EMC Course Notes 2024

Printed Circuit Board Layout

Printed Circuit Board Layout

Some circuits are fabricated on tiny silicon wafers and others consist of various components connected by wires. However, the circuits that are often the center of an EMC engineer's attention are those that are laid out on fiberglass epoxy boards. Printed circuit boards like the one illustrated in Figure 11.1 can be found in nearly all electronic systems. Circuit *components* with metal pins are connected by copper *traces*. *Surface mount technology* (SMT) components are glued to the top and/or bottom of a board. *Pin-in-hole* (or *through-hole*) components are held to the board by their pins, which extend through the board and are soldered to traces on the opposite side.

Single-layer boards have all the traces routed on one side of the board. Two-layer boards have traces on both sides. Multi-layer boards have several layers of copper traces separated by layers of fiberglass epoxy (or a similar dielectric). The number of layers is usually even. Four-layer boards are very common in low-cost products. Boards with dozens of layers are sometimes used to connect densely populated boards with high component pin counts.

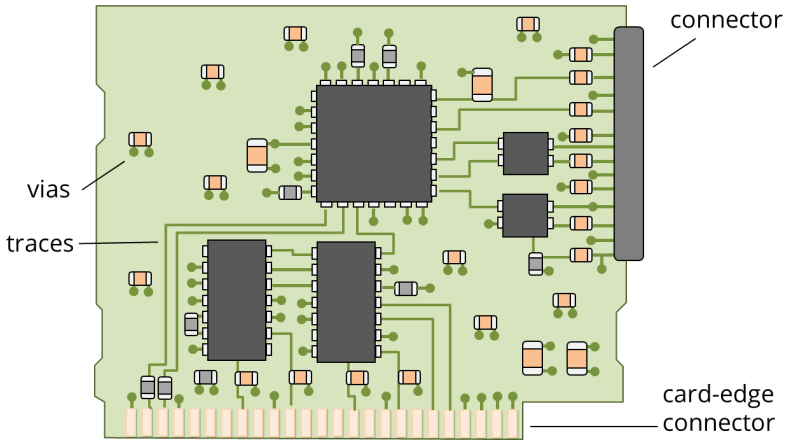

Figure 11.1. A printed circuit board.

Multi-layer boards often have entire layers with solid copper planes dedicated to the distribution of power to the components on the board. These planes provide a low-impedance path for the DC power currents, as well as a return-current path for single-ended signals.

The placement of components and the routing of traces usually play a crucial role in determining the electromagnetic compatibility of products employing printed circuit boards. Well-laid-out boards will not radiate significantly, and they do a good job of minimizing currents and fields that might couple noise to cables or other objects off the board. They also are configured to minimize opportunities for external currents or fields to couple interfering signals on to the board.

Strategies for Laying Out Printed Circuit Boards

Many board designers employ a list of guidelines to help them place components and route traces. For example, a typical guideline might be "minimize the length of all traces carrying high-speed digital signals." That might be good advice in general, but long traces carrying high-speed digital signals do not necessarily cause a problem. High-speed digital signals

are routinely routed across large boards, or even between boards, without any EMC or signal integrity problems. So, exactly how important is it to minimize the length of these traces?

A designer may not be familiar with the reason for the guideline or may not fully understand the consequences of violating the guideline for a particular application. And often, boards laid out by people trying to adhere to a long list of EMC guidelines are worse than boards laid out by people with no EMC guidance at all.

For example, suppose someone is laying out a high-speed multi-layer printed circuit board and needs to route a trace carrying a high-frequency signal from a digital component to an analog amplifier. They want to minimize the chance of having an EMC problem, so they search the web for EMC design guidelines and find three guidelines that seem to pertain to their situation:

1. minimize the length of high-speed traces,
2. gap any solid planes between analog and digital circuits, and
3. never let a high-speed trace cross over a gap in the signal return plane.

They envision three possible routing strategies as illustrated in Figure 11.2. The first option routes the trace directly between the two components but leaves the plane between them solid. The second option gaps the plane but routes the trace over the gap. The third option routes the trace around the gap. Each of these alternatives violates one of the guidelines.

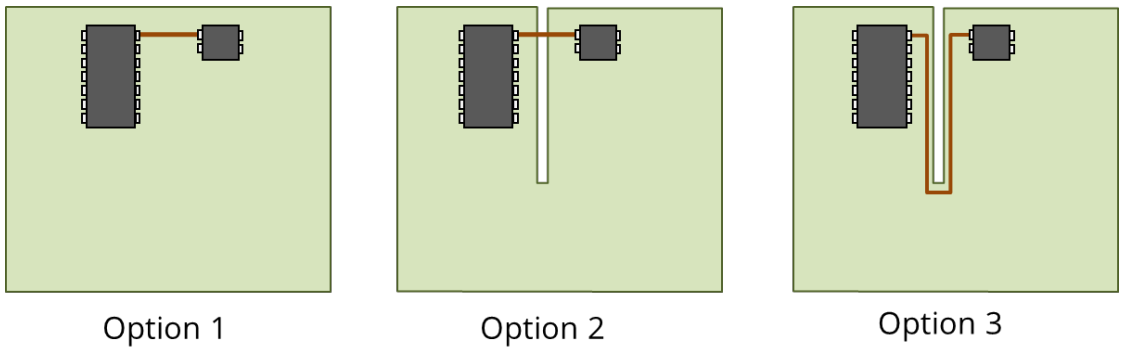

Figure 11.2. Which is the best trace-routing alternative?

Is each alternative equally good because it satisfies 2 of 3 guidelines? Are they all bad because they all violate at least one guideline? Making the right choice can be the difference between a board that meets all requirements and a board that has severe radiated emissions or immunity problems. In this example, Option 1 would almost always be the correct choice. The guideline to gap ground planes between analog and digital components is usually terrible advice, but it is still found in some texts and many internet sources.

There are no EMC design rules that can be applied in every situation. Blindly adhering to a list of EMC design rules is not likely to result in an EMC compliant design. The best strategy for designing and laying out EMC compliant printed circuit boards involves identifying and quantifying possible sources, victims and coupling paths.

In this chapter, we will explore steps that every EMC engineer should follow when laying out a printed circuit board or reviewing an existing board design. These steps include:

- Identifying potential EMI sources and victims,
- Identifying potential antennas/ports,
- Quantifying possible coupling that could impact EMC compliance.

By understanding how various design choices can make a difference in an EMC compliance test, component placement and trace routing decisions become more clear.

Identifying Potential EMI Sources and Victims

A typical circuit board may have dozens, hundreds or even thousands of circuits. Each circuit is a potential source of energy that might eventually be coupled unintentionally to other circuits or devices. Each circuit is also a potential victim of unintentionally coupled noise. However, some circuits are much more likely than others to be a noise source and other circuits are much more likely to be victims. EMC engineers (and board designers) should be able to recognize those circuits that are potentially good sources and those that are potentially most susceptible. Circuits of particular interest are discussed below.

Digital Clock and Signal Circuits

Synchronous digital circuits employ a system clock. Clock signals are constantly switching and have narrow band harmonics. They are often among the most energetic signals on a printed circuit board. For this reason, it is common to see narrow band radiated emission peaks at harmonics of the clock frequency, as illustrated in Figure 11.3.

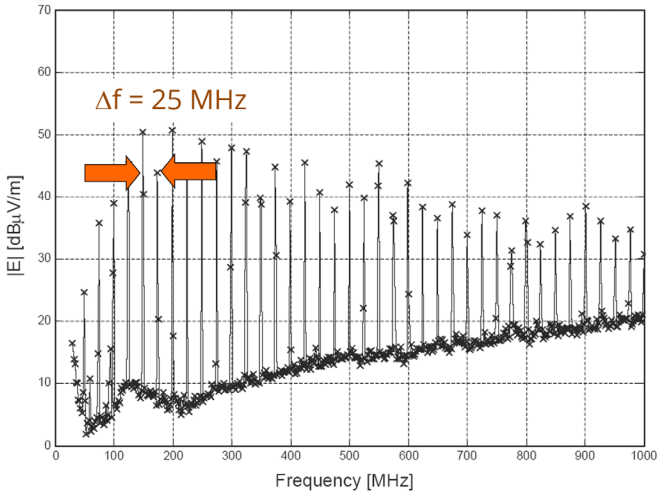

Figure 11.3. Radiated emissions from a product with a 25-MHz clock.

In this figure, the radiated emissions are dominated by harmonics of a 25-MHz clock. The noise floor from 200-1000 MHz is the thermal noise of the spectrum analyzer used to make the measurement (corrected to reflect the antenna factor). To make this product compliant with the FCC or CISPR Class B radiated emission specification, the amplitude of the source

harmonics could be decreased, the unintentional “antenna” could be made less efficient, or the source-antenna coupling path could be attenuated.

Digital signals are not necessarily as periodic as clock signals and their random nature results in noise that may be more broadband. Digital signals that toggle more often can result in radiation emission profiles similar to clock signals. An example of this would be the least significant bit on a microprocessor address bus, since stepping through consecutive addresses can cause this signal to toggle at the clock frequency. The exact form and strength of the radiation from digital signals depends on many factors including the software that is running, and the encoding scheme employed.

Analog Signals

Analog signals can be broadband or narrowband, high frequency or low frequency. If your board employs analog signals, it is a good idea to be familiar with how these signals look in both the time and frequency domains. Narrowband, high-frequency analog signals can be strong emissions sources. Fortunately, it’s relatively easy to control the path of these signal currents because they will take the path of least inductance and minimize their own loop areas. Low-frequency analog signals can be more problematic because the signal current paths are less well contained. These signals may be vulnerable to interference from other low-frequency sources sharing the same current path. For both high- and low-frequency analog signals, it’s important to follow the current and evaluate the coupling to and from the signal path.

Power Switching Circuits

Switch-mode power supplies, DC-DC converters, and power inverters generate different voltages by rapidly switching the current into an inductor or transformer on and off. Typical switching frequencies are in the 10-100 kHz range for high power inverters and 1-10 MHz range for low power DC-DC converters. The spikes of current generated by this switching can put noise on the power distribution nets and couple to other devices and circuits on the board. Although this noise is relatively periodic, producing narrow band harmonics, it can appear as broadband noise in a conducted or radiated emissions test when the switching frequency is lower than the resolution bandwidth of the measurement.

The small hump in the noise floor around 120 MHz in Figure 11.3 is due to power switching noise. In this product, the switching noise is negligible relative to the clock noise. However, in other products the power switching noise can dominate. For example, some DC-DC converters switching at MHz frequencies with sub-nanosecond transition times create broadband noise at GHz frequencies that can impact the performance of on-board radio receivers.

Power switching noise can be reduced by slowing down the transition time of the switching circuit. However, transition times that are greater than about 1% of the switching period can noticeably reduce the efficiency of the power supply. Field-coupled noise from switching power circuits is typically mitigated by keeping the switch nodes and switching current loops away from nearby circuits that might carry the noise away from the immediate area. Conducted coupling is generally controlled by carefully selected and located filter components.

DC Power Distribution and Low-Speed Digital Signals

DC power and low-speed digital signals do not usually have enough power at radiated emission frequencies to be troublesome. Nevertheless, DC power and low-speed digital signal conductors often carry strong high-frequency currents that can be a significant source of radiated emissions problems. The source of these currents is typically the high-speed components connected to these conductors. In some cases, the clock-frequency currents on nominally low-frequency traces can exceed the clock-frequency currents on the clock traces.

Figure 11.4 shows a map of the near magnetic field above a memory module. The near magnetic field provides an indication of the currents flowing in the lead frame of the component package. The frequency of the measurement is the third harmonic of the clock frequency. Note that more current is being drawn from the DC power supply pins than is being drawn from the signal pins.

