




FIGURE 11-1. Transient power-supply current (A) without and (B) with a decoupling capacitor.

When a logic gate switches, a current transient dI occurs in the power distribution system as shown in Fig. 11-1A. This current transient flows through both the ground and power traces. The transient current flowing through the power and ground inductances produces a noise voltage that appears between the V_{cc} and ground terminals of the logic gate. In addition, the transient current flowing around a large loop makes an efficient loop antenna.

The magnitude of the power supply voltage transient can be reduced by decreasing the inductances L_p and L_g , and/or by decreasing the rate of change of the current (dI/dt) that flows through these inductances. The inductance can be reduced, but not eliminated, by using power and ground planes or grids, as discussed in Chapter 10. The loop area and inductance can both be minimized by supplying the transient current from another source, such as a capacitor or capacitors, located near the logic gate as shown in Fig. 11-1B. The noise voltage across the logic gate is then a function of the capacitor C_d and the traces between it and the gate (represented by L_{p2} and L_{g2} in Fig. 11-1B). The number of capacitors, their type, value, and placement with respect to the IC are important in determining their effectiveness.

Decoupling capacitors, therefore, serve two purposes. First, they provide a source of charge close to the IC so that when the IC switches, the decoupling capacitors can supply the required transient currents through a low-impedance path. If the capacitor cannot supply the required current, then the magnitude of the voltage bus will dip, and the IC may not function properly. The second purpose of the decoupling capacitor is to provide a low ac impedance between the power and ground rails, which effectively shorts out (or at least minimizes) the noise injected back into the power/ground system by the IC.

11.2 TRANSIENT POWER SUPPLY CURRENTS

Two different power supply transient currents occur when a digital IC switches. These currents are shown in Fig. 11-2. First, when the logic gate switches from




FIGURE 11-2. Transient currents produced when a CMOS logic gate switches.

low to high, a transient current I_L is needed to charge the load capacitance C_L . This current only occurs in logic gates connected to outputs. The second current results from the totem pole output structure of the IC, and it occurs on both the high-to-low and the low-to-high transition. Part way through the switching cycle, both transistors are partially on, which produces a low impedance across the power supply and draws a transient current I_d as shown in Fig. 11-2. This dynamic internal current occurs even in gates that are not connected to external loads.

These transient power supply currents are drawn every time the logic circuit changes state. Which of these two currents predominate depends on the type of IC. For digital ICs with a lot of input/output (I/O) drivers, such as clock drivers, buffers, or bus controllers, the transient load current I_L will predominate. For large application-specific integrated circuits (ASICs), microprocessors, or other devices with a large amount of internal processing, the dynamic internal current I_d will predominate. The magnitude of these currents can be estimated by the following means.

11.2.1 Transient Load Current

The magnitude of the transient load current can be determined by considering the effect of the capacitive load C_L on the IC, as shown in Fig 11-2. The transient current through the load capacitance, during the switching cycle, will be $I_L = (C_L V_{cc}) / t_r$. For more than one load connected to the IC, multiply C_L




FIGURE 11-3. Isosceles triangular wave.

by the number of loads n , thus giving the magnitude of the transient load current as follows:

$$I_L = \frac{nC_L V_{cc}}{t_r}, \quad (11-1)$$

where C_L is the capacitance of one load, n is the number of loads on the IC, V_{cc} is the IC supply voltage, and t_r is the rise time of the output waveform. A typical complementary metal-oxide semiconductor (CMOS) gate, which acts as a load, has an input capacitance of 7 to 12 pF.

Experience has shown that the waveshape of the transient current pulse can be approximated by a triangular wave (see Fig. 11-3) with a peak amplitude determined by Eq. 11-1, and a pulse width equal to twice the rise time t_r (Archambeault, 2002 or Radu et al., 1998). This transient load current will normally only occur on the low-to-high transition because, as was indicated in Table 10-4, the power supply or decoupling capacitor is the source of the transient load current only on the low-to-high transition.*

For example, a device operating from a 5-V power supply with a 1-ns rise time output pulse, driving 10 CMOS loads with 10 pF of input capacitance each will have a transient I/O current of 500 mA.

11.2.2 Dynamic Internal Current

It is often harder to obtain estimates of the shoot-through current than the transient load current. Many IC data sheets provide no information on

* The one exception to this is a stripline located adjacent to both a power and ground plane; in which case, the decoupling capacitor will be the source of the current on both the high-to-low and the low-to-high transitions.

the transient power supply currents drawn by the IC. Some CMOS device data sheets, however, provide information on the internal power dissipation, by providing the value of an equivalent power dissipation capacitance C_{pd} . The power dissipation capacitance can be thought of as an equivalent internal capacitance that can be used to estimate the magnitude of the shoot-through current; this approach is similar to the way the load capacitance C_L was used to determine the magnitude of the transient load current. This capacitance is often specified as so much capacitance per gate. The total power dissipation capacitance will then be the capacitance per gate C_{pd} multiplied by the number of gates switching simultaneously (that may be difficult to determine). Other IC data sheets actually list a value for the dynamic power supply current I_{ccd} , which is specified in units of A/MHz.

The magnitude of the dynamic internal current I_d can then be estimated by using one of the following two equations:

$$I_d = I_{ccd} f_0 \quad (11-2a)$$

or

$$I_d = \frac{nC_{pd} V_{cc}}{t_r}, \quad (11-2b)$$

where I_{ccd} is the dynamic power supply current, f_0 is the clock frequency, C_{pd} is the power dissipation capacitance per gate, n is the number of gates switching, V_{cc} is the supply voltage, and t_r is the switching time of the gate.

