

Effective Strategies for Choosing and Locating Printed Circuit Board Decoupling Capacitors

Todd Hubing

Electromagnetic Compatibility Laboratory
University of Missouri-Rolla
Rolla, Missouri USA

Abstract—Power bus decoupling is an important part of digital printed circuit board design. Effective strategies for implementing power bus decoupling depend on the board construction as well as the demands that the circuits place on the power distribution network. This paper discusses strategies for locating and mounting decoupling capacitors in various situations as well as methods for estimating the total amount of decoupling capacitance required.

Keywords - simultaneous switching noise, power bus noise; decoupling capacitors; connection inductance; design guidelines

I. INTRODUCTION

Decoupling capacitors are an important part of any digital printed circuit board design. These capacitors are connected between power and power return conductors to help stabilize the voltage delivered to active digital devices. When there is a sudden change in the current drawn by a particular device, decoupling capacitors provide a local source of charge so that current can be supplied quickly without causing the voltage on the power pins to dip suddenly. Boards with inadequate decoupling suffer from excessive power bus noise, which may cause signal integrity or EMC problems affecting the reliability of the product.

Figure 1 provides a simple schematic to illustrate the concept of printed circuit board decoupling. When an active device on a printed circuit board exhibits a sudden change in the amount of current it draws, this change in current causes some of the supply voltage to be dropped across the connection inductance. This drop in voltage is observed by all the devices on the board, not just the switching device. If the drop in voltage is too great, the circuits on the board may fail to operate properly. It is also possible for voltage fluctuations on the power bus to cause conducted or radiated electromagnetic interference problems.

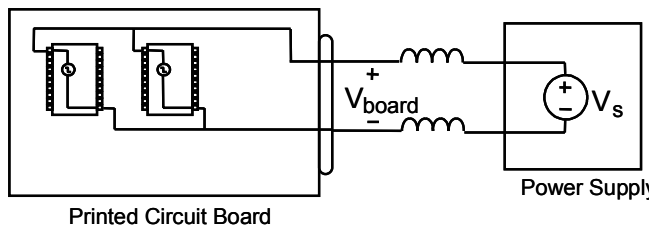


Figure 1. Voltage drop caused by current drawn through power supply connection inductance.

One method of countering this voltage fluctuation is to provide capacitance between the power and ground conductors near the active devices. This capacitance provides a local source of charge that can respond to sudden changes in the current demand while stabilizing the voltage on the power bus.

II. BACKGROUND

Printed circuit board decoupling is as old as printed circuit board technology itself. Early boards with one metallization layer generally contained a large decoupling capacitor near the point where power came on to the board. Larger devices that were likely to exhibit a sudden change in the amount of current they required often had their own decoupling capacitors across their power pins.

The earliest printed circuit boards were built for analog circuits. Analog circuit designers were very familiar with noise sources and filtering, so power bus decoupling was nothing new to them. Later, with the introduction of digital circuits, power distribution was no longer an integral part of the circuit design. Digital circuits required a supply voltage that was sufficiently stable, but the overall quality of the product was no longer tied to specific levels of power bus noise. As the number of digital printed circuit boards increased, printed circuit board decoupling became an afterthought. It was generally possible to build boards with a nominal amount of decoupling that operated flawlessly.

As the speed of digital circuits increased, more of the energy in power bus transients was at high frequencies. This forced designers to pay more attention to the decoupling capacitors they used and specifically to the equivalent series inductance of these capacitors and their connections to the board. A common practice in the 1970s and 80s was to put two decoupling capacitors on high-speed digital devices. One was a large-valued capacitor to filter the lower frequencies and the other was a smaller-valued capacitor (with a lower equivalent series inductance) to filter the higher frequencies.

At that time, decoupling capacitors were widely regarded to be ineffective at frequencies above their *self-resonant frequency* (the frequency at which their self-capacitance resonated with their connection inductance). Clayton Paul was one of the first to show that capacitors were generally effective well beyond their self resonant frequency and that their effectiveness at high frequencies depended entirely on their connection inductance [1]. As the speed of digital circuits has continued to rise, the importance of understanding and

controlling decoupling capacitor inductance has become even more important. Today, effective decoupling strategies have as much to do with achieving appropriate inductances as they do with providing appropriate capacitances.

