacitors sufficiently close to the power load. Greater physical separation

increases the inductance L_2 of the current path from the capacitors to the load, also contradicting requirement (5.11). The available component technology therefore imposes a tradeoff between the high capacity and low parasitic inductance of a capacitor.

Gate switching times of a few tens of picoseconds are common in contemporary integrated circuits, creating high transients in the power load current. Only on-chip decoupling capacitors have a sufficiently low parasitic inductance to maintain a low impedance power distribution system at high frequencies. Placing a sufficiently large on-chip decoupling capacitance, as determined by (5.15), requires a die area many times greater than the area of the load circuit. Therefore, while technically feasible, the single-tier decoupling solution is prohibitively expensive. A more efficient approach to the problem, a multi-stage hierarchical placement of decoupling capacitors, is described in the following section.

5.5 Hierarchical placement of decoupling capacitance

Low impedance, high frequency power distribution systems are realized in a cost effective way by using a hierarchy of decoupling capacitors. The capacitors are placed in several stages: on the board, package, and circuit die. The impedance characteristics of a power distribution system with several stages of decoupling capacitors are described in this section. The evolution of the system output impedance is described as the decoupling stages are consecutively placed across the network. Arranging the decoupling capacitors in several stages eliminates the need to satisfy both the high capacitance and low inductance requirements, as expressed by (5.11) and (5.15), in the same decoupling capacitor stage.

Board decoupling capacitors

Consider a power distribution system with decoupling capacitors placed on the board, as shown in Fig. 5.10. The circuit is analogous to the network shown in Fig. 5.7. Similar to the single-tier decoupling scheme, the board decoupling capacitance should satisfy the following condition in order to meet the target specification at low frequencies,

$$C_{\mathrm{b}} > \frac{L_{\mathrm{r}}}{Z_0(R_{\mathrm{r}} + R_{\mathrm{b}}^C)} \,. \tag{5.16}$$

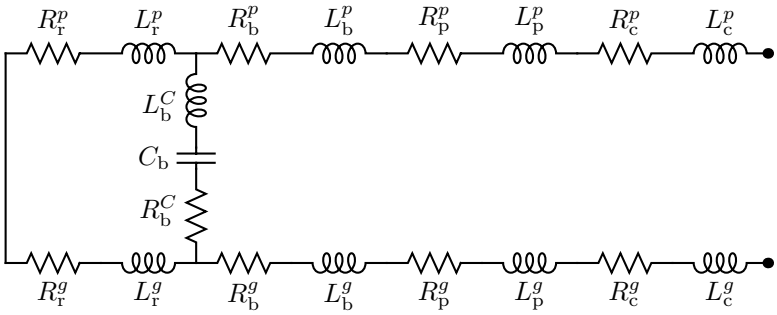

Fig. 5.10. A circuit model of a power distribution system with a board decoupling capacitance.

$L_{\rm r} = L_{\rm r}^p + L_{\rm r}^g$ and $R_{\rm r} = R_{\rm r}^p + R_{\rm r}^g$ are the resistance and inductance, respectively, of the network upstream of the board capacitance, including the output inductance of the voltage regulator and the impedance of the current path between the voltage regulator and the board capacitors, as defined by (5.1). At frequencies greater than $f_{\rm res}^{\rm B}$, the power current flows through the board decoupling capacitors, bypassing the voltage regulator. Condition (5.11), however, is not satisfied due to the relatively high inductance $L_b + L_p + L_c$ resulting from the large separation from the load and the high parasitic series inductance of the board capacitors $L_{\rm b}^C$. Above the frequency $f_{R_{\rm B}} = \frac{1}{2\pi}\frac{1}{R_{\rm B}C_{\rm b}}$, the impedance of the power distribution system is

$$Z_{\rm B} = R_{\rm B} + j\omega L_{\rm B}, \tag{5.17}$$

where $L_{\rm B} = L_{\rm b}^C + L_{\rm b} + L_{\rm p} + L_{\rm c}$ and $R_{\rm B} = R_{\rm b}^C + R_{\rm b} + R_{\rm p} + R_{\rm c}$. Although the high frequency inductance of the network is reduced from $L_{\rm tot}$ to $L_{\rm B}$, the impedance exceeds the target magnitude Z_0 above the frequency $f_{\rm max}^{\rm B} = \frac{1}{2\pi}\frac{Z_0}{L_{\rm B}}$, as shown in Fig. 5.11.

Package decoupling capacitors

The excessive high frequency inductance $L_{\rm B}$ of the power distribution system with board decoupling capacitors is further reduced by placing an additional decoupling capacitance physically closer to the power load, *i.e.*, on the integrated circuit package. The circuit model of a power distribution system with board and package decoupling capacitors is shown in Fig. 5.12. The impedance of the network with board and package decoupling capacitors is illustrated in Fig. 5.13. The package decoupling capacitance decreases the network inductance $L_{\rm B}$ in

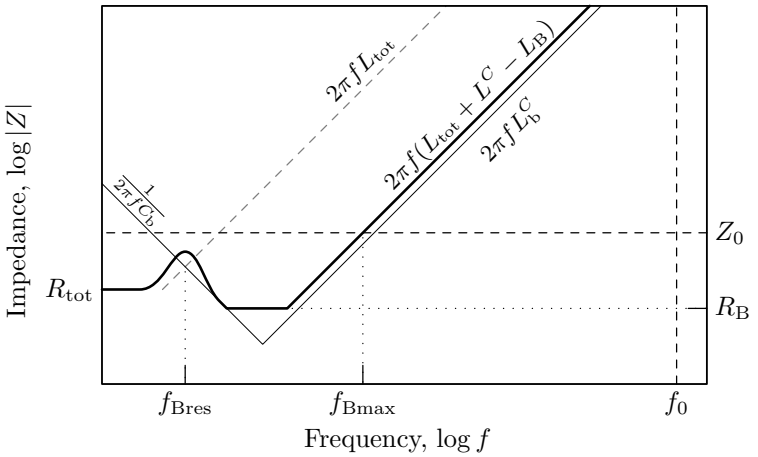

Fig. 5.11. Impedance of the power distribution system with the board decoupling capacitance shown in Fig. 5.7. The impedance characteristic is shown with a black line. The impedance of the power distribution system with no decoupling capacitors is shown, for comparison, with a gray dashed line. The impedance of the board decoupling capacitance is shown with a thin solid line.