The waveshape of this dynamic current pulse can also be approximated by a triangular wave associated with each edge of the clock, and having a peak amplitude determined by Eq. 11-2a or 11-2b. The pulse width is equal to twice the rise time t_r (Archambeault, 2002, or Radu et al., 1998). Because the dynamic internal current occurs on every transition of the clock (low to high and high to low), the dynamic internal current will contain harmonics that are multiples of twice the clock frequency.

11.2.3 Fourier Spectrum of the Transient Current

For an isosceles-triangular wave as shown in Fig. 11-3, the amplitude of the current in the n th harmonic is given by (Jordan, 1985, p. 7-10)

$$I_n = \frac{2It_r}{T} \left(\frac{\sin\left(\frac{n\pi t_r}{T}\right)}{\frac{n\pi t_r}{T}} \right)^2, \quad (11-3)$$

where I is the amplitude of the triangular wave, t_r is the rise (fall) time, T is the period, and n is the harmonic number. Figure 11-4 shows a log-log plot of




FIGURE 11-4. Envelope of Fourier spectrum of an isosceles triangular wave.

the envelope of the harmonics for a wave with a t_r/T ratio of 0.1. As can be observed, above a frequency of $1/\pi t_r$, the harmonics fall off with frequency at a 40 dB per decade rate, and both odd and even harmonics are present.

Table 11-1 lists the percentage of current present in the first six harmonics versus the t_r/T ratio. Note that as the rise time becomes faster (smaller t_r/T), the percent current in the fundamental decreases, and the current has a slower falloff with frequency (harmonic number). For a t_r/T ratio of 0.1 or less, the fundamental frequency will contain less than 20% of the current.

TABLE 11-1. Percent Current in the First Six Harmonics versus the t_r/T Ratio for an Isosceles-Triangle Waveform.

t_r/T	Harmonic					
	1	2	3	4	5	6
0.05	10%	10%	9%	9%	8%	7%
0.1	19%	18%	15%	12%	8%	5%
0.2	35%	23%	10%	2%	0%	<1%
0.3	44%	15%	2%	2%	3%	<1%




FIGURE 11-5. The waveshape of the total transient current drawn by an IC.

The $1/\pi t_r$ frequency will occur at a frequency equal to k times the fundamental frequency, where k (which may be a noninteger) is equal to

$$k = \left(\frac{1}{\pi}\right) \left(\frac{T}{t_r}\right). \quad (11-4)$$

The heavy line in Table 11-1 indicates between which harmonics the $1/\pi t_r$ frequency falls.

11.2.4 Total Transient Current

The total transient current drawn by the IC is the sum of the transient load current I_L and the dynamic internal current I_d . Figure 11-5 shows the waveshape of the total transient current. The transient load current spikes occur at the clock frequency, and the dynamic internal currents spikes occur at twice the clock frequency.

Archambeault (2002, pp. 127–129) has demonstrated a good correlation between measured and predicted V_{cc} -to-ground noise voltage using the current waveshape shown in Fig. 11-5.

11.3 DECOUPLING CAPACITORS

Effective power supply decoupling has become increasingly more difficult to achieve as the result of increasing clock frequencies and faster rise times. Ineffective decoupling can lead to excessive power bus noise as well as to excessive radiated emission.

Many designers are decoupling digital logic ICs by placing a single 0.1- or 0.01- μF capacitor adjacent to the IC. This method is the same one that has been used on digital logic ICs for the last 50 years, so it still must be the correct approach—right? After all, how much has IC technology changed in the last 50 years? I think it is interesting that this method has worked for as many years as

it has. No one should be surprised that, possibly, we are now to the point that a new approach to digital IC decoupling is required.

It is important to understand that decoupling is *not* the process of placing a capacitor adjacent to an IC to supply the transient switching current as shown in Fig. 11-6A; rather it is the process of placing an L - C network adjacent to an IC to supply the transient switching current as shown in Fig. 11-6B. All decoupling capacitors have inductance in series with them. Therefore, the decoupling network is a series resonant circuit. As can be observed from Fig. 11-7, the inductance comes from three sources, as follows:

1. The capacitor itself
2. The interconnecting PCB traces and vias
3. The lead frame inside the IC




FIGURE 11-6. (A) Ideal decoupling network and (B) actual decoupling network.




FIGURE 11-7. Equivalent circuit for decoupling capacitor connected to an integrated circuit, for a DIP package with power and ground on diagonally opposite pins of the IC.

Although Fig. 11-7 is for a dual in-line package (DIP), the total inductance value is still representative of that obtained with other IC packages. Figure 11-7 assumes that the capacitor is placed at one end of the DIP, close to one of the pins (power or ground), and about an inch away from the other pin.

The internal inductance of a surface mount technology (SMT) capacitor itself is typically 1–2 nH, the interconnecting PCB traces and vias add 5 to 20 nH or more according to the layout, and the internal lead frame of the IC may have 3 to 15 nH of inductance according to the type of IC package. The inductances of the interconnecting PCB traces, however, are usually the only parameters under the system designer's control.

It is extremely important therefore to minimize the inductance of the traces between the IC and the decoupling capacitor. The printed circuit traces should be as short as possible and should be placed as close together as possible to minimize their loop area. Notice also from Fig. 11-7 that the decoupling capacitor contributes the least amount of inductance. Therefore, it is not the major problem.

From the above, we observe that the total inductance can vary from a low of about 10 nH to a high of 40 nH. Typically it is in the range of 15 to 30 nH. It is this inductance that limits the effectiveness of the decoupling network. It is very important to remember this fact—we are placing an L - C network between the power and ground, *not* a capacitor!

Because of this combination of capacitance and inductance, the decoupling network will, at some frequency, become resonant. At the resonant frequency, the magnitude of the inductive reactance is equal to the magnitude of the capacitive reactance and the network is a very low-impedance and an effective bypass. Above the resonant frequency, the circuit becomes inductive and its impedance increases with frequency.