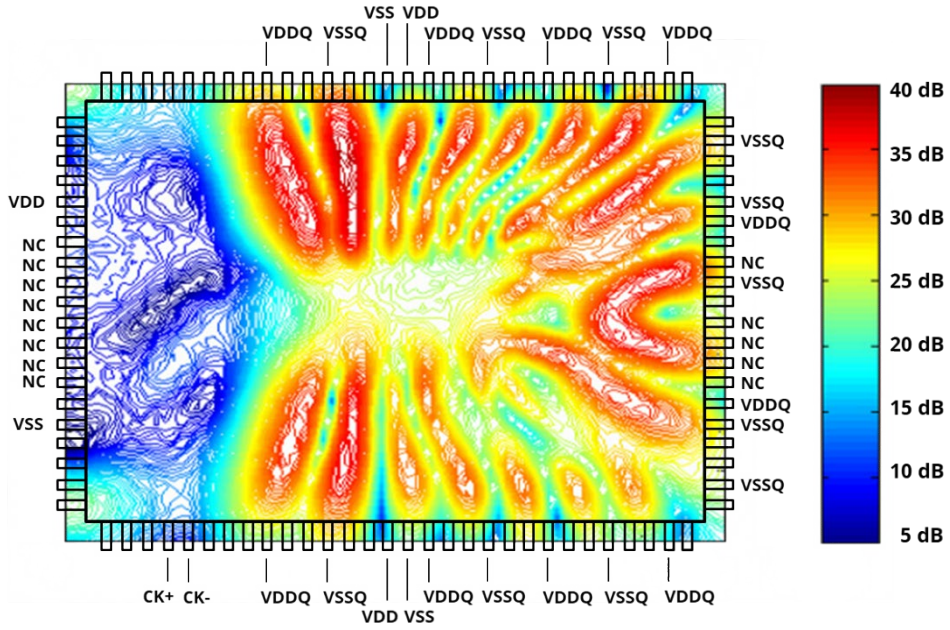

Figure 11.4. Near magnetic field above a packaged integrated circuit.

Figure 11.5 shows a similar plot of the near magnetic fields above a microcontroller implemented in a field programmable gate array (FPGA). The field was measured at the fundamental clock frequency. In this case, the strongest current is on the clock trace. However, currents of similar amplitude are flowing in some of the pins across the top. These are address lines that were not switched at all during the measurement.

Why do high-frequency currents and voltages appear on low-frequency data lines? There are several ways that this can happen depending on the design and layout of the integrated circuit (IC). Some ICs do a good job of containing their internally generated noise and others do not. A poor design can cause high-frequency voltage fluctuations on every input and output trace connected to the IC. And even good designs tend to put significant amounts of current on some of the I/O pins at harmonics of the internal clock frequencies.

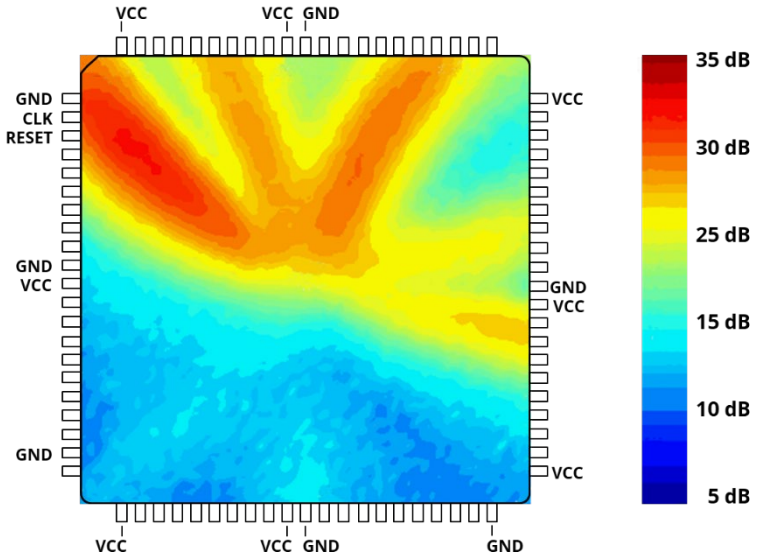

Figure 11.5. Near magnetic field above a microprocessor.

When laying out a printed circuit board with an IC that is clocked internally at a high frequency, it is a good idea to treat every pin on that IC as if it were a high-frequency source with the same characteristics as the internal clock. Otherwise, some of the power or low-speed digital traces connected to the IC could provide a conduit for high-frequency noise coming from the IC.

Wide-bandwidth, High-gain Amplifiers

It is also important to identify potential victim circuits to anticipate potential EM immunity problems with the board layout. Almost all of the circuits on a typical board are potentially vulnerable to large transients or strong RF noise; but wide-bandwidth, high-gain amplifier circuits deserve particular attention. This includes circuits that contain op-amps, even if the surrounding circuitry limits the nominal bandwidth and/or gain. Amplifiers with a wide bandwidth are generally sensitive to noise in a wide range of frequencies, including frequencies outside the nominal signal bandwidth. High-gain amplifiers can be sensitive to small amounts of coupled noise, and moderate amount of noise can saturate the amplifier resulting in the generation of strong signal harmonics.

Identifying Antennas/Ports

To fail a conducted or radiated emissions test, signals or noise generated on the board must be coupled off the board. Likewise, to fail a conducted or radiated immunity test, signals or noise originating off the board must be coupled onto the board. Identifying the unintentional antennas or “ports” that facilitate the coupling of noise to and from the circuit board is a key part of any EMC design review. To meet radiated emissions and immunity requirements, it is important to identify the structures that will ultimately be the unintentional antennas. This includes structures that may not be located on the board, but are driven by the board (e.g., cables and connections to the enclosure). To meet conducted emissions requirements, and immunity requirements, it is important to identify the ports (e.g., connectors or access points) that will bring interference onto or off-from the circuit board.

Unintentional Antennas

As described in the section on electromagnetic radiation in Chapter 3, there are three conditions that must be met for most unintentional antennas to radiate effectively:

- The antenna must have two conducting parts,
- neither part can be electrically small, and
- something must induce a voltage between the two conducting parts.

Printed circuit boards tend to be electrically small at frequencies below about 100 MHz ($\lambda > 3$ meters). This implies that any efficient antenna parts must be found off the board. Generally, at low frequencies the only viable antenna parts are the attached cables and/or a metallic chassis. If a printed circuit board is laid out in a manner that minimizes the possibility of inducing a voltage between any two of these antenna parts, then it is much less likely to have a radiated emission or radiated susceptibility problem.

Figure 11.6 shows two printed circuit board layouts. In both layouts, the connector and chassis connections represent possible efficient antenna parts. It's important to hold all the antenna parts to the same potential. As little as 1 mV driving one antenna part relative to another could cause a radiated emissions failure. In Layout #1, Connector 1 is located on the opposite side of the board from Connector 2 and the chassis connections. It would be difficult to keep the circuits on the board from driving one connector relative to the other.

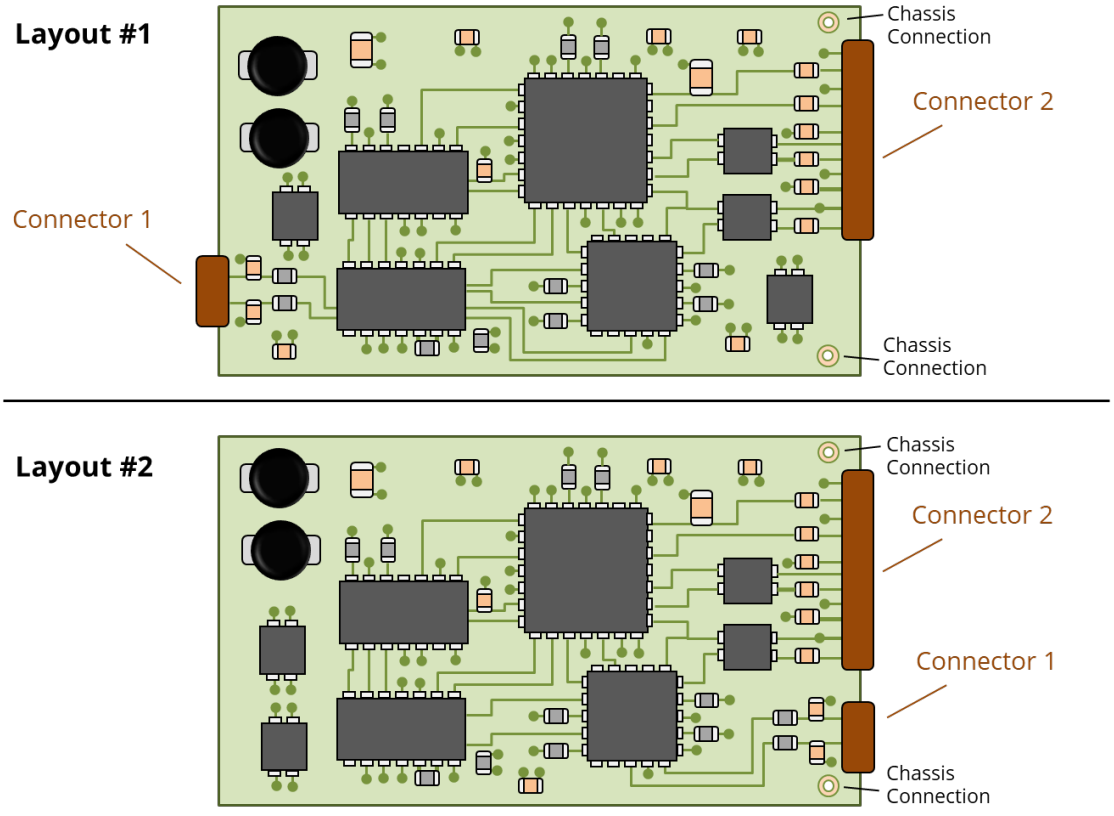

Figure 11.6. Two printed circuit board layouts.

Layout #2 is much less likely to have radiated emissions or immunity problems, because it would be much more difficult to develop a significant voltage between any two conductors capable of serving as an efficient antenna. In fact, locating all the I/O connectors on one side of a board is often an important first step in guaranteeing emissions and immunity compliance. This is especially true if the board does not have a metal chassis or enclosure to establish a high-frequency reference for all external connections.

At frequencies above 100 MHz, wavelengths are shorter, and it becomes more likely that objects mounted on the board (or the board itself) can serve as efficient antenna parts. Nevertheless, even at frequencies up to several GHz, these antenna parts should be relatively easy to spot. For example, at 1 GHz the wavelength in free space is 30 cm. A quarter wavelength is 7.5 cm. Therefore, an efficient antenna part will have to be at least several centimeters long and be driven relative to something that is just as large or larger.

Recall that differential currents (currents with nearby return paths) are relatively inefficient radiation sources. This means that a trace that is right beside or above its current return path is not a good antenna part. So, if one half of our antenna is a metallic plane on the board, the other half must stick up and away from the plane. This helps to make these antenna parts readily identifiable even at high frequencies. Table 11.1 lists common antenna parts found on printed circuit boards above and below 100 MHz.

Table 11.1. Printed circuit board objects that may or may not be parts of a good antenna.

Good Antenna Parts		Poor Antenna Parts	
< 100 MHz	**> 100 MHz**	**< 100 MHz**	**> 100 MHz**
cables or cable shields	tall components or heatsinks	integrated circuits	integrated circuits
metal enclosure or chassis	seams in shielding enclosures	microstrip or stripline traces	microstrip or stripline traces
	sparsely populated power planes	anything not big	densely populated power planes

EMC Ground on a Circuit Board

The concept of a printed circuit board EMC ground was introduced in Chapter 8. The importance of identifying a circuit board's EMC ground in the early stages of the layout cannot be overstated. Everything that is electrically large and metal on or near the circuit board should be bonded to the EMC ground (or kept far from components that could couple noise to or from it). This is a critical part of ensuring EMC compliance and it cannot be accomplished without first identifying which locations on the board constitute the EMC ground.

As indicated in Chapter 8, the circuit board's EMC ground is typically the region of the return plane directly beneath and near an external connection. As Figure 11.7 illustrates, bonding external connections to the EMC ground at high frequencies can prevent them from being driven by a voltage capable of causing radiated emissions.

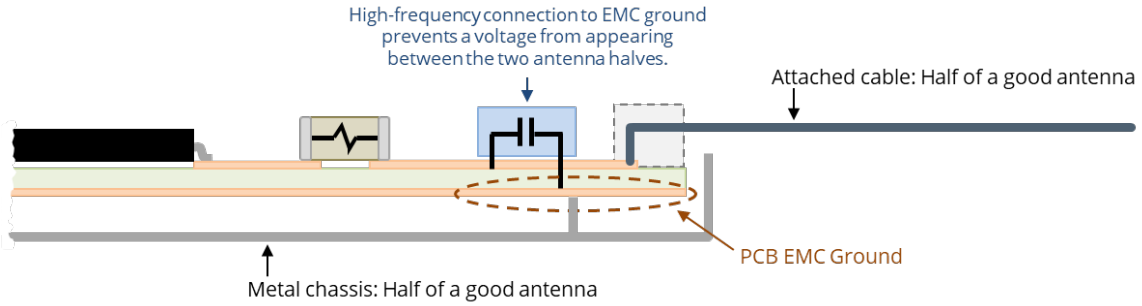

Figure 11.7. Bonding to the board's EMC ground to control emissions.