III. DECOUPLING GUIDELINES

Anyone who has ever searched the literature looking for information on printed circuit board decoupling is sure to have come across conflicting advice. Papers have been published suggesting that capacitors should always be located near the active devices they are meant to decouple. Other papers indicate that capacitor location is not critical. Some advise that decoupling capacitors should be connected through short traces directly to the active device. Others advise that traces should never be used with decoupling capacitors. With little effort, it is possible to find conflicting advice related to capacitor values, capacitor loss, capacitor mounting and many other aspects of good printed circuit board decoupling strategies.

A primary reason for the discrepancies in the literature is that printed circuit boards are used in a wide variety of applications that have different design goals. The best power distribution strategy for a 50-layer all digital board with a thousand interconnections will not be the best strategy for a single-layer, high-volume, mixed-signal board. Both boards may be state-of-the-art, high-frequency designs; but the constraints on the power distribution will be very different.

Generally speaking, it is convenient to divide printed circuit boards into three categories for the purposes of defining power bus decoupling strategies: *boards without power planes*, *boards with closely spaced power planes*, and *boards with power planes that are not closely spaced*. In this section we'll briefly describe decoupling design rules and strategies for each type of board.

A. Boards without Power Planes

In many ways, printed circuit board decoupling is more straight-forward and easier to accomplish effectively on boards that do not have a power plane. This is especially true if the board in question has one or more solid ground planes. When power is distributed on traces, the designer has a great deal of control over the power bus noise seen by each device. The relative isolation between the active devices on the board allows the designer to decouple each device individually ensuring that the device's high-frequency current demands are met without permitting unacceptable voltage swings at the power input pins.

In general, the nominal value of a decoupling capacitor determines the total amount of charge available to the active device while the connection inductance between the active device and the capacitor is what limits the ability of the capacitor to respond to sudden demands for current. It is necessary to have enough capacitance available and also to ensure that the connection inductance will allow the capacitor to respond quickly enough to meet the needs of the active device.

We will discuss methods of quantifying the amount of capacitance necessary in the next section. Generally accepted

guidelines for locating and connecting decoupling capacitors on boards without power planes are provided below:

- Provide at least one "local" decoupling capacitor for each active device and at least one larger "bulk" decoupling capacitor for each voltage distributed on the board.
- Local decoupling capacitors should be connected between the voltage and ground pins of the active device. The area of the loop formed by the capacitor/device connection should be minimized.
- Bulk decoupling capacitors should be located near the point where a voltage comes on to the board. If the voltage is generated on the board, the bulk decoupling should be near the location where it is generated.
- Bulk decoupling capacitors should be sized to meet the transient current needs of the entire board (see following section).
- As a general rule, two local decoupling capacitors with the same nominal value are better than one capacitor with twice the nominal value. Two capacitors have a lower overall connection inductance [1] and provide better high-frequency filtering to the rest of the power bus [2].

B. Boards with Closely Spaced Power Planes

Decoupling a printed circuit board with closely spaced power planes is much different than decoupling a board without power planes. The planes themselves contribute a decoupling capacitance (by nature of their close proximity) that becomes very important at high frequencies. Also, the small loop area and wide conductors associated with the planes causes the contribution of the planes to the overall connection inductance to become negligible relative to the inductance of the connections between the capacitors and the planes. This means that the location of a decoupling capacitor on a board has little to do with the inductance of its connection to an active device. In the first few nanoseconds of a transition, typically sized boards with a plane spacing of 10 mils or less, draw virtually all of their current from the inter-plane capacitance regardless of where (or how well) surface mount decoupling capacitors are connected [3].

