

15 BYPASS CAPS AND DECOUPLING SYSTEMS

There are two views regarding bypass capacitor decoupling strategies today. They appear to be quite different on the surface, and they lead to some differences in results when taken to their logical conclusions. I tend to call them (a) the traditional approach, viewing bypass capacitors almost as little storage batteries (what I sometimes refer to as "little storage buckets of electrons"), and (b) the power system impedance approach, where the focus is on achieving a power distribution system impedance level below some target maximum at all relevant frequencies. We look at both of them here, and I'll highlight the differences in conclusions between them.

TRADITIONAL APPROACH

Figure 15-1 shows a representative logic circuit. If the input goes low (logical zero) Q1 turns on and pulls the base of Q2 low, thereby also forcing the base of Q4 low. Q3 turns on and the output goes high (logical one). Thus, this logic device is an inverting gate. That is, the output goes in the opposite direction from the input.

In a typical application, logic gates are connected together, as shown in Figure 15-2, with the output from one connected to the input of another. Thus, when the input of the first gate goes low, for example, its output will go high, the input of the second gate then goes high, and the output of the second gate goes low.

Logic circuits are designed so that the signal level of their logical states falls within a known and expected range. That is how the second gate knows that it is receiving either a logical one or a logical zero from the first gate. But what if a logical zero value from the first gate was a high enough voltage level that the sec-

INTEL-1010
U.S. Patent 7,151,430

ond gate either could not "decide" what its level was, or it mistook it for a logical one? A logical system error would then result and the system would likely fail.

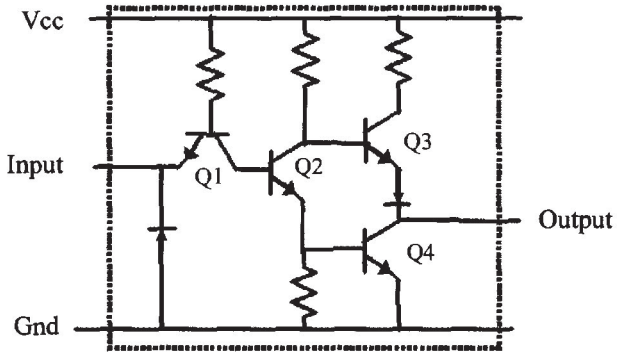


Figure 15-1 Typical logic gate.

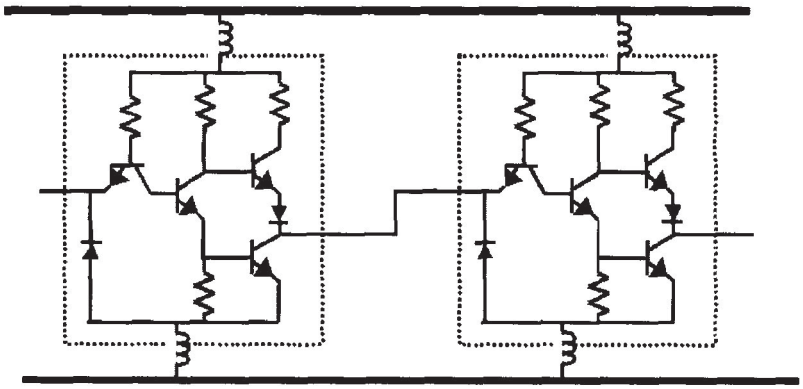


Figure 15-2 Two logic gates connected in series.

Circuits are designed with a wide enough tolerance range so that the devices themselves are unlikely to cause this to happen. But there are external factors that might cause something like this to occur. Consider the situation illustrated in Figure 15-3. Two logic gates are connected in series. There is a common power supply for them, Vcc. There is some level of inductance in the path between the power supply and the devices. (Remember, this was one of the possible contributors

INTEL-1010
U.S. Patent 7,151,430

to common mode problems identified in Chapter 9.) Several factors contribute to this inductance: the planes, the pads, the vias and escapes connected to the pads, and so on. And remember that current always flows in a closed loop.

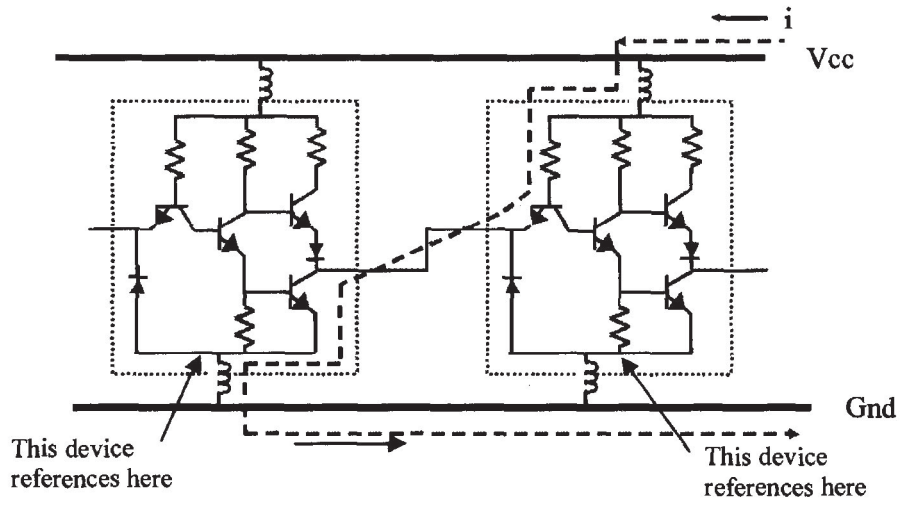


Figure 15-3 Part of current path through first gate when its output goes low.

So when the output of the first device goes low, there is a current path that loops around from Vcc through the power input pin on the second device, through the device to the output of the first device, then through the ground pin, and back to Vcc. In particular, this current flows through the stray inductance shown at the bottom of the device. The transient switching voltage across this inductance is given by the expression V = L x di/dt (Equation 4-4), where di/dt is the change of current divided by the change in time (where the change in time is the rise time or, in this case, the fall time) of the logical transition. If the fall time is very short, we can have a significant voltage across this inductor even if the inductance is small. Furthermore, if this is a large, several-hundred-pin ball grid array (BGA), there might be a substantial amount of transient current flowing though this inductor.

Now the output signal from this device is the sum of the normal logical zero voltage *plus* the voltage generated across the inductor. If the voltage across the inductor is large enough, this signal might look like a logical one to the next device in

INTEL-1010
U.S. Patent 7,151,430

the string, and a logic error would occur. Figure 15-4 illustrates how this might happen. As the signal falls from its high to low state, a voltage proportional to the fall time of the signal appears across the inductor. This voltage *adds* to the signal voltage. The result is the dotted curve shown in the figure. The voltage increment across the inductor, which causes the voltage level associated with this device to rise, or "bounce" up to a level higher than it would normally be, is sometimes referred to as *ground bounce.*

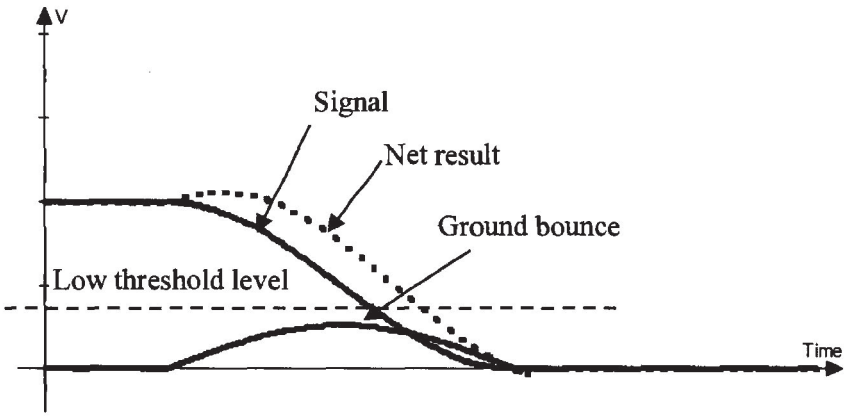


Figure 15-4 Effect of ground bounce on a falling signal transition.

The horizontal line in the figure represents a low threshold. The effect of the ground bounce can be interpreted one (or both) of two ways:

1. At a given point in time after the switch, the signal is higher than it would otherwise be by the amount of the ground bounce.

2. The effective propagation time of the signal slows down; that is, the signal crosses the threshold later than it otherwise would.

Remember that the left side of Figure 15-3 represents only a single gate. There may be a large number of gates all switching at the same time, each adding an increment of current through the stray inductance. If there are a large number of gates, the resulting signal error or delay can prevent the proper operation of the circuit. The issue is this: How do we prevent this from happening?

Figure 15-5 illustrates one way. Place a capacitor around the device that connects between the device and the stray inductances that can contribute to the ground bounce. Figure 15-5 illustrates where. Now, when there is a short-term transient surge of current, the required electrons (remember, current is the flow of electrons) are stored conveniently in the capacitor rather than having to transition from the power supply. This is where the term "little storage buckets of electrons" comes from.

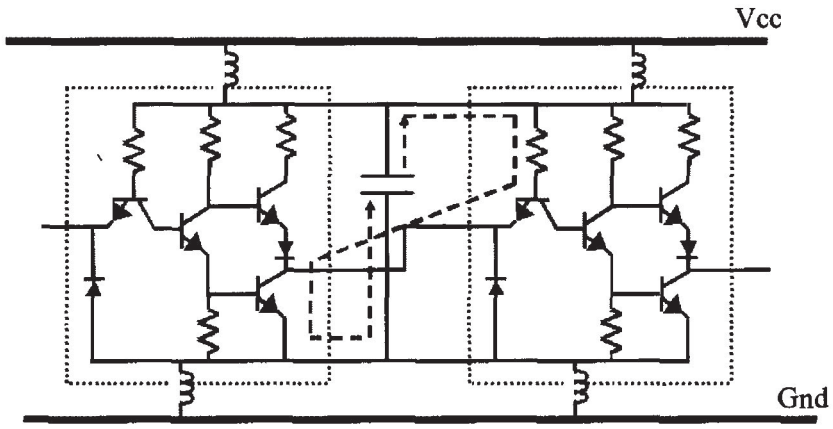


Figure 15-5 A properly placed capacitor stores and then delivers the short-term transient charge needed for switching.

