IMPEDANCE CHARACTERISTICS OF DECOUPLING CAPACITORS IN MULTI-POWER DISTRIBUTION SYSTEMS

Mikhail Popovich and Eby G. Friedman

Department of Electrical and Computer Engineering
University of Rochester
Rochester, New York 14627, USA
Email: nhlover@ece.rochester.edu, friedman@ece.rochester.edu

ABSTRACT

To decrease power consumption without affecting circuit speed, multiple power supply voltages are often used in modern high performance ICs such as microprocessors. To maintain the impedance of a power distribution system below a specified level, multiple decoupling capacitors are placed at different levels of the power grid hierarchy. The system of decoupling capacitors used in power distribution systems with multiple power supplies is the focus of this paper. The dependence of the impedance on the power distribution system parameters is investigated. An antiresonance phenomenon is intuitively explained in this paper. Design techniques to cancel and shift the antiresonant spikes out of range of the operating frequencies are presented.

1. INTRODUCTION

Power dissipation has become a critical design issue in high performance microprocessors as well as battery powered and wireless electronics, multimedia and digital signal processors, and high speed networking. An effective way to reduce power consumption is to lower the supply voltage. Reducing the supply voltage, however, increases the circuit delay [1]. The increased delay can be compensated by changing the critical paths with behavioral transformations such as parallelization or pipelining [2]. The resulting circuit consumes less power while satisfying global throughput constraints at the cost of increased circuit area.

Recently, the use of multiple on-chip supply voltages has become common practice [3]. This strategy has the advantage of allowing modules along the critical paths to operate with the highest available voltage level (in order to satisfy target timing constraints) while permitting modules along the non-critical paths to use a lower voltage (thereby reducing energy consumption). In this manner, the energy consumed is decreased without affecting circuit speed. This scheme tends to result in smaller area as compared to parallel architectures. The problem of using multiple supply voltages to reduce the power requirements has been investigated in the area of high level synthesis for low power [4,5]. While it is possible to use many supply voltages, in practice such a scenario is expensive. Practically, the availability of a small number of voltage supplies (two or three) is reasonable.

This research was supported in part by the Semiconductor Research Corporation under Contract No. 2003-TJ-1068, the DARPA/ITO under AFRL Contract F29601-00-K-0182, the National Science Foundation under Contract No. CCR-0304574, Fulbright Program under Grant No. 87481764, grants from the New York State Office of Science, Technology & Academic Research to the Center for Advanced Technology - Electronic Imaging Systems and to the Microelectronics Design Center, and by grants from Xerox Corporation, IBM Corporation, Intel Corporation, Lucent Technologies Corporation, and Eastman Kodak Company.

The design of the power distribution system has become an increasingly difficult challenge in modern CMOS circuits [6]. As CMOS technologies are scaled, the power supply voltage is lowered. As clock rates rise and more functions are integrated on-chip, the power consumed has greatly increased. Assuming that only a small fraction of the power supply voltage (about 5%) is permitted as ripple voltage (noise), a target impedance for an example power distribution system is [7]

$$Z_{target} = \frac{V_{dd} \times r}{I} = \frac{1.8 \text{ volts} \times 5\%}{50 \text{ ampers}} = 0.002 \text{ ohms}, \qquad (1)$$

where V_{dd} is the power supply voltage, r is the allowable ripple voltage, and I is the current. With general scaling theory [8], the current I is increasing and the power supply voltage is decreasing. The impedance of a power distribution system should therefore be decreased to satisfy power noise constraints. The target impedance of a power distribution system is falling at an alarming rate, a factor of five per computer generation [9]. The specific impedance must be satisfied not only at DC, but also at all frequencies where current transients exist [10].

A power distribution network is a complex multi-level system. The design of a power distribution system with multiple supply voltages is the primary focus of this paper. The influence of a second supply voltage on a system of decoupling capacitors is also investigated. The paper is organized as follows. The impedance of a power distribution system with multiple supply voltages is described in Section 2. A case study of the dependence of the impedance on certain power distribution system parameters is presented in Section 3. Some specific conclusions are summarized in Section 4.

2. IMPEDANCE OF A POWER DISTRIBUTION SYSTEM WITH MULTIPLE SUPPLY VOLTAGES

The impedance of a power distribution network is an important issue in modern high performance ICs such as microprocessors. The impedance should be maintained below a target level to guarantee the power and signal integrity of a system. The impedance of a power distribution system with multiple power supplies is described in subsection 2.1. The antiresonance of capacitors connected in parallel is addressed in subsection 2.2. The dependence of the impedance on the power distribution system is investigated in subsection 2.3.