The EMC ground also plays an important role in meeting immunity requirements. When properly identified, the PCB EMC ground is the place on the board where injected noise currents must eventually go. Knowing where this is allows the board designer to direct these currents from the point where they are injected to the place where they need to go without flowing across the rest of the board.

As illustrated in Figure 11.8, the same capacitor used to control emissions provides a good path for injected high-frequency noise currents. Transient protection in the same location can be used, if needed, to redirect high-voltage transients.

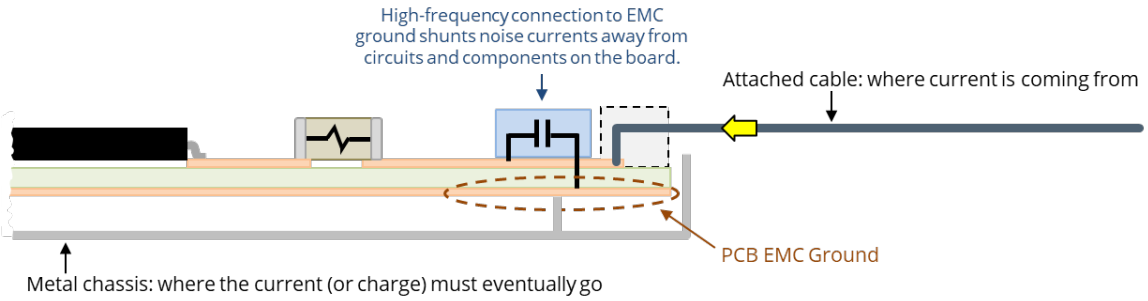

Figure 11.8. Bonding to the board's EMC ground to meet immunity requirements.

The examples in Figures 11.7 and 11.8 show a board mounted in a metal chassis. However, the PCB EMC ground is just as important in boards mounted in plastic enclosures. In this case, incoming noise currents usually need to be routed to the largest available metal structure. This is typically the board's ground plane. The procedure for identifying PCB EMC grounds is more fully described in Chapter 8.

Identifying and Quantifying the EM Coupling

Once we have identified the potential sources or victims and the potential antennas/ports, a good board layout is simply a matter of controlling the coupling between the two. Earlier, we learned that there are 4 possible electromagnetic coupling mechanisms:

- Conducted coupling,
- Electric field coupling,
- Magnetic field coupling, and
- Radiation.

On a circuit board, the coupling between a source and its antenna/port will not be radiation coupling, so there are only three coupling mechanisms that we need to consider.

Conducted Coupling

Conducted coupling will only occur if the source is directly connected to and drives one good antenna part relative to another. Two examples of conducted coupling are illustrated in Figure 11.9. Figure 11.9(a) shows a single-ended signal being transmitted on an unshielded twisted wire pair. A common-mode voltage is developed between the unbalanced driver and the balanced cable, as described in Chapter 7. This voltage drives the cable relative to the board and chassis connections.

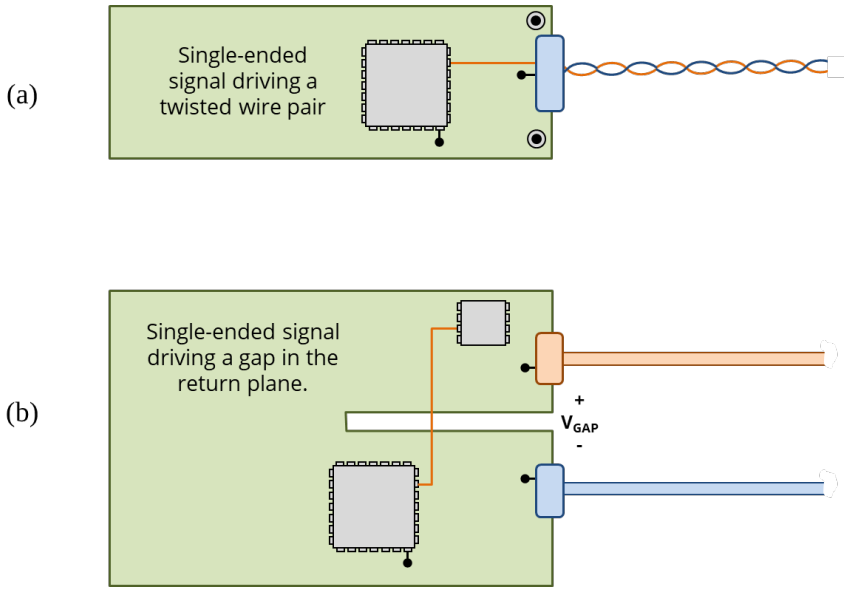

Figure 11.9. Two examples of conducted coupling driving an antenna structure.

Figure 11.9(b) illustrates a single-ended signal trace crossing over a gap in the return plane. The impedance in the path of the return current generates a voltage across the gap. This voltage drives cables referenced to one side of the gap relative to cables referenced to the other side.

Conducted coupling tends to be easy to spot once the source and the antenna parts have been identified. By simply paying attention to the high-frequency current paths, conducted coupling between intentional signals and unintentional antennas can usually be avoided.

Electric Field Coupling

An example of electric field coupling that can produce radiated emissions is illustrated in Figure 11.10. A large heatsink is mounted over an integrated circuit. As the IC pulls current from the power and ground at high frequencies, it takes on a voltage relative to the plane below it. At frequencies where the heatsink is not electrically small, it is potentially an effective antenna part. Electric field coupling from the surface of the IC drives the heatsink relative to the board planes. As little as 1 mV between the heatsink and board at the structure's resonant frequency can cause a radiated emissions failure.

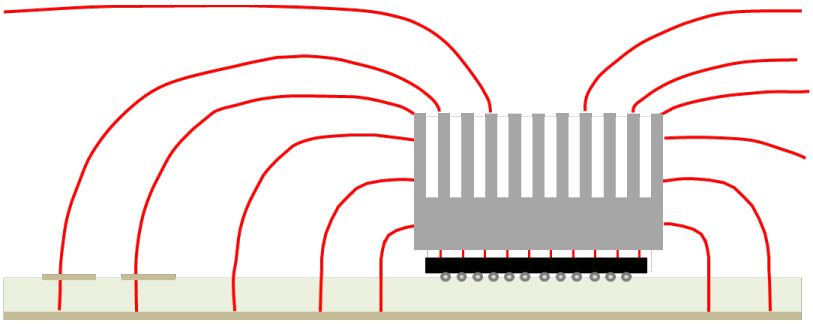

Figure 11.10. Integrated circuit package coupling to a heatsink.

Another way that the structure in Figure 11.10 can cause radiated emissions is for the field from the heatsink to couple to nearby circuit board traces. These traces can carry noise to other areas of the board. If the traces connect to wires that leave the board, those wires can be driven relative to the board and chassis grounds.

When heatsinks extend horizontally beyond the surface of the IC, it's sometimes necessary to route traces directly beneath the heatsink. In this case, the electric field coupling to traces can be very strong, as much as half the noise voltage on the IC itself. If this noise cannot be tolerated, traces should be routed on a layer beneath the uppermost circuit board plane.

Note that grounding the heatsink to the uppermost plane would reduce its voltage relative to the plane. This would reduce any electric field coupling from the heatsink. Unfortunately, it is often challenging to establish a low-impedance bond between a heatsink and a circuit board plane at high (e.g., GHz) frequencies. In many cases, it is better to minimize the size of the heatsink and keep unnecessary traces and components a safe distance away.

Of course, electric field coupling doesn't have to involve heatsinks. Any component or trace with a high-frequency voltage relative to the board ground can be the source of electric field coupling. The switch nodes of DC-to-DC converters and motor drivers are common examples of circuit board traces that can strongly couple to nearby traces or components and carry that noise to unintentional antennas. A worst-case analysis of this coupling can determine whether it is likely to couple enough noise to be problem. If so, the coupling can be reduced by layout changes or (in some cases) using board-level electric field shielding.

Magnetic Field Coupling

Figure 11.11 illustrates a very common magnetic field coupling problem in many printed circuit board designs. An otherwise well-designed board has connectors attached to each side. A circuit consisting of a single microstrip trace driven at one end and terminated at the other end is located between the two connectors.

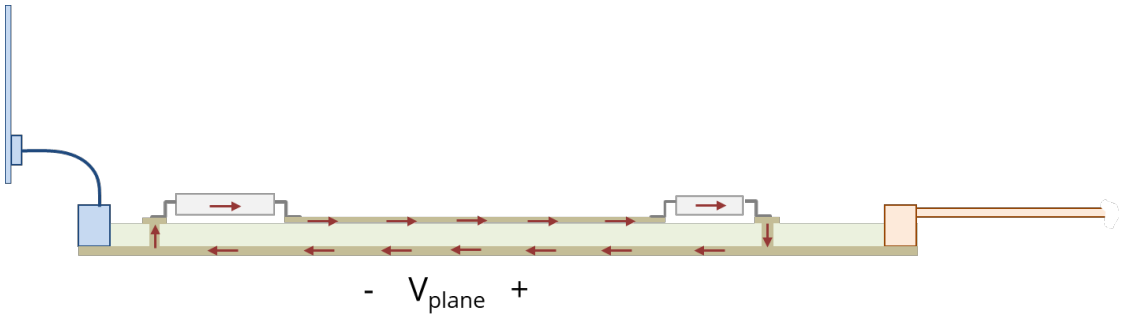

Figure 11.11. Example of magnetic field coupling on a circuit board.

Microstrip traces are not efficient radiated emission sources, so the only possible antenna parts in this design are the two connectors and their attached cables. One might expect the two antenna parts to be at the same potential because they are connected to each other with a wide copper plane. However, an important requirement for a "ground" conductor is that it does not carry power or signal currents capable of producing an objectionable voltage difference.

The "ground" plane in this design does carry signal currents. The current flowing in the plane generates a magnetic flux that wraps around the plane. The flux wrapping the plane produces a voltage difference that can drive things referenced to the plane on opposite sides. This voltage difference is difficult to quantify precisely because it depends on many factors including the position and orientation of the cables. Nevertheless, a general rule of thumb is that tens of milliamps at tens of megahertz will effectively induce millivolts of voltage. Increasing the trace width or trace height increases this voltage, while stripline traces with currents returning above and below the trace produce no voltage at all.

A good way to avoid this problem is to make all the external connections on one edge of the board. Then, provided no high-speed traces are routed near and parallel-to the edge, the plane on that edge will provide a good reference ground.

Another option when high-speed circuitry must be located between connectors is to mount the board on a metal chassis that makes a good connection to the board ground plane on both sides. This effectively shorts out the voltage induced across the plane and allows currents to circulate harmlessly on the chassis.

Crosstalk to an I/O Trace

Although, strictly speaking, it is not an independent coupling mechanism, another common problem that occurs with printed circuit board layouts is inadvertent crosstalk between signal traces and I/O traces capable of carrying that noise off the board. An example of this is illustrated in Figure 11.12. A moderately high-speed trace is routed alongside another trace that attaches to a connector. Voltages and/or currents coupled from one trace to the other (via electric or magnetic fields) can be propagated down the I/O trace and off the board. This coupling effectively drives the cable relative to the board and chassis ground.

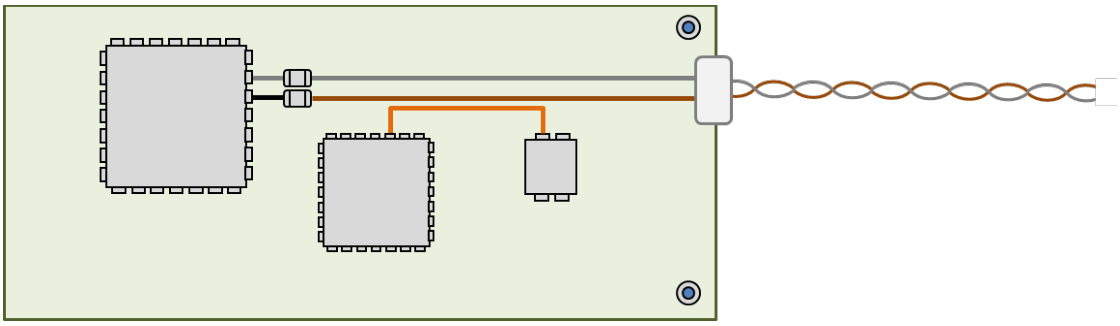

Figure 11.12. Crosstalk to an I/O trace.