Conceptually, this answer is pretty simple. It does raise a few issues, however. For example, what size capacitor do we use? Do we use more than one capacitor? Where do we place the capacitors? How do we connect them to the circuit board? Figure 15-6 helps us consider some of these questions.

Size and Quantity Looking in from the left side of Figure 15-6, let the first curve represent the required amount of charge needed to complete the logical transition. The area under the curve represents the required charge, and the slope of the curve represents when that charge is required. The furthest curve in (far right) represents what the power system can supply. It can supply a very large amount of charge, but probably not very quickly. The first problem is that power systems usually have a large amount of inductance associated with them, so their response is slowed by that inductance. The second problem is that power systems might be placed some distance away from the need. If electrons can travel at 6 inches in a

INTEL-1010
U.S. Patent 7,151,430

nanosecond, and if the device switches in a nanosecond, then all the electrons needed during the switching phase need to be located within 6 inches of the device. Half of them need to be within .5 ns (3 inches) away. One-quarter of them need to be within 1.5 inches, and so on.

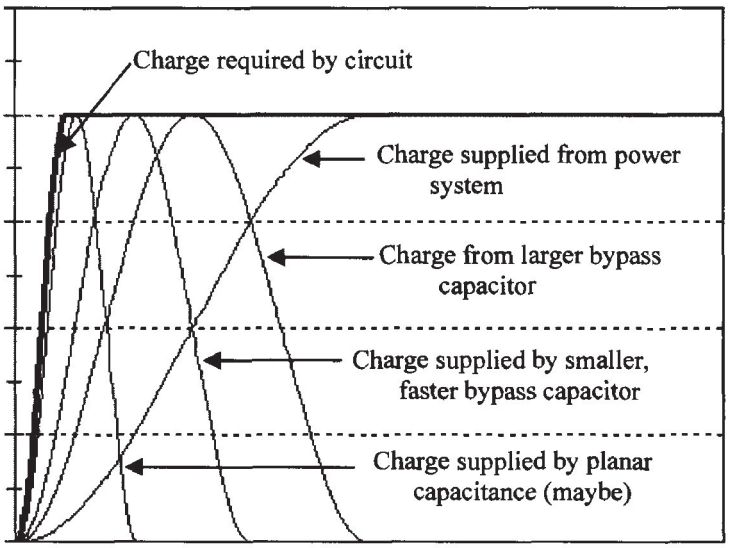


Figure 15-6 Different capacitor types play different roles in supporting transient switching requirements.

The curve next to the power system curve represents what a larger bypass capacitor might be able to provide. It can provide a fair amount of charge, and the inductance associated with it is probably smaller than that of the power system, but the inductance is still not low enough to allow the capacitor to respond quickly enough to meet all the needs of the device.

Adjacent to that curve is another curve for a smaller capacitor. In general, smaller capacitors have less inductance than larger capacitors do, so they can respond faster. But they also store less charge, so they might not be able, by themselves, to handle the requirement.

INTEL-1010
U.S. Patent 7,151,430

Finally, planar capacitance, if it exists, might be able to supply charge very quickly, since there is very little inductance associated with a plane. There are, however, several issues associated with planar capacitance:

- It does not exist by accident. It must be designed in.
- The charge is distributed across the plane, not localized.
- The charge on the plane is subject to the same 6-inch rule we just described.

Therefore, planar capacitance, if designed in, can supply charge very quickly, but the amount of charge it can supply is limited.

Finally, it is possible (as shown) that all of these sources of charge together might not be able to meet the requirement. If that should be the case, the device simply cannot switch as fast as it wants to and the system slows down.

The conclusion from all this is that often more than one solution may be required. We may need a larger (say .1 μF) capacitor for bulk charge and a smaller (say .01 or even .001 μF) capacitor for fast response. We may find that we cannot reliably meet the switching needs at all without designing in some planar capacitance. That is why the engineer you are working for may require two (or more) decoupling capacitors for each decoupling requirement on the board. That is why a great many boards now being designed for high-speed applications have planar capacitance designed in (see Chapter 16).

Placement It should be clear from the discussion related to ground bounce that one of the purposes for using bypass decoupling capacitors is to minimize the inductance in the current path associated with the switching device. The way to do this is to place the capacitor as close as possible to the switching requirement. (Indeed, some people would argue that it is not possible to get too close.) The question is this: If there are two capacitors, which one is placed closer to the device?

This answer is straightforward and pretty noncontroversial. Place the smaller, faster one closer to the requirement. It is the one that will be providing the first increment of charge, so it should be placed closest to where the charge is needed.

But since the capacitor is probably smaller than the device, should it be placed closer to the power pin or the ground pin? This question can spark an argument that can go on for days without a conclusion. Engineers can have very strong

INTEL-1010
U.S. Patent 7,151,430

feelings about this. After listening to all the comments on e-mail forums whenever this topic is raised, my opinion (and why) is as follows:

The objective of the decoupling capacitor is to stabilize the reference voltage (so devices have a stable reference for logical ones and zeros).

- Therefore the capacitor should be placed closest to the reference voltage.
- The reference voltage is ground (zero volts) for both TTL and ECL logic devices, so the capacitor should be placed closest to the ground pin for these logic families.
- The reference voltage for CMOS devices is typically halfway between the lower and higher supply voltages, so the capacitor in these circuits should be placed an equal distance between the power and ground pins on the devices.

Connection Another source of controversy is how to connect the capacitor to the circuit. One school of thought is to connect everything directly to the power and ground planes. Another school of thought is to run traces between the device pins (power and ground) and the capacitor, and then connect the capacitor and the device to the planes with vias at that point (the capacitor). The argument for this second position is that, by definition, there is switching noise between the device and the capacitor, and we want to keep that noise off the planes so it does not interfere with other devices. The argument for the former is that connecting to the planes provides the lowest inductance path, and therefore the fastest switching. Even more recently, some people have argued that vias can have more inductance than traces, and that in fact the traces are the lowest inductance paths.

I know of no definitive studies or results that address this point. After listening to all the arguments and looking at the results of all the boards we have designed, my personal conclusion is that it just doesn't matter very much which way you connect the capacitors to the circuit. What you *do* want to do is minimize the inductance at the pads and vias, however they are connected.

INTEL-1010
U.S. Patent 7,151,430

Summary Summarizing the traditional approach:

- Use two or more capacitors per requirement, one for bulk and one for speed.
- Select low-inductance capacitor types.
- Use low-inductance pad and via designs.
- Consider designing in planar capacitance for even faster response.
- Place the faster capacitor closest to the reference voltage.

You might recognize that virtually all application notes from semiconductor suppliers (if they address decoupling at all) suggest variations on this approach. Almost none of them suggest decoupling strategies that are based on the power system impedance approach. This is not particularly surprising since their primary interest is the performance of their particular device. Generally, they have little knowledge about the rest of the system design.

POWER SYSTEM IMPEDANCE APPROACH

Recently, there has been a growing interest in looking at the decoupling issue from a total power system standpoint. The argument goes like this: The traditional approach focuses on the individual IC components and thus is wasteful. What it really achieves by default, and what could be achieved more economically if we addressed it head on, is a (target) low impedance for the power system at all relevant frequencies.

Ideal Response As we place bypass capacitors around the board, by default we are creating a power system impedance curve that is "high" at DC but low at all other frequencies. By focusing on the individual ICs we are not optimizing the system. What if, instead, we focused on the system? Then the needs of the individual ICs would be taken care of by default.

The objective, therefore, is to develop a power system bypass decoupling strategy with the ideal characteristic of infinite impedance at DC and zero (or at least very low) impedance at all other frequencies. This type of response is suggested in Figure 15-7. The ideal, of course, can't be obtained, but the issue is how close we can come to the ideal.

INTEL-1010
U.S. Patent 7,151,430

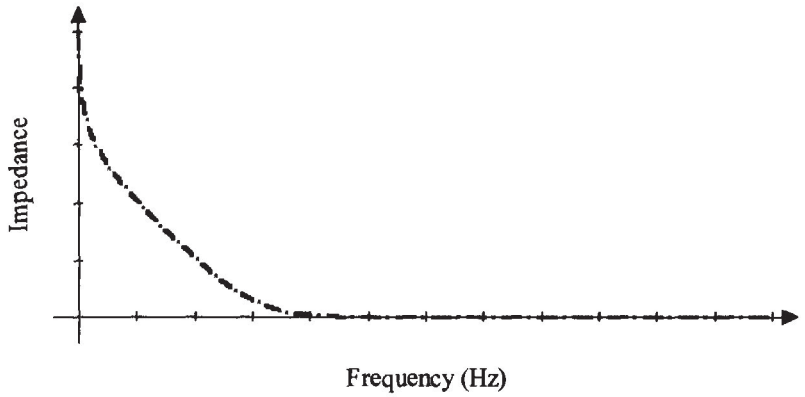


Figure 15-7 Ideal response curve for the power distribution system.

Capacitor Response Let's start with a simple .01 μF capacitor. The impedance, Z, of an ideal capacitor is given by the expression (see Equation 6-1):

$$Z = 1/\omega C = 1/2\pi fC$$
where f = the frequency in Hz

This relationship is shown by the dotted line in Figure 15-8. It is a straight line (on logarithmic scales) and follows the familiar pattern of having very high impedance at very low frequencies and very low impedance at very high frequencies.

Inductance Effects The problem is that a real capacitor doesn't act like this at all. All real capacitors have some inductance associated with them. Some of that inductance is inherent in the structure of the device itself and some is associated with how the device is mounted onto the board. Good SMT devices might have 5 to 10 nH inductance associated with them, so a "real" .01 μF capacitor actually looks more like that shown in Figure 15-9.