The resonant frequency of a series L - C circuit is

$$f_r = \frac{1}{2\pi\sqrt{LC}} \quad (11-5)$$

Figure 11-8 shows the impedance versus frequency of various decoupling capacitors when in series with 30 nH of inductance. The top horizontal scale also shows the equivalent digital logic rise time that corresponds to the frequency shown on the bottom horizontal scale. The relationship between rise time and frequency is $t_r = 1/\pi f$.

The point where the capacitive reactance is equal to the inductive reactance is the resonant frequency of that combination of capacitor and inductor. The impedance at resonance decreases to a low value (equal to just the series resistance of the network), because the capacitive reactance cancels the inductive reactance and leaves only the resistance.

As can be observed in Fig. 11-8, the commonly used 0.1- μ F capacitor resonates with 30 nH of inductance at 3 MHz, and a 0.01- μ F capacitor resonates




FIGURE 11-8. Impedance of different value decoupling capacitors in series with 30 nH of inductance.

at 9 MHz, which is not impressive in this day of 100+ MHz clocks. Above about 50 MHz, the impedance of the decoupling network is dominated by the 30 nH of inductance regardless of what value capacitor is used. If a good layout was achieved and the inductance was half of the 30 nH shown in Fig. 11-8, the resonant frequency would only increase by the $\sqrt{2}$, or a factor of 1.41. Therefore, a 0.1- μ F capacitor will resonate at about 4 MHz, and a 0.01- μ F capacitor will resonate at about 13 MHz when in series with 15 nH of inductance.

Figure 11-8 clearly demonstrates that *placing a single capacitor adjacent to an IC is not an effective way to decouple digital logic at frequencies above 50 MHz, regardless of the capacitor value or placement.* Figure 11-8 also shows that for frequencies less than about 50 MHz, it is possible to tune the decoupling network, by choosing different value capacitors to place the resonant dip at a problem frequency. But this same approach will not work above 50 MHz.

Figure 11-9 shows measured values of the V_{cc} -to-ground noise voltage versus frequency on a board using 0.01- μ F decoupling capacitors. The measurements were made between the power and ground pins of the IC, with and without decoupling capacitors on the board. In the frequency range of 20 to 70 MHz,




FIGURE 11-9. Measured V_{cc} -to-ground noise voltage versus frequency for a printed circuit board with and without decoupling capacitors.

the presence of the decoupling capacitors significantly decreased the magnitude of the V_{cc} -to-ground noise voltage. However, in the frequency range of 70 to 120 MHz, the noise voltage was the same, with and without the decoupling capacitors, this result indicates that the capacitors were ineffective in this frequency range, which confirms the results predicted by Fig. 11-8.

11.4 EFFECTIVE DECOUPLING STRATEGIES

Contamination of the V_{cc} bus by clock harmonics that result from ineffective decoupling can cause signal integrity problems and can be the cause of excessive radiated emission. Possible solutions to the high-speed decoupling problem are as follows:

1. Slow down the rise time
2. Decrease the current transient
3. Decrease the inductance in series with the capacitor
4. Use multiple capacitors

The first two approaches listed above go against the advancement in technology and will not provide a long-term solution to the problem. Decreasing the inductance in series with the decoupling capacitor is desirable and should always be done, to whatever extent possible; however, this by itself will not solve the high-speed decoupling problem. Even if the total inductance in series with the capacitor was somehow reduced to only 1 nH, which is an unlikely possibility, a decoupling network with a 0.01- μ F capacitor would have a resonant frequency of only 50 MHz. Therefore, it is not possible to move the resonant frequency of a single capacitor decoupling network up to a frequency of a few hundred megahertz using any realistic value of capacitance.

Table 11-2 lists the resonant frequency of various values of capacitors when in series with 5, 10, 15, 20, and 30 nH of inductance.

The conclusion that we can draw from Table 11-2, is that in most cases, it is impossible to raise the resonant frequency of the decoupling L - C network to be above the clock frequencies commonly used in most digital electronics today, no less above the frequency of the harmonics. With only 10 nH of inductance, a 1000-pF capacitor will only have a resonant frequency of 50 MHz.

TABLE 11-2. Resonant Frequency (in MHz) of Various Value Capacitors When in Series With the Listed Value of Inductance.

Capacitor (μ F)	5 nH	10 nH	15 nH	20 nH	30 nH
1	2.3	1.6	1.3	1	0.9
0.1	7.1	5	4.1	3.6	3
0.01	22.5	16	13	11	9
0.001	71.2	50	41	36	30

At frequencies below the decoupling network resonance, the two most important considerations are (1) to have sufficient capacitance (see Eq. 11-12) to provide the required transient current, and (2) to provide an impedance that is low enough to short out the noise current generated by the IC (see Section 11.4.5). However, above resonance the most important consideration is to have a low enough inductance, such that the decoupling network is still a low impedance and shorts out the noise current. Therefore, a decoupling network can still be effective above resonance, if a way can be found to lower its inductance sufficiently.

11.4.1 Multiple Decoupling Capacitors

No single decoupling capacitor network will provide a low enough inductance. Therefore, the real solution to the high-frequency decoupling problem lies in the use of multiple decoupling capacitors. Three approaches have been proposed. They are as follows:

1. The use of multiple capacitors all of the same value.
2. The use of multiple capacitors of two different values.
3. The use of multiple capacitors of many different values, usually spaced a decade apart. For example, a 1 μ F, 0.1 μ F, 0.01 μ F, 0.001 μ F, 100 pF, etc.

11.4.2 Multiple Capacitors of the Same Value

When a number of identical L - C networks are connected in parallel as shown in Fig. 11-10, the total capacitance C_t becomes

$$C_t = nC, \quad (11-6a)$$


where C is the capacitance of one of the networks and n is the number of networks in parallel.

