UltraCAD Design, Inc.

DESIGN NOTE

MAINTAINING CLEAN POWER

EMI and noise problems can manifest themselves in subtle ways. One thing most people instinctively understand, though, is the need for a "clean" power and ground system. But problems can manifest themselves there in subtle ways, too.This Technical Note looks at the power and ground system from three standpoints: power and ground isolation, capacitive decoupling, and the potential for distributed capacitance planes.

Power Supply Isolation:

Let's start with the recognition that power and ground systems are usually inherently "clean." At least, that would seem to be a necessary condition for good system performance. So we start with a power supply system that has been designed to be adequate for the system in which it is going to be used.

If the power system is inherently "clean," how then does noise get into it? And that is the insight that is necessary for good circuit layout. ***The purpose is NOT to keep noise on the power planes from getting into your circuit, the purpose is to keep circuit noise from getting onto the power planes.*** Once this concept is recognized, PCB layout can be viewed differently and more effectively.

There have been a few good articles and presentations given on this topic. One of the better presentations was given by Reinhold Henke of VoicePro Systems at the PCB Design Conference in Burlingame, CA. in March, 1993 (Reprinted in their Proceedings.) Parts of this Section borrow heavily on some of the points he made.

In high speed, high performance circuits, the noise of interest comes from current transients caused by rapid (fast rise time) switching. These current transients flow on traces and planes. The traces and planes have some finite inductance. Thus, a current pulse with a fast rise time flowing through an inductance creates a voltage transient. The voltage transient affects noise margins for logic devices and radiates from antennas to create EMI problems.

So the design rules are fairly fundamental:

1. Minimize these voltage transients
2. Control where they flow
3. Keep them off the power and ground planes and from passing between power supply systems.

Minimize Transients With Bypass Caps:

At a 1993 Design Conference, one of the speakers pointed out that at very fast rise times (.5 to 1.0 nsec) the highest frequency components are so high (300 to 500 Mhz) that even the best bypass capacitors look like inductors. He then made an outrageous extension of that point to conclude that bypass capacitors are ineffective in today's logic systems and can be ignored by engineers and designers alike! This section explains why bypass capacitors are used and what their limitations are. Our hope is that this discussion will clarify, in non-technical terms, an area that has been confused by significant amounts of misinformation through the years.

Switching Transients: Consider what happens when the signal levels change at the inputs of a single IC logic chip. The change in state of the output levels can cause a corresponding change in the current requirements through the chip at the V^+ and ground pins. Remember that current is the flow of electrons. Therefore, the change in current requirement is really a change in the "flow rate" of electrons. The "speed" at which this flow rate changes is measured by the rise time of the chip.

Therefore, for example, for a chip with a 1.0 nsec rise time, the flow rate of electrons (the current) must change from the old level to the new level in about 1.0 nsec. That's pretty fast!

For cases where the current requirement increases, the additional current is supplied through the power and ground planes. In a normal, steady state analysis, this poses no problem whatsoever. But what happens in the first few nanoseconds? The power and ground planes have considerable

11502 NE 20th, Bellevue, WA. 98004 Phone: (206) 450-9708 FAX: (206) 450-9790

inductance associated with them, and they simply cannot respond with additional charge (electron flow) that fast. So we have two mutually exclusive requirements:

(a) a logic chip that requires additional electron flow (current) in about a nanosecond or so, and

(b) a source (power and ground plane) that simply cannot supply electron flow that fast.

What happens under these circumstances is not always predictable. Some of the possible results are:

Timing Errors: The chip simply may slow down and switch only as fast as charge can be supplied. This can result in logic timing problems and logic errors due to logic state uncertainty.

Noise Errors: Current supply may not be sufficient to support output levels. So, output levels may fall out of specification and noise margins significantly decrease, resulting in logic level uncertainty and logic errors.