This problem is usually obvious once you see it. However, on a board with hundreds or thousands of traces, this situation arises more often than it should. If the board layout software can't check for I/O traces that are routed in the vicinity of high-speed traces, then this should be done manually. The same also applies to I/O traces routed in the vicinity of traces connected to vulnerable inputs, since the easiest way for radiated noise to get onto a board is through the I/O.

Printed Circuit Board Design Review

Now that we know a little more about noise sources, antennas and coupling mechanisms on printed circuit boards, let's look at some of the steps that can be taken to systematically review a printed circuit board design. We'll assume we have circuit schematics and a preliminary board layout available. We may also need component datasheets and physical descriptions of the various signals on the board and the I/O. Finally, we'll need to understand the product's function, grounding strategy and EMC requirements.

The design review steps provided here are intended to provide general guidance concerning things to look for and basic methods for determining when something must be changed. They are presented in a specific order that generally works well but may not be optimum in all situations. It's more of a checklist rather than a set of step-by-step instructions.

1. Identify connections to the ground structure and the locations of all I/O.

An essential part of making good EMC design decisions is understanding where your board's EMC ground is located. Basically, all the I/O will be bonded to (or referenced to) this ground at high frequencies. Any I/O that can't be bonded to ground, must be filtered, shielded, or isolated. Alternatively, we must be prepared to demonstrate that the ungrounded I/O cannot possibly contribute to an EMC compliance problem.

If the board carries high-speed signals and does not have a metal chassis, it's important to put all the connectors on one edge or in one corner. If high-speed circuitry must be located between connectors, then there must be a plan in place to prevent one cable from being driven relative to another.

2. Identify ground plane layer(s) and the extent of the plane.

Ideally, the board has enough solid ground (current-return) planes to ensure that high-frequency return currents come back directly above and/or below the signal trace. The

planes may not need to extend to areas of the board where they are not utilized, but these planes should never be gapped for any reason.

3. Identify ground traces or fills on other layers.

A board with solid ground planes should usually not have ground traces or ground fill on other layers. When a connection needs to be made to ground, it should be made with a via connecting directly to the plane. Ground traces and ground fill add unnecessary inductance to the ground plane connection. Sometimes the presence of a ground trace or fill indicates that two or more components are sharing a connection to the ground plane. This can result in common-impedance coupling. Ground traces and ground fills serve a useful purpose on boards with 1 or 2 layers, but they should not be used on most multi-layer boards with solid ground planes.

4. Identify power distribution layer(s) and the type of power distribution.

If the board has active components that have three or more power pins at a given voltage, that voltage should probably be distributed on a closely spaced (i.e., < 0.25 mm) pair of planes. High-frequency decoupling capacitors should be mounted on the side of the board closest to the plane pair. If the active components have only one or two power pins, power can be routed on planes or traces. The decoupling should be appropriate for the power distribution geometry as described in the *Power Bus Decoupling* section in this chapter.

5. Identify the I/O nets.

With board layout viewers, it's easy to highlight each net attached to a connector pin. These are the nets that will carry noise emissions off the board and/or bring interfering signals or transients onto the board. Ideally, these nets terminate in a series component near the connector. Follow each I/O net and look for possible ways that noise might couple to or from the trace. If there is a potential coupling problem, do a worst-case coupling calculation to determine if a layout change is required.

If the I/O net connects directly to an IC, make sure signals or noise from the IC can't drive the I/O hard enough to fail a compliance test (e.g., with more than 1 mV of common-mode voltage difference). Remember that harmonics of an IC's internal clock can be very strong on any of its input or output pins. If necessary, provide filtering and/or transient protection on the I/O net.

6. Identify filter components.

Filters play an important role in EMC compliance. Look for filter components to ensure that they are well-mounted to minimize parasitics and prevent inadvertent coupling that would bypass the filter. Make sure that higher order filters are designed appropriately and won't ring at their cut-off frequency. If the filter has an inductor, determine whether it needs to be shielded. If the filter uses ferrites, check the datasheet to be sure they are appropriate for that application. If the filter employs a common-mode choke, make sure it is filtering a balanced pair of signal or power conductors.

7. Identify traces requiring transition-time control with series resistance.

By default, it is a good idea to assume every digital signal trace driving a capacitive load needs a series resistor to control its transition time. Then, remove the resistor from any trace where it is clearly not required. The most common reason for removing the transition

time control resistor is when you can follow the trace current from end to end and see clearly that there is no coupling path that could carry signal noise to a victim circuit or to an unintentional antenna. If there is any doubt about the coupling, do a worst-case coupling calculation.

The transition time of signals on controlled impedance traces with matched terminations should be controlled using logic with a controlled slew-rate. If that logic is not available, the signal should be filtered, or the traces should be well-isolated from anything that could couple noise in or out of the signal path.

8. Identify critical components.

Some components on a circuit board require special attention during the design review. For example, inductors are generally associated with filters, power converters, or RF circuits. Each of these applications requires a little extra attention during layout. Identifying the inductor locations helps to focus attention on regions where the board layout can be critical.

Analog amplifiers (particularly op-amps) are another type of component with special layout concerns. Op-amps inherently have high gain and a wide bandwidth. Small amounts of noise coupled to the input or feedback loop of an op-amp can disrupt the entire circuit. Ensure that the loop areas associated with the input and feedback components are small. Make sure the traces are routed over a solid ground plane and do a worst-case calculation of the coupled E-fields and H-fields that would occur during an immunity test.

Transient protection components are also critical. Make sure they are chosen appropriately for the application. Also, be sure they connect to the correct "ground." Components protecting the board should be connected to the circuit board EMC ground. Components protecting specific IC inputs should be connected to the signal current-return conductor associated with that input. For fast transients, the inductance of the connection is critical. Be sure the time it takes for the transient protection to turn on is shorter than the time it takes for a component to be damaged.

Phase-locked loops (e.g., as used in clock multipliers) are very sensitive to small amounts of power bus noise. Some ICs that employ multipliers have certain power inputs that require filtered power. Often, this is accomplished with a ferrite on the power input. It's generally a good idea to follow the recommendations of the IC manufacturer, but never put a ferrite on the ground connection.

When ferrites appear in series with traces that are not carrying DC power, check to ensure that a ferrite is the right choice. For capacitive loads, resistors are usually a much better option. If the ferrite is on a ground conductor, or a single-ended signal-return conductor, it should probably be removed.

9. Identify power inverter circuits/devices.

Locate the switching voltage nodes and minimize their surface area. Keep them away from any conductors that might carry the switching noise away from the immediate area. Identify the switching current loops and look for magnetic field coupling to nearby circuits that could carry the switching noise away.

Watch for unintended coupling from the components as well as the traces. Do a worst-case calculation of this coupling and use shielded components if necessary.

Power converters that have come out in the past few years tend to be much smaller and quieter than older converters. Low to medium power DC-to-DC converters should be switching at frequencies in the MHz range. They should not have isolated power grounds. Many of the converters available today have EMI reducing features such as controlled slew-rates, spread-spectrum switching and optimized pin layouts. Take advantage of these features and don't stick with older components just because they have been good enough in the past.

Finally, if the power converter is in the vicinity of an intentional receiving antenna (e.g., GPS or Wi-Fi), board-level shielding may be required to reduce coupling from the converter that could impact the signal-to-noise ratio of the receiver.

10. Identify currents and coupling for each EMC compliance test.

As a final check of the board layout, it's a good idea to review each EMC compliance test requirement and visualize the impact of the layout on that test. For example, if the board will be tested for conducted emissions, identify the most likely noise sources and follow the EM coupling paths that could bring that noise to the LISN. Do a quick, worst-case calculation for any coupling that looks like it could cause a compliance problem. Ensure that the power input filter will have sufficient insertion loss at the required frequencies.

Similarly, for radiated emissions, identify the most likely noise sources and evaluate the worst-case coupling to potential antenna structures. Don't try to calculate radiated field strengths but ensure that no two antenna parts can be driven by a millivolt or more at any radiated emissions test frequency. For emissions requirements more stringent than CISPR 32 Class B, this voltage difference might need to be as low as 100 microvolts.

If the board will be subjected to a bulk current injection (BCI) test, follow the current as it comes in on the harness and makes its way to the metal table-top. The last segment of this path may be the parasitic capacitance between the table-top and the device's chassis or circuit board planes. Be sure there is a low impedance path between the points where the BCI currents enter and exit the board. Watch for any EM coupling that might occur to vulnerable components or circuits. Evaluate the worst-case coupling to ensure that the board will not have any trouble complying with the requirements of this test.

For radiated immunity tests, follow the same procedure used for the bulk current injection test. In addition, if the board is not in a metal enclosure, evaluate the worst-case coupling that could occur directly from the fields to vulnerable circuits.

If the board will be subjected to electrical fast transients, model the source to determine the worst-case current waveform. Follow the current from the harness to the place where it will leave the board. Watch for coupling that might result in a failure to comply with the test requirement. Add transient protection or filtering if necessary.

If the board is in a metal enclosure, the ESD currents are most likely to come in on connector pins. Follow the current on any I/O nets that will be subjected to an ESD test. Use blocking resistors or transient protection circuits to protect any components that may not survive a direct discharge. Make sure the ESD current has a conducted path to the metal enclosure without being pulled too far onto the board.

If the board is in a plastic enclosure, look for possible arc paths through seams or openings in the enclosure. If the path cannot be blocked, be sure the ESD current can get to the board's EMC ground without being pulled onto the board.

Boards that are not in a shielded enclosure will experience strong field coupling during the initial part of the transient. Make sure that short, low-energy voltage spikes coupled to any of the circuits will not cause an unacceptable ESD failure. Use resistors to slow the response time of high-impedance digital inputs. Verify that any high-speed data errors will be detected and corrected.

By visualizing each EMC test and the path of the relevant conduction and displacement currents, potential compliance issues often become obvious. This allows board designers to address them before the first prototypes are built and tested. And even in situations where compliance vs. non-compliance is not clear, the visualization exercise allows the designer to anticipate possible problems. In these situations, mounting pads and clearances can be incorporated in the design to allow for filtering, transient protection or shielding if it is necessary.

Power Bus Decoupling

A sudden change in the amount of current drawn by an electronic device can cause a momentary dip or spike in the voltage on the power distribution bus. This voltage variation can impact the operation of other devices on the same power bus. The goal of power bus decoupling is to reduce power bus noise coupling between active circuits.

In many ways, power bus decoupling is similar to power bus filtering. However, most forms of power bus filtering impede the flow of high-frequency currents on the power bus. The focus of power bus decoupling is to meet the high-frequency current demands of active devices without allowing the power bus voltage to vary significantly.

Consider the printed circuit board illustrated in Figure 11.13(a). A low impedance, 3.3-volt source supplies power to two active components on a printed circuit board. Assuming the resistance of the traces and connections is negligible, the voltage appearing at the pins of both components is approximately 3.3 volts. However, a sudden change in the current drawn by one of the components can cause the voltage delivered to the board to momentarily drop due to the connection inductance.