2.1. Impedance of a Power Distribution System

A model of the impedance of a power distribution system with two supply voltages is shown in Fig. 1. The impedance seen from the

load of the power supply V_{dd1} is applicable for the load of the power supply V_{dd2} if Z_1 is substituted for Z_3.

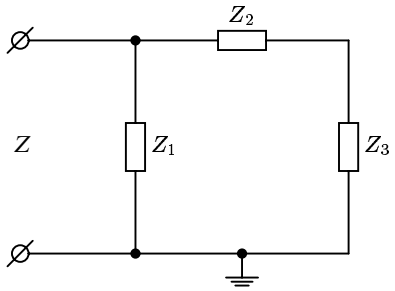

Figure 1: Impedance of a power distribution system with two supply voltages seen from the load of the power supply V_{dd1}.

The impedance of the power distribution system shown in Fig. 1 can be modeled as

$$Z = \frac{Z_1 Z_2 + Z_1 Z_3}{Z_1 + Z_2 + Z_3}. \tag{2}$$

Decoupling capacitors have traditionally been modeled as a series RLC network [11]. A schematic representation of a power distribution network with two supply voltages where a decoupling capacitor is represented by an RLC series network is shown in Fig. 2. In this case, the impedance of the power distribution network is

$$Z = \frac{a_4 s^4 + a_3 s^3 + a_2 s^2 + a_1 s + a_0}{b_3 s^3 + b_2 s^2 + b_1 s}, \tag{3}$$

$$\begin{aligned}
a_4 &= L_1(L_2 + L_3),\\
a_3 &= R_1 L_2 + R_2 L_1 + R_1 L_3 + R_3 L_1,\\
a_2 &= R_1 R_2 + R_1 R_3 + \frac{L_1}{C_2} + \frac{L_2}{C_1} + \frac{L_1}{C_3} + \frac{L_3}{C_1},\\
a_1 &= \frac{R_1}{C_3} + \frac{R_3}{C_1} + \frac{R_1}{C_2} + \frac{R_2}{C_1},\\
a_0 &= \frac{C_2 + C_3}{C_1 C_2 C_3},\\
b_3 &= L_1 + L_2 + L_3,\\
b_2 &= R_1 + R_2 + R_3,\\
b_1 &= \frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3},
\end{aligned}$$

and $s = j\omega$ is a complex frequency.

The minimum power distribution system impedance is limited by the Effective Series Resistance (ESR) of the decoupling capacitors. For on-chip applications, the ESR includes the parasitic resistance of the decoupling capacitor and the resistance of the power distribution network connecting a decoupling capacitor to a load. The resistance of the on-chip power distribution network is greater than the parasitic resistance of the on-chip decoupling capacitors. For on-chip applications, therefore, the ESR is represented by the resistance of the power delivery system. Conversely, for printed circuit board (PCB) applications, the resistance of the decoupling capacitors dominates the resistance of the power delivery system. In this case, the ESR is primarily the resistance of the decoupling capacitors. In order to achieve a target impedance as described by (1), multiple decoupling capacitors are placed at different levels of the power grid hierarchy [6].

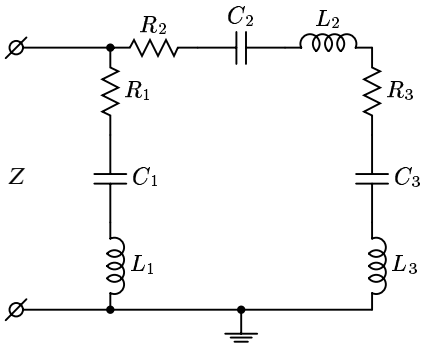

Figure 2: Impedance of power distribution system with two supply voltages and the decoupling capacitors represented as series RLC networks.