The following guidelines apply to boards with closely spaced power planes:

- Choose the largest nominal capacitance available in a given package size. However, do not use capacitors that have a nominal capacitance less than the parallel plate capacitance that naturally occurs between the power and power-return planes. A board made with FR-4 material containing one pair of power distribution planes spaced 0.25 mm (10 mils) apart has an inter-plane capacitance of approximately 16 pF/cm².
- The location of the local decoupling capacitors is not critical because their performance is dominated by the inductance of their connection to the planes. At the frequencies where they are effective they can be

located anywhere within the general vicinity of the active devices [3].

- The effective connection inductance is approximately inversely proportional to the number of decoupling capacitors. Therefore, high-speed circuit boards often have many local decoupling capacitors for every active device on the board.
- Connection inductance is determined by the loop area formed by the capacitor body, mounting pads, traces and vias [4].
- Never use traces attached to a decoupling capacitor pad. Locate vias within or adjacent to the pad, preferably as close together as possible.
- If there is no room for the via adjacent to the pad, then move the whole capacitor. Capacitor location doesn't matter, but connection inductance is critical.
- Mount all of the local decoupling capacitors on the face of the board nearest to the planes. Connection inductance is nearly proportional to the distance from the planes.

C. Boards with Power Planes Spaced > 0.5 mm

If the power and ground planes on a printed circuit board are greater than about 0.5 mm, then the inductance between the planes is no longer negligible. This applies to most 4-layer boards built with a 1-mm core and multi-layer boards that have signal layers in between the power and ground layers.

It is possible to use this inductance to our advantage if we locate decoupling capacitors very near the active device they are meant to decouple. As illustrated in Figure 2, if the vias of an active device and a decoupling capacitor are sufficiently close, the mutual inductance between the two vias will cause current drawn into the active device to be pulled from the

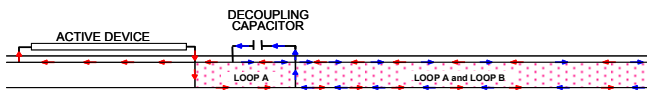


Figure 2. Mutual inductance due to magnetic flux coupling between the planes.

decoupling capacitor before it can be drawn from the planes [5, 6, 7]. This means that the voltage between the planes will not fluctuate so much provided there is a decoupling capacitor located very close to the device.

In order to take advantage of this phenomenon, it is important for the via of the active device that carries current between the planes to be close to the via from the capacitor connecting to the furthest plane. It is also important that the mutual inductance between the planes be greater than the connection inductance above the planes. In order to meet this condition, it is essential to minimize the decoupling capacitor connection inductance to the planes.

The following decoupling guidelines apply to boards with power planes spaced greater than about 0.5 mm apart:

- The total value of the bulk decoupling is determined by the transient power requirements of the active devices on the board (see following section).
- Local decoupling capacitors are intended to be effective at higher frequencies. The inductance of their connection to the power distribution planes is far more critical than their nominal capacitance.
- Choose the largest nominal capacitance available in a given package size. Nominal capacitance values are not nearly as critical as connection inductance.
- The location of the local decoupling capacitors is critical. Local decoupling capacitors should be located as close as possible to the power or ground pins of the active device they are decoupling. To determine which pin the decoupling capacitor should be located next to, first determine which power distribution plane is farthest from the active device. Local decoupling should be provided near the pins that connect to the most distant plane. For example, if the components are above Layer 1 on a 4-layer board and Layers 2 and 3 are Vcc and GND respectively, then the decoupling capacitors should be located near the GND pins of the active device. If there are any active devices below Layer 4 on this board, then decoupling for these devices should be located next to the Vcc pins.
- Orient the local decoupling capacitor so that the pin connected to the most distant plane is nearest the active device's pin connecting to the most distant plane. For example, in the board above if the active device is above Layer 1, and the decoupling capacitor is below Layer 4, then the Vcc pin of the decoupling capacitor should be located near the GND pin of the active device.
- If the decoupling capacitor can be located near enough to the active device to share the same via, this is optimal. However under no circumstances should a trace be used between the decoupling capacitor mounting pads and the vias. Decoupling capacitors should have vias located in or adjacent to the mounting pads to minimize their connection inductance.
- To minimize connection inductance: Never use traces on decoupling capacitors, locate the via adjacent to the mounting pad, locate the two capacitor vias as close together as possible, and mount all of the local decoupling capacitors on the face of the board nearest to the planes.