For example, assume the connection between the power source and the board was made using a 30-cm long twisted wire pair with an inductance per unit length of 5 nH/cm. If the current drawn by the board changes from 1.00 amperes to 1.15 amperes in 20 ns, the voltage seen by both components on the board would drop momentarily. The average voltage during this transition would be,

$$
\begin{aligned}
V &= V_{initial} - L\frac{\partial I}{\partial t} \\
&\approx V_{initial} - L\frac{\Delta I}{\Delta t} \\
&\approx 3.3\text{v} - \left(150\times10^{-9}\text{H}\right)\frac{0.15\text{A}}{20\times10^{-9}\text{s}} \\
&\approx 2.2 \text{ volts.}
\end{aligned}
\tag{11.1}
$$

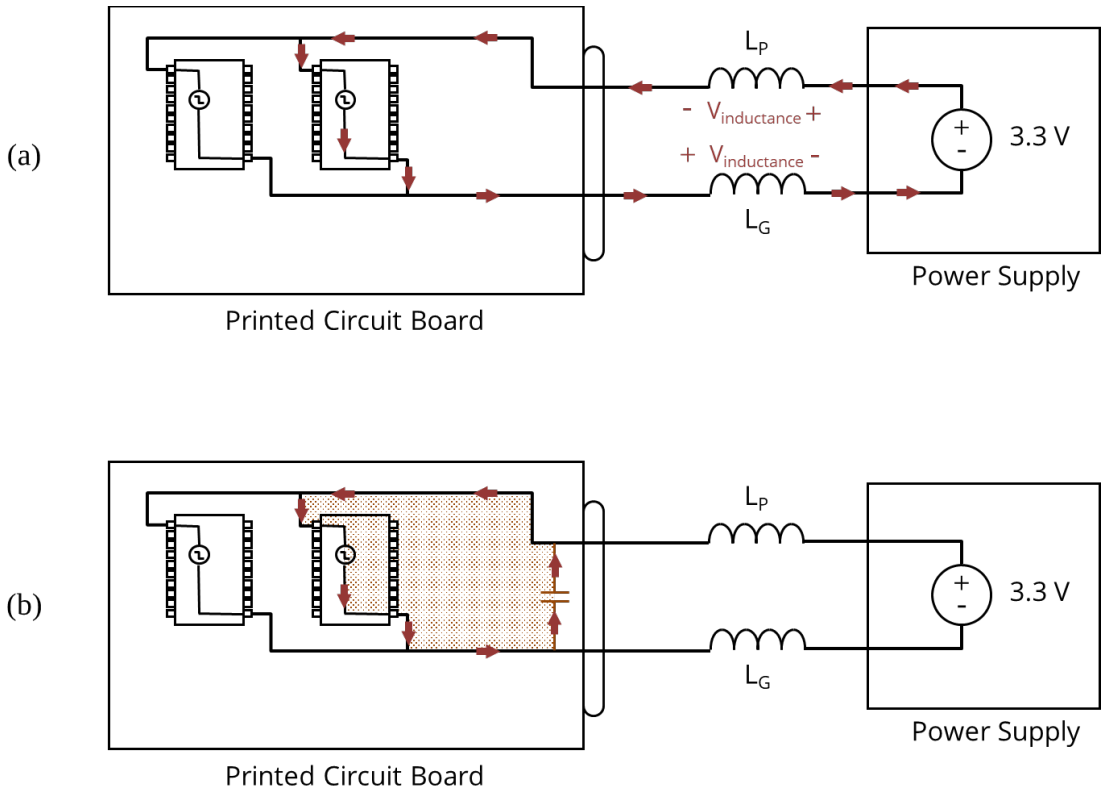

Figure 11.13 A sudden change in the current drawn by one component can affect the voltage on a printed circuit board power bus.

A larger or more sudden current change would cause the voltage to drop even more. This voltage transient would be observed by every component on the board and could result in data errors or, in severe cases, damage to the components.

One solution to this problem is to place a capacitor on the board in parallel with the power source as shown in Figure 11.13(b). The capacitor stores charge on the board and supplies current in response to any change in voltage across its terminals. Therefore, the capacitor helps to stabilize the voltage on the board. Capacitors with both terminals connected to different reference voltages, like the capacitor in Figure 11.13(b), are called *decoupling capacitors*, because they reduce common impedance coupling between components sharing the same power bus.

Power Bus Impedance

The ratio of the power bus voltage change to the change in current supplied by the power bus is the power bus impedance,

$$Z_{powerbus} = \frac{\Delta V}{\Delta I}. \tag{11.2}$$

If we know how much current we need to supply to a board at any given frequency, and we know how much noise voltage we are willing to tolerate, we can calculate the maximum allowable power bus impedance at that frequency. For example, if we plot the power bus impedance of the circuit in Figure 11.13(a) as a function of frequency, we see that the impedance increases with frequency due to the power bus inductance, as shown by the solid line in Figure 11.14.

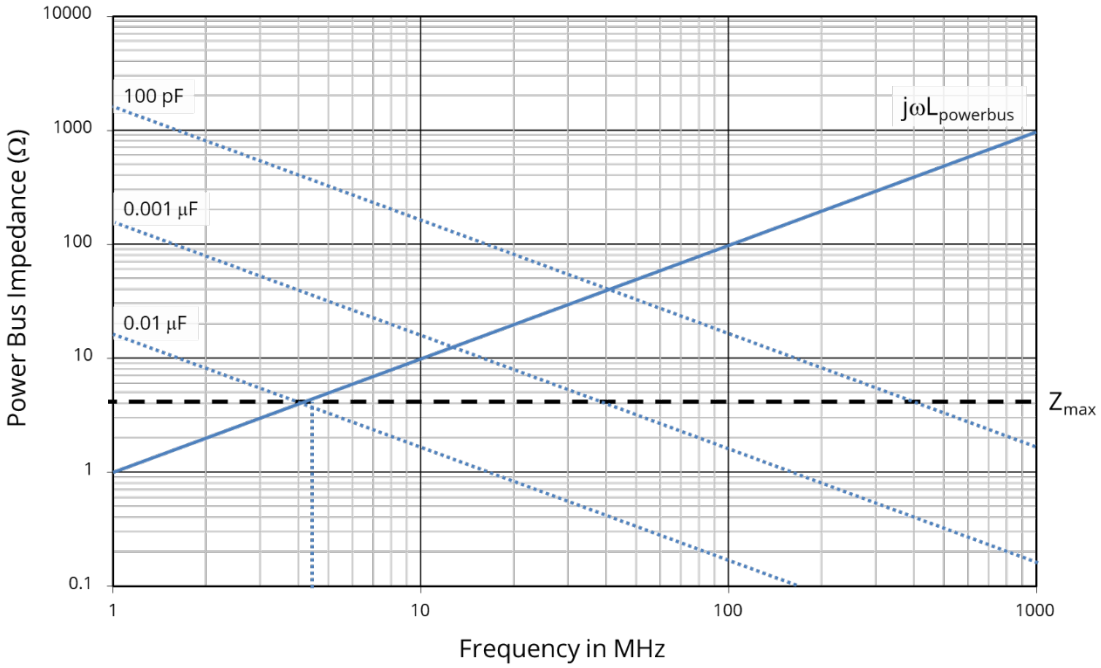

Figure 11.14 Power bus impedance.

If we were willing to tolerate no more than 0.033 volts of power bus noise, but we needed to be able to supply at least 8 mA of current at any frequency up to 25 MHz, the maximum allowable power bus impedance would be,

$$Z_{\max} = \frac{0.033 \text{ volts}}{0.008 \text{ amps}} = 4.1\,\Omega. \tag{11.3}$$

The maximum allowable power bus impedance is indicated by the horizontal dashed line in Figure 11.14. Note that power bus impedance exceeds the maximum allowable impedance at frequencies greater than,

$$f_0 = \frac{Z_{\max}}{2\pi L_{powerbus}} = \frac{4.1\,\Omega}{2\pi(150 \text{ nH})} = 4.4 \text{ MHz}. \tag{11.4}$$

In other words, drawing 8 mA of current from the power bus at any frequency above 4.4 MHz would cause the power bus noise to exceed the specification.

Suppose that we add a decoupling capacitor to the circuit as indicated in Figure 11.13(b). The impedance of the power bus now consists of the bus inductance and the decoupling capacitance in parallel. The impedance of the capacitor for several possible capacitance values is indicated by the dotted lines in Figure 11.14.

As indicated in the figure, this application requires a capacitor with a value of at least 0.01 μF to meet the power bus noise specification. Larger values of capacitance are even better, but smaller values do not supply enough current at the lower frequencies.

At high frequencies, the dotted lines in Figure 11.14 suggest that the capacitors are more than adequate. However, the connection between the active device and the decoupling capacitor also has a loop area as indicated in Figure 11.15(a). We need to model the connection inductance associated with the decoupling capacitor to adequately characterize the high-frequency behavior of the power bus. In this case, approximating the connection as a square 3-cm x 3-cm loop, we estimate the inductance of the connection to be approximately L_{global} = 40 nH. The subscript *global* indicates that this capacitor has approximately the same inductance to all the active devices on the board. The impedance of the connection inductance is plotted in Figure 11.16. Since this inductance is in series with the active device, the overall power bus impedance is determined by adding this impedance to the impedance of the rest of the bus.

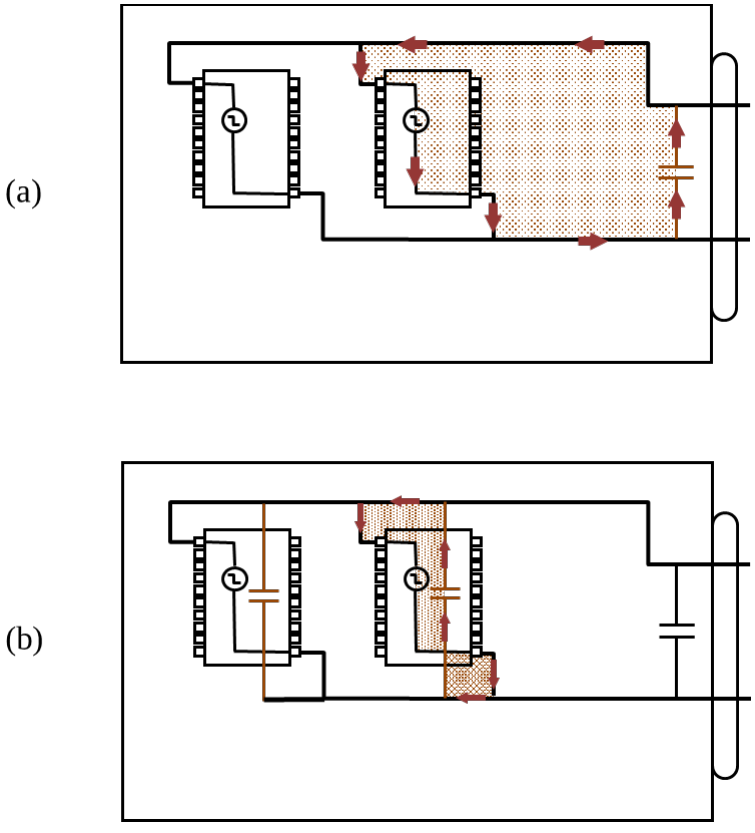

Figure 11.15. The loop area associated with a decoupling capacitor connection is responsible for connection inductance.

Note that the effect of adding the global decoupling capacitor to this circuit was to increase the maximum effective frequency of the power bus from f_0 = 4.4 MHz to f_1 = 16 MHz. Also note that the new maximum frequency is determined entirely by the connection inductance of the global decoupling capacitor. Changing the value of this capacitor without changing the connection inductance will have no effect on the bandwidth of the decoupling.

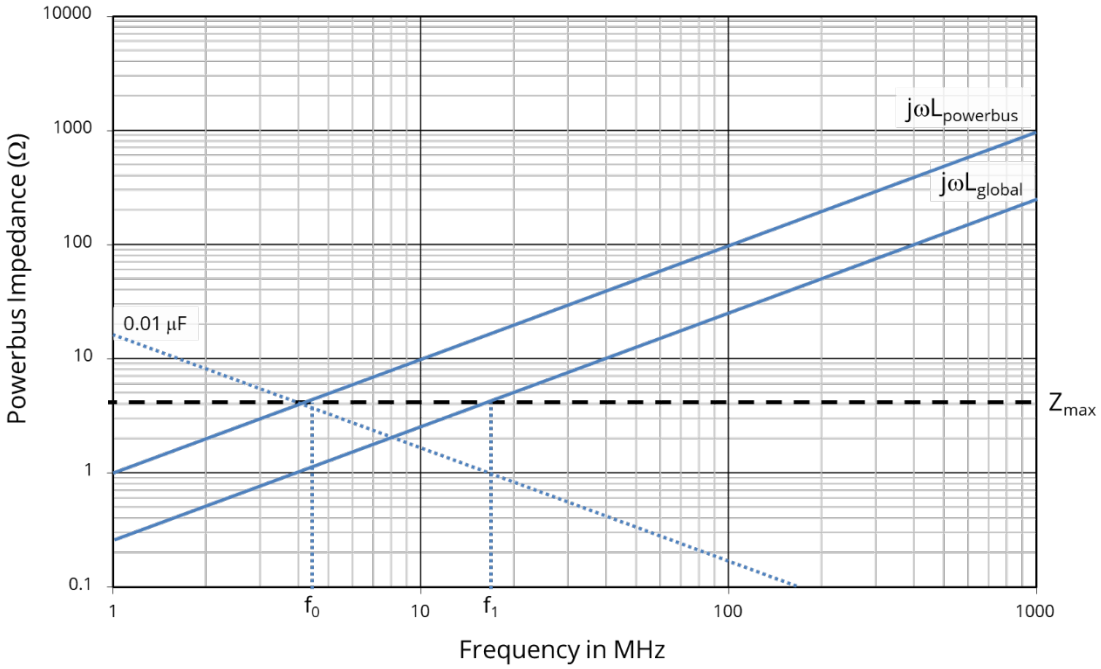

Figure 11.16. Power Bus impedance with local decoupling capacitor.

