

Bypass Capacitors

An Interview With Todd Hubing

Douglas Brooks, President

UltraCAD Design, Inc

www.ultracad.com

Positions I have taken in the past regarding bypass caps have generated some considerable controversy. In addition, Todd Hubing, et al, wrote an article (Footnote 1) on plane capacitance that has, unfortunately, been widely misinterpreted.

Todd is an Assistant Professor in Electrical Engineering at the University of Missouri-Rolla. He was the keynote speaker at PCB Design Conference East last fall. After the conference, he and I entered into an e-mail exchange regarding his research and opinions regarding bypass caps. My e-mails tended to be in the form of statements about which I asked Todd to agree or disagree, and then to comment on. The exchange is reprinted below in an "interview" format. It has been edited for length, clarity and continuity, but not for content. Todd's responses, in particular, have not been edited.

DB: It was a pleasure meeting you in Boston. I have several questions related to bypass caps and would like to make several 'statements' and ask whether you agree or disagree with them.

Your paper gives evidence that at very high frequencies, the planes provide enough capacitance to support device switching. But in digital circuits, a switched signal contains numerous harmonics that can be, and will be, well below the highest frequency component related to the rise time of the signal. While the planes may be able to support the highest frequency requirements, can they can support ALL the harmonics present in a digital signal?. Would you support the following statements? (a) The planes are not a complete substitute for bypass caps in a practical digital circuit. (b) Bypass caps play an important role in providing charge to switching devices at all but the highest frequency harmonics.

TH: Yes, certainly I would agree that bypass caps provide most of the current at the frequencies where they are effective. This is evident from the data which shows that at low frequencies the power bus impedance is highly dependent on (in fact nearly equal to) the impedance of the added decoupling capacitors. Removing these caps can severely degrade the signal waveform.

DB: Your paper also provides evidence that the placement of caps does not have much, if any, impact on the impedance of the board. But that may not be the same issue as

providing charge to an IC at the time of switching. Would you agree with this statement? A bypass cap that is used for providing charge to support the lower frequency harmonic switching requirements of a digital signal should be placed as close as possible to the device being supported.

TH: I would not agree with that statement if the boards in question have a power/ground plane spacing of 10 mils or less. At the frequencies where added decoupling capacitors are effective, the time delay associated with moving charge across the board is inconsequential. We have demonstrated this with time and frequency domain measurements on a number of boards.

Boards with a power/ground spacing greater than 20 mils or so behave quite differently. The inductance of the planes cannot be neglected. We described this in a paper presented at the 1995 IEEE EMC Symposium (Footnote 2). It is important to place decoupling capacitors near the chip they are decoupling in a board like this. Also, the decoupling capacitors are typically effective at frequencies up to 1GHz or higher.

DB: So, if we are bypassing frequencies that are too low for the planes to be effective AND if the inductance of the path along the plane between the cap and the device begins to limit the cap's effectiveness, then "closer is better".

TH: I agree with the statement. Although, I might add that the inductance of the planes between the device and the cap is never an issue (practically speaking) for boards with a 10 mil or less spacing between planes.

DB: Is it possible to avoid the, 'less than 10 mil spacing -- more than 10 mil spacing' qualifier that sometimes comes up in these discussions?.

TH: I don't think it is possible to avoid the 'less than 10 mil --- more than 10 mil' qualifier. Boards with a greater than ~30 mil spacing between planes behave quite differently than boards with ~10-mil or less spacing. The critical spacing depends on the size of the board, but the ≤ 10 mil or ≥ 30 mil guideline appears to be appropriate for common board sizes. Components on a boards with 30 mil or greater spacing may never see the board capacitance (at any frequency) due to the mutual inductance of the vias between the planes.

DB: OK, but what about the mutual inductance of the via between the device and the plane-pair? Is its effect different, or is it part of the consideration?

TH: The magnetic flux coupling between the planes is much stronger than the magnetic flux coupling above the planes. This is hard to explain without drawing a picture, but essentially the time varying magnetic flux between the planes causes charge to be pulled out of the nearest decoupling cap BEFORE it can be drawn from the planes if the spacing between the planes is greater than ~30 mils. Our paper in the 1995 EMC Symposium Proceedings has an illustration that shows this more clearly. There is also measured data from test boards and real products in this paper that illustrate this effect.

DB: Equation 9 in your paper shows that the bandwidth of a board increases by the square root of the number of bypass caps. Your data supports this and figures 6, 8, and 9 also show this. Would you agree that adding more bypass caps increases that bandwidth of a board up until the point where the bandwidth is so broad that no further gains are relevant?

TH: At frequencies where a board with closely spaced planes becomes resonant, the added decoupling capacitors do not contribute significantly to reducing the board's impedance. This could be considered an upper limit on the bandwidth of the decoupling. The main point of the paper was that even at frequencies well below board resonance, the added decoupling capacitors tend to be ineffective. It is not that these capacitors are poor sources of charge, but that the internal planes become a much better source of charge at high frequencies. For a given set of planes, larger numbers of decoupling capacitors have a lower effective series inductance and therefore a higher effective bandwidth. Every time we double the number of caps, we halve the inductance and get 40% more bandwidth. However, as a practical matter we quickly reach a point of diminishing returns and added decoupling on boards with 10 mil or less power/ground plane spacing rarely has an effect at frequencies above 100 MHz.

DB: Let me come at this another way. Since the power distribution system of a PCB can look like a capacitor, there is a frequency at which the available capacitance from the planes is more efficient than that which can be obtained by adding a capacitor to the board.

TH: True for boards with 10 mil or less spacing.

DB: Increasing the number of capacitors on a board increases the bandwidth of the power distribution system up until the bandwidth reaches this fre-

quency. At that point, adding more capacitors does not make any additional incremental improvement. BUT, the capacitors added up to that point have improved the performance of the PCB.