The impedance of this capacitor, as a function of frequency, is given by Equation 15-1:

$$Z = \left|2\pi fL - \frac{1}{2\pi fC}\right| = \left|\frac{(2\pi f)^2 LC - 1}{2\pi fC}\right| \qquad [15\text{-}1]$$

INTEL-1010
U.S. Patent 7,151,430

This is the solid curve plotted in Figure 15-8. Note that it follows the curve of an ideal capacitor for lower frequencies, but then as the influence of the inductor comes into play the curve starts to fall more sharply, until it reaches a minimum.

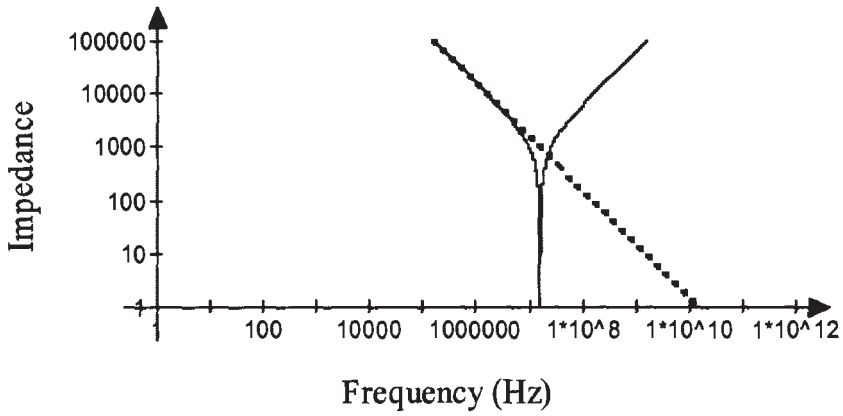


Figure 15-8 Frequency response curve of an ideal and real bypass capacitor.

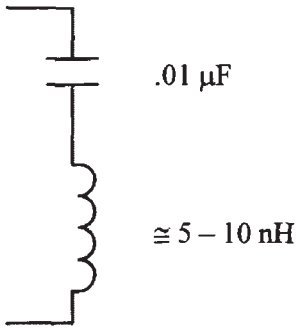


Figure 15-9 Real bypass capacitor.

Then it starts rising rapidly. This minimum occurs when the numerator of the impedance formula goes to zero (refer to Equation 6-13), which happens, in this case (C = .01 μF and L = 10 nH), when:

$$f = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{(.01)10^{-6}(10)10^{-9}}} = 16 \text{ MHz}$$

We call this minimum the resonance point, or the *self-resonant frequency* of the capacitor. It is the point where the impedance of this LC circuit goes to zero. Not just to a very low value; not just close to zero; it goes to (in the absence of any resistance) absolute zero.

There are two things to recognize in Figure 15-8. The first is the existence of a series resonant point, as just discussed. The second is the behavior of the impedance curve above resonance. We refer to the shape of the curve as *downward sloping to the right* before resonance and *upward sloping to the right* past the resonant point. The significance of these terms is that capacitive circuits are characterized by downward-sloping curves and inductive circuits are characterized by upward-sloping curves. So the impedance curve is inductive past the resonant point. (The formal meaning of that is that the voltage phase shift is positive past the resonant point.)

Some people argue that the effectiveness of a bypass capacitor disappears above resonance; that is, above resonance it acts like an inductor and is therefore useless as a capacitor. This is not necessarily true. Yes, the circuit is inductive above the resonant point, but the overall impedance is still less than that for an ideal capacitor for at least some frequency range above resonance. This is because the capacitor and inductor combined can have (and do have) a lower impedance than the pure capacitor alone until some frequency above resonance where the inductive effect finally and fully takes over.

Multiple Capacitors Now let's start doing something of more interest. Let's look at 200 of these .01 μF capacitors all connected in parallel. This is something like what we do on a circuit board when we add a large number of equal capacitors around the various devices. The combined effect is as if we had a single capacitor with a capacitance of 2.0 μF and a lead inductance of only .05 nH. The response for this set of capacitors is shown in Figure 15-10, along with the curve for a single ideal .01 μF capacitor and a single .01 μF capacitor with associated inductance. Note the set of 200 capacitors still has the same effective resonant frequency. Since each of the individual capacitors has the same self-resonant frequency, the set must have the same self-resonant frequency. (Actually the set has 200 individual self-resonant frequencies, all at the same point.)

INTEL-1010
U.S. Patent 7,151,430

Now there are two things to note about Figure 15-10. The first is that at every frequency the impedance of the 200 capacitors is significantly lower than that of an individual capacitor. And, there is a significant range over which that impedance is relatively low. Remembering that our objective is to develop a strategy for achieving a low impedance across a wide frequency range, this curve already suggests part of the approach to getting there—add more capacitors.

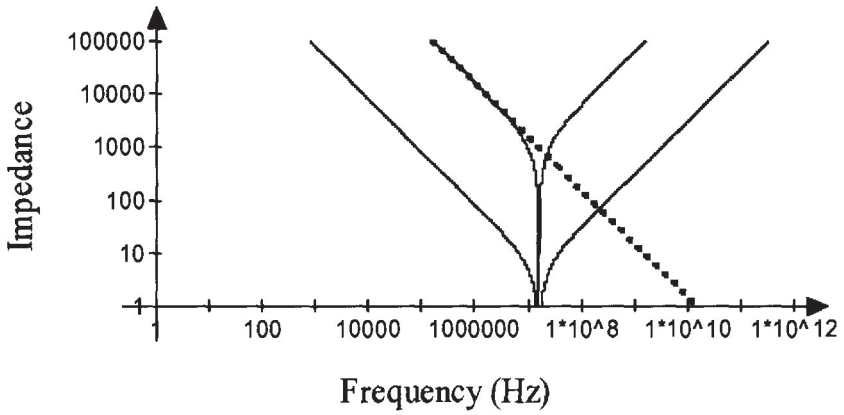


Figure 15-10 Response when we add multiple values of a capacitor.

The second is that there is a significant frequency range (from about 16 MHz to about 200 MHz) where the new impedance curve is inductive, yet the impedance is still well below that for the ideal individual capacitor. Just because the impedance curve is inductive doesn't mean that it isn't meeting the objective of providing low impedance for the power system.

Additional Capacitor Value We can carry our strategy a step further and add additional capacitors of a different value. For example, lets add 20 each 50 μF capacitors. This will provide low impedance at the lower frequencies. These larger capacitors, however, are likely to have larger associated inductance with them. Assume that inductance to be 20 nH per capacitor. The combined effect of these capacitors is shown on the curve in Figure 15-11, along with the effect of the 200 each .01 μF capacitors we have already looked at. We are now beginning to develop a frequency response curve that really begins to look like our objective. If we combined the best parts of the two individual curves, the impedance level would be less than 1 Ω from 100 Hz to approximately 3 GHz.

Planar Capacitance We could, if we wanted to, incorporate into the board stackup some planar capacitance. We do that by placing a power and ground plane adjacent to each other very closely. The smaller the distance between the two planes, the greater the capacitance. A significant advantage to planar capacitance is that there is very little inductance associated with it. We can often assume that it is a pure capacitor over the range of interest.

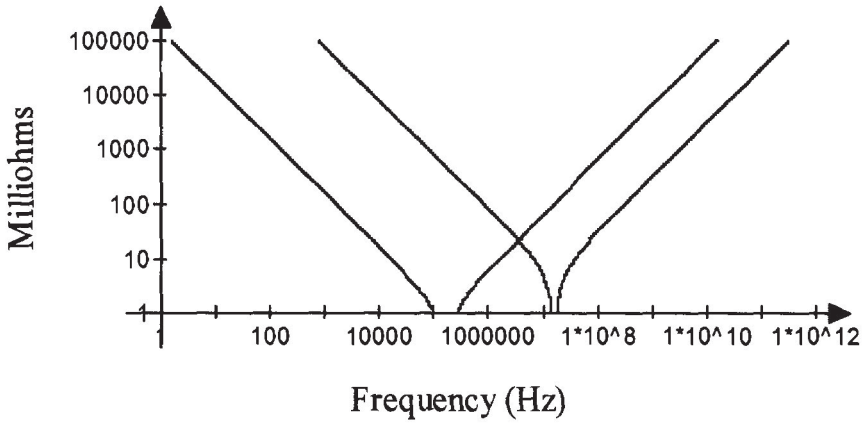


Figure 15-11 Frequency response of two different sets of capacitor values.

A reasonably good formula for approximating planar capacitance is given by Equation 3-10:

$$C = \frac{.2248 \times \varepsilon_r \times A \times (n-1)}{d}$$

where C is the capacitance in pF
ε_r is the relative dielectric coefficient
$\varepsilon_r = 1$ for air
$\varepsilon_r \approx 4$ for FR4

INTEL-1010
U.S. Patent 7,151,430

A = area of the overlapping planes, in^2
d = distance between planes, in inches
n = number of plates

In using this formula, be sure to do the following:

- Be careful to observe the units you are using.
- Account for only the overlapping area of the planes.
- Adjust for such things as holes and voids.

A reasonable number for the amount of capacitance that can be achieved with planar capacitance is in the range of 100 to 300 pF per square inch.

Continuing on with our illustration, lets add 950 pF of planar capacitance to the other capacitors we have assumed. Figure 15-12a shows the result. It would be tempting to look at only the best part of each curve in Figure 15-12a to get something that looks like Figure 15-12b. In this figure, the overall impedance is less than 200 mΩ at all frequencies above 1 kHz. If our objective was to define a set of capacitors that provided a power system frequency response curve less than 200 mΩ at all relevant frequencies, this set would appear to meet the target.

Appearances can be deceiving, however. In Figure 15-12a, look at each point where a downward-sloping line intersects an upward-sloping line. The downward-sloping line is capacitive. The upward-sloping line is inductive. At the point where they cross, their impedances (really their reactances) are equal in magnitude but opposite in sign. This is characteristic of the point where a parallel LC circuit goes into resonance. And at that parallel resonance point, the impedance rises infinitely high. So there is really an infinite impedance peak (pole) at the points where each pair of upward-sloping and downward-sloping curves cross.