If it is necessary to provide effective power bus decoupling at higher frequencies, we must reduce the connection inductance associated with the decoupling capacitor. One effective way to accomplish this is to provide *local* decoupling capacitors near each active device as shown in Figure 11.15(b). The local decoupling capacitor does not need to have a large nominal value, because the low-frequency current demand is met by the power supply and the global decoupling capacitor. The value of the local decoupling capacitor only needs to be large enough to supply current to a single active device at frequencies where the global decoupling is no longer effective.

Since the current drawn by each individual device is generally less than the current drawn by the entire board, a new value of Z_{max} is required to calculate the minimum value of the local decoupling capacitor. In this case, if we assume the component being decoupled draws half the current that the board draws (40 mA), then

$$Z_{\text{device\_max}} = \frac{0.033 \text{ volts}}{0.004 \text{ amps}} = 8.2\ \Omega\,. \tag{11.4}$$

The global decoupling was adequate up to 16 MHz, so this is the frequency at which the local decoupling must start to be effective. Therefore, the minimum value of the local decoupling capacitance must be,

$$C_{local} \geq \frac{1}{2\pi f_1 Z_{device\,\max}} = \frac{1}{2\pi\left(16\times10^{6}\right)\left(8.2\right)} = 0.0012\,\mu\text{F}. \tag{11.5}$$

The highest frequency at which the local decoupling capacitor is effective is determined by its connection inductance. For example, if the inductance of the connection to the local decoupling capacitor is 15 nH, the decoupling is adequate up to,

$$f_2 = \frac{Z_{device\,\max}}{2\pi L_{local}} = \frac{8.2}{2\pi\left(15\times10^{-9}\right)} = 87\text{ MHz}, \tag{11.6}$$

which is good enough to meet our design goal of providing decoupling up to 25 MHz. If the device drew significant current at higher frequencies, it would be necessary to further reduce the impedance of the connection to the decoupling capacitor possibly by distributing the power on solid planes with multiple decoupling capacitors connected to the planes. The bus impedance seen by the device with the local decoupling is plotted in Figure 11.17.

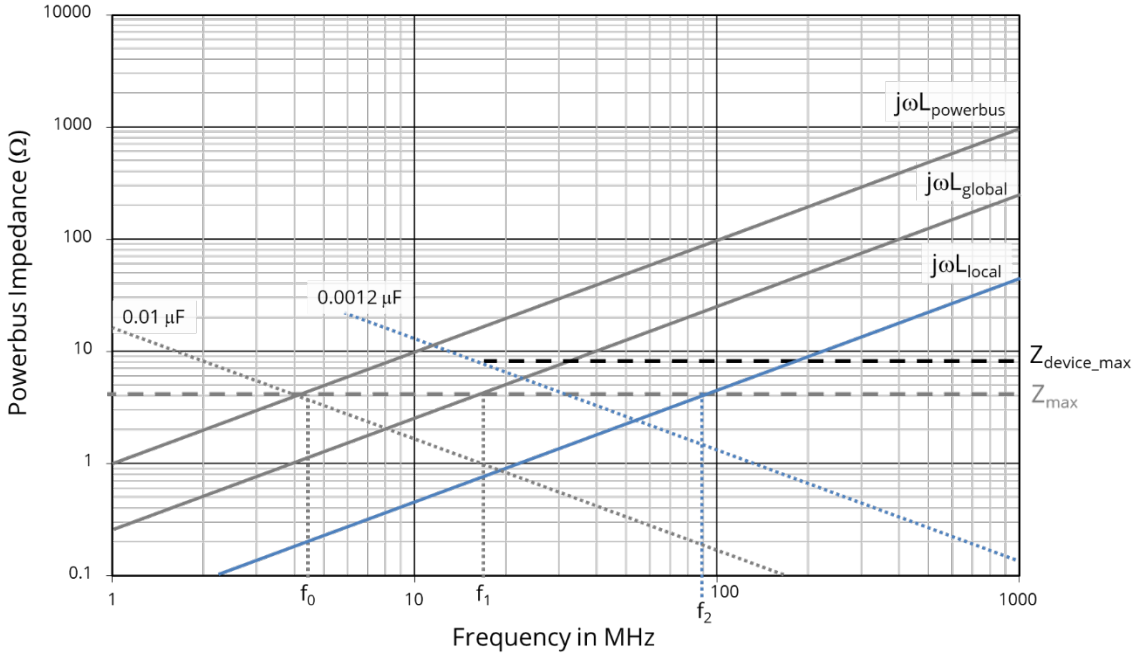

Figure 11.17. Power Bus impedance with one global and one local decoupling capacitor.

Decoupling PCBs with No Power Planes

The goal of printed circuit board decoupling is to provide sufficient current to the active devices without experiencing an unacceptable variation in the voltage on the power distribution bus. 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 provide high-frequency decoupling to each device individually.

If the current requirements of an active device are well-understood, calculations like those shown in the previous sections can be made to determine exactly how much capacitance is required and how much connection inductance can be tolerated. However,

it's important to remember that the calculated capacitances are minimum values, and the calculated inductances are maximum values. The decoupling will work just as well or better if there is more than enough capacitance. It will also work just as well or better if the capacitors are connected through a lower inductance.

Given that we rarely know the precise current demands of an active device as a function of frequency, power bus decoupling tends to be an exercise in designing for the worst-case. We can estimate the worst-case current demands at any given frequency. We can calculate the minimum capacitance necessary to provide that current at any given frequency. And, we can calculate the maximum inductance that would allow that much current to flow for a given voltage variation at that frequency.

Ultimately, circuit board decoupling comes down to getting enough capacitance connected to the device through a sufficiently low inductance. If the power distribution is routed on traces, high-frequency decoupling will be provided by a local capacitor mounted near the device in a manner that minimizes the inductance of the loop formed by the power inputs and the capacitor.

Figure 11.18 shows various examples of good local decoupling capacitor connections to boards without power planes. The three examples on the left of the figure illustrate cases where both the power and power-return are routed on traces. If there is a power-return (e.g., GND) plane, then the lowest-inductance connection will always be made by connecting both the active device and the decoupling capacitor directly to the plane rather than using a power-return trace. This is illustrated in the two examples on the right side of the figure. Note that the existence of the power-return plane allows a much lower inductance connection to be made.

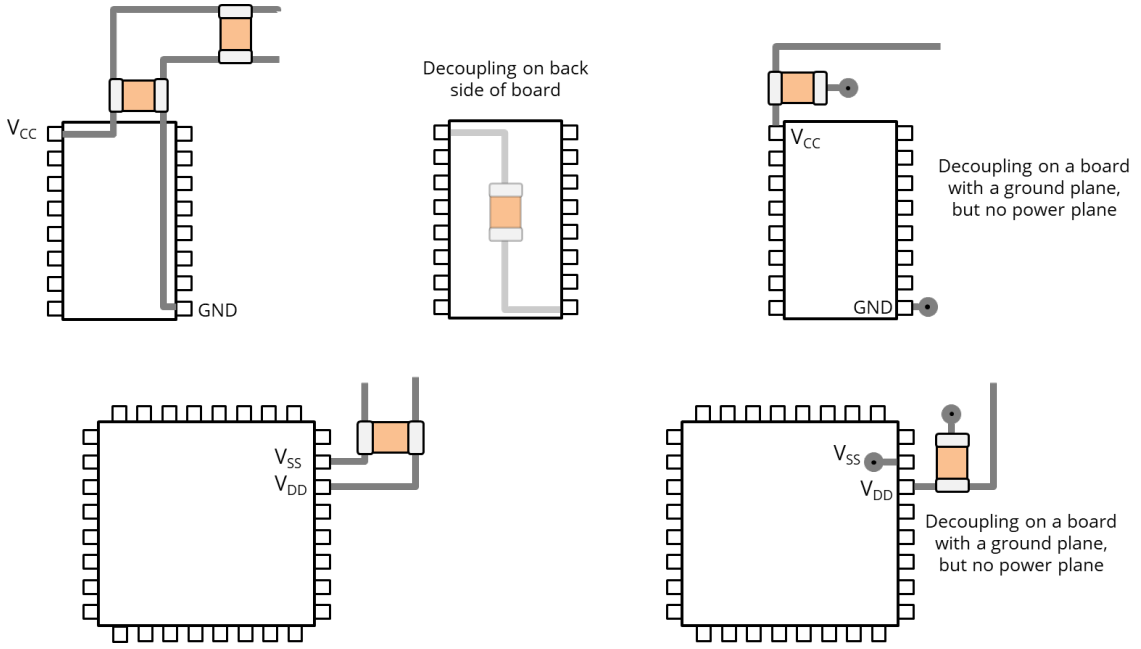

Figure 11.18. Examples of decoupling capacitor connections on a board without a power plane.

General guidelines for locating and connecting decoupling capacitors on boards without a power plane include:

- Provide at least one local decoupling capacitor for each active device and at least one larger global decoupling capacitor for each voltage distributed on the board.
- Local decoupling capacitors should be connected between the power and power-return pins of the active device. The inductance of the loop formed by the capacitor/device connection should be minimized.
- Local decoupling capacitors typically have nominal values from 0.01 µF to 4.7 µF depending on package size and maximum voltage requirements. It's fine to use values larger than necessary if the capacitors meet all other requirements.
- Global decoupling capacitors are normally located near the point where the power comes on to the board. If the voltage is generated on the board, the global decoupling should usually be near the location where it is generated.
- Global decoupling capacitors should be sized to meet the low-frequency or transient current needs of the entire board. Typically, on boards without power planes, global decoupling capacitors have values equal to 1 - 10 times the sum of the values of the local decoupling capacitors connected to the same bus.

Sometimes it is advantageous to connect two local decoupling capacitors to the same power input. This might be done to get a larger overall capacitance, or to provide better filtering to the rest of the power bus. If two local decoupling capacitors are connected to the same power input, they should generally have the same nominal value. Otherwise, they may exhibit a strong parallel resonance in the middle of the frequency range where they are designed to operate. To illustrate this, consider the impedance of the two capacitors in Figure 11.19. A 0.01-µF capacitor and a 10-µF capacitor are used to decouple an active device. Each has 5 nH of connection inductance. The transfer impedance (i.e., the ratio of the power bus voltage to the current drawn by the active device) exhibits a peak at about 22 MHz. This is the parallel resonance between the inductance of the larger-valued capacitor and the capacitance of the smaller-valued capacitor. In this circuit, a pulse of current drawn by the active device would cause the voltage on the power bus to ring at this frequency.

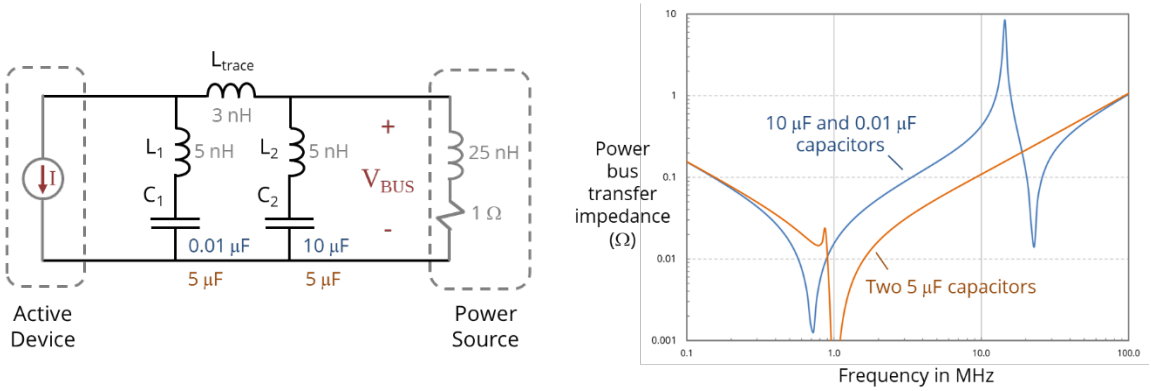

Figure 11.19. Impedance of two capacitors with unequal values compared to two capacitors with the same value.