As described in [11], the ESR of the decoupling capacitors does not change the location of the poles and zeros of the power distribution system impedance. Only the damping factor of the RLC system formed by the decoupling capacitor is affected. Representing a decoupling capacitor with a series LC network, the impedance of the power distribution system with dual power supply voltages is

$$Z = \frac{a_4 s^4 + a_2 s^2 + a_0}{b_3 s^3 + b_1 s}, \tag{4}$$

$$\begin{aligned}
a_4 &= L_1(L_2 + L_3),\\
a_2 &= \frac{L_1}{C_2} + \frac{L_2}{C_1} + \frac{L_1}{C_3} + \frac{L_3}{C_1},\\
a_0 &= \frac{C_2 + C_3}{C_1 C_2 C_3},\\
b_3 &= L_1 + L_2 + L_3,\\
b_1 &= \frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3}.
\end{aligned}$$

2.2. Antiresonance of Parallel Capacitors

To maintain the impedance of a power distribution system below a specified level, multiple decoupling capacitors are placed in parallel at different levels of the power grid hierarchy. The ESR affects the quality factor of the RLC system resonance frequency by acting as a damping element. The influence of the ESR on the impedance is therefore ignored. If all of the parameters of the circuit shown in Fig. 2 are equal, the impedance of the power distribution system can be described as a series RLC circuit. Expression (4) has four zeros and three poles. Two zeros are located at the same frequency as the pole when all of the parameters of the circuit are equal. The pole is thereby canceled for this special case and the circuit behaves as a series RLC circuit with one resonant frequency.

If the parameters of the power distribution system are not equal, the zeros of (4) are not paired. In this case, the pole is not canceled by a zero. For instance, in the case of two capacitors connected in parallel as shown in Fig. 3, in the frequency range from f_1 to f_2, the impedance of the capacitor C_1 has become inductive whereas the impedance of the capacitor C_2 remains capacitive. In this case, an LC tank will produce a peak at a resonant frequency located between f_1 and f_2. Such a phenomenon is called *antiresonance* [9].

The location of the antiresonant spike depends on the ratio of the Effective Series Inductance (ESL) of the decoupling capacitors. Depending upon the parasitic inductance, the peak impedance

caused by the decoupling capacitor is shifted to a different frequency, as shown in Fig. 3. For instance, if the parasitic inductance of C_1 is greater than the parasitic inductance of C_2, the antiresonance will appear at a frequency ranging from f_1 to f_2, *i.e.*, before the self-resonant frequency f_2 of the capacitor C_2. If the parasitic inductance of C_1 is lower than the parasitic inductance of C_2, the antiresonance will appear at a frequency ranging from f_2 to f_3, *i.e.*, after the self-resonant frequency of the capacitor C_2. The ESL of the decoupling capacitors, therefore, determines the frequency of the antiresonant spike of the system.

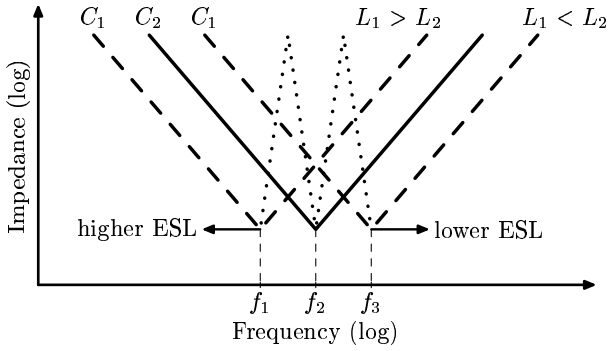

Figure 3: Antiresonance of the two capacitors connected in parallel, $C_2 = C_1$. Two antiresonant spikes appear between frequencies f_1 and f_2 and f_2 and f_3 (dotted lines).

2.3. Dependence of Impedance on Power Distribution System Parameters

In practical applications, a capacitor C_2 placed between V_{dd1} and V_{dd2} exists either as a parasitic capacitance or as a decoupling capacitor. Intuitively, from Fig. 2, by increasing C_2, the greater part of the decoupling capacitor C_3 is connected in parallel with C_1. Alternatively, by decreasing the impedance Z_2, the greater part of Z_3 is connected in parallel with Z_1, reducing the impedance of the power distribution system as seen from the load of the power supply V_{dd1}. The value of a parasitic capacitance is typically much smaller than a decoupling capacitor such as C_1 and C_3. The decoupling capacitor C_2 can be chosen to be equal to or greater than C_1 and C_3. Depending upon the placement of the decoupling capacitors, the ESL can vary from 50 nH at the power supply to almost negligible values on-chip.

Antiresonance is highly undesirable because at a particular frequency, the impedance of a power distribution network can become unacceptably high. To cancel the antiresonance at a given frequency, a smaller decoupling capacitor is placed in parallel, shifting the antiresonance spike to a higher frequency. This procedure is repeated until the antiresonance spike appears at a frequency out of range of the system operating frequency, as shown in Fig. 4.

Another technique for shifting the antiresonance spike to a higher frequency is to decrease the ESL of the decoupling capacitor. The dependence of the impedance of a power distribution system on the ESL is discussed below.