The actual frequency response curve is shown in Figure 15-13. The formulas and calculations for this type of problem are quite complex. We at UltraCAD wrote our own calculator specifically to work with this type of problem and to generate graphical results for analysis. The result from that calculator is shown in Figure 15-13. Note, in particular, that the axes are scaled slightly differently between Figures 15-12 and 15-13, accounting for part of the visual difference between them, even though the same circuit is being graphed. Note also that the graphs should extend infinitely high at their peaks and infinitely low at their minimums. They don't because of the graphical limitations associated with the program, not the algebra.

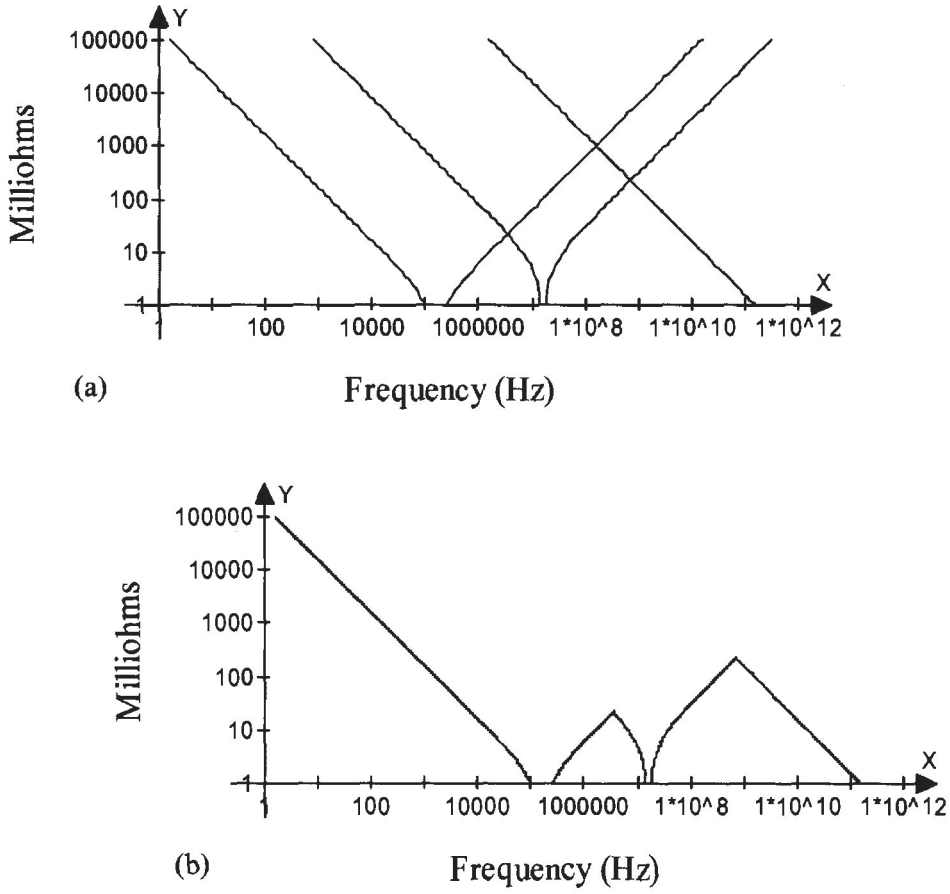


Figure 15-12 Power system frequency curve developed from 20 each 50 μF caps, 200 each .01 μF caps, and 950 pF in planar capacitance.

The implications of this type of curve can be very serious. Most stray frequencies that happened to find their way onto the power system would be shunted to power or ground through an impedance represented by the curve, less than .2 Ω at most frequencies. But any stray noise frequencies (harmonics) that just happened to be near the peaks of the curve shown in Figure 15-13 would *not* be shunted to the power system. They would be free to wander over the system with no attenuation, and would ultimately radiate out off the board. If a board whose power system impedance curve looked like Figure 15-13 were tested for EMI radiation, it would

INTEL-1010
U.S. Patent 7,151,430

likely fail. We would expect the strongest radiation frequencies to be near the peaks of the curve.

That is why we don't want any peaks in our impedance curves.

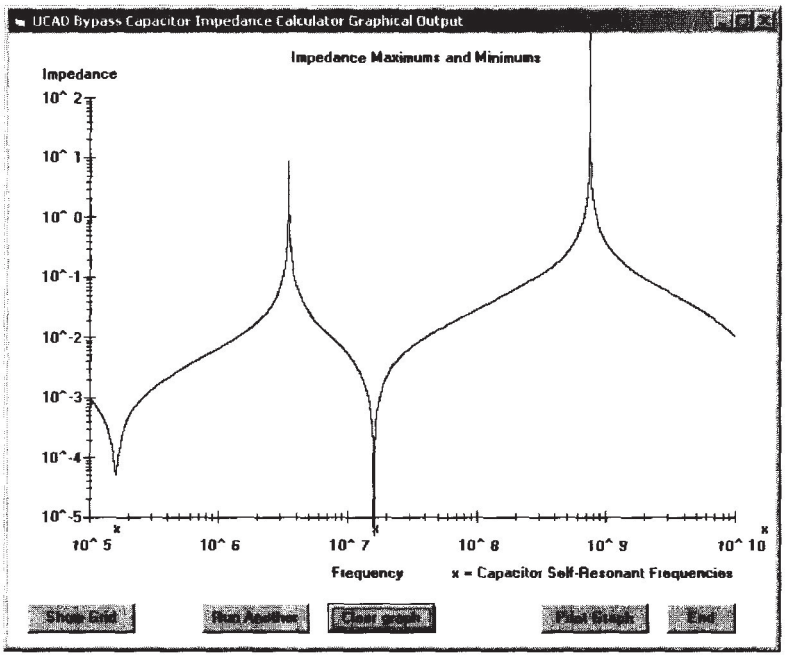


Figure 15-13 Actual response curve of the selected capacitors.

Equivalent Series Resistance (ESR) Well, we know the real world can't be as bad as that shown in Figure 15-13, and it's not. The reason is that bypass capacitors are not pure LC circuits. The equivalent circuit for a typical capacitor actually looks like Figure 15-14. There is a resistance in series with the capacitor and the associated stray inductance. We call this resistance ESR. It is usually very low and is often ignored in most circuit applications.[1]

It turns out ESR has an important role to play in bypass capacitor applications. Most engineers understand that it effectively lowers the peaks and raises the

[1]The equivalent circuit for a capacitor could also include a parallel "leakage" resistance around the capacitor. This equivalent parallel resistance is very high and is ignored in most circuits.

troughs in the power system impedance curve. The peak of the impedance curve is inversely proportional to the ESR value; that is, $Z \approx 1/ESR$. Figure 15-15 illustrates the effect ESR can have on the power system impedance curve shown in Figure 15-13.

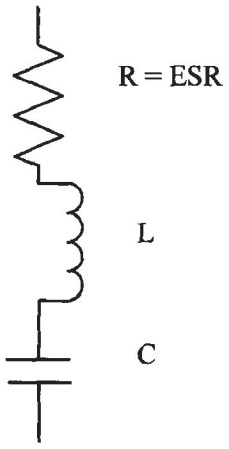


Figure 15-14 A real bypass capacitor includes both inductance and series resistance.

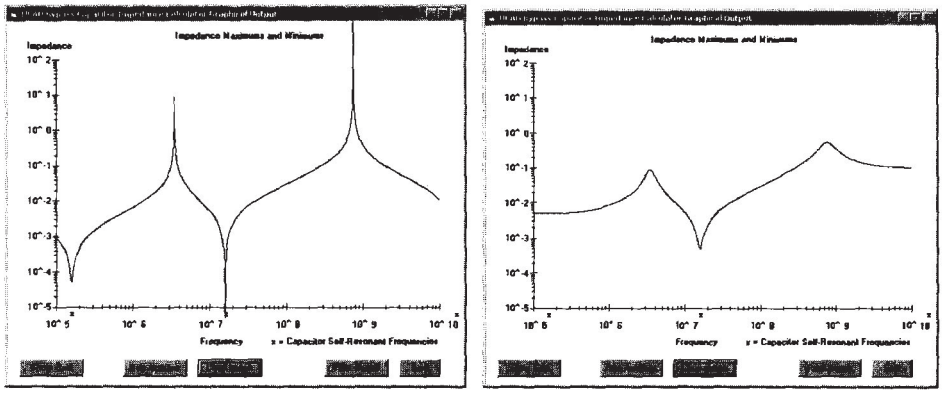


Figure 15-15 Lower ESR (left) will increase the peaks and lower the troughs of the impedance curve; higher ESR (right) will tend to smooth the response.

What a great many engineers do not understand, however, is the impact that ESR has on the overall system impedance curve when a significant number of by-

INTEL-1010
U.S. Patent 7,151,430

pass capacitors are connected in parallel. For example, here are three myths that many engineers firmly believe but are simply not true when ESR begins to influence a circuit.

- *Myth 1:* The system series resonant points (the minimum values of the curve, or the troughs) are at the self-resonant points of the individual capacitors. Truth: When two (or more) capacitors are used with individual series resonant points close together, the system series resonant points are *between* the individual capacitor self-resonant points.

- *Myth 2:* The minimum impedance of the system equals ESR. Truth: The minimum value of the impedance function is *less than* ESR, and it gets smaller as the self-resonant frequencies of the capacitors used in the system get closer together.

- *Myth 3:* Using lower ESR capacitors leads to a better system impedance response curve than using higher ESR capacitors does. Truth: Moderate ESR is better than lower ESR from a system impedance standpoint.

The following analysis demonstrates these points.

Self-Resonant Frequencies Assume a simple capacitor with capacitance C, inductance L, and ESR equal to R. The inductance should be considered from the practical sense—not only the inherent inductance associated with the capacitor physical structure itself, but also the PCB pads and attachment process, and so on. The impedance through this capacitor is given by Equation 15-2:

$$Z = R + j\,\omega L + 1/j\,\omega C, \text{ or}$$
$$Z = R + j(\omega L - 1/\omega C) \qquad [15\text{-}2]$$

where ω is the angular frequency, $\omega = 2\pi f$.