If more than one local decoupling capacitor is employed, they should have similar nominal capacitance values. The orange curve in Figure 11.19 shows the transfer impedance for the

same circuit when the two capacitors both have a nominal value of 5 μF. In this case, their self-resonances are at nearly the same frequency and the parallel resonance between them is significantly damped.

Note that below about 300 kHz, the transfer impedance of the circuits in Figure 11.19 is the same. At low frequencies, the total value of the connected capacitance determines the response. Above 50 MHz, the both circuits exhibit the same transfer impedance. At high frequencies (above the self-resonant frequency of both capacitors), the connection inductance dominates, and the nominal capacitance values are no longer relevant.

Decoupling PCBs with Closely-Spaced Power/Power-Return Planes

There is a limit to how low the connection inductance can be when power is distributed to components on traces (generally several nanohenries or more). Many of the active devices found on circuit boards today can draw amperes of current at relatively high frequencies. These devices require very low power bus impedances that can only be achieved by distributing the power on closely-spaced circuit board planes. Closely-spaced planes not only permit low-inductance connections; they also provide their own inter-plane capacitance, which can provide excellent decoupling at very high frequencies.

Figure 11.20 illustrates the cross-section of a board with power planes containing an active device and three decoupling capacitors. Note that for any given connection between the active device and a capacitor, the loop area between the planes is small relative to the area associated with the connections to the planes. Also, the planes are wide relative to the connecting via diameters, so the magnetic flux wrapping the planes is weak. The value of the connection inductance is dominated by the connection between the active device and the planes and the connections between the capacitors and the planes. In this example, the capacitor on the far left is connected through vias that are near each other and has the lowest connection inductance to the planes. This capacitor will provide more decoupling current to the active device than the other capacitors even though it is located farther away.

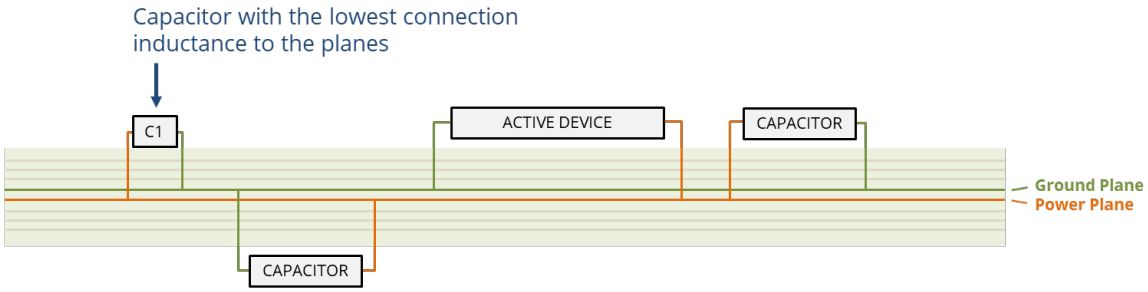

Figure 11.20. Cross-section of a multi-layer PCB with an active device and 3 decoupling capacitors.

Figure 11.21 shows a lumped-element model for the power distribution impedance on a board with closely-spaced power and ground planes. If the planes are very close (e.g., less than 0.5 mm), the plane inductance is negligible. And at frequencies where the maximum dimensions of the board are much less than a wavelength, the plane impedance can be modeled with a single capacitor. Capacitors mounted on the surface of the board have a connection inductance as indicated in the figure. While this model is only valid at frequencies well below the first plane resonance, it generally serves as an excellent model at all frequencies where the decoupling capacitors are effective.

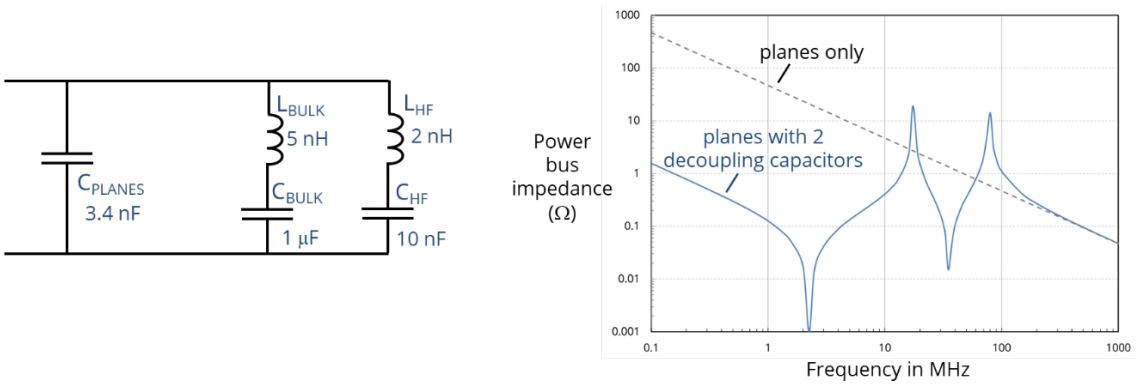

Figure 11.21. Lumped-element model and power bus impedance for power distribution on a board with closely spaced power and ground planes.

The circuit in Figure 11.21 models two capacitors connected to a pair of closely spaced planes. One of them, C_{BULK}, is a large capacitor with a high nominal value intended to provide the capacitance needed for low-frequency decoupling. The other, C_{HF}, has a small package size intended to facilitate the low-inductance connection to the planes needed for high-frequency decoupling.

At low frequencies, the inductances can be neglected, and the impedance of the power bus is approximately equal to,

$$Z_{powerbus} = 1 \Big/ j\omega(C_{BULK} + C_{HF} + C_{PLANES}) \quad . \tag{11.7}$$

All the capacitors on the board help to decouple the power bus, although the larger-valued capacitors are the most effective. At higher frequencies, the larger inductances become important. For example, the bulk capacitor represented in Figure 11.21 resonates with its own connection inductance at,

$$f_{self-resonance} = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{(5\,\text{nH})(1.0\,\mu\text{F})}} = 2.25\ \text{MHz}. \tag{11.8}$$

The power bus impedance is very low near the self-resonant frequency, but above that the bulk capacitor becomes an inductance, and impedance rises with increasing frequency. Eventually, the inductance of the bulk capacitor exhibits a parallel resonance with the remaining capacitors. At this parallel resonance, the impedance of the power bus peaks. As the frequency continues to increase, the smaller-valued capacitor with smaller connection inductance eventually goes self-resonant producing another null in the power bus impedance. But that null will be followed by another peak when both capacitors look like inductors and form a parallel resonance with the plane capacitance. At the parallel resonant peaks, the impedance of the power bus can be very high, and the board will tend to ring at these frequencies if there is not sufficient loss to dampen these resonances.

The decoupling bandwidth can be improved by adding more capacitors. Two identical capacitors in parallel have twice the capacitance and half the inductance of one capacitor. Ten times the number of capacitors will have ten times the capacitance and one tenth of the inductance. Figure 11.22 illustrates the impact of increasing the number of decoupling

capacitors. With 20 decoupling capacitors instead of 2, the impedance is reduced and the frequency range over which the capacitors are effective is extended.

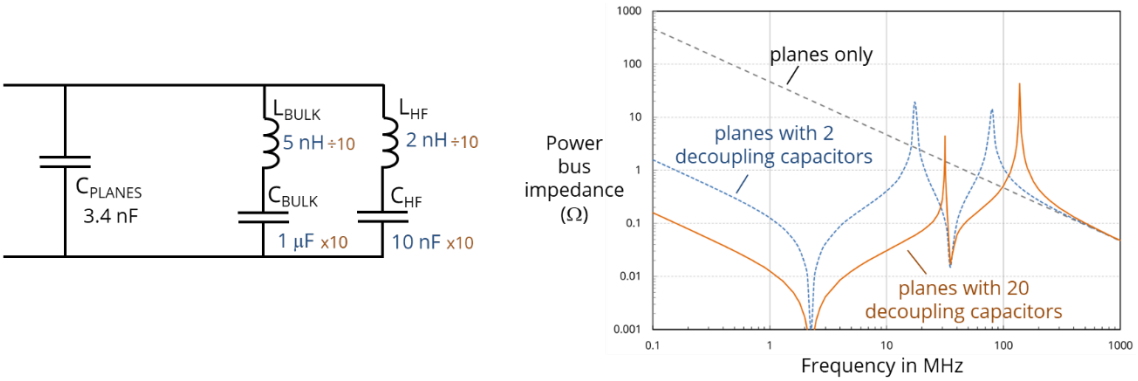

Figure 11.22. Impact that increasing the number of capacitors by a factor of 10 has on the power bus impedance of a board with closely spaced power and ground planes.

The peaks in the response shown in Figure 11.22 can be a concern for boards with relatively few components. Current is not easily drawn from the power bus at these frequencies, and a transient current draw could cause the bus voltage to oscillate. One way of preventing this is to choose capacitors with a range of values, so that their series resonances are spread across the entire frequency range of interest. For example, instead of using 10 high-frequency decoupling capacitors with a value of 10 nF, use 1 nF, 2 nF, 4.7 nF, 6.8 nF, 10 nF, 20 nF, 47 nF, 68 nF, 100 nF and 200 nF. This spaces the nulls in the response roughly equally across the log-frequency spectrum. The peaks occur between the nulls, so they are also spread out. Both the nulls and the peaks are more damped with this arrangement.

While spacing these resonances can be effective, it is not usually necessary or desirable if the board is heavily populated. The resistances in the power distribution due to the power draw of the active devices, the ESR of the decoupling capacitors, and other naturally occurring losses tend to provide enough damping so that these power bus resonances are not a significant concern. For example, Figure 11.23 shows the power bus impedance of the circuits in Figures 11.21 and 11.22 when the active devices provide 1 Ω of loading on the power bus. This would be the case, for example, if a 3.3-volt board drew 3.3 amps of DC current.

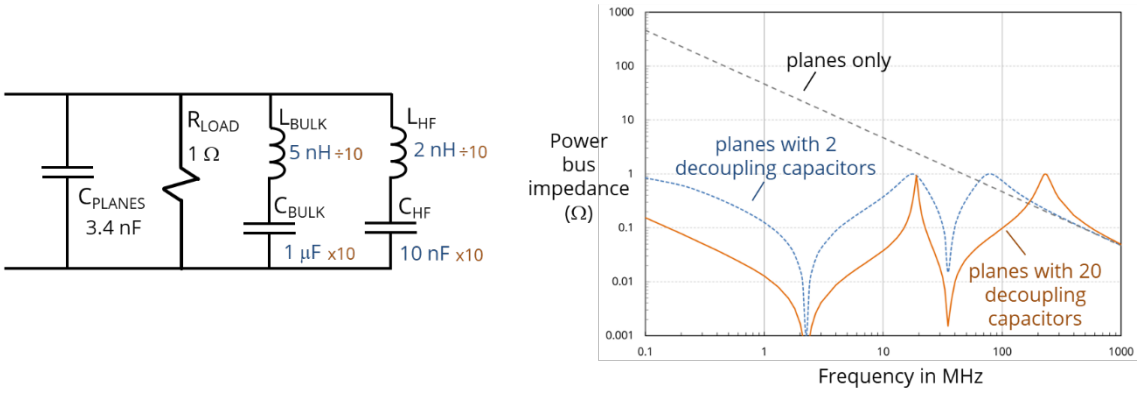

Figure 11.23. Impact of load resistance on the power bus impedance of a board with closely spaced power and ground planes.

Power Plane Resonances

Note that, for the example in Figure 11.23, the decoupling capacitors provided no benefit at frequencies above 200 MHz. When boards have closely spaced power and ground planes, the plane capacitance becomes a good source of high-frequency current. Surface mounted decoupling capacitors with 1-2 nH of connection inductance are unable to provide enough current to compete with the planes at hundreds of megahertz. Because of this, simple circuit models that neglect the plane inductance work well over the entire range of frequencies where the surface mounted decoupling capacitors are effective.