TH: True for boards with 10 mil or less spacing (at least in theory). Adding more caps continues to increase the bandwidth of the decoupling until the board itself becomes resonant. In practice however, the number of decoupling caps required is an impractical number. Doubling the number of caps provides a 40% increase in bandwidth. Once you have 40 or 50 caps on the board, another dozen doesn't have much effect. Yet for many designs 40 or 50 is not nearly enough to reach that theoretical upper limit. On boards with a 30 mil or greater spacing, the caps located nearest the noise source or the point of the noise measurement are the most critical no matter how many other caps are present on the board.

DB: You comment that "Once you have 40 or 50 caps on the board, another dozen doesn't have much effect." I will simply observe that on one board we designed for a customer, the engineer had (count 'em) 2,100 bypass caps on the board!! That is extreme, but 1,000 is not unusual on a large board. So, the square root of 2100 begins to get you up there.

TH: Wow! I didn't realize that anybody was putting this many decoupling caps on one board. I don't believe that this is a very good design strategy. Large numbers of decoupling caps increase the cost and decrease the MBTF of a board. It's more effective to utilize the interplane capacitance of the board for high-frequency decoupling when you need it.

DB: Whether IC power and ground pins should be connected to the planes or to the bypass caps by traces has become a hotly debated topic. The argument for connecting IC's to the plane is lower inductance. The argument for connecting the IC to the cap is to keep noise off the plane. Would you support this argument? The preferred wiring of bypass caps is from the IC to the cap, and then to the plane, until the fastest frequency requirement (presumably associated with rise time) becomes high enough that it cannot be supplied through the inductance of the trace, at which point it is better to attach the IC directly to the plane.

TH: You have summarized the situation quite nicely. A very good theoretical argument can be made in favor of connecting power and ground pins to the cap first and then the planes. We have seen some products decoupled this way which were operating at fairly high speeds and they did not appear to have any problem supplying adequate charge to the devices. On the other hand, this approach can be difficult to implement on devices that have many power and ground

pins. It is not yet clear that the benefits of this approach justify the design problems it creates. We're still looking at this issue though (on a product that is in our lab right now).

DB: Then, regarding the placement of bypass caps in this case, would you agree that "closer is better".

TH: When the planes are 30 mils or more apart, closer is better. There are no situations that I can think of where "closer is worse".

DB: Some authors, Howard Johnson (Footnote 3) for example, advocate placing bypass caps at strategic places on a board to provide a return path for a return signal and thereby reducing loop area and therefore EMI emissions. One place he has specifically advocated this is near connectors. Would you agree?

TH: I agree with Howard Johnson's contention that the decoupling capacitors play an important role in returning some of the signal current. What was not known at the time the book was written however, is that on a board with closely spaced power/ground planes most of the current in the 100 KHz - 100 MHz frequency range returns through the decoupling caps with the lowest inductance connection even if these caps are not physically located near the circuit.

DB: Expanding on this point; I don't want to put words in Howard Johnson's mouth here, but I think he might make the following argument (Footnote 4): Suppose the signal loop is out the signal pin, down a trace to another device, to ground, and back to the ground pin of the first device. Now the return path will want to be directly under the signal trace (path of lowest loop area.) Suppose the POWER plane is directly under the signal trace, so the return path is on the power plane. The question, then, is how does it get from there to the two ground pins? The answer (I believe Howie would say) is through nearby bypass caps. If the caps are close, the loop area will be smaller than if they are further away. Therefore, would you agree that if we want to limit loop area, "closer is better?"

TH: We have recently built and measured a number of test boards to investigate the issue of "where do return currents flow when you put them on one plane, but another plane is closer to the signal path?" For closely spaced planes, the return currents do not seek the nearest decoupling cap to get from one plane to the other. They utilize the interplane capacitance in the vicinity of the signal vias.

This appears to be true even if the planes are not so closely spaced. Our test boards had an interplane spacing of 43 mils and yet above ~30 MHz virtually all of the current returned on the plane closest to the trace even though this plane was not connected to anything.

We're still looking at this. We'd like to come up with a simple model that could be used to predict where the currents will flow for a given trace/plane geometry and a given frequency.

Footnotes:

1. Hubing, Drewniak, Van Doren, and Hochanson, "Power Bus Decoupling on Multilayer Printed Circuit Boards," IEEE Transactions on Electromagnetic Compatibility, Vol 37, No 2, May, 1995, pp. 155-166. All my references are to this paper.
2. Hubing, Van Doren, Sha, Drewniak, and Wilhelm, "An Experimental Investigation of 4-Layer Printed Circuit Board Decoupling, Proceedings of the 1995 IEEE International Symposium on Electromagnetic Compatibility, August, 1995, pp. 308-312.
3. Dr. Howard Johnson, President, Signal Consulting, Inc. and author of "High-Speed Digital Design: A Handbook of Black Magic," Prentice-hall, 1993.
4. E-mail dated 3/4/97, Subject: "Re: decoupling/bypass capacitors at connectors" sent over the Signal Integrity E-Mail Forum. To subscribe to this forum, send an e-mail to 'si-admin@silab.eng.sun.com' with the word 'subscribe' as the subject.

If you found this article interesting, you might also like the following references:

"Maintaining Clean Power, Part One," Brookspeak, Printed Circuit Design Magazine, April, 1997, also available as one of UltraCAD's Technical Notes (t007.pdf) (See www.ultracadm.com and follow the links to "Technical Notes")

"Ground Bounce" Parts I and II, Brookspeak, Printed Circuit Design Magazine, August and September, 1997, also available from our web site. (See www.ultracadm.com and follow the links to "Brookspeak columns and articles.")