EMI Radiation: There will be inductance on the plane and at the power and ground pins of the IC. One way to increase current flow is to increase voltage across the inductance. So the voltage at the input pin of the IC may drop below that of V^+, thus "forcing" more current flow. This voltage transient will be very brief (maybe only a very few nsecs), but it can be very destructive. It is one source of EMI noise radiation. This EMI radiation can cause noise spikes in adjacent traces and can result in subsequent FCC compliance testing problems.

Can we calculate these effects? Not with any formulas that normal people know how to use! Can we measure them? Maybe, with good test equipment. Can we deal with them and even control them? Yes.

Conceptual Solution: Since the problem is an inability of the power and ground system to supply enough charge (source of electrons) to meet the initial current requirement during the first few nsecs of switching, the natural answer is to provide a supply of charge close to where it is needed. That's what a bypass capacitor does. Conceptually, this is a straightforward answer.

Practical Solutions: The practical problem with this approach, at least with today's very fast circuits, is that:

(a) capacitors large enough to supply sufficient charge have too much lead inductance to be able to respond fast enough,

(b) capacitors with low enough lead inductance to respond quickly cannot supply enough total charge, and

(c) at today's very fast switching speeds (sub-nsec) NO bypass capacitor may be able to respond fast enough to completely supply the chip's initial switching require-ments.

Consider the traces illustrated in Figure 1. Trace (a) illustrates the transient current requirement for the chip as logic levels change state. Trace (b) illustrates conceptually the fastest response that the power and ground planes can provide. Trace (c) illustrates what a standard bypass capacitor can provide. It can only supply a "burst" of charge, not a long term supply, and it has its own response time. Trace (d) illustrates that a smaller, faster bypass capacitor can supply charge faster (because its lead inductance is typically smaller), but it stores much less charge and can't meet the short term requirements.

Therefore it seems clear that not only are bypass capacitors absolutely necessary in high speed switching circuits for supplying "local" charge, but more than one size and type of bypass capacitor may be necessary to meet all the switching requirements.

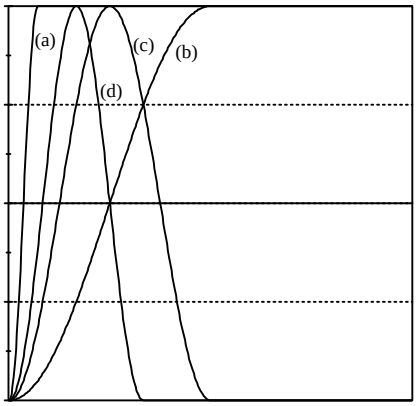

(a) Desired Response
If system could respond

(b) Response possible
from power planes

(c) Additional charge
supplied from larger bypass cap

(d) Charge available
from smaller, faster bypass cap

Figure 1

Furthermore, since the constraining factor in supplying charge to high speed switching circuits is lead inductance, it is important to address this directly. Lead inductance is primarily related to two factors: physical construction of the capacitor itself and its placement in the circuit. The physical construction of bypass caps and their selection is the province of the engineer. But there are some important considerations to think about regarding their placement.

Bypass Cap Placement: As stated above, one of the design objectives is to keep noise transients OFF the power and ground planes. Thus, normal trace routing, illustrated in Figure 2(a), causes the charge that flows between the bypass cap and the IC leads to flow across the planes. Reinhold Henke uses the analogy of a "voice" in describing this. If the "voice" gets on a power or ground plane, it can be "heard" by other circuits. Therefore, you want to keep transient noise signals ("voices") off the planes. Figure 2(b) is BETTER because none of this transient flow occurs on the planes. There *is*, however, a transient voltage fluctuation along the trace between the bypass cap and the IC pins.