Resonance occurs, by definition, when the j term is zero (see Equation 6-13) as shown in Equation 15-3:

$$\omega L = 1/\omega C \qquad [15\text{-}3]$$
$$\omega^2 = 1/LC$$
$$\omega = 1/\sqrt{LC}$$

The impedance through the capacitor at resonance is R.

Effects of Multiple Capacitors Assume we have n identical parallel capacitors, as earlier. The equivalent circuit of the n identical capacitors is the single capacitor whose values are

$$C = nC$$
$$L = L/n$$
$$R = R/n$$

The impedance of this system is now given by Equation 15-4:

$$Z = R/n + j(\omega L/n - 1/\omega nC) \qquad [15\text{-}4]$$

The resonant frequency of this system is, again, where the j term goes to zero, or where

$$\omega L/n = 1/\omega nC$$

which results in exactly the same self-resonant frequency as before. Paralleling capacitors does not change the self-resonant frequency, but it effectively increases the capacitance, reduces the inductance, and reduces the ESR compared to a single capacitor. The resulting impedance response curve tends to "flatten out" or broaden compared to a single capacitor, as shown in Figure 15-10.

Historically, on circuit boards, circuit designers have used a large number of bypass capacitors of "the same" value. The advantage of this process has been the increased C and the reduced L and R that result. But the capacitor values are typically not exactly equal. Each one has a tolerance.

Parallel Capacitors Take the case of two parallel capacitors of different value, shown in Figure 15-16. Let R1 = R2 = R to simplify the arithmetic. (This assumption does little harm and greatly helps the intuition.) Let us also assume that:

$$C1 > C2$$
$$L1 > L2$$

which means that fr1 (the self-resonant frequency of C1) is lower than fr2.

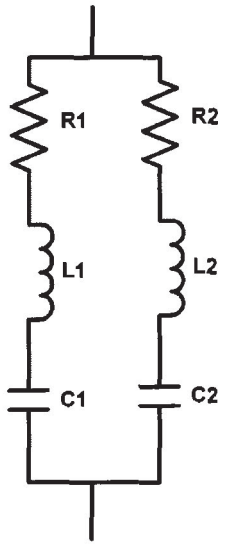


Figure 15-16 Two capacitors in parallel

Now, the combined impedance through this parallel pair of capacitors is given by Equation 15-5:

$$X1 = \omega L1 - 1/\omega C1$$
$$X2 = \omega L2 - 1/\omega C2$$
$$Z1 = R + jX1$$
$$Z2 = R + jX2$$

$$Z = \frac{1}{\frac{1}{Z1} + \frac{1}{Z2}} = \frac{(R + jX1)(R + jX2)}{2R + j(X1 + X2)} \qquad [15\text{-}5]$$

From this, we can derive the real and imaginary terms of the impedance expression (multiply numerator and denominator by (2R – j(X1+ X2)) as shown in Equations 15-6 and 15-7:

$$Re(Z) = \frac{R\left[2\left(R^2 - X1X2\right) + (X1 + X2)^2\right]}{4R^2 + (X1 + X2)^2} \qquad [15\text{-}6]$$

$$\mathrm{Im}(Z) = \frac{(X1 + X2)\left(R^2 + X1X2\right)}{4R^2 + (X1 + X2)^2} \qquad [15\text{-}7]$$

Further, we derive that the magnitude and phase of the impedance term are given by Equations 15-8 and 15-9:

$$|Z| = \sqrt{\mathrm{Re}(Z)^2 + \mathrm{Im}(Z)^2} \qquad [15\text{-}8]$$

$$\theta = \tan^{-1}\left(\frac{\mathrm{Im}(Z)}{\mathrm{Re}(Z)}\right) \qquad [15\text{-}9]$$

The curve of impedance as a function of frequency is shown in Figure 15-17. It is instructive to look at this curve and the real and imaginary terms of the impedance expression formula together.

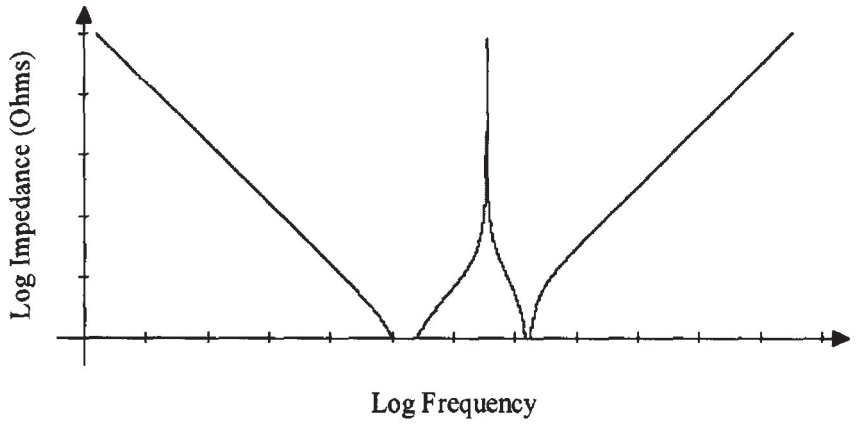


Figure 15-17 Two capacitors in parallel.

Resonance Resonance occurs (by definition) when the imaginary term (Equation 15-7) is zero. This is also the point at which the phase angle is zero. The impedance at that point is simply the real part of the impedance expression.

INTEL-1010
U.S. Patent 7,151,430

The imaginary term for Z goes to zero under the two conditions given by Equations 15-10 and 15-11:

$$X1 = -X2 \quad [15\text{-}10]$$
$$R^2 = -X1X2 \quad [15\text{-}11]$$

The first condition represents the "pole" between the self-resonant frequencies of the two capacitors. X1 equals –X2 when the reactance term of C1 is inductive (+) and increasing, the reactance term for C2 is capacitive (–) and decreasing, and where the two reactance terms are equal. This is the *anti-resonance* point that occurs at a frequency between fr1 and fr2.

Assuming R is small, the second condition can only occur where either X1 or X2 is small. X1 is small near fr1 and X2 is small near fr2. X1 and X2 must be of opposite sign, since R^2 must be positive. Therefore, these system resonant points must be *between* fr1 and fr2, and they must not be equal to fr1 or fr2 (unless, in the limit, R = 0).

Conclusion The system resonant frequencies are not necessarily the same as the capacitor self-resonant frequencies unless ESR is zero (so much for Myth 1).

It can further be shown that at this point, where the imaginary term is zero and $R^2 = -X1X2$, the real term, and thus the impedance itself, simply reduces to R.

Impedance at fr1 At fr1, the self-resonant frequency of C1, X1 = 0. It can be shown that the phase of the impedance function at this point is

$$\Theta = \tan^{-1}[RX2/(2R^2 + X2^2)]$$

If X1 = 0, then X2 must be negative (capacitive, under the conditions we have been assuming) so

$$\Theta < 0$$

Only in the limit where R = 0 does Θ go to zero.

The *magnitude* of impedance at the point where X1 = 0 is shown by Equation 15-12:

$$|Z| = R\sqrt{\frac{R^2 + X2^2}{4R^2 + X2^2}} \qquad [15\text{-}12$$

This is *less than* R for any positive (i.e., R > 0) value of R. In the limit, it is equal to R for R << X2 and equal to R/2 if R >> X2.

The results are exactly symmetrical if we are looking at fr2, the point where X2 = 0.

Conclusion The minimum value for the impedance function is at a frequency other than the system self-resonant frequency and also at a frequency other than the self-resonant frequency of the capacitor. It is less than ESR (R) when two capacitors are connected in parallel. Further, the minimum value declines as X2 gets smaller, or as the self resonant frequencies of the capacitors are moved closer together, or as the number of capacitors increases.

Impedance at Anti-resonance If we let X1 = –X2, then Im(Z) goes to zero, by definition. This is the anti-resonant point between fr1 and fr2. At this point, the system impedance (Z) is given by Equation 15-13:

$$Z = \frac{R}{2} + \frac{X^2}{2R} \qquad [15\text{-}13]$$

For small values of R, this is inversely proportional to R and can be a very large number if R << X. This is why there is concern about very high impedances at the anti-resonant point. If R, on the other hand, is only in the range of .1 or .01, then this number might be more manageable. But note that if R >> X, then at the anti-resonant point Z ≅ R/2. This is less than the impedance at the system resonant point. Do not overlook the significance of this result. The impedance at the anti-resonant *peak* can be *less* than the impedance at the system resonant point!

So consider this: If Z equals (approximately) R at the minimum (the system resonant point), under what conditions is Z also equal to R at the maximum (the system anti-resonant point)? Remember R equals the capacitor's ESR. Under those conditions, the impedance curve will be (at least approximately) flat. It turns out that Z at the anti-resonant frequency exactly equals R if:

$$R = X1 = -X2$$

Therefore, we can achieve a (relatively) flat impedance response curve if we position our capacitor values such that, at the anti-resonant points, X1 = –X2 = ESR.

This has a very significant consequence. As ESR gets smaller, then, for a flat impedance response, X1 and X2 must be smaller at the anti-resonant points. This means that fr1 and fr2 must be closer together. And *this* means that as ESR gets smaller, it requires more capacitors to achieve a relatively flat impedance response.

Conclusion Low values of ESR require more capacitors to achieve a smooth response. Moderate values of ESR lead to smoother response curves with more manageable anti-resonant peaks (so much for Myth 3).

General Case Analysis As we add more values for C, the algebra associated with these kinds of analyses gets very difficult. We at UltraCAD wrote our own program so we could look at various capacitor configurations and see what happens in a more real-world situation (see Appendix H). The following examples illustrate some of the results we achieved using that calculator.[2]

Figure 15-18 shows the impedance curves for three different sets of bypass capacitors used to design a power system. The sets included 100, 150, or 200 individual bypass capacitors, respectively. The capacitor values were carefully selected so that their self-resonant frequencies were spread evenly over a defined frequency range (evident in the figure). The graphical output covers a frequency range from about 1.0 MHz to about 1,000 MHz. To this extent, the example is contrived, but it illustrates what could be achieved under the most controlled of conditions. It is representative of what would happen over more general conditions. All capacitors had the same ESR, .01 Ω.