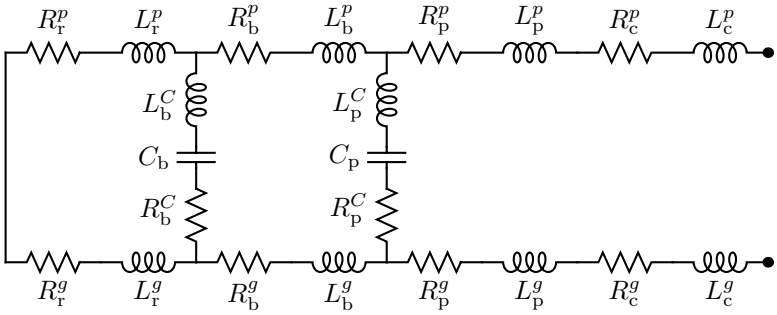

Fig. 5.12. A circuit model of a power distribution system with board and package decoupling capacitances.

a fashion similar to the board decoupling capacitance. A significant difference is that for the package decoupling capacitance the capacity requirement (5.15) is relaxed to

$$C_{\mathrm{p}} > \frac{L_{\mathrm{b}}^{'}}{Z_0(R_{\mathrm{b}}^{'} + R_{\mathrm{p}}^{C})} , \tag{5.18}$$

where $L_{\mathrm{b}}^{'} = L_{\mathrm{b}} + L_{\mathrm{b}}^{C}$ and $R_{\mathrm{b}}^{'} = R_{\mathrm{b}} + R_{\mathrm{b}}^{C}$ are the effective inductance and resistance in parallel with the package decoupling capacitance. The

upstream inductance $L_{\mathrm{b}}^{'}$ as seen by the package capacitance at high frequencies (*i.e.*, $f > f_{R_{\mathrm{B}}}$) is significantly lower than the upstream inductance L_{r} as seen by the board capacitors. The package capacitance requirement (5.18) is therefore significantly less stringent than the board capacitance requirement (5.16). The lower capacitance require-

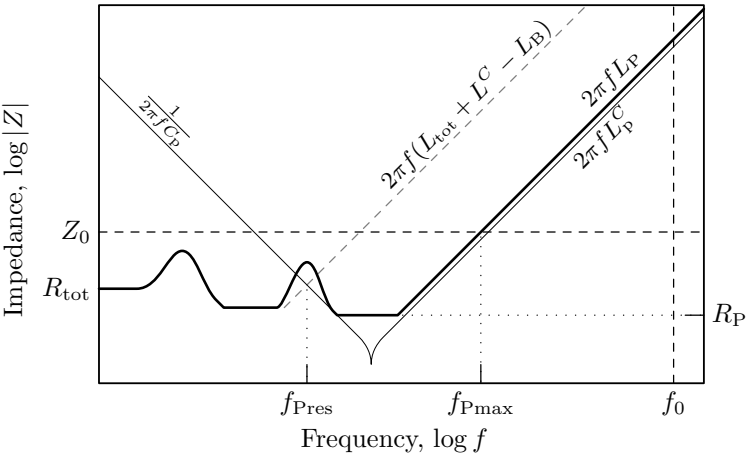

Fig. 5.13. Impedance of the power distribution system with board and package decoupling capacitances as shown in Fig. 5.7. The impedance characteristic is shown with a black line. The impedance of the power distribution system with only the board decoupling capacitance is shown with a dashed gray line for comparison. The impedance of the package decoupling capacitance is shown with a thin solid line.

ment is satisfied by using several small form factor capacitors mounted onto a package. As shown in Fig. 5.13, the effect of the package decoupling capacitors on the system impedance is analogous to the effect of the board capacitors, but occurs at a higher frequency. With package decoupling capacitors, the impedance of the power distribution system above a frequency $f_{R_{\mathrm{P}}} = \frac{1}{2\pi}\frac{1}{R_{\mathrm{P}}C_{\mathrm{p}}}$ becomes

$$Z_{\mathrm{P}} = R_{\mathrm{P}} + j\omega L_{\mathrm{P}}, \tag{5.19}$$

where $L_{\mathrm{P}} = L_{\mathrm{p}}^{C} + L_{\mathrm{p}} + L_{\mathrm{c}}$ and $R_{\mathrm{P}} = R_{\mathrm{p}}^{C} + R_{\mathrm{p}} + R_{\mathrm{c}}$. Including a package capacitance therefore lowers the high frequency inductance from L_{B} to L_{P}. The reduced inductance L_{P}, however, is not sufficiently low to satisfy (5.11) in high speed circuits. The impedance of the power system exceeds the target magnitude Z_0 at frequencies above $f_{\mathrm{max}}^{\mathrm{P}} = \frac{1}{2\pi}\frac{Z_0}{L_{\mathrm{P}}}$.

On-chip decoupling capacitors

On-chip decoupling capacitors are added to the power distribution system in order to extend the frequency range of the low impedance characteristics to f_0. A circuit model characterizing the impedance of a power distribution system with board, package, and on-chip decoupling capacitors is shown in Fig. 5.14. The impedance characteristics of a

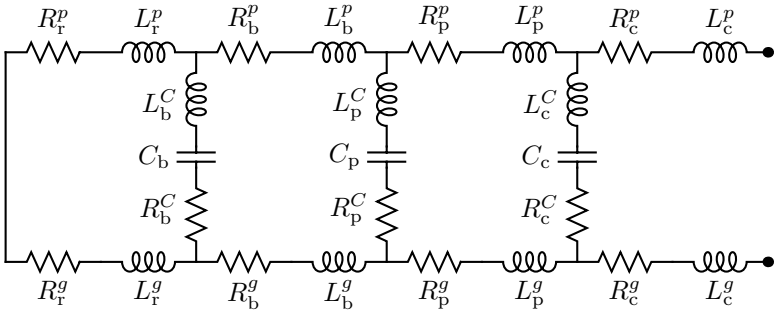

Fig. 5.14. A circuit network characterizing the impedance of a power distribution system with board, package, and on-chip decoupling capacitances.

power distribution system with all three types of decoupling capacitors are illustrated in Fig. 5.15. To ensure that the network impedance does not exceed Z_0 due to the inductance L_{P} of the current loop terminated by the package capacitance, the on-chip decoupling capacitance should satisfy

$$C_{\mathrm{c}} > \frac{L_{\mathrm{p}}^{'}}{Z_0(R_{\mathrm{p}}^{'} + R_{\mathrm{c}}^{C})} , \tag{5.20}$$

where $L_{\mathrm{p}}^{'} = L_{\mathrm{p}} + L_{\mathrm{p}}^{C}$ is the effective inductance in parallel with the package decoupling capacitance. The capacity requirement for the on-chip capacitance is further reduced as compared to the package capacitance due to lower effective inductances, $L_{\mathrm{p}} < L_{\mathrm{b}}$ and $L_{\mathrm{p}}^{'} < L_{\mathrm{b}}^{'}$.