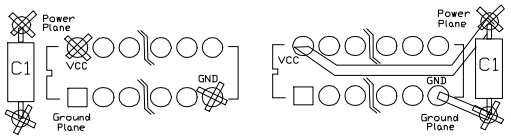

Figure 2

That is why the cap is placed near the ground pin of the IC. In that way, the transient noise voltage is minimized between the the ground plane and the IC ground. It is important to note that the IC signals are referenced to the IC ground --- which therefore should be kept as close as possible to the system ground at the ground plane. The relative noise on the signal leads, compared to other devices on the board, will be approximately equal to the noise signal between the IC ground pin the the plane. On the other hand, the signals are relatively less sensitive to the noise on the V_{CC} line of the IC.

Note that the transient noise that can occur between the bypass capacitor ground trace and the IC ground pin is analogous to the ground bounce that can occur inside the IC itself (caused by the inductance of the wire between the chip and the package ground pin. The consideration given to corner pin vs center pin grounding of devices underscores the importance of this type of trace routing.

Figure 2, therefore, illustrates the priorities associated with this type of routing:

First:	Keep noise off the planes
Second:	Minimize the noise at the ground pin
Third:	Minimize the noise at the V_{CC} pin

It should also be noted that Figure 2 applies to **voltage** controlled logic systems, where signals are typically referenced to ground. In **current** controlled logic systems (such as ECL) the signals are typically referenced to the highest supply. That is why current controlled logic systems normally use negative voltage power supplies and the highest voltage (ground) is used as the reference plane. In such systems, the shortest trace leads would be to the highest voltage (ground) pin (not to a negative supply pin).

Summary Design Rule Considerations:

What does all this mean for designs and design rules, then?

(a) Bypass capacitors serve a very necessary purpose and are required in today's fast logic circuits.

(b) It is possible that more than one source of transient charge (bypass capacitor) is required. Some designers place larger bypass capacitors near a group of ICs, and other, smaller, and faster ones close to each IC. It is worth noting that there is not necessarily a one-to-one association between capacitors and ICs. No IC can know that it "owns" one bypass cap and not another!

(c) The bypass capacitor closest to the IC should have the lowest possible lead inductance (so that it can supply charge as fast as possible). It should also be as close to the IC as is physically possible. That is, the printed circuit board trace lengths must be absolutely minimized. Therefore, surface mount ceramic capacitors placed on the bottom of the board directly across the supply pins usually represent the optimal placement.

(d) Since there must be at least some finite lead length, the trace lengths associated with the ground pins should be the shortest. This is because the internal IC signals are referenced to ground and are more sensitive to voltage transients at the ground pin than they are to transients at the V^+ pin. (*Note that ECL signal levels are referenced to the highest voltage. That is why ECL circuits have "negative" supply voltages. Therefore, the shortest bypass capacitor traces should be at the higher voltage --- ground --- reference plane for ECL circuits.*)

(e) Finally, since there must be some, finite lead lengths, and since these must have some, finite (though very small) inductance, some EMI radiation is inevitable at the power and ground pins. Good designers recognize this and approach their designs from the standpoints of (1) recognizing it, (2) minimizing it, and (3) protecting the rest of their circuit and the environment from it.

Distributed Capacitance Planes:

Some of our customers have asked about our opinion of the relatively new distributed capacitance planes that can be used in PC board fabrication. Some board manufacturers have licensed this technology from Zycon and offer it as an option to their customers.

Generally, *UltraCAD* does not promote or otherwise comment on other brands or companies. But we see this as a technological issue that may have relevance to our customers. Therefore, we have decided to offer these observations and opinions in hopes that they will help clarify some issues for our customers.

The distributed capacitance plane is basically a power/ground pair of planes separated by a very thin dielectric material. Typically, instead of designing a board with a single ground plane and a single power plane, one would design the board with two pairs of these planes. The plane-pairs provide a source of distributed capacitance that may provide advantages over a board without this structure.