The total inductance L_t of n identical L - C networks in parallel is

$$L_t = \frac{L}{n}, \quad (11-6b)$$

where L is the inductance of one of the networks. By referring to Eq. 10-8, we can conclude that Eq. 11-6b will only be correct if the mutual inductance between the inductances of the individual networks are negligible compared with their self-inductance. Therefore, the L - C networks must be physically separated from each other.

Notice from Eqs. 11-6a and 11-6b that when identical L - C networks are paralleled, the capacitance (the good parameter) is multiplied up by the number of networks, and the inductance (the bad parameter) is divided down by the

FIGURE 11-10. Three identical L - C networks in parallel.

number of networks. Therefore, the inductance of the network can be reduced to any desirable value just by using a sufficient number of L - C networks in parallel.

The requirements for effectively paralleling multiple L - C networks are as follows:

1. Make all the capacitors the same value; this way they share the current equally.
2. Each capacitor must feed the IC through a different inductance than the others. Therefore, they cannot be located together, because of the mutual inductance that would occur. Spread them out.

Rather than just considering the resonant frequency, a better approach to analyzing the effectiveness of a decoupling network is to consider the IC as a noise current generator as shown in Fig. 11-11. The decoupling network can then be designed to be a low impedance across the frequency band of interest, to short out the noise current and to prevent it from contaminating the power bus. A maximum allowable value for the decoupling network impedance, which is sometimes referred to as a *target impedance*, can be specified, for example, 100 to 200 $m\Omega$, and the impedance of the chosen decoupling network can be calculated to determine over what frequency range the impedance is less than the target impedance. See Section 11.4.5 for more information on determining target impedance.

Figure 11-12 is a plot of the impedance versus frequency for various numbers (1, 8, 64, 512) of identical L - C networks in parallel. In all cases, the *total* capacitance equals 0.1 μF and the inductance in series with each capacitor is 15 nH.

Although the high-frequency impedance is reduced dramatically by the use of multiple capacitor networks, the low-frequency impedance is not—only the frequency of the resonant dips shift. Actually, at low frequencies, where the single capacitor network was resonant, the impedance is actually higher when multiple capacitors are used than when only one capacitor is used. This is because the total capacitance, in this case 0.1 μF , is not large enough to be a low impedance at these frequencies.




FIGURE 11-11. The IC as a noise current generator.




FIGURE 11-12. Impedance versus frequency for a decoupling network consisting of from 1 to 512 equal-value capacitors when in series with 15 nH of inductance. In all cases, the total capacitance equals 0.1 μF .

Figure 11-13 is similar to Fig. 11-12, except in this case the multiple capacitors add up to a *total* capacitance of 1 μF (instead of 0.1 μF). As a result, the low-frequency impedance is significantly reduced. For the case of 512 capacitors, Fig. 11-13 shows that the impedance of the L - C decoupling network is less than 0.2 Ω from 1 MHz to 1000 MHz. With 64 capacitors, the impedance is less than 0.5 Ω from less than 1 MHz to 350 MHz. Notice




FIGURE 11-13. Impedance versus frequency for a decoupling network consisting of from 1 to 512 equal-value capacitors when in series with 15 nH of inductance. In all cases, the total capacitance equals 1.0 μ F.

also that the resonant frequency of the decoupling networks in Fig. 11-13 are lower than the resonant frequency of the networks in Fig. 11-12, but the decoupling effectiveness is improved (lower impedance over a wider frequency range). Therefore, *the use of a large number of equal value capacitors is an effective way to provide a low-impedance decoupling network that is effective over a large frequency range.* This strategy can be very effectively used when decoupling a large IC.

11.4.3 Multiple Capacitors of Two Different Values

Two different decoupling capacitor values are sometimes recommended, based on the theory that the large value capacitor will provide effective low-frequency decoupling, and the small value capacitor will provide effective high-frequency decoupling. If two different capacitor values are used, there will be two distinct resonant dips as shown in Fig. 11-14, which is good.

Although the above is true, when capacitors of different values are placed in parallel, a potential problem exists with respect to the parallel resonance (sometimes referred to as an “antiresonance”) that occurs between the two networks. Figure 11-14 shows a plot of the impedance of a 0.1- μ F capacitor in





FIGURE 11-14. Impedance versus frequency for a decoupling network consisting of a 0.1 and a 0.01- μ F capacitor, both in series with 15 nH of inductance.

parallel with a 0.01- μ F capacitor, which both have 15 nH of inductance in series with them. Figure 11-14 clearly shows the two resonant dips produced by the two different value capacitors in series with 15 nH of inductance, one at 4.1 MHz and the other at about 13 MHz. Notice, however, that an impedance spike occurs at about 9 MHz, which is bad. This effect, as described by Paul (1992), is caused by the parallel, or antiresonance, between the two networks.

Why this occurs can best be understood by referring to Fig 11-15A, which shows two L - C decoupling networks with different value capacitors connected between V_{cc} and ground. Assume that $C_1 \gg C_2$ and $L_1 = L_2$. Let f_{r1} be the frequency where the capacitor C_1 is resonant with inductor L_1 , and let f_{r2} be the frequency where capacitor C_2 is resonant with inductor L_2 .

Below the frequency at which either of the networks is resonant ($f < f_{r1}$), both networks look capacitive, and the total capacitance is equal to the sum of the two capacitors, which for all practical purposes is just equal to the large capacitor. Therefore, the small capacitor has no, or very little, effect on the performance of the decoupling network.

Above the frequency where both networks are resonant ($f > f_{r2}$), both networks look inductive and the total inductance is equal to the two inductors in parallel, or one half the inductance. This improves the decoupling at frequencies above f_{r2} .