At higher frequencies, where the maximum dimensions of the plane pair are not small relative to a wavelength, the power planes can no longer be modeled as a simple capacitance. Planes with large surface areas and few connected components can behave like cavity resonators. At their resonant frequencies, they can radiate like patch antennas.

For rectangular planes, the first cavity resonance occurs when the largest dimension of the plane pair is a half wavelength (in the cavity dielectric). Various online calculators, such as the one shown in Figure 11.24, can be used to determine the frequencies associated with higher order modes. There are also numerical modeling tools available that can calculate the resonances associated with plane pairs of arbitrary shape.

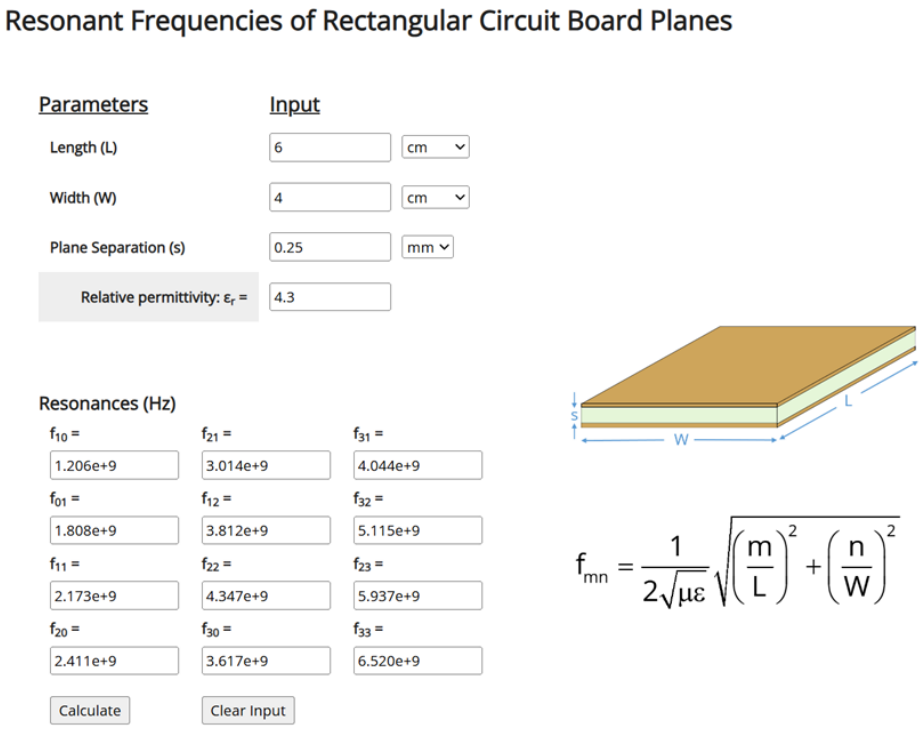

Figure 11.24. An example of an online tool for calculating the resonances of rectangular circuit board plane pairs on the LearnEMC website.

If the power plane pairs are very closely spaced (e.g., less than 0.1 mm apart), the copper loss in the planes damps the resonances significantly. Similarly, if the planes are not so closely spaced, but have a high-density of attached components, the component losses damp the resonances. Because of this, power plane resonances are rarely a significant issue

on heavily populated circuit boards. On the other hand, plane resonances should be anticipated when large areas of the plane pair extend into unpopulated regions of the board.

Decoupling PCBs with Widely Spaced Power Planes

Often, printed circuit boards will have solid power and ground planes, but the spacing between these planes will be 1 mm or greater. For these boards, the inductance associated with the planes can no longer be neglected and the model in Figure 11.21 does not apply. The plane inductance inhibits the ability of the planes to provide charge quickly. On the other hand, by careful placement of the decoupling capacitors, this inductance can be used to help pull current from the decoupling capacitors while reducing the power bus noise on the planes.

Figure 11.25 helps to illustrate how this works. In this example, an active device and a decoupling capacitor are located on the top side of the board. The inductance between the planes is dominated by the magnetic flux that wraps the vias connected to the lower plane (in this case the power plane). The flux wrapping each via adds in the region between the two vias. However, the flux wrapping the vias cancels outside this region. The flux between the vias represents an inductance that is impeding the ability of the decoupling capacitor to supply current. The flux outside this region represents a mutual inductance that helps to pull charge from the decoupling capacitor.

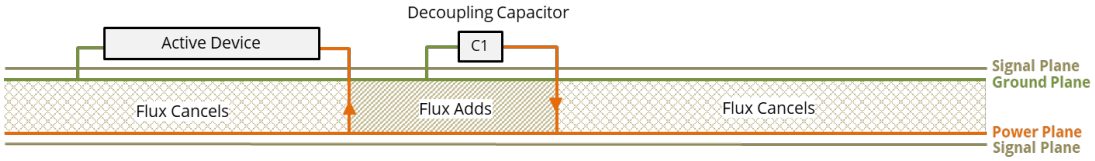

Figure 11.25. Current flow in a board with >1-mm plane spacing.

If the via from the active device is located near the via from the decoupling capacitor, the mutual inductance is stronger. If the mutual inductance due to the loop area between the planes is greater than the self inductance associated with the loop formed by the decoupling capacitor connection outside the planes, then more charge will be pulled from the decoupling capacitor than from the planes.

For widely spaced planes, when the two vias are very close to each other, the flux cancelation between the planes is high. This mutual inductance tends to force current to be drawn from the capacitor instead of the planes. This is exactly what we want if the goal is to keep noise currents off the planes.

When the active device and the local decoupling capacitor are on the same side of the board, the mutual inductance is maximized by sharing a via connection to the most distant plane as illustrated in Figure 11.26a. On the other hand, if the decoupling capacitor and active device are on opposite sides of the board, the vias pulling current between the planes are not connected to the same plane. This case is illustrated in Figure 11.26b where the capacitor's power via needs to be located near the active devices power-return (or GND) via.

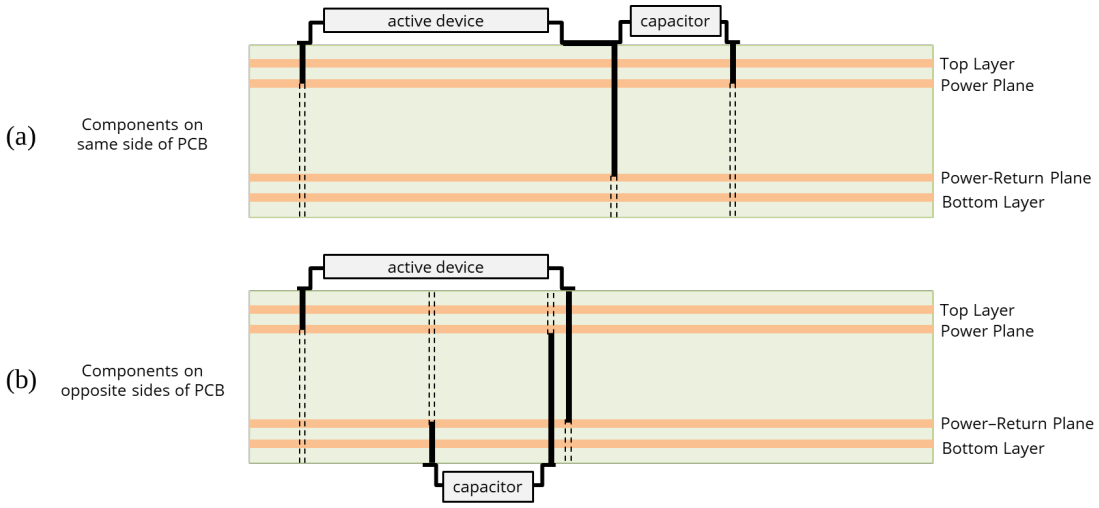

Figure 11.26. Locating capacitors to maximize mutual inductance in a board with >1-mm plane spacing.

If the mutual inductance between the planes is greater than the capacitor's connection inductance above the planes, most of the current supplied to the active device will be pulled from the capacitor. When implemented correctly, this type of connection maximizes the current available to the active device and minimizes the voltage induced in the power planes.

Connecting Capacitors to Boards with Power Planes

Whether the planes are closely-spaced or widely-spaced, the connection inductance between the decoupling capacitors and the planes determines how well they will work at high frequencies. Minimizing this inductance is essential. On boards with power planes, the main parameter determining the connection inductance is the loop area illustrated in the top half of Figure 11.27. This area is bounded by the path of the current coming out of the planes flowing through the capacitor and back to the planes.

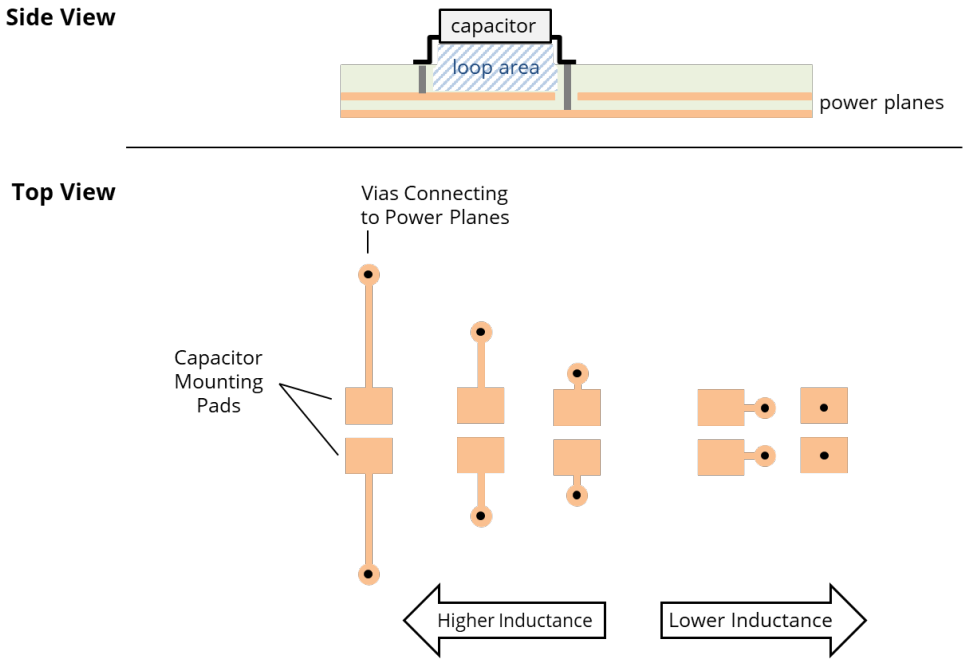

Figure 11.27. Locating the vias connecting to the planes close to each other reduces the connection inductance.

The distance from the power planes to the capacitor is a critical factor in determining this loop area. However, the distance between the connecting vias is equally important. Viewed from above, it's essential that the power and power-return vias be located as close to each other as possible. Typically, this means choosing an SMT capacitor with the smallest possible package size and locating the connecting vias in or adjacent to the mounting pads as indicated in bottom right side of Figure 11.27.

Note that vias located at opposite ends of the pads can have twice the distance between them as vias located to the side of the pads. This results in approximately twice the inductance and, for global decoupling capacitors, twice as many capacitors would be required to achieve the same high-frequency performance.

In most situations, small SMT capacitors (e.g., 0402) with vias in or to the side of the pads represent the most cost-effective and space-effective decoupling. In certain situations, particularly when the plane area is very limited, the use of special low-inductance decoupling capacitor packages may be warranted. Low-inductance package options include reverse-aspect-ratio capacitors (e.g., 0204) that are wider than they are long. These capacitors allow two via connections to each pad and minimize the distance between power and power-return vias. Another option is capacitors that have four or more terminals that facilitate interspersed power and power-return connections.

Power Circuit Layout

Many circuit boards employ switching power circuits to convert a supplied power voltage to one or more different voltages required by the active components on the board. The high-frequency switching in these circuits creates noise that may show up in conducted or radiated emissions tests. The goal of the board designer is to contain this switching noise

and not allow it to couple to traces or components that might ultimately carry this noise off the board.

Figure 11.28 illustrates how different power converters might be connected to produce the different voltages required for circuits on the board. There are many power-converter design options that employ different circuit topologies, different switching components, and different switching frequencies and modulations. The proper layout for reducing conducted and radiated emissions largely depends on the inverter topology. However, the layout is also influenced by requirements related to other factors such as heat dissipation, current carrying capacity, proximity to other circuits, and galvanic isolation requirements.