What is apparent from the figure is that the impedance curve improves as the number of capacitors increases. More important, each curve is *everywhere* better than the curve above it. That is, the peaks of the 150-capacitor curve are lower than the troughs of the 100-capacitor curve, even though the ESR is the same in both cases. This alone would seem to disprove Myth 2.

[2]See also Brooks, "ESR and Bypass Capacitor Self-Resonant Behavior: How to Select Bypass Caps," available at www.ultracad.com.

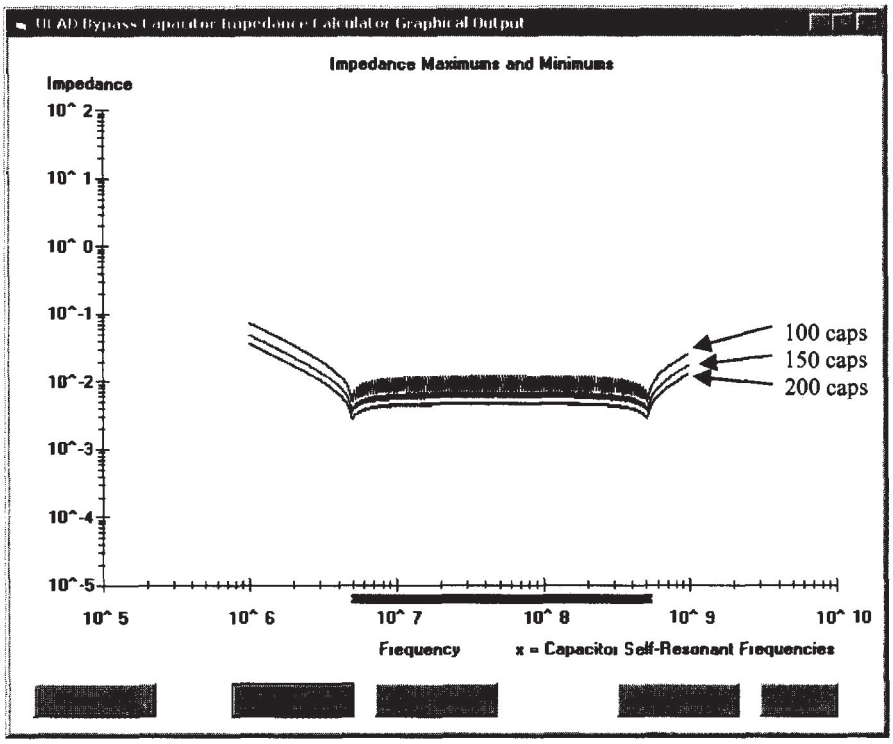


Figure 15-18 Power system impedance response improves with increasing numbers of capacitors.

Figure 15-19 shows the same 150-capacitor curve as Figure 15-18, with ESR equal to .01 Ω, along with a second curve for the same 150 capacitors with ESR of .001 Ω. Although the .001-Ω ESR curve has lower troughs than the other curve, it also has higher peaks. One would be hard-pressed to argue that the lower ESR curve was the better power system impedance response curve. So much for Myth 3.

Consequences So what does all this show us? First, it shows that it is possible to design a power system impedance response curve for any system that can meet an arbitrary target impedance at all frequencies. This type of analysis also suggests that this can probably be done with a smaller number of capacitors, overall, than we would use using the traditional approach to bypass capacitor decoupling. A further advantage of this approach is that, at least conceptually, we could calculate the exact shape and value of the response curve at all relevant frequencies. Obviously, it is not the board designer who would design this kind of decoupling

strategy; it is the circuit and system design engineer. But this design approach would lead to slightly different conclusions than would the traditional approach.

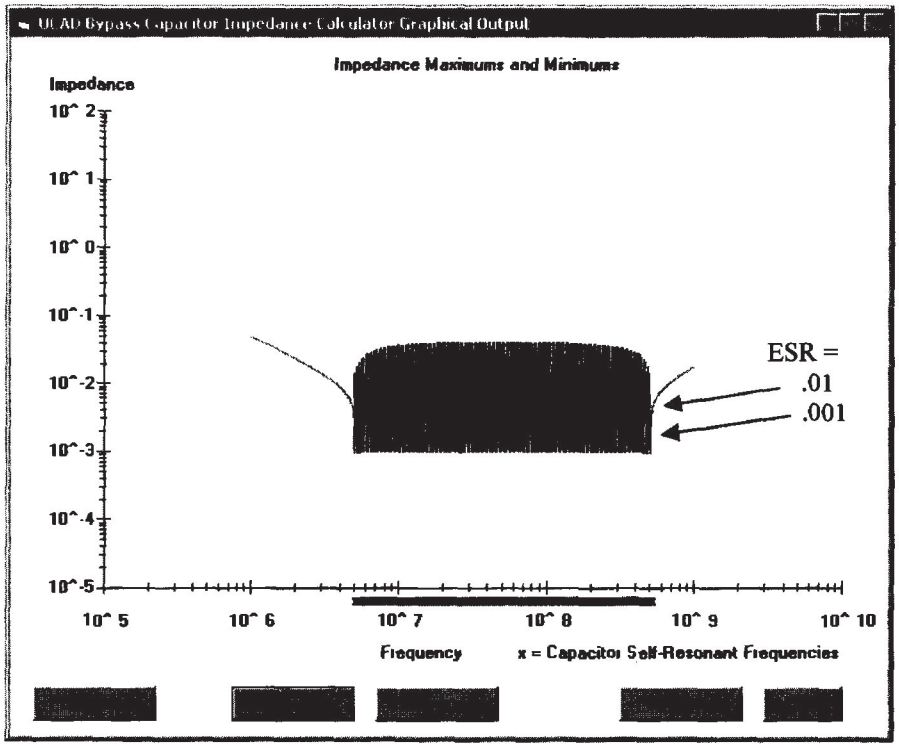


Figure 15-19 Power system impedance response improves with more moderate ESR.

Those conclusions would include the following:

- You must pick the capacitor type carefully, both from the standpoint of capacitance value and of self-resonant frequency.
- Intentionally designed-in planar capacitance is almost a necessity.
- It is best to have a wide range of values and self-resonant frequencies (taking into consideration both the inductance associated with the capacitors and that associated with the pads and the assem-

bly processes mounting them on the boards), rather than to have smaller clusters of identical values.

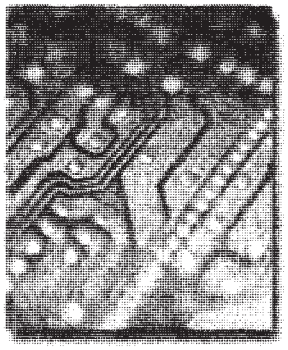
- Moderate ESR is better than low ESR.
- Capacitor placement is not particularly important, except that the electrons still must get (traveling at 6 in/ns) to where they are needed in a timely manner.

SUMMARY

The jury is still out on the relative merits of these two approaches to bypass capacitor decoupling. Personally, I believe they both have a lot to offer. Overall, I like the logic behind the power system impedance approach, and the results of the quantitative analyses are hard to ignore. Intuitively, however, it is hard to ignore that (especially large) logic devices require a lot of charge support at the moment of switching. The traditional approach argues that you can't place charge supporting capacitors too close to the requirement, something I also believe in.

I believe the very best approach to capacitive decoupling is a combination of these two approaches. Fundamentally, use the power system impedance approach, but then make sure there is enough decoupling at every pin to meet each device's requirement.

INTEL-1010
U.S. Patent 7,151,430



16 POWER SYSTEMS

*D*esigners are sometimes confused about how to deal with multiple power planes on printed circuit boards and how to deal with their separation. We often talk about the horizontal separation between planes as a *split,* and there are some design rules that may apply when designing in the area of splits.

Before we talk about how to deal with splits between planes, it is instructive to talk about why we have separate power planes in the first place—to accommodate separate power supplies. One reason we might have separate power supplies is because we might have different voltages in our circuits. The issue of separate voltages might thus be a reasonable place to start.

POWER SUPPLY VOLTAGES

On any given board, or in any given application, we may find DC power supply voltages of 5 V, 3.3 V, 2.5 V, 1.8 V, and even 1.2 V. Older designs have included 12-V supplies, and certain types of circuits may have special design requirements up to thousands of volts.

There are a variety reasons these different voltages exist. One reason relates simply to the semiconductor technology used in the fabrication of the IC making up the circuit. Newer technologies and newer physical semiconductor structures often have inherently lower saturation and switching voltages, leading to lower overall supply voltage requirements for their operation.

A primary driving force for lower voltage relates to the power required for the circuit's operation. If the same circuit function can be achieved with a lower

INTEL-1010
U.S. Patent 7,151,430

voltage supply, a power saving is usually achieved. Power is the product of voltage times current. Therefore, the power reduction can be directly proportional to the voltage reduction. However, if the voltage reduction also leads to a lower current requirement (e.g., through Ohm's Law), the power reduction can be proportional to the square of the voltage reduction. This has obvious implications for battery-operated circuits and the thermal management of larger, more complex circuits.

Another driving force for lower operating voltages is rise time. If, for example, a technology can switch (dV/dt) a 5-V signal in 2 ns, perhaps it can be modified to switch a 2.5-V signal in half that time, with no other technological changes required.

Finally, different circuit functions may require special voltages. These functions might include such things as transmitter stages, high voltages required by cathode ray tubes and laser printers, switching requirements for electromagnetic relays and machine tools, and audio speaker requirements.

Different voltages require separate regulation circuits. This factor would define a minimum number of regulated power supplies, one for each individual voltage. However, there may be additional, multiple regulated supply sections for any given voltage. For example, we may have different regulated supplies for the analog and digital sections of a board, even if they both have the same overall power supply voltage requirement.

We may also have different regulated supplies for different or sequential stages of the same circuit. For example, if there are different circuit paths for the individual R, G, and B components of a video signal, each path might have its own regulated power supply, even though the individual power requirements are identical. In extreme cases, some engineers design regulated supplies for each stage of a signal path through a circuit.