FIGURES 11-15. (A) Decoupling network with two different value capacitors, and (B) the equivalent circuit of A for a frequency f where $f_{r1} < f < f_{r2}$

At a frequency between the resonances of the two networks ($f_{r1} < f < f_{r2}$), however, the network with the larger capacitor has become inductive, and the network with the smaller capacitor is still capacitive. The equivalent circuit of the two networks is therefore a capacitor in parallel with an inductor as shown in Fig. 11-15B, or a parallel resonant circuit, the impedance of which is large at resonance, which produces the resonant spike shown in Fig. 11-14.

The exact shape, amplitude, and location of the parallel resonant peak will vary as a function of the ratio of the two capacitor values, the equivalent series resistance (ESR) of the capacitors, and the PCB layout. If the two capacitor values are within a two-to-one ratio, the amplitude of the resonant peak will be reduced to an acceptable value, because it will occur in the resonant dip. For example, even the same nominal value capacitors will have different values, because of their tolerance. The resonant peak produced between two capacitors that have a 20%, or even 50%, tolerance will not cause a problem. The major problem with the antiresonance occurs when the capacitors have values that are an order of magnitude or more apart.

Therefore, we can conclude that when two different value capacitors are used in each of the decoupling networks that:

- The small capacitor network will have *no effect* on the decoupling performance at frequencies below which the large capacitor network is resonant.
- The decoupling will *be improved* at frequencies above which both capacitor networks are resonant, because the inductance will decrease.
- The decoupling will actually *be worse* at some frequencies between these two resonance frequencies, because of the impedance spike caused by the parallel resonant network, which is bad.

For example, Bruce Archambeault (2001) provides information on the decoupling effectiveness for various decoupling strategies by using a network analyzer to measure the power to ground plane impedance of a test PCB. For the case of two different value decoupling capacitors, he concludes:

There is no noticeable improvement in the high frequency decoupling performance when a second capacitor value is added. In fact the decoupling performance is worse in the frequency range where much of the typical noise energy exists (50-200 MHz).

Archambeault's data show that at some frequencies, between the resonant frequencies of the two decoupling networks, the noise increased by as much as 25 dB, in the frequency range of 50 to 200 MHz when two different capacitor values were used, as compared with the results when all the capacitors were of the same value.

11.4.4 Multiple Capacitors of Many Different Values

Capacitors of many different values (typically spaced a decade apart) are sometimes also recommended, based on the theory that the multiple impedance dips produced by the resonances of the different value capacitors is advantageous, because it will provide a low impedance at many frequencies.

However, when decoupling capacitors of many different values are used, additional impedance spikes are also produced. Figure 11-16 shows the impedance plot for the case when four different value decoupling capacitors are used, all in series with 15 nH of inductance. The solid line plot is for the case of a 0.1- μ F, a 0.01- μ F, a 0.001- μ F, and a 100-pF capacitor all in parallel. Each capacitor is in series with 15 nH of inductance. As can be observed four resonant dips exist, one for each value of capacitor. However, the antiresonance also produces three resonant spikes in the impedance plot. If some clock harmonics should fall on or near the frequency of these spikes, the power to ground noise will actually increase. Some people consider this approach as the equivalent of playing "Russian Roulette," because one must hope that none of the clock harmonics fall on or near any of the resonant spikes. Notice also, as shown in Fig. 11-16, that the amplitude of the resonant peaks increases with frequency, whereas the impedance of the resonant dips remains the same.




FIGURE 11-16. Impedance versus frequency for a decoupling network consisting of a 0.1, a 0.01, and a 0.001 μ F plus a 100-pF capacitor (solid line), and a network that consists of four 0.1- μ F capacitors (dashed line). For both networks, the capacitors are in series with 15 nH of inductance.

If, however, instead of four different value capacitors, four 0.1- μ F capacitors had been used, the impedance plot would be as shown by the dashed line in Fig. 11-16. Both plots are for the same number of capacitors. The lower impedance at low frequency is the result of the greater total capacitance provided by the four parallel 0.1- μ F capacitors. The absence of the resonant spikes is the result of all of the capacitors being of the same value. At frequencies above 200 MHz, the results are the same for the case of the four 0.1- μ F capacitors and the four different value capacitors. This results from the fact that at these frequencies only the inductance in series with the capacitors matters, and in both cases four inductances are in parallel. However, below 200 MHz the results are better for the case of the four 0.1- μ F capacitors, except at the few frequencies that correspond to the resonant dips in the multiple-value capacitor case. Which impedance plot would you rather have for your product?

My preference for effective high-frequency decoupling, therefore, is to use multiple capacitors, all of the same value. This approach works well and has fewer pitfalls than either of the schemes using multiple capacitors of different values.

11.4.5 Target Impedance

To be an effective decoupling network, the impedance of the network must be kept below some target value over the frequency range of interest. If this can be accomplished, then the location of the resonant frequency does not matter. If the target impedance were 200 m Ω , then for the case of the 64-capacitor configuration, which was depicted in Fig. 11-12, the range over which the impedance would have been below 200 m Ω is from about 8 MHz to 130 MHz.

However, using a target impedance that is constant across the frequency range is overly restrictive and not necessary. From Eq. 11-3 and Fig. 11-4, we know that the amplitude of the harmonics of a triangular wave fall off at a rate of 40 dB per decade above a frequency of $1/(\pi t_r)$. Therefore, the target impedance can rise above this frequency without increasing the noise voltage. If the target impedance is allowed to rise at a rate of 20 dB per decade above this corner frequency (as shown by the solid line in Fig. 11-17), then the noise will still decrease at a 20 dB per decade rate beyond this frequency. This approach greatly simplifies the decoupling network design and minimizes the required number of capacitors.