IV. HOW MUCH DECOUPLING IS REQUIRED?

There are essentially three basic approaches that are commonly used to determine how much decoupling capacitance is necessary for a given board design. These approaches are described below.

A. Simple guidelines

Despite the importance of providing adequate printed circuit board decoupling, most board designers use very simple design guidelines to determine how much bulk and local decoupling to use. Generally these design guidelines are based on past experience with other board designs or component application notes that may or may not be relevant to a particular board design. These guidelines rarely take into account parameters such as current drawn by the components, allowable noise voltages or the board stack-up. Instead, a guideline is likely to read something like, "... include one 0.01 μF local decoupling capacitor for each V_{CC} pin of every active component on the board plus 1 bulk decoupling capacitor with a value equal to 5 times the sum of the local decoupling capacitance." Simple guidelines such as this may work well for one particular application, while being totally inappropriate for others. Without accounting for the power requirements of the circuit design in question or recognizing that the location and mounting inductance will affect the amount of decoupling required, it is unlikely that optimum decoupling will be achieved based on this type of design guideline.

B. Maximum Impedance Approach

A more systematic approach to printed circuit board decoupling is to determine the maximum allowable impedance that the power distribution network is allowed to have and then choose the decoupling capacitance to meet this goal [8, 9]. The maximum allowable impedance is generally determined by dividing the maximum allowable power bus noise voltage by the maximum current drawn by the active device(s) being decoupled. For example, if a particular component can draw as much as 2 amps of current and the power bus noise voltage must be kept below 300 mV, then the maximum power bus impedance would be,

$$Z_{max} = \frac{300 \text{ mV}}{2 \text{ A}} = 150 \text{ m}\Omega. \quad (1)$$

A given decoupling capacitor will have an impedance that decreases with frequency at lower frequencies and then increases at higher frequencies where its connection inductance dominates. Decoupling capacitors are chosen to ensure that there is enough capacitance to provide an adequately low impedance at the low frequencies and that it is connected with a low enough inductance to meet the high-frequency impedance requirements. Figure 3 illustrates this concept for a multi-layer board. In this case, because the power planes are closely spaced, all of the capacitors connected to the board appear to be in parallel with the inter-plane capacitance of the board. At the lowest frequencies, the power supply has a source impedance that is below the maximum impedance. Above ω_1 , the inductance of the connection to the power supply causes the power bus impedance to be too high, so a bulk decoupling capacitor is chosen with enough capacitance to be effective at ω_1 . The impedance of the bulk decoupling capacitor is low enough to meet the power bus impedance requirements between ω_1 and ω_2 . A local decoupling capacitor with a smaller connection inductance is chosen to provide adequate

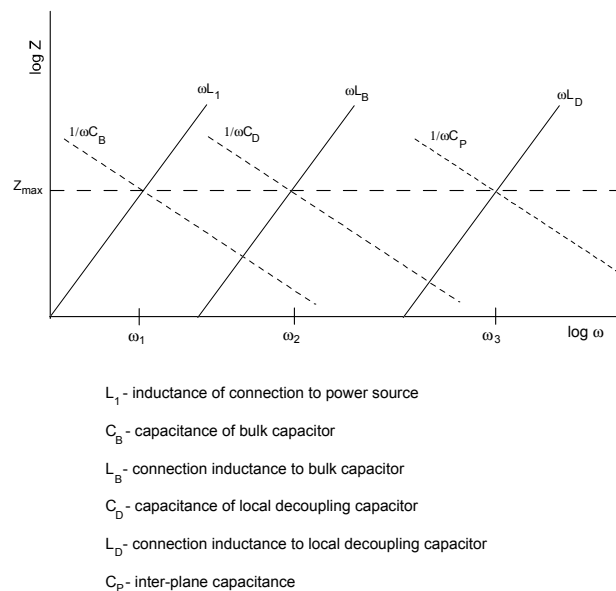


Figure 3. Power bus impedance of a multi-layer board with 2 decoupling capacitors.

decoupling between ω_2 and ω_3 . At frequencies higher than ω_3 , the inter-plane capacitance is adequate to ensure that the impedance target is met.