The reason for multiple regulated supplies is usually noise isolation. All signals flow in a closed loop. For every signal flowing down a trace, there is a return signal coming back. The return signal is usually on the plane closest to the trace, and (at least in fast rise-time circuits) is usually as close as possible to the trace.

Although return current density is greatest underneath the trace, current gradients spread out on the plane beyond the edges of the trace (see Figure 16-1).

INTEL-1010
U.S. Patent 7,151,430



Figure 16-1 Current distribution of a return signal spreads out on the plane underneath the signal trace.

Here is a point that is important to understand: Current is the flow of electrons. Figure 16-1 shows the variation in current density near a trace. It necessarily, therefore, illustrates the *electron density* in the same region. Electrons have charge. Charge density is voltage. Therefore, there is also a *voltage density* (actually a voltage gradient) that occurs near a trace.

Variations in voltage gradients can occur for several reasons. One is high-frequency return current, as suggested in Figure 16-1. Another is simply the flow across the plane of current from the regulated power supply required for charging (and discharging) the various bypass (and planar) capacitances associated with the circuit. Voltage gradients across planes of as much as 250 mV are not at all uncommon on circuit boards.

These (changing) voltage gradients constitute noise, and it is instructive to note that this noise is not coming from somewhere totally outside our circuit. In fact, it is generated by the circuit itself. One guideline for good power supply management is not to focus on preventing noise on the planes from getting into our circuits. It is, instead, to *keep the noise that is generated by our circuits from getting onto the planes in the first place.*

One reason engineers use different regulated power supply regions for different circuits and different stages of the same circuit is to try to prevent noise generated by one (part of a) circuit to interact with, and interfere with, signals in another (part of the) circuit.

INTEL-1010
U.S. Patent 7,151,430

NEED FOR POWER PLANES

We have now covered several topics in this book about why power and ground planes are required in high-speed design applications. Here is a summary of some of those important reasons.

Impedance control (Chapters 10, 11, 14): If we want to control trace impedance as a strategy for the control of reflections (using proper trace termination techniques), then good, solid, continuous planes are almost always required. It is very difficult to control trace impedance without the use of planes.

Loop areas (Chapter 9): Loop area can be visualized as the area defined by the path of the signal (traveling down a trace) and its return current. When the return signal is on the plane immediately under the trace, loop area is minimized. Since EMI is directly related to loop area, EMI is minimized when good, solid, continuous planes exist under traces.

Crosstalk (Chapters 12, 13): The two most practical ways to control crosstalk are (a) separation between traces and (b) closeness of the traces to their reference planes. Crosstalk is inversely proportional to the square of the distance between the traces and their reference planes.

Need for Plane Pairs

Decoupling (Chapter 15): The capacitance formed by the proximity of two planes placed close together can be very important and beneficial in circuit decoupling at very high frequencies, where bypass capacitors and their associated mounting and lead inductance begin to have problems.

EMI (Chapter 9): Planar capacitance can be effective in controlling EMI radiations caused by both differential mode and common mode noise signals.

STRATEGIES FOR DESIGNING WITH PLANES

For all these reasons, the use of planes is very important and beneficial in PCB design. But then some relevant questions are these: How many planes should I use? What should be on them? Where should they be placed?

INTEL-1010
U.S. Patent 7,151,430

For example, every different power supply voltage is typically distributed on its own plane. It is logical that each different regulated supply of the same voltage also gets its own plane; otherwise some of the regulated supply sections would simply be shorted out. Very often, however, all the different regulated supplies are *referenced* to the same voltage potential—zero volts. Is it possible, then, to have only one ground plane (at zero volts) that can service all the individual regulated power supplies, or do we need a separate ground plane for each individual regulated supply?

In looking at the question of whether each power plane needs its own, separate individual zero-voltage reference plane, we have to look at why the separate power supply exists at all. For example, suppose the overriding issue is power dissipation. Then it may be perfectly fine for two supplies to reference the same ground plane. In some cases, a circuit simply involves two types of ICs with different supply requirements, and a single reference (zero-voltage) plane may work well here, too. And, if we have mixed signals (as in an A/D circuit) but all the circuits are known to be quiet, so noise is not a concern, a single reference plane might be adequate.

As noted earlier, however, one reason for using different regulated power supplies is for noise control—for example, keeping transmitter noise out of a receiver section or keeping digital noise out of an analog section. All the reasons for using planes have a single common denominator—noise control. So it is almost axiomatic in high-speed designs that noise control is the predominant design issue, and our design strategies are undertaken with noise control as an important focus.

Now, if circuit noise is our primary focus, then we almost always need to define separate reference (ground or zero-voltage) planes for each individual regulated supply section. Consider the implications for not doing so. Assume, in a transmitter section, we send a signal down the trace shown in Figure 16-2a. The return signal, although primarily underneath the trace, extends for some distance beyond the edges of the trace. If any receiver circuitry is anywhere near this return gradient, some of this transmitter noise may couple into the receiver section. In addition, there are current gradients that flow from the transmitter-regulated power supply across the ground plane while charging (and discharging) any bypass capacitors that may exist. These currents may also couple into receiver circuitry. The whole point of having a separate receiver power supply section is to try to isolate the two sections. A single ground (reference plane) tends to work against this objec-

INTEL-1010
U.S. Patent 7,151,430

tive. So, good design practice is to have separate ground planes for the receiver and transmitter sections (Figure 16-2b).

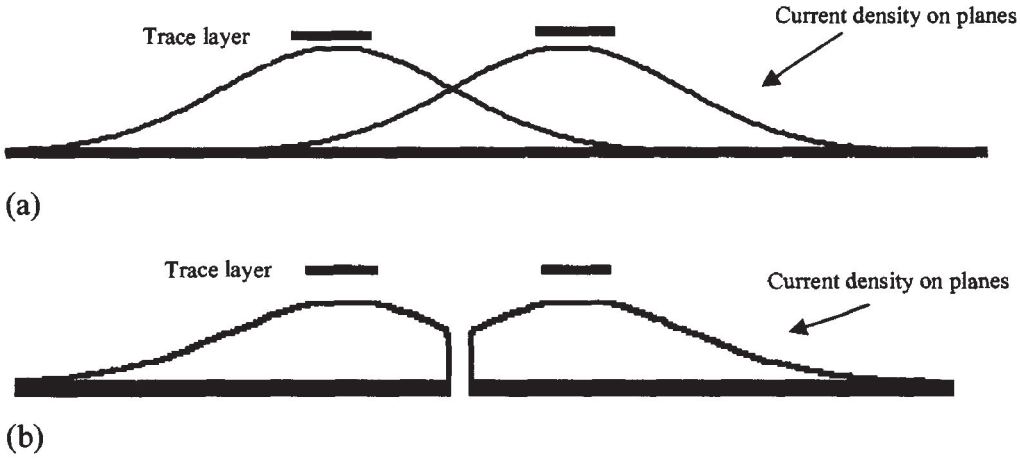


Figure 16-2 Current gradients from one trace can extend across the plane and underneath another unrelated trace (a). Separate ground planes (b) can constrain these current gradients.

Similarly, a digital signal flowing down a trace generates a return gradient on the reference plane under the trace. Such a signal may cause a noise signal to be injected into a nearby analog trace. Again, for this reason, good design practice dictates using separate analog and digital ground planes.

In fact, whenever we use separate regulated supplies for noise control purposes, it is good design strategy to use a separate reference plane for each one, even though all the reference planes are at the same nominal voltage potential—zero volts. Not doing so tends to defeat the very purpose of having separate regulated supplies in the first place.

In summary, separate regulated power supply stages can exist for a variety of reasons. Very often they exist for the purpose of controlling circuit noise, which, in turn, almost always requires the use of planes. If each different regulated power supply section requires its own separate distribution plane, then each should have its own individual (zero-voltage) reference plane, also.

INTEL-1010
U.S. Patent 7,151,430

SOME DESIGN RULES

Now, based on the preceding, let's assume we have several regulated power supplies (associated with separate circuits on the board), each distributed on its own plane, and each with its own, identifiable, separate reference (zero-voltage) plane. What PCB design strategies and rules are appropriate?

Connecting Reference Planes Together Ultimately, almost all regulated power supplies reference to the same thing, commonly a zero reference (zero-voltage) point in the system. Typically, if one puts an Ohmmeter between the various reference planes, we find that they are all connected together. It is usually important, however, that when reference planes are connected together, the connection is made at a *single point.*

Suppose, for example, this were not true. Suppose we had both a digital and an analog section on our board, and the digital and analog reference planes were connected at two (or more) points. There are conditions under which certain return signals could travel on both planes (not just on the plane under the trace), thereby negating the separation we were trying to accomplish in the first place. Even worse, it is possible (under certain conditions) for noise signals to circulate in a loop around through both reference planes, crossing between them at the two connection points. Such current loops are called *ground loops*. Their origin (cause) is often obscure, but their effects usually are not. The effects include noise problems, EMI radiation problems, and in extreme cases power dissipation and heat problems.

Control over ground loops is relatively simple. If there is only a single connection point for the reference planes, there is no loop over which a signal can travel.

But this poses a new interesting question: Where should that single point be? On some systems, particularly where there are mother and daughter boards, the planes are separately routed back in a "star" fashion to a single point, usually at the primary power supply for the system. If there are multiple regulated power "islands" on a board, these might be connected to each other at single points with zero-ohm resistors (really just a jumper) or with ferrite beads. The beads, of course, are used to block higher frequency components while still allowing a DC connection between the planes. Some argue, however, that if we have been really effective in isolating our power requirements, there are no higher frequency components to filter, and therefore a zero-ohm resistor is all that is necessary.

INTEL-1010
U.S. Patent 7,151,430

In A/D circuits, it is common to have digital power and ground planes on the digital side of the IC and analog power and ground planes on the analog side of the IC. We often connect the analog and digital grounds together at a single point directly under the IC (or at least very close to the IC) with zero-ohm resistors or with ferrite beads.