Using this approach, one can easily estimate the number of decoupling capacitors needed to provide effective high-frequency decoupling. The minimum number of capacitors n is equal to

$$n = \frac{2L}{Z_t t_r}, \quad (11-7)$$

where L is the inductance in series with each capacitor, Z_t is the low-frequency target impedance, and t_r is the switching (rise/fall) time of the logic. Using at




FIGURE 11-17. Target impedance (solid line). Impedance of a decoupling network consisting of 64, 0.01- μ F capacitors, each in series with 10 nH of inductance (dashed line).

least this number of capacitors will keep the high-frequency impedance of the decoupling network at or below the target impedance.

The key to an optimum decoupling design is to know what inductance to use in Eq. 11-7. As far as the power-to-ground noise on the IC chip itself is concerned, the total inductance (decoupling capacitor, PCB traces, and IC lead frame) must be considered. However, nothing can be done at the PCB level about the inductance of the IC lead frame. Besides, when we measure the power-to-ground noise on an IC we measure it at the pins of the IC, not at the chip. In addition, as far as contaminating the PCB power bus with noise, it is the voltage at the IC-to-PCB interface (i.e., the IC pins) that matters, not the noise voltage at the chip itself.

Therefore, the objective of decoupling should be to minimize the V_{cc} -to-ground noise voltage at the pins of the IC. To accomplish this, one can neglect the internal inductance of the IC. Therefore, we only have to consider the decoupling capacitor inductance and the PCB trace (including via) inductance. The internal inductance of a good SMT capacitor will be 1.5 nH or less. The PCB trace inductance will be about 10 nH/in, and the inductance of a through-via on a 0.062-in-thick PCB will be about 0.8 nH.

The low-frequency target impedance is often determined by considering the magnitude of the total transient current and the allowable variation of the

supply voltage. Therefore, we can write

$$Z_t = \frac{k dV}{dI}, \quad (11-8)$$

where dV is the allowable power supply transient voltage variation, dI is the amplitude of the transient power supply current drawn by the IC, and k is a correction factor discussed in the next paragraph.

Let us assume that our concern is the V_{cc} -to-ground noise spikes and their effect on the operation of the circuit. We can determine from Table 11-1 that no more than about 50% of the current is contained in the frequencies below the $1/\pi t_r$ frequency. Therefore, the low-frequency target impedance could be increased by a factor of two, because the dI used in Eq. 11-8 was the total transient current. Therefore, in this case we can use $k = 2$ in Eq. 11-8. This approach will limit the total noise voltage spike to less than the target impedance times the total transient current.

For example, consider the case of decoupling a large IC having a rise/fall time of 2 ns operating from a 5-V supply, where it is desirable to keep the maximum power supply deviation to no more than 5% of the supply voltage. Assume also that the total transient current drawn by the IC is 2.5 A. From Eq. 11-8 with $k = 2$, we get a low-frequency target impedance of 200 m Ω . Assume that each decoupling capacitor has 10 nH of inductance in series with it. The target impedance will then be 200 m Ω up to the $1/(\pi t_r)$ frequency, which for a 2-ns rise time is equal to 159 MHz, above which the target impedance can be allowed to increase at a 20 dB per decade rate. This target impedance is shown as the solid line in Fig. 11-17.

From Eq. 11-7 we determine that 50 capacitors would be necessary to meet the 200-m Ω target impedance at 159 MHz. If the lowest frequency of interest were 2 MHz, the total decoupling capacitance necessary to meet the 200-m Ω target impedance, at 2 MHz, would be 400 nF. Dividing 400 nF by 50, gives 8000 pF for the minimum value of each capacitor. The total capacitance must also meet the transient current criteria of Eq. 11-12, which in this case it does. If the designers decide to use 64, 0.01- μ F capacitors, the impedance versus frequency plot will be as shown by the dashed line in Figure 11-17. The decoupling network will then have an impedance below the target impedance at all frequencies above 1.3 MHz.

11.4.6 Embedded PCB Capacitance

Taking the concept of using a large number of equal value capacitors to its limit, one can conclude that the ideal decoupling configuration is an infinite number of infinitesimal capacitors (i.e., instead of using discrete capacitors, one should use distributed capacitance).

Could this result possibly be achieved by taking advantage of the interplane capacitance between the power and ground planes of a PCB? To be effective above about 50 MHz, a power-ground plane capacitance of about 1000 pF/in² would be required. Standard layer spacings of 0.005 to 0.010 in, however, provide capacitances that are 1/5 to 1/10 of this value.

Therefore, if advantage is to be taken of the concept of distributed capacitance, then a new PCB structure, that has additional capacitance needs to be developed. This increase in distributed capacitance could possibly be achieved in one of the following two ways:

1. Reduce the layer spacing
2. Increase the dielectric constant of the PCB material

11.4.6.1 Effective Decoupling Capacitance Area. The effective decoupling capacitance, however, is not the capacitance of the total power-ground-plane pair. Because of the finite velocity of propagation, of electromagnetic energy the effective interplane capacitance is only the capacitance of an area that is located within a finite radius r of the IC being decoupled, as shown in Fig. 11-18. Charge that is stored in the interplane capacitance farther away than this cannot be moved fast enough to arrive at the IC during the switching transient.

The velocity of propagation of electromagnetic energy in a dielectric was given in Eq. 5-14 and is repeated here:

$$v = \frac{c}{\sqrt{\epsilon_r}}, \quad (11-9)$$




FIGURE 11-18. Area of effective interplane decoupling capacitance.

where c is the speed of light (12 in/ns) and ϵ_r is the relative dielectric constant of the material.

The radius of the effective capacitance area is equal to

$$r = vt = \frac{12t}{\sqrt{\epsilon_r}}, \quad (11-10)$$

where r is the radius in inches, t is the time required to move the charge in nanoseconds, and ϵ_r is the relative dielectric constant of the material. If we want to move charge on an FR-4 epoxy-glass PCB ($\epsilon_r = 4.5$) in 0.5 ns, then r would be equal to approximately 3 in.