The plot in Figure 3 does not indicate the series LC resonances that occur at each "V" in the curve or the parallel LC resonances that occur at each peak. For a low-loss board, the parallel LC resonances cause each peak to extend well above Z_{max} resulting in ringing on the power bus that significantly exceeds the maximum allowable voltage. However, in most practical, densely populated board designs, these peaks are well damped and need not be a factor when calculating the amount of decoupling necessary.

Some board designers prefer to choose decoupling capacitors with different values in order to spread the nulls and peaks in the board response over a wide range of frequencies [9]. Other designers prefer to use just one value of decoupling capacitor in order to get the maximum possible amount of low-frequency decoupling capacitance for a given number of low-inductance capacitor connections. (This is sometimes referred to as the "deep-V" approach.) Either approach can be used to determine how much decoupling is required for a particular board design.

C. Capacitance Ratio Approach

A significant drawback of the maximum impedance approach is that the current requirements of the board as a function of frequency are not usually provided or easy to determine. Typically, the peak currents used to determine Z_{max} are calculated in the time domain. This can result in over-designed boards however, since the current requirements in one frequency band may be quite different than the current requirements in another.

The capacitance ratio approach avoids a direct calculation of the board currents, by recognizing that CMOS circuits draw

currents proportional to device capacitances on the board.

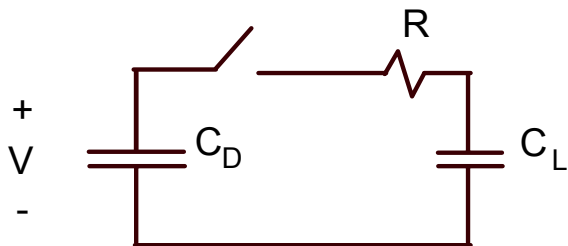


Figure 4. Once capacitor charging another.

Consider the circuit in Figure 4, which consists of two capacitors and a switch connected in series. The capacitor on the right, C_L , represents a CMOS load capacitance that will draw current from the power bus when it is charged. The capacitor on the left, C_D , represents a decoupling capacitor attempting to maintain a constant voltage on the power bus. At time $t=0$, the switch closes charging C_L and causing the power bus voltage to decrease slightly. A higher ratio of C_D to C_L , results in a smaller change in the power bus voltage.

To quantify this, we note that the energy stored in the system is initially equal to,

$$E_{initial} = \frac{1}{2} C_D V_i^2 \quad (2)$$

where V_i is the initial voltage across C_D . The energy dissipated in R in order to fully charge C_L is,

$$E_{dissipated} = \frac{1}{2} C_L V_f^2 \quad (3)$$

where V_f is the final voltage across C_L and C_D . Subtracting (3) from (2) and setting the result equal to the total stored energy in the final state yields,

$$\frac{1}{2} C_D V_i^2 - \frac{1}{2} C_L V_f^2 = \frac{1}{2} (C_L + C_D) V_f^2 \quad (4)$$

Solving (4) for C_D we see that,

$$C_D = \frac{2C_L}{\left(\frac{V_i}{V_f}\right)^2 - 1} \approx C_L \frac{V_i}{\Delta V} \quad \text{for } \Delta V \ll V_i \quad (5)$$

where ΔV is the change in voltage $V_i - V_f$. This equation suggests that total decoupling capacitance should be set to a value that is equal to the total device capacitance times the power bus voltage divided by the maximum power bus noise.