To determine the location of the antiresonant spikes, the roots of the denominator of (4) are evaluated. One pole is located at $\omega = 0$. Two other poles are located at frequencies,

$$\omega = \pm\sqrt{\frac{C_3+C_1C_3/C_2+C_1}{C_1C_3(L_1+L_2+L_3)}}. \qquad (5)$$

To shift the poles to a higher frequency, the ESL of the decoupling capacitors must be decreased. If the ESL of the decoupling capacitors is close to zero, the impedance of a power delivery network will not produce overshoots over a wide range of operating frequencies. Expression (5) shows that by minimizing the decoupling capacitor C_2 between the two supply voltages, the operating frequency of the overshoot-free impedance of a power delivery network can be increased.

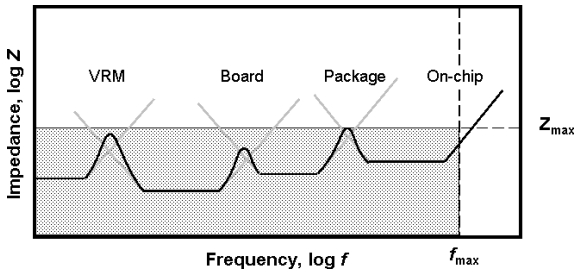

Figure 4: Impedance of the power distribution system as a function of frequency. Decoupling capacitors are placed at different hierarchical levels to shift an antiresonant spike above the maximum operating frequency of the system.

As discussed above, the location of the antiresonant spike is determined by the ESL ratio of the decoupling capacitors. The magnitude of the antiresonance is determined by the total ESL of C_1, C_2, and C_3. By lowering the system inductance, the quality factor is decreased. The peaks become wider in frequency and lower in magnitude. The amplitude of the antiresonant spikes can be decreased by lowering the ESL of all of the decoupling capacitors within the power distribution system. When the parasitic inductance of C_2 is similar in magnitude to the other decoupling capacitors, from (3), the poles and zeros do not cancel, affecting the circuit behavior. The zero at the resonant frequency of a system (the minimum value of the impedance) decreases the antiresonant spike. The closer the location of an antiresonant spike to the resonant frequency of a system, the greater the influence of a zero on the antiresonance behavior. From a circuit perspective, the more similar the ESL of each capacitor, the smaller the amplitude of the antiresonant spike. Decreasing the inductance of the decoupling capacitors has the same effect as increasing the resistance. Increasing the parasitic resistance of a decoupling capacitor is limited by the target impedance of the power distribution system and is therefore not recommended. Decreasing the inductance of a power distribution system is highly desirable and, if properly designed, the inductance of a power distribution system can be significantly reduced [12].

3. CASE STUDY

The dependence of the impedance on the power distribution system parameters is described in this section to quantitatively illustrate the concepts presented in Section 2. An on-chip power distribution system is assumed in this example. The total budget of on-chip decoupling capacitance is distributed among the low voltage power supply ($C_1 = 150$ nF), high voltage power supply ($C_3 = 50$ nF), and the capacitance placed between the two power supplies ($C_2 = 10$ nF). The ESR and ESL of the power distribution network are chosen to be equal to 0.01 ohms and 100 pH, respectively. The target impedance is 0.1 ohms.

For typical values of an example power distribution system, an antiresonant spike is produced at approximately 100 MHz with a magnitude greater than the target impedance, as shown in Fig. 5. According to (5), to shift the antiresonant spike to a higher frequency, the capacitor C_2 should be decreased. As C_2 is decreased to 3 nF, the antiresonant spike appears at a higher frequency, approximately 178 MHz, and is of greater magnitude. To further decrease the impedance of a power distribution system with multiple power supply voltages, the total ESL of the decoupling capacitors should be decreased. As the total ESL of the system is decreased to 10 pH, the impedance of the power distribution system is below the target impedance over a wide frequency range, from approximately 10 MHz to 2.2 GHz.