The effective capacitance of a power-ground-plane pair is

$$C = \frac{\epsilon A_e}{s}, \quad (11-11)$$

where ϵ is the dielectric constant, s is the spacing of the planes, and A_e the effective area of the planes.

The capacitance is proportional to the effective area (which is determined by the velocity of propagation and the rise/fall time of the IC), the spacing between the planes, and the dielectric constant. From Eq. 11-11 we can observe that *decreasing the layer spacing is an effective way of increasing the board's effective decoupling capacitance.*

However, if the dielectric constant of the PCB material is increased, to increase the capacitance, Eq. 11-9 tells us that the velocity of propagation will decrease. Thus, if the dielectric constant is increased, then the radius of the effective area decreases by the square root of ϵ , and the area of the effective capacitance circle decreases by the square of this, or proportional to ϵ . Therefore, the area A_e decreases proportional to the increase in ϵ . Hence from Eq. 11-11 we conclude that *changing the dielectric constant of the PCB has no effect on the effective decoupling capacitance*, because ϵ increases as much as the effective area A_e decreases.

Therefore, the effective interplane decoupling capacitance remains the same regardless of the dielectric constant of the PCB material, and the only effective way to increase the effective capacitance is to place the planes closer together. Hence, no advantage is gained on a large, high-speed digital logic PCB, by using exotic (read expensive) high dielectric materials to increase the embedded decoupling capacitance, and the only effective way to increase the effective decoupling capacitance is to place the power and ground planes closer together, or to use multiple power and ground planes.*

*In the case of a PCB with an area less than the area of the effective interplane decoupling capacitance circle, the whole board area will contribute to the effective decoupling capacitance, and in that case increasing the dielectric constant of the material will increase the decoupling capacitance. The ϵ_r of the high dielectric constant material must be used to calculate the propagation delay used to determine the radius of the effective decoupling capacitance circle.

11.4.6.2 Practical Implementation of Embedded Capacitance. In 1989–1990 Zycon, (Zycon since became Hadco and now is Sanmina), developed a special PCB laminate with a 2-mil spacing between layers using standard FR-4 epoxy glass as the dielectric. This laminate known as ZBC-2000[®] provides 500 pF/in² of interplane capacitance.* By using two sets of power and ground planes made from this laminate in a PCB, one can obtain the desired 1000 pF/in² of capacitance.

Although Sanmina (Howard and Lucas, 1992) and Unisys (Sisler, 1991) have patents on the 2-mil thick embedded capacitance printed circuit boards, the technology is readily available and multisourced. Sanmina's trade name for this technology is Buried Capacitance.[®] Although this technology has been available for over 15 years, it is just now becoming popular. Conversion to an embedded capacitance PCB is simple, because no new artwork is required. Only the layer stackup is changed. It is therefore very easy to try out this technology by having two sets of prototype boards made, one the standard way and the second using the Buried Capacitance[®] layers, and then the performance of the two can be directly compared. The most common layer stackup is shown in Fig. 11-19.

Other layer stackups are also possible using the embedded capacitance approach, as discussed in Section 16.4.2. Because of the use of two power-ground-plane pairs, a four-layer board becomes a six-layer board, and a six-layer board becomes an eight-layer board when embedded capacitance is used. Although the ZBC-2000[®] laminate is slightly more expensive than standard




FIGURE 11-19. Stackup of a typical embedded capacitance printed circuit board.

* ZBC-2000 is a registered trademark of Sanmina-SCI, San Jose, CA.




FIGURE 11-20. V_{cc} -to-ground noise voltage from 0 to 200 MHz on (A) a standard printed circuit board, and (B) on an embedded capacitance printed circuit board. (© 1991 UP Media Group)

PCB laminate, the primary cost increase comes from having to print the two extra PCB layers.

Figure 11-20 shows the V_{cc} -to-ground noise voltage from 1 to 200 MHz on an embedded capacitance PCB using the stackup of Fig. 11-19 and on an otherwise identical standard (nonembedded capacitance) board that has only one power and one ground plane (Sisler, 1991). Comparing the two plots in Fig. 11-20 shows that the embedded capacitance board has much reduced noise at all frequencies above 30 MHz and virtually no measurable noise above 60 MHz. For the data shown in Fig. 11-20, the standard board had 135 decoupling capacitors on it, whereas the embedded capacitance board had no decoupling capacitors on it. Both boards contained the same bulk decoupling capacitors. (see Section 11.8)

The decoupling, however, was worse on the embedded capacitance board at frequencies below 30 MHz, in particular at 20 MHz. This is the result of the embedded capacitance board not having sufficient total capacitance to be effective at these low frequencies. This result is similar to that shown in Fig. 11-12, where increasing the number of capacitors reduced the high-frequency impedance of the power distribution network but did not reduce the low-frequency impedance. To reduce the low-frequency impedance, additional capacitance is required as demonstrated by comparing Figs. 11-12 and 11-13.

For the case shown in Fig. 11-20B, adding back four of the original 135 decoupling capacitors (thereby increasing the total capacitance) reduced the low-frequency noise voltage on the embedded capacitance board to what it was on the standard board.

Subsequent measurements, by others, on embedded capacitance printed circuit boards have shown them to provide effective decoupling at frequencies as high as 5 GHz, with few or no discrete decoupling capacitors on the board.

I am not aware of anyone who has made measurements above 5 GHz, but there is no reason to suspect that embedded capacitance approach would not be effective at frequencies even greater than this.