Advantages of hierarchical decoupling

Hierarchical placement of decoupling capacitors exploits the trade-off between the capacity and the parasitic series inductance of a capacitor to achieve an economically effective solution. The total decoupling capacitance of a hierarchical scheme $C_{\mathrm{b}} + C_{\mathrm{p}} + C_{\mathrm{c}}$ is larger than the total decoupling capacitance of a single-tier solution C as determined

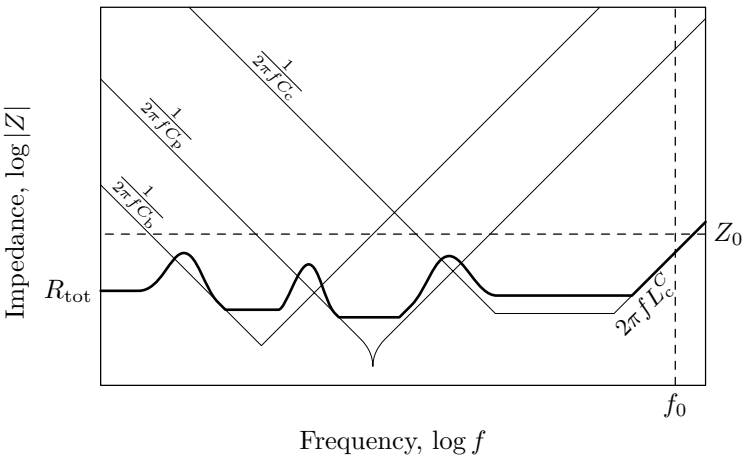

Fig. 5.15. Impedance of a power distribution system with board, package, and on-chip decoupling capacitances, as shown in Fig. 5.7. The impedance is shown with a black line. The impedance characteristic is within the target specification outlined by the dashed lines. The impedance characteristics of the decoupling capacitances are shown with thin solid lines.

by (5.15). The advantage of the hierarchical scheme is that the inductance limit (5.11) is imposed only on the final stage of decoupling capacitors which constitute a small fraction of the total decoupling capacitance. The constraints on the physical dimensions and parasitic impedance of the capacitors in the remaining stages are dramatically reduced, permitting the use of cost efficient electrolytic and ceramic capacitors.

The decoupling capacitance at each tier is effective within a limited frequency range, as determined by the capacitance and inductance of the capacitor. The range of effectiveness of the board, package, and on-chip decoupling capacitances overlaps each other, as shown in Fig. 5.15, spanning an entire frequency region of interest from DC to f_0 (the maximum operating frequency).

As described in Section 5.4, a decoupling capacitor lowers the high frequency impedance by allowing the power current to bypass the inductive interconnect upstream of the capacitor. In a power distribution system with several stages of decoupling capacitors, the inductive interconnect is excluded from the power current loop in several steps, as illustrated in Fig. 5.16. At near-DC frequencies, the power current loop extends through an entire power supply system to the power source,

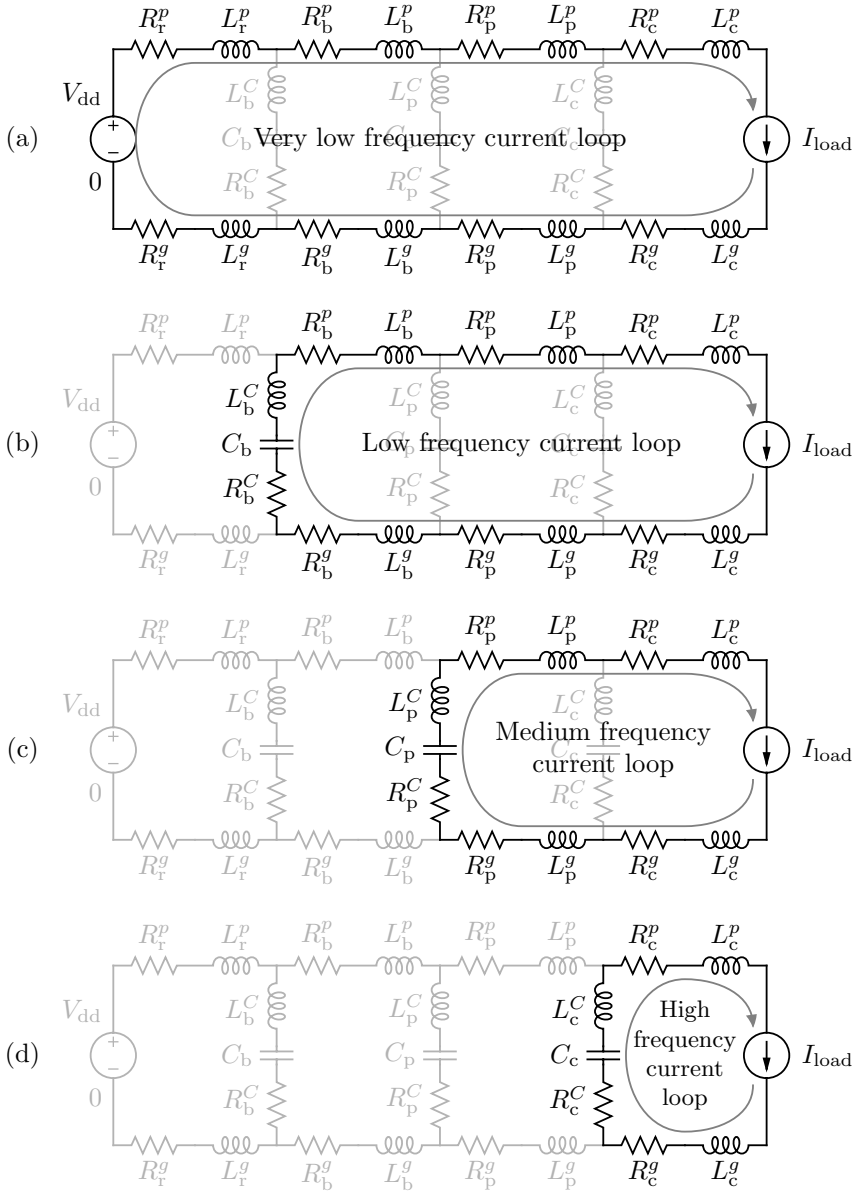

Fig. 5.16. Variation of the power current path with frequency. (a) At low frequencies the current loop extends through an entire system to the power source; as frequencies increase, the current loop is terminated by the board, package, and on-chip decoupling capacitors, as shown in (b), (c), and (d), respectively.

as shown in Fig. 5.16(a). As frequencies increase, the power current is shunted by the board decoupling capacitors, shortening the current loop as shown in Fig. 5.16(b). At higher frequencies, the package capacitors terminate the current loop, further reducing the loop size, as depicted in Fig. 5.16(c). Finally, at the highest frequencies, the power current is terminated by the on-chip capacitors, as illustrated in Fig. 5.16(d). The higher the frequency, the shorter the distance between the shunting capacitor and the load and the smaller the size of the current loop. At transitional frequencies where the bulk of the current is shifted from one decoupling stage to the next, both decoupling stages carry significant current, giving rise to resonant behavior.

Each stage of decoupling capacitors determines the impedance characteristics over a limited range of frequencies, where the bulk of the power current flows through the stage. Outside of this frequency range, the stage capacitors have an insignificant influence on the impedance characteristics of the system. The lower frequency impedance characteristics are therefore determined by the upstream stages of decoupling capacitors, while the impedance characteristics at the higher frequencies are determined by the capacitors closer to the load. For example, the board capacitors determine the low frequency impedance characteristics but do not affect the high frequency response of the system, which is determined by the on-chip capacitors.