For example, suppose we have a board with one device that drives 32 CMOS loads (5 pF per load) and has a power dissipation capacitance, C_{PD} , of 380 pF. If the power bus voltage is 3.3 volts and the maximum allowable power bus noise voltage is 300 mV, then the amount of decoupling capacitance required for this board is,

$$C_d \approx [380 \text{ pF} + (32 \times 5 \text{ pF})] \frac{3.3 \text{ V}}{0.3 \text{ V}} = 0.01 \mu\text{F} \quad (6)$$

Since this approach does not account for the inductance of the capacitor connections, it is up to the designer to ensure that the capacitance is available to the device at all of the frequencies of interest. One method of doing this is to give each mounted decoupling capacitor an effective capacitance value that is its nominal value minus a term related to its connection inductance and frequency or transition time [10].

V. CONCLUSIONS

Power bus decoupling is an important part of any printed circuit board design. Strategies for effective decoupling depend on many factors including the board layout, board stackup, current requirements of the active devices and power bus noise requirements. Three general approaches have been described depending on whether or not the board has power planes and how closely these planes are spaced. Minimizing connection inductance is an important aspect of all three approaches.

While the three approaches described will help to determine the best way to locate and mount decoupling capacitors, they do not address the issue of how much decoupling is enough for a given application. Two methods were described for answering the question, "How much is enough?" While neither method is likely to give you an exact answer to this question, either method is preferable to relying on luck or rules-of-thumb based on experiences with older-technology products. A proactive, informed strategy for power bus decoupling, is much more likely to result in an effective board design that meets or exceeds all of its requirements.

REFERENCES

- [1] C. R. Paul, "Effectiveness of Multiple Decoupling Capacitors," IEEE Trans. on Electromagnetic Compatibility, vol. 34, no. 2, May 1992, pp. 130-133.
- [2] T. Zeeff, T. Hubing, T. Van Doren and D. Pommerenke, "Analysis of Simple Two-Capacitor Low-Pass Filters," IEEE Transactions on Electromagnetic Compatibility, vol. 45, no. 4, Nov. 2003, pp. 595-601.
- [3] T. H. Hubing, J. L. Drewniak, T. P. Van Doren, and D. Hockanson, "Power Bus Decoupling on Multilayer Printed Circuit Boards," IEEE Transactions on Electromagnetic Compatibility, vol. 37, no. 2, May 1995, pp. 155-166.
- [4] T. Zeeff, "Estimating the Connection Inductance of a Decoupling Capacitor," UMR EMC Laboratory Technical Report TR01-1-030, University of Missouri-Rolla, June 2001.
- [5] T. H. Hubing, T. P. Van Doren, F. Sha, J. L. Drewniak, and M. Wilhelm, "An Experimental Investigation of 4-Layer Printed Circuit Board Decoupling," Proceedings of the 1995 IEEE International Symposium on Electromagnetic Compatibility, Atlanta, GA, August 1995, pp. 308-312.

- [6] J. Fan, J. Drewniak, J. Knighten, N. Smith, A. Orlandi, T. Van Doren, T. Hubing and R. DuBroff, "Quantifying SMT Decoupling Capacitor Placement in DC Power-Bus Design for Multilayer PCBs," *IEEE Trans. on Electromagnetic Compatibility*, vol. EMC-43, no. 4, November 2001, pp. 588-599.
- [7] J. Fan, W. Cui, J. Drewniak, T. Van Doren, J. Knighten, "Estimating the Noise Mitigation Effect of Local Decoupling in Printed Circuit Boards," *IEEE Trans. on Advanced Packaging*, vol. 25, no. 2, May 2002, pp. 154-165.
- [8] H. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, 1993, Section 8.2.
- [9] L. Smith, R. Anderson, D. Forehand, T. Pelc and T. Roy, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," *IEEE Transactions on Advanced Packaging*, vol. 22, no. 3, Aug. 1999, pp. 284-291.
- [10] T. Van Doren, "Power Bus Noise Estimation Algorithm," UMR EMC Laboratory Technical Report TR99-3-024, University of Missouri-Rolla, May 1999.