Overlapping Planes If we have separate regulated power supplies with their own reference planes, it is good design practice not to let unrelated portions of the planes overlap. For example, do not let a portion of an analog power plane overlap a portion of a digital ground plane (see Figure 16-3). Remember, a capacitor in its simplest form is simply two conducting surfaces separated by a dielectric. The area over which two planes overlap forms a small capacitor. It may be a very small capacitor, to be sure. Nevertheless, any capacitance provides a path over which noise may travel from one regulated supply to another, working to defeat the purpose for which the separation existed in the first place.

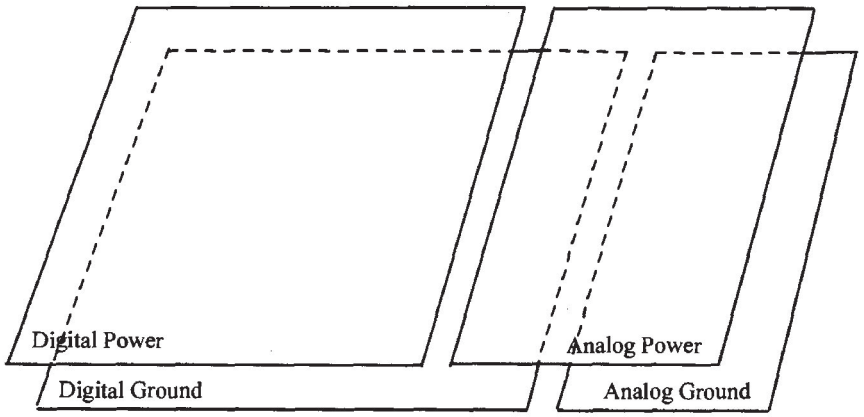


Figure 16-3 Overlapping planes can cause capacitive coupling between unrelated planes.

An important part of the PCB placement process is to place components in such a way that their common regulated supplies (and grounds) can be efficiently grouped together and not overlap other circuits.

INTEL-1010
U.S. Patent 7,151,430

Decoupling to Wrong Plane We use bypass capacitors to decouple our circuits—that is, to connect the power and ground planes together (from an AC standpoint) at specific points on the board. It is probably obvious that we don't want to drop a bypass capacitor from one power plane to an unrelated reference (ground) plane. Again, the reason is that we can (and almost certainly will, in this case) couple noise from one regulated supply section into the other. Unfortunately, this mistake can occur accidentally, sometimes fairly easily. As designers, we must be careful to check our engineers' net lists to ensure that this mistake hasn't happened. Even worse, we must guard against making this mistake ourselves. It is very embarrassing when this happens. There is simply no good answer to the question, "Why did you connect it that way?"

Signals Crossing Separations Remember that signals reference to their (usually) nearest plane, be it power or ground. If we have a separation between two reference planes, we never want to route a trace across that separation. Figure 16-4 illustrates an example of routing a digital trace across an analog plane. There are four primary problems that can result from doing this, all of them bad. They parallel three of the reasons we mentioned earlier for using planes in the first place.

1. Good impedance control requires continuous control over geometry, and a continuous return path underneath the trace. If the trace crosses the boundary between two planes, the return signal cannot "jump" the gap. (Don't assume you can solve this problem with a bypass cap. Doing so would couple two unrelated planes together.) This causes an impedance discontinuity and a reflection, and therefore a potential noise problem at that point.

2. If the return signal can't jump the gap, it must find some other path. This almost certainly increases the loop area for the signal and therefore the potential for EMI.

3. If the return signal does somehow find its way onto the analog plane, we now have digital noise on the analog plane, defeating the purpose of the separation in the first place.

4. Suppose two traces cross a separation between two planes. Since their return currents cannot jump the gap, these return currents must find another path. Even though the signal traces are separated from each other, their return signal paths might not be separated, and their *return signals* might crosstalk. Thus, when

INTEL-1010
U.S. Patent 7,151,430

signals cross plane splits, crosstalk may result even though there is no apparent cause or reason. This type of crosstalk can be very hard to diagnose.

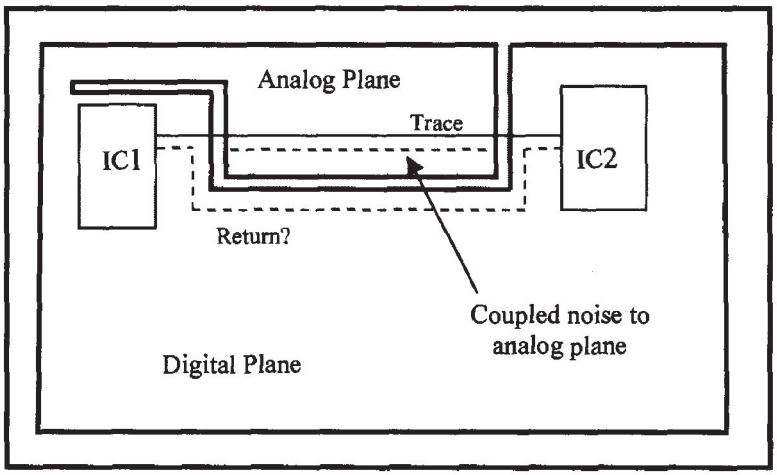


Figure 16-4 Never route a high-speed trace over an unrelated plane.

The primary way to avoid the problem of traces crossing splits in planes is to be careful in the layout and placement stages of the design. Group circuits by regulated supply voltage and then by signal flow. If a particular trace must extend into another regulated supply area, perhaps little "islands" or "peninsulas" (on both the power and ground planes) can be created for the traces involved.

If you cannot find a suitable way of routing traces without violating this design principle, then the circuit design engineer should be consulted. The engineer is the one responsible for deciding whether the design guideline can be relaxed in this particular case, whether the layout should be changed, or whether (in extreme cases) a different component selection is called for.

INTEL-1010
U.S. Patent 7,151,430

STACKUPS

We summarized earlier why planes and plane pairs are important in high-speed design. From those rules we can derive two basic rules that help us define what makes a good board stackup.

Rule 1: Every high-speed trace must be referenced to a continuous plane. The reasons include loop area control, impedance control, and crosstalk control.

Rule 2: Every power supply should have a parallel power–ground capacitive plane pair. The reasons include help in capacitive decoupling circuits and for EMI control.

A good board stackup meets these two requirements. Any stackup that does not meet *both* these requirements may suffer high-speed problems as a result.

Figure 16-5 was adapted from one created by Rick Hartley.[1] It shows illustrations of good and poor stackups from a high-speed design standpoint. Every "good" design is characterized by every trace being referenced to a nearby plane and every power supply having a power–ground planar pair. Every "poor" design is missing at least one of these two elements.

One conclusion that is drawn from Figure 16-5 is that it is hard to have an efficient six-layer design. That is true. Six-layer designs with four trace layers cannot meet both criteria for effective high-speed design layout at the same time. Therefore, if six-layer boards are going to be used with four trace layers in high-speed designs, we need to understand and accept that some compromises are being made.

[1]Rick Hartley, "Controlling Radiated EMI Through PCB Stackup," *Printed Circuit Design*, July, 2000, p. 16.

INTEL-1010
U.S. Patent 7,151,430

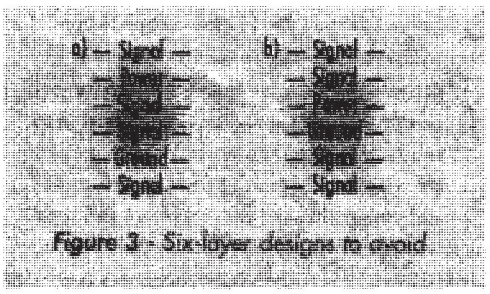


Figure 3 - Six-layer designs to avoid

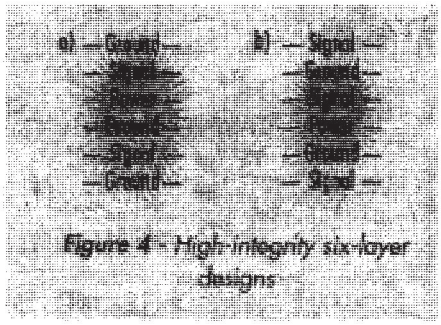


Figure 4 - High-integrity six-layer designs

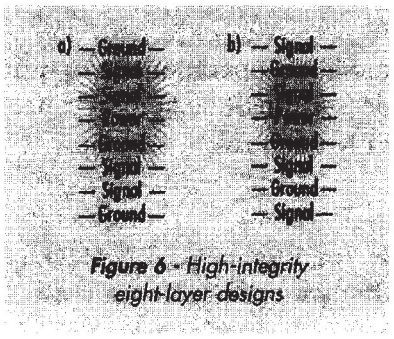


Figure 6 - High-integrity eight-layer designs

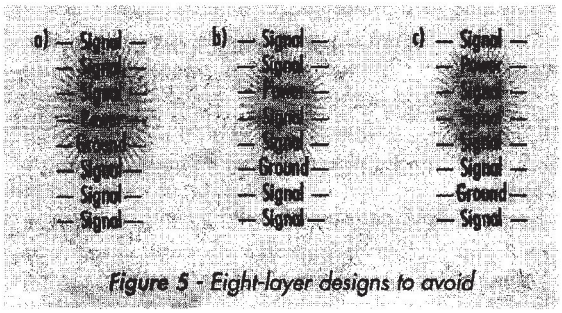


Figure 5 - Eight-layer designs to avoid

Figure 16-5 Examples of good and not-so-good board stackups (courtesy Rick Hartley and *Printed Circuit Design* magazine).

CONCLUSION

It is common to have different regulated supply voltages on a circuit board. If circuit noise control is not a major issue, then it is possible for all these supplies to reference a common (zero-volt) ground. If circuit noise—especially high-speed circuit noise—is an issue, however, then the regulated supply voltages are commonly distributed on their own individual planes. In this case it is normal for each regulated supply "plane" to have its own individual (zero-voltage) reference plane. These planes are usually constructed as plane pairs on the board. The boundaries defining individual planes are often called *splits*.

Certain, relatively simple design rules are followed when we have split planes:

INTEL-1010
U.S. Patent 7,151,430