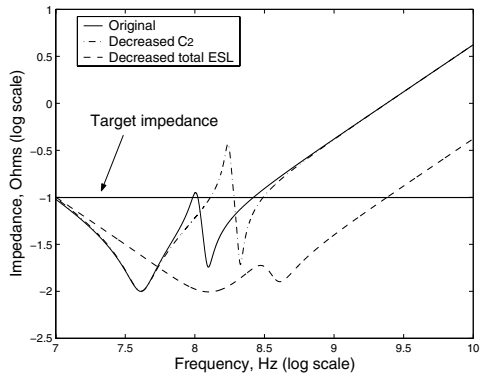

Figure 5: The impedance of a power distribution system with dual power supply voltages as a function of frequency, $R_1 = R_2 = R_3 = 10\,\mathrm{m}\Omega$, $C_1 = 150\,\mathrm{nF}$, $C_2 = 10\,\mathrm{nF}$, $C_3 = 50\,\mathrm{nF}$, and $L_1 = L_2 = L_3 = 100\,\mathrm{pH}$. The impedance of the example power distribution network produces an antiresonant spike with a magnitude greater than the target impedance (the solid line). The antiresonant spike is shifted to a higher frequency with a larger magnitude by decreasing C_2 to 3 nF (the dashed-dotted line). By decreasing the total ESL of the system to 10 pH, the impedance can be maintained below the target impedance over a wide frequency range, from approximately 10 MHz to 2.2 GHZ (the dashed line).

The design of a power distribution system with multiple power supply voltages is a complex task and requires many iterative steps. In general, to maintain the impedance of a power delivery system below a target level, the proper combination of design parameters needs to be determined. In on-chip applications, ESL and C_2 can be arbitrarily chosen. At the board level, ESR and C_2 can be adjusted to satisfy target impedance specifications. At the package level, ESL, C_2, and ESR are design parameters of the system. Usually, the total decoupling capacitance is constrained by the technology and application. In certain cases, it is possible to increase the decoupling capacitance. From (4), note that by increasing the decoupling capacitance, the overall impedance of a power distribution system with multiple power supply voltages can be significantly decreased.

4. CONCLUSIONS

It has become common practice to use multiple on-chip power supply voltages to reduce power dissipation without degrading system speed. To maintain the impedance of a power distribution system below a specified impedance, multiple decoupling capacitors are placed at different levels of the power grid hierarchy. The decoupling capacitors are placed both with progressively decreasing value to shift the peak antiresonance beyond the maximum operating frequency and with increasing ESR to control the damping characteristics and broaden the frequency range of the antiresonant spikes, ensuring that the amplitude of the output impedance remains below the target impedance. Another strategy is to limit the magnitude of the antiresonant spikes by reducing the ESL of all of the decoupling capacitors.

5. REFERENCES

[1] J. Mermet and W. Nebel, *Low Power Design in Deep Submicron Electronics*, Kluwer Academic Publishers, 1997.

[2] A. Chandrakasan, W. J. Bowhill, and F. Fox, *Design of High-Performance Microprocessor Circuits*, Wiley-IEEE Press, 2000.

[3] K. Usami and M. Horowitz, "Clustered Voltage Scaling Technique for Low-Power Design," *Proceedings of the IEEE International Workshop on Low Power Design*, pp. 3–8, April 1995.

[4] J.-M. Chang and M. Pedram, "Energy Minimization Using Multiple Supply Voltages," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 5, No. 4, pp. 425–435, December 1997.

[5] S. Raje and M. Sarrafzadeh, "Variable Voltage Scheduling," *Proceedings of the ACM International Symposium on Low Power Design*, pp. 9–14, April 1995.

[6] A. V. Mezhiba and E. G. Friedman, *Power Distribution Networks in High Speed Integrated Circuits*, Kluwer Academic Publishers, 2004.

[7] L. D. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc, and T. Roy, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," *IEEE Transactions on Advanced Packaging*, Vol. 22, No. 3, pp. 284–291, August 1999.

[8] S. R. Nassif and O. Fakhouri, "Technology Trends in Power-Grid–Induced Noise," *Proceedings of the ACM International Workshop on System Level Interconnect Prediction*, pp. 55–59, April 2002.

[9] L. D. Smith, "Packaging and Power Distribution Design Considerations for a Sun Microsystems Desktop Workstation," *Proceedings of the Electrical Performance of Electronic Packaging Conference*, pp. 19–22, October 1997.

[10] W. D. Becker *et al.*, "Modeling, Simulation and Measurement of Mid-Frequency Simultaneous Switching Noise in Computer Systems," *IEEE Transactions on Components, Packaging, and Manufacturing Technology, Part B*, Vol. 21, Issue 2, pp. 157–163, May 1998.

[11] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*, Addison-Wesley, 1990.

[12] A. V. Mezhiba and E. G. Friedman, "Inductive Properties of High-Performance Power Distribution Grids," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 10, Issue 6, pp. 762–776, December 2002.