5.6 Resonance in power distribution networks

Using decoupling capacitors is a powerful technique to reduce the impedance of a power distribution system within a significant range of frequencies, as discussed in preceding sections. A decoupling capacitor, however, *increases* the network impedance in the vicinity of the resonant frequency f_{res}. Controlling the impedance behavior near the resonant frequency is therefore essential to effectively use decoupling capacitors. A relatively high quality factor, $Q \gtrsim 3$, is assumed in the expression for the peak impedance of the parallel resonant circuit described in Section 5.4. Maintaining low impedance characteristics in power distribution systems necessitates minimizing the quality factor of all of the resonant modes; relatively low values of the quality factor are therefore common in power distribution systems. The impedance characteristics of parallel resonant circuits with a low quality factor, $Q \approx 1$, are the focus of this section.

The purpose of the decoupling capacitance is to exclude the high impedance upstream network from the load current path, as discussed in Section 5.4. The decoupling capacitor and the inductive upstream network form a parallel resonant circuit with a significant resistance in both the inductive and capacitive branches. An equivalent circuit diagram is shown in Fig. 5.17.

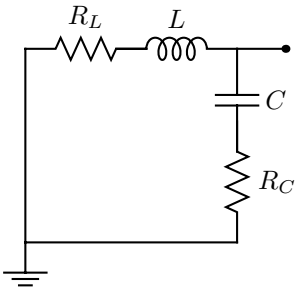

Fig. 5.17. A parallel resonant circuit with a significant parasitic resistance in both branches. L and R_L represent the effective impedance of the upstream network; C and R_C represent the impedance characteristics of the decoupling capacitor.

Near the resonant frequency, the impedance of the upstream network exhibits an inductive-resistive nature, $R_L + j\omega L$. The impedance of the decoupling capacitor is well described near the resonant frequency as $R_C + j\omega C$. The effective series inductance of the capacitor is significantly lower than the upstream network inductance (otherwise the capacitor would be ineffective, as discussed in Section 5.4) and can be typically neglected near the resonant frequency. The impedance characteristics of the capacitor and upstream network are schematically illustrated in Fig. 5.18.

The impedance of the tank circuit shown in Fig. 5.17 is

$$\begin{aligned} Z_{\text{tank}} &= (R_L + j\omega L) \,\Big\|\, \left(R_C + \frac{1}{j\omega C}\right) \\ &= \frac{(R_L + j\omega L)\left(R_C + \frac{1}{j\omega C}\right)}{(R_L + j\omega L) + \left(R_C + \frac{1}{j\omega C}\right)} \\ &= R_C + \frac{s(L - CR_C^2) + (R_L - R_C)}{s^2 LC + sC(R_C + R_L) + 1}\,. \end{aligned} \tag{5.21}$$

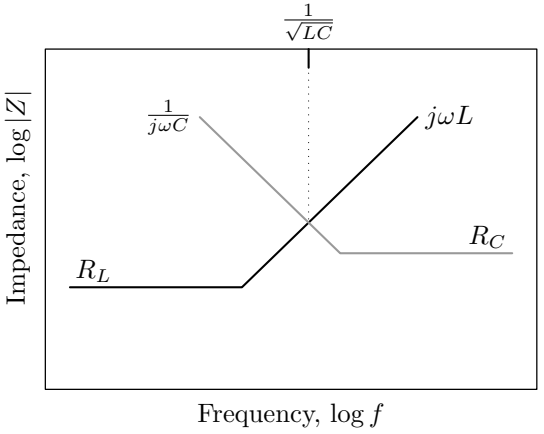

Fig. 5.18. An asymptotic plot of the impedance magnitude of the inductive and capacitive branches forming the tank circuit shown in Fig. 5.17.

The poles of the system described by (5.21) are determined by the characteristic equation,

$$s^2 + 2\zeta\omega_n + \omega_n^2 = 0, \tag{5.22}$$

$$\text{where} \qquad \omega_n = \frac{1}{\sqrt{LC}} \tag{5.23}$$

$$\text{and} \qquad \zeta = \frac{R_L + R_C}{2}\sqrt{\frac{C}{L}}\,. \tag{5.24}$$

Note that the magnitude of the circuit impedance varies from R_L at low frequencies ($\omega \ll \frac{1}{\sqrt{LC}}$) to R_C at high frequencies ($\omega \gg \frac{1}{\sqrt{LC}}$). It is desirable that the impedance variations near the resonant frequency do not significantly increase the maximum impedance of the network, *i.e.*, the variations do not exceed or only marginally exceed $\max(R_L, R_C)$, the maximum resistive impedance. Different constraints can be imposed on the parameters of the tank circuit to ensure this behavior [109]. Several cases of these constraints are described below.

Case I: The tank circuit has a monotonic (*i.e.*, overshoot- and oscillation-free) response to a step current excitation if the circuit damping is critical or greater: $\zeta \leqslant 1$. Using (5.24), sufficient damping is achieved when

$$R_L + R_C \geqslant 2R_0 = 2\sqrt{\frac{L}{C}}\,, \tag{5.25}$$

where $R_0 = \sqrt{\frac{L}{C}}$ is the characteristic impedance of the tank circuit. Therefore, if the network impedance $R_L + j\omega L$ is to be reduced to the target impedance Z_0 at high frequencies (meaning that $R_C = Z_0$), the required decoupling capacitance is

$$C \geqslant \frac{L}{Z_0^2} \frac{4}{(1 + R_L/Z_0)^2} . \tag{5.26}$$

Case II: Alternatively, the monotonicity of the impedance variation with frequency can be ensured. It can be demonstrated [109] that the magnitude of the impedance of the tank circuit varies monotonically from R_L at low frequencies, $\omega \ll \omega_n$, to R_C at high frequencies, $\omega \gg \omega_n$, if

$$R_L R_C \geqslant R_0^2 = \frac{L}{C} . \tag{5.27}$$

The required decoupling capacitance in this case is

$$C \geqslant \frac{L}{Z_0^2} \frac{1}{R_L/Z_0} . \tag{5.28}$$

Note that condition (5.27) is stronger than condition (5.25), *i.e.*, satisfaction of (5.27) ensures satisfaction of (5.25). Correspondingly, the capacitance requirement described by (5.28) is always greater or equal to the capacitance requirement described by (5.26).

Case III: The capacitance requirement determined in Case II, as described by (5.28), increases rapidly when either R_L or R_C is small. In this situation, condition (5.28) is overly conservative and restrictive. The capacitive requirement is significantly relaxed if a small peak in the impedance characteristics is allowed. To restrict the magnitude of the impedance peak to approximately 1% (in terms of the resistive baseline), the following condition must be satisfied [109],

$$\left(b_2 r^2 + b_1 r + b_0\right) R_{\max}^2 \geqslant R_0^2 = \frac{L}{C} , \tag{5.29}$$

where $R_{\max} = \max(R_L, R_C)$, $r = \min\left(\frac{R_L}{R_C}, \frac{R_C}{R_L}\right)$, $b_0 = 0.4831$, $b_1 = 0.4907$, and $b_2 = -0.0139$. The decoupling capacitance requirement is relaxed in this case to

$$C \geqslant \frac{L}{Z_0^2} \frac{1}{b_2 r^2 + b_1 r + b_0} . \tag{5.30}$$

Case III is the least constrained as compared to Cases I and II.

The design space of the resistances R_L and R_C as determined by (5.25), (5.27), and (5.29) is illustrated in Fig. 5.19 in terms of the circuit characteristic impedance $R_0 = \sqrt{\frac{L}{C}}$. In the unshaded region near the origin, the impedance exhibits significant resonant behavior. The rest of the space is partitioned into three regions. The more constrained regions are shaded in progressively darker tones. In the lightest shaded region, only the Case III condition is satisfied. The Case I and III conditions are satisfied in the adjacent darker region. All three conditions are satisfied in the darkest region.

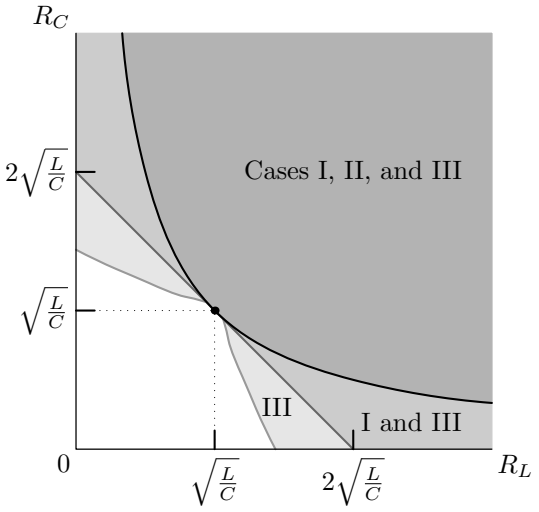

Fig. 5.19. Design space of the resistances R_L and R_C for the tank circuit shown in Fig. 5.17. All three design constraints, as determined by (5.25), (5.27), and (5.29), are satisfied in the darkest area. The constraints of Cases I and III are satisfied in the medium gray area, while only the Case III constraint is satisfied in the lightly shaded area.

Alternatively, conditions (5.25), (5.27), and (5.29) can be used to determine the decoupling capacitance requirement as a function of the circuit inductance L and resistances R_L and R_C. The normalized capacitance $C\frac{R_L^2}{L}$ versus the normalized resistance R_C/R_L is shown in Fig. 5.20. The shading is analogous to the scheme used in Fig. 5.19. The darkest region boundary is determined by condition (5.28), the boundary of the intermediate region is determined by condition (5.26), and the boundary of the lightest region is determined by condition

(5.30). Similar to the resistance design space illustrated in Fig. 5.19, Case II is the most restrictive, while Case III is the least restrictive. The capacitance requirements decrease in all three cases as the normalized resistance $\frac{R_C}{R_L}$ increases.

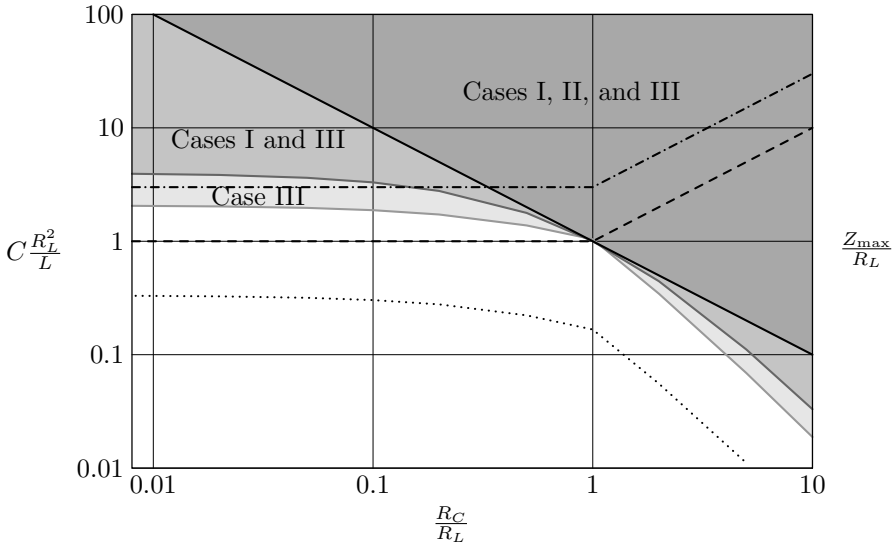

Fig. 5.20. Decoupling capacitance requirements as determined by Cases I, II, and III, *i.e.*, by (5.26), (5.28), and (5.30), respectively. Case II requires the greatest amount of capacitance while Case III requires the lowest capacitance.

The normalized maximum impedance of the network $Z_0 = \frac{R_{\max}}{R_L} = \frac{\max(R_L,R_C)}{R_L}$ is also shown in Fig. 5.20 by a dashed line. For $\frac{R_C}{R_L} \leqslant 1$, the maximum network impedance is $Z_0 = R_L \left(\frac{Z_0}{R_L} = 1\right)$. Increasing resistance R_C to R_L therefore reduces the required decoupling capacitance without increasing the maximum impedance of the network. Increasing R_C beyond R_L further reduces the capacitance requirement, but at a cost of increasing the maximum impedance. For $\frac{R_C}{R_L} \geqslant 1$, the maximum impedance becomes $Z_0 = R_C \left(\frac{Z_0}{R_L} = \frac{R_C}{R_L}\right)$.

The requirement for the decoupling capacitance can be further reduced if the resonant impedance is designed to significantly exceed $R_{\max}$. For comparison, the capacitance requirement as determined by (5.15) for the case $Z_0 = 3R_{\max}$ ($Q \approx 3$) is also shown in Fig. 5.20 by the dotted line. This significantly lower requirement, as compared to

Cases I, II, and III, however, comes at a cost of a higher maximum impedance, as shown by the dashed line with dots.

5.7 Full impedance compensation

A special case of both (5.26) and (5.28) is the condition where

$$R_L = R_C = R_0 = \sqrt{\frac{L}{C}}\,, \tag{5.31}$$

as shown in Figures 5.19 and 5.20. Under condition (5.31), the zeros of the tank circuit impedance (5.21) cancel the poles and the impedance becomes purely resistive and independent of frequency,

$$Z_{\text{tank}}(\omega) = R_0. \tag{5.32}$$

This specific case is henceforth referred to as fully compensated impedance. As shown in Fig. 5.20, choosing the capacitive branch resistance R_C in accordance with (5.31) results in the lowest decoupling capacitance requirements for a given maximum circuit impedance.

The impedance of the inductive and capacitive branches of the tank circuit under condition (5.31) is illustrated in Fig. 5.21. The condition of full compensation (5.31) is equivalent to two conditions: $R_L = R_C$, *i.e.*, the impedance at the lower frequencies is matched to the impedance at the higher frequencies, and $\frac{L}{R_L} = R_C C$, *i.e.*, the time constants of the inductor and capacitor currents are matched, as shown in Fig. 5.21.

A constant, purely resistive impedance is achieved across the entire frequency range of interest, as illustrated in Fig. 5.22, if each decoupling stage is fully compensated. The resistance and capacitance of the decoupling capacitors in a fully compensated system are completely determined by the impedance characteristics of the power and ground interconnect and the location of the capacitors. The overall capacitance and resistance of the board decoupling capacitors are, according to (5.31),

$$R_{\text{b}}^{C} = R_{\text{r}}, \tag{5.33}$$

$$C_{\text{b}} = \frac{L_{\text{r}}}{R_{\text{r}}^2}\,. \tag{5.34}$$

The inductance of the upstream part of the power distribution system as seen at the terminals of the package capacitors is $L_{\text{b}}^{C} + L_{\text{b}}$. The capacitance and resistance of the package capacitors are

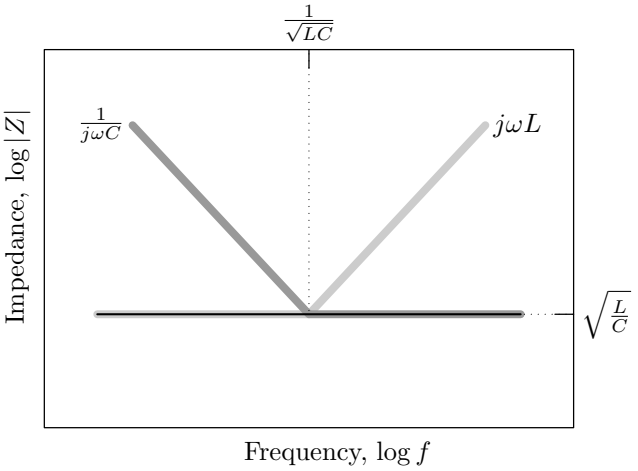

Fig. 5.21. Asymptotic impedance of the inductive (light gray line) and capacitive (dark gray line) branches of a tank circuit under the condition of full compensation (5.31). The overall impedance of the tank circuit is purely resistive and does not vary with frequency, as shown by the black line.

$$R_{\mathrm{p}}^{C} = R_{\mathrm{b}}^{C} + R_{\mathrm{b}} = R_{\mathrm{r}} + R_{\mathrm{b}}, \tag{5.35}$$

$$C_{\mathrm{p}} = \frac{L_{\mathrm{b}}^{C} + L_{\mathrm{b}}}{(R_{\mathrm{r}} + R_{\mathrm{b}})^2}\,. \tag{5.36}$$

Analogously, the resistance and capacitance of the on-chip decoupling capacitors are

$$R_{\mathrm{c}}^{C} = R_{\mathrm{p}}^{C} + R_{\mathrm{p}} = R_{\mathrm{r}} + R_{\mathrm{b}} + R_{\mathrm{p}}, \tag{5.37}$$

$$C_{\mathrm{c}} = \frac{L_{\mathrm{p}}^{C} + L_{\mathrm{p}}}{(R_{\mathrm{r}} + R_{\mathrm{b}} + R_{\mathrm{p}})^2}\,. \tag{5.38}$$

Note that the effective series resistance of the decoupling capacitors increases toward the load, $R_{\mathrm{b}}^{C} < R_{\mathrm{p}}^{C} < R_{\mathrm{c}}^{C}$, such that the resistance of the load current loop, no matter which decoupling capacitor terminates this loop, remains the same as the total resistance R_{tot} of the system without decoupling capacitors,

$$\begin{aligned} R_{\mathrm{b}}^{C} + R_{\mathrm{b}} + R_{\mathrm{p}} + R_{\mathrm{c}} &= \\ R_{\mathrm{p}}^{C} + R_{\mathrm{p}} + R_{\mathrm{c}} &= \\ R_{\mathrm{c}}^{C} + R_{\mathrm{c}} &= R_{\mathrm{tot}}. \end{aligned} \tag{5.39}$$

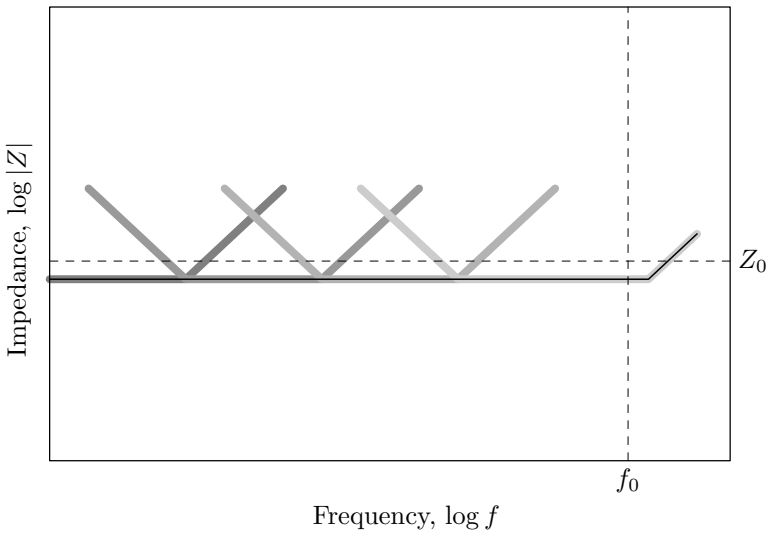

Fig. 5.22. Impedance diagram of a power distribution system with full compensation at each decoupling stage. The impedance of the decoupling stages is shown in shades of gray. The resulting system impedance is purely resistive and constant over the frequency range of interest, as shown by the black line.

If the resistance of the decoupling capacitors is reduced below the values determined by (5.33), (5.35), and (5.37), the resonance behavior degrades the impedance characteristics of the power distribution system.

5.8 Case study

A case study of a power distribution system for a high performance integrated circuit is presented in this section. This case study is intended to provide a practical perspective to the analytic description of the impedance characteristics of the power distribution systems developed in the previous section.

Consider a microprocessor consuming 60 watts from a 1.2 volt power supply. The average power current of the microprocessor is 50 amperes. The maximum frequency of interest is assumed to be 20 GHz, corresponding to the shortest gate switching time of approximately 17 ps. The objective is to limit the power supply variation to approximately 8% of the nominal 1.2 volt power supply level under a 50 ampere load. This objective results in a target impedance specification of

Table 5.1. Parameters of a case study power distribution system

Circuit parameter	Initial system	Near-critical damping	Fully compensated
R_r	1 mΩ	1 mΩ	1 mΩ
L_r	10 nH	10 nH	10 nH
C_b	5 mF	5 mF	10 mF
R_b^C	0.1 mΩ	1 mΩ	1 mΩ
L_b^C	0.3 nH	0.3 nH	0.3 nH
R_b	0.3 mΩ	0.3 mΩ	0.3 mΩ
L_b	0.2 nH	0.2 nH	0.2 nH
C_p	250 μF	250 μF	296 μF
R_p^C	0.2 mΩ	1.5 mΩ	1.3 mΩ
L_p^C	1 pH	1 pH	1 pH
R_p	0.1 mΩ	0.1 mΩ	0.1 mΩ
L_p	1 pH	1 pH	1 pH
C_c	500 nF	500 nF	1020 nF
R_c^C	0.4 mΩ	1.5 mΩ	1.4 mΩ
L_c^C	1 fH	1 fH	1 fH
R_c	0.05 mΩ	0.05 mΩ	0.05 mΩ
L_c	4 fH	4 fH	4 fH

0.08×1.2 volts/50 amperes = 2 milliohms over the frequency range from DC to 20 GHz. Three stages of decoupling capacitors are assumed. The parameters of the initial version of the power distribution system are displayed in Table 5.1.

The impedance characteristics of the overall power distribution system are shown in Fig. 5.23. The resonant modes of this system are significantly underdamped. The resulting impedance peaks exceed the target impedance specifications. The impedance characteristics improve significantly if the damping of the resonant mode is increased to the near-critical level. The greater damping can be achieved by only manipulating the effective series resistance of the decoupling capacitors. The impedance characteristics approach the target specifications, as shown in Fig. 5.23, as the resistance of the board, package, and on-chip capacitors is increased from initial values of 0.1, 0.2, and 0.4 milliohms to 1, 1.5, and 1.5 milliohms, respectively. Further improvements in the system impedance characteristics require increasing the decoupling capacitance. Fully compensating each decoupling stage renders the impedance purely resistive. The magnitude of the fully compensated

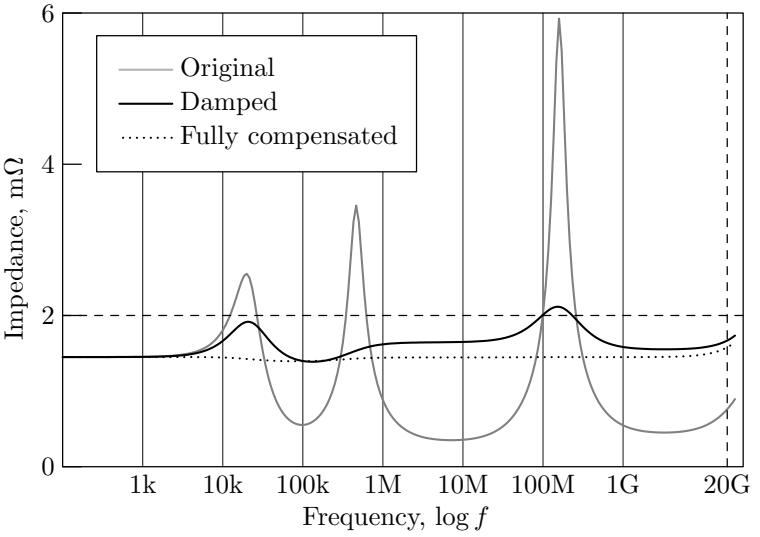

Fig. 5.23. Impedance characteristics of a power distribution system case study. The initial system is significantly underdamped. The resonant impedance peaks exceed the target impedance specification, as shown with a gray line. As the effective series resistance of the decoupling capacitors is increased, the damping of the resonant modes also increases, reducing the magnitude of the peak impedance. The damped system impedance effectively satisfies the target specifications, as shown with a solid black line. The impedance characteristics can be further improved if the impedance of each decoupling stage is fully compensated, as shown by the dotted line.

impedance equals the total resistance of the power distribution system, 1.45 milliohms $(1 + 0.3 + 0.1 + 0.05)$, as described in Section 5.7.

The resonant frequencies of the case study are representative of typical resonant frequencies encountered in power distribution systems. The board decoupling stage resonates in the kilohertz range of frequencies, while the frequency of the resonant peak due to the package decoupling stage is in the low megahertz frequencies [108], [110]. The frequency of the chip capacitor resonance, often referred to as the chip-package resonance as this effect includes the inductance of the package interconnect, typically varies from tens of megahertz to low hundreds of megahertz [105], [111], [112].

5.9 Design considerations

The power current requirements of high performance integrated circuits are rapidly growing with technology scaling. These requirements necessitate a significant reduction in the output impedance of the power distribution system over a wider range of frequencies for new generations of circuits. Design approaches to improve the impedance characteristics of next generation power distribution systems are discussed in this section.

As described in previous sections, the inductance of the power distribution interconnect structures is efficiently excluded from the path of high frequency current by the hierarchical placement of decoupling capacitors. The inductance of the power and ground interconnect, however, greatly affects the decoupling capacitance requirements. As indicated by (5.16), (5.18), and (5.20), a lower inductance current path connecting the individual stages of decoupling capacitors relaxes the requirements placed on the decoupling capacitance at each stage of the power distribution network. Lowering the interconnect inductance decreases both the overall cost of the decoupling capacitors and the effective impedance of the power distribution system. It is therefore desirable to reduce the inductance of the power distribution interconnect.

As indicated by (5.18) and (5.20), the lower bound on the capacitance at each decoupling stage is determined by the effective inductance of the upstream current path $L^C + L^{\text{int}}$, which consists of the inductance of the previous stage decoupling capacitors L^C and the inductance of the interconnect connecting the two stages L^{int}. The impedance characteristics are thereby improved by lowering both the effective series inductance of the decoupling capacitors, as described in Section 5.9.1, and the interconnect inductance, as described in Section 5.9.2.

5.9.1 Inductance of the decoupling capacitors

The series inductance of the decoupling capacitance can be decreased by using a larger number of lower capacity capacitors to realize a specific decoupling capacitance rather than using fewer capacitors of greater capacity. Assume that a specific type of decoupling capacitor is used to realize a decoupling capacitance C, as shown in Fig. 5.7. Each capacitor has a capacity C_1 and a series inductance L_1^C. Placing $N_1 = \frac{C}{C_1}$ capacitors in parallel realizes the desired capacitance C with

an effective series inductance $L^{C_1} = \frac{L_1^C}{N_1}$. Alternatively, using capacitors of lower capacity C_2 requires a larger number of capacitors $N_2 = \frac{C}{C_2}$ to realize the same capacitance C, but results in a lower overall inductance of the capacitor bank $L^{C_2} = \frac{L_2^C}{N_2}$. The efficacy of this approach is enhanced if the lower capacity component C_2 has a smaller form factor than the components of capacity C_1 and therefore has a significantly smaller series inductance L_2^C. Using a greater number of capacitors, however, requires additional board area and incurs higher component and assembly costs. Furthermore, the efficacy of the technique is diminished if the larger area required by the increased number of capacitors necessitates placing some of the capacitors at a greater distance from the load, increasing the inductance of the downstream current path L_2.

The series inductance of the decoupling capacitance L^C, however, constitutes only a portion of the overall inductance $L_2 + L^C$ of the current path between two stages of the decoupling capacitance, referring again to the circuit model shown in Fig. 5.7. Once the capacitor series inductance L^C is much lower than L_2, any further reduction of L^C has an insignificant effect on the overall impedance. The inductance of the current path between the decoupling capacitance and the load therefore imposes an upper limit on the frequency range of the capacitor efficiency. A reduction of the interconnect inductance is therefore necessary to improve the efficiency of the decoupling capacitors.

5.9.2 Interconnect inductance

Techniques for reducing interconnect inductance can be divided into *extensive* and *intensive* techniques. Extensive techniques lower inductance by using additional interconnect resources to form additional parallel current paths. Intensive techniques lower the inductance of existing current paths by modifying the structure of the power and ground interconnect so as to minimize the area of the current loop. The inductance of a power distribution system can be extensively decreased by allocating more metal layers on a printed circuit board and circuit package for power and ground distribution, increasing the number of package power and ground pins (solder balls in the case of ball grid array mounting) connecting the package to the board, and increasing the number of power and ground solder bumps or bonding wires connecting the die to the package. Extensive methods are often constrained by technological and cost considerations, such as the number and thickness

of the metal and isolation layers in a printed circuit board, the total number of pins in a specific package, and the die bonding technology. Choosing a solution with greater interconnect capacity typically incurs a significant cost penalty.

Intensive methods reduce the interconnect inductance without incurring higher manufacturing costs. These methods alter the interconnect structure in order to decrease the area of the power current loop. For example, the inductance of a current path formed by two square parallel power and ground planes is proportional to the separation of the planes[3] h,

$$L_{\text{plane}} = \mu_0 h. \tag{5.40}$$

Thus, parallel power and ground planes separated by 1 mil (25.4 μm) have an inductance of 32 pH per square. A smaller separation between the power and ground planes results in a proportionally smaller inductance of the power current loop. Reducing the thickness of the dielectric between the power and ground planes at the board and package levels is an effective means to reduce the impedance of the power distribution network [113], [114].

The same strategy of placing the power and ground paths in close proximity can be exploited to lower the effective inductance of the "vertical conductors," *i.e.*, the conductors connecting the parallel power and ground planes or planar networks. Examples of such connections are pin grid arrays and ball grid arrays connecting a package to a board, a flip-chip area array of solder balls connecting a die to a package, and an array of vias connecting the power and ground planes within a package. In regular pin and ball grid arrays, this strategy leads to a so-called checkerboard pattern [115], as shown in Fig. 5.24.

5.10 Limitations of the one-dimensional circuit model

The one-dimensional lumped circuit model shown in Fig. 5.3 has been used to describe the frequency dependent impedance characteristics of

[3] Equation (5.40) neglects the internal inductance of the metal planes, *i.e.*, the inductance associated with the magnetic flux within the planes. Omission of the internal inductance is a good approximation where the thickness of the planes is much smaller than the separation between the planes and at high signal frequencies, where current flow is restricted to the surface of the planes due to a pronounced skin (proximity) effect.

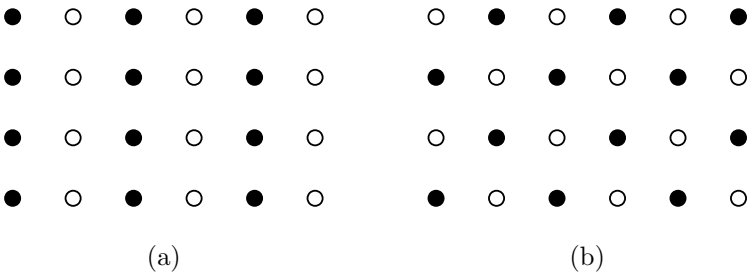

Fig. 5.24. Placement of area array connections for low inductance. The inductance of the area array interconnect strongly depends on the pattern of placement of the power and ground connections. The inductance of pattern (a) is relatively high, while pattern (b) has the lowest inductance. The power and ground connections are shown in black and white, respectively.

power distribution systems. This model captures the essential characteristics of power distribution systems over a wide range of frequencies. Due to the relative simplicity, the model is an effective vehicle for demonstrating the primary issues in the design of power distribution systems and related design challenges and tradeoffs. The use of this model for other purposes, however, is limited due to certain simplifications present in the model. Several of these limitations are summarized below.

A capacitor at each decoupling stage is modeled by a single *RLC* circuit. In practice, however, the decoupling capacitance on the board and package is realized by a large number of discrete capacitors. Each capacitor has a distinct physical location relative to the load. The parasitic inductance of the current path between a specific capacitor and the power load is somewhat different for each of the capacitors. The frequency of the resonant impedance peaks of an individual capacitor therefore varies depending upon the location of the capacitor. The overall impedance peak profile of a group of capacitors is therefore spread, resulting in a lower and wider resonant peak.

More significantly, it is common to use two different types of capacitors for board level decoupling. High capacity electrolytic capacitors, typically ranging from a few hundreds to a thousand microfarads, are used to obtain the high capacitance necessary to decouple the high inductive impedance of a voltage regulator. Electrolytic capacitors, however, have a relatively high series inductance of several nanohenrys. Ceramic capacitors with a lower capacity, typically tens of microfarads,

and a lower parasitic inductance, a nanohenry or less, are used to extend the frequency range of the low impedance region to higher frequencies. Effectively, there are two decoupling stages on the board, with a very small parasitic impedance between the two stages.

Using lumped circuit elements implies that the wavelength of the signals of interest is much larger than the physical dimensions of the circuit structure, permitting an accurate representation of the impedance characteristics with a few lumped elements connected in series. The current transitions at the power load are measured in tens of picoseconds, translating to thousands of micrometers of signal wavelength. This wavelength is much smaller than the size of a power distribution system, typically of several inches, making the use of a lumped model at first glance unjustified. The properties of power distribution systems, however, support a lumped circuit representation over a wider range of conditions as suggested by the aforementioned simple size criterion.

The design of a power distribution system is intended to restrict the flow of the power current as close to the load as possible. Due to the hierarchy of the decoupling capacitors, the higher the current frequency, the shorter the current loop, confining the current flow closer to the load. As seen from the terminals of the power source, a power distribution system is a multi-stage low pass filter, each stage having a progressively lower cut-off frequency. The spectral content of the current in the on-board power distribution system is limited to several megahertz. The corresponding wavelength is much larger than the typical system dimensions of a few inches, making a lumped model sufficiently accurate. The spectral content of the power current within the package network extends into the hundreds of megahertz frequency range. The signal wavelength remains sufficiently large to use a lumped model; however, a more detailed network may be required rather than a one-dimensional model to accurately characterize the network impedances.

A one-dimensional model is inadequate to describe an on-chip power distribution network. The on-chip power distribution network is the most challenging element of the power delivery system design problem. The board and package level power distribution networks consist of several metal planes and thousands of vias, pins, and traces as well as several dozens of decoupling capacitors. In comparison, the on-chip power distribution network in a high complexity integrated circuit typically consists of millions of line segments, dramatically exacerbating the

complexity of the design and analysis process. The design and analysis of on-chip power distribution networks is the focus of the Chapter 7.

5.11 Summary

The impedance characteristics of power distribution systems with multiple stages of decoupling capacitances have been described in this chapter. These impedance characteristics can be briefly summarized as follows.

- The significant inductance of the power and ground interconnect is the primary obstacle to achieving a low output impedance power distribution system
- The hierarchical placement of decoupling capacitors achieves a low output impedance in a cost effective manner by terminating the power current loop progressively closer to the load as the frequency increases
- The capacitance and effective series inductance determine the frequency range where the decoupling capacitor is effective
- Resonant circuits are formed within the power distribution networks due to the placement of the decoupling capacitors, increasing the output impedance near the resonant signal frequencies
- The effective series resistance of the decoupling capacitors is a critical factor in controlling the resonant phenomena
- The lower the inductance of the power interconnect and decoupling capacitors, the lower the decoupling capacitance necessary to achieve the target impedance characteristics