Embedded capacitance boards have other advantages also. The power and ground plane inductances are significantly reduced because of the close spacing of the planes, as well as because of the use of multiple power and ground planes. Resonance problems are also minimized because the embedded capacitance power/ground plane sandwich tends to be a low-Q structure, as a result of its low L/C ratio (see Eq. 10-3). In addition, by removing 90% or more of the decoupling capacitors and their associated vias, board routing is greatly simplified, and in many cases the board can be reduced in size. Also, in some applications, the removal of most of the discrete decoupling capacitors can eliminate the need for surface mounting components on both sides of the board.

Embedded capacitance boards tend to have similar power-to-ground noise voltages at all locations on the board, whereas standard boards, which use discrete decoupling capacitors, have power-to-ground noise voltages that are dependent on the location where the measurement is taken.

I am convinced that in the future, some form of embedded capacitance printed circuit boards will become the standard in the industry, because they provide a simple, effective way to provide high-frequency power supply decoupling using a minimum number of discrete capacitors. This transition to the common use of embedded capacitance may be accelerated when the Sanmina and Unisys patents expire in a few years. Theoretically, distributed capacitance decoupling becomes more effective as the frequency is increased, whereas discrete capacitor decoupling networks become less effective as frequency is increased.

11.4.7 Power Supply Isolation

The only true solution to the high-frequency decoupling problem is the use of a multiplicity of capacitors either discrete or embedded into the PCB.

Another approach, although it does not solve the fundamental decoupling problem, can and has been used to minimize the detrimental side effects of poor decoupling. The objective is to prevent the noisy power plane, which is the result of ineffective decoupling, from contaminating the rest of the PCB.

This objective can be accomplished by isolating, or segmenting, the noisy portion of the power plane from the rest of the PCB's power plane, and by feeding power to the isolated plane through a π -filter, as shown in Fig. 11-21. Note that only the power plane is split, not the ground plane. This approach does not reduce the noise or improve the decoupling effectiveness on the isolated power plane, but it prevents this noise from contaminating the main power plane.

This approach is most effective when only a small number of circuits operate at high frequencies. The high-frequency circuit, or circuits, are then isolated




FIGURE 11-21. Filter used to feed power to an isolated power plane.




FIGURE 11-22. Example of an isolated power plane.

from the remainder of the circuits and powered from an isolated power plane as shown in Fig. 11-22. For example, if a microprocessor is the only chip operating at a high clock frequency, then the microprocessor and its clock oscillator might be powered from the isolated power plane. Another example would be to segregate the clock oscillator and clock drivers from the remainder of the circuitry and power them from an isolated power plane. Multiple isolated power planes could also be used for different circuits.

Figure 11-23 shows the V_{cc} -to-ground noise voltage, from 20 to 120 MHz, on a PCB with an isolated plane powering the microprocessor and its oscillator circuit. Figure 11-23A shows the V_{cc} -to-ground noise voltage on the isolated plane that powers the microprocessor, and Fig. 11-23B shows the V_{cc} -to-ground noise voltage on the main power plane. In Fig. 11-23A, we observe a large amount of clock noise resulting from ineffective decoupling of the microprocessor. Comparing Figs. 11-23A and 11-23B, we observe that the clock




FIGURE 11-23. V_{cc} -to-ground noise voltage from 20 to 120 MHz on (A) an isolated power plane, and (B) on the main power plane.

noise on the main V_{cc} plane is virtually eliminated above 60 MHz, and it is much-reduced below 60 MHz when compared with the noise on the isolated power plane.

Note that when this approach is used, no signal traces can cross over the split in the power plane, on an adjacent layer, as this would interrupt the return current path as discussed in Section 16.3. All signal traces connecting to the ICs powered from the isolated plane must be routed on layers adjacent to a solid plane (power or ground). Therefore, this approach has routing restrictions, whereas the distributed capacitance approach discussed previously performs better and actually simplifies routing as the result of the eliminated capacitors and their associated vias.

The isolated power plane approach is sometimes implemented by printing the isolated plane on one of the signal layers, instead of splitting the actual power plane. In this case, the isolated power plane is often referred to as a power island or "power puddle." The results are the same with either the split power plane or the power island approach.

11.5 THE EFFECT OF DECOUPLING ON RADIATED EMISSIONS

The transient power supply currents drawn when digital logic ICs switch produce radiated emissions by the following three different mechanisms:

1. The transient power supply currents flowing around the loop between the IC and the decoupling capacitor will cause the loop to radiate.
2. The transient power supply currents flowing through the impedance of the ground will produce a ground noise voltage that will excite any cables connected to the system, which causes them to radiate.

3. The V_{cc} -to-ground noise voltage couples, via the power bus, to other ICs (e.g., I/O drivers) and eventually ends up on power and signal cables, which causes them to radiate.

Effective decoupling can minimize all three of the above radiation mechanisms.

Figure 11-24A shows an IC and its decoupling capacitor C_1 . The radiation from the decoupling loop is proportional to the loop area times the magnitude of the current in the loop (see Eq. 12-2). The closer the capacitor is to the IC and the smaller the loop area, the less the radiation.

In addition, the transient power supply current also produces a voltage drop in the ground, which is shown as V_g in Fig. 11-24A. This ground noise voltage will excite any cables connected to the system, which causes common-mode radiation from the cables as shown in Fig. 12-2 and discussed in Section 12.2.

Figure 11-24B shows the same IC as in Fig. 11-24A, but this time it has two identical decoupling capacitors C_1 and C_2 ; one capacitor is located on each side of the IC. Because the two capacitors are equal in value and are located symmetrically with respect to the IC, they will each provide half of the current required by the IC when it switches. Therefore, the radiation from the loop that contains capacitor C_1 will be cut in half, or reduced by 6 dB. The current in the loop that contains C_2 will contain the other half of the transient current, and it will also produce an equal amount of radiation.




FIGURE 11-24. Current in the IC to decoupling capacitor loop, for the case of (A) one decoupling capacitor, and (B) two identical decoupling capacitors.