Our board manufacturer friends have explained the primary tradeoffs regarding this structure to us in these terms:

The added cost of using distributed planes is approximately the same as adding two more layers to the circuit board. But the inherent capacitance of the planes, then, can be used to eliminate up to as much as 80% of the bypass capacitors on the board. This reduction in parts count can have two additional potential benefits:

1. It may mean that SMT parts (principally caps) that had to be mounted on the back side of a board can now be moved to the front or eliminated altogether, resulting in significant manufacturing assembly savings.
2. It may also mean that the reduced parts count allows for routing economies that, in turn, result in a reduced layer count.

Therefore, the cost of the extra layer pair can be (perhaps more than) gained back through reduced parts cost, manufacturing savings, and potential layout savings.

UltraCAD takes no position on these claims. They are legitimate issues that an engineer should consider on a case by case basis in evaluating the benefits of the distributed plane technology for his or her particular design.

What we **do** see as a significant potential benefit is the possibility for improved EMI performance with distributed planes. In the Section above, we pointed out that a major problem in quieting boards where very fast logic circuits were used was the inability of even small bypass caps to provide charge to a switching device fast enough.This is because even the smallest bypass caps still have some finite lead inductance. Even very small inductances become important at the high frequency components of very fast rise times.

It would seem that it would be hard to find a source of charge for a switching device with lower inductance than a distributed capacitance plane directly under the device and connected directly to its power and ground pins. The amount of available charge may be very small, and the plane may not be a substitute for the use of bypass capacitors for the larger amounts of stored charge the circuit requires. But the plane may be just what is needed to provide the initial charge required at the very point of switching that lead inductance prevents ordinary bypass caps from providing.

For this reason, we believe a very significant potential benefit of the distributed plane technology is its promise for very quiet circuit switching with respect to EMI emissions. In systems where FCC compliance issues are important, this potential benefit, alone, may offset any additional costs associated with the technology.

(In all fairness we should add that some of the board fabricators we respect question if there are any benefits to distributed capacitance planes, and they worry about dielectric breakdown.)

Power Supply Boundaries:

Usually, the basic power supply either is placed near a corner of the board or is off the board altogether with power supplied through a connector. In the latter case, there may be, and should be, some sort of filtering at the point of entry. In all cases, there is a defined point where the power supply ends and where the power distribution system for the circuit begins. THAT is the singular point where the power supply should tie into the planes.

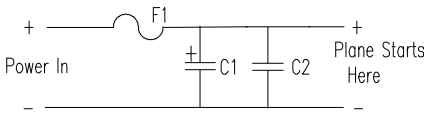

Figure 3

For example, in Figure 3 the typical power supply filtering ends at C2 and the power and ground planes start at that point. C1 has a large capacitance for bulk filtering and C2 is a small capacitor for high frequency filtering. (Some people may add a third, very small value capacitor for very high frequency filtering.) Figure 4 shows alternative ways traces might be routed on the board. Figure 4(a) illustrates normal routing practices,

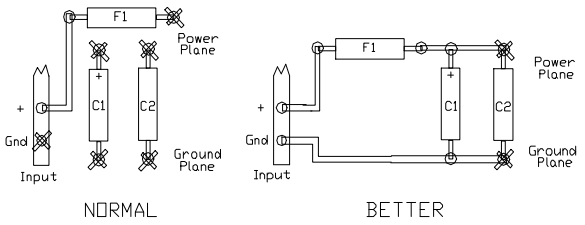

(a) (b)

Figure 4

where power ties to the planes at several places instead of a single point, as illustrated in Figure 4(b). Thus, in Figure 4(a), any current flowing between the basic power supply and C1 and C2, or between C1 and C2 themselves, flows on the plane. This is not the case in Figure 4(b). Again using Henke's voice analogy: If there are "voices" on the plane, they can be "heard" by the other circuits. So you want to keep those "voices" off the plane; and in doing so, be sure to use wide, low impedance traces between the input connector and the components before they tie to the plane.

It is important to remember once again that the filtering here works BOTH WAYS. It not only keeps power supply noise off the board, but it keeps board noise off the main power supply and keeps it from appearing on other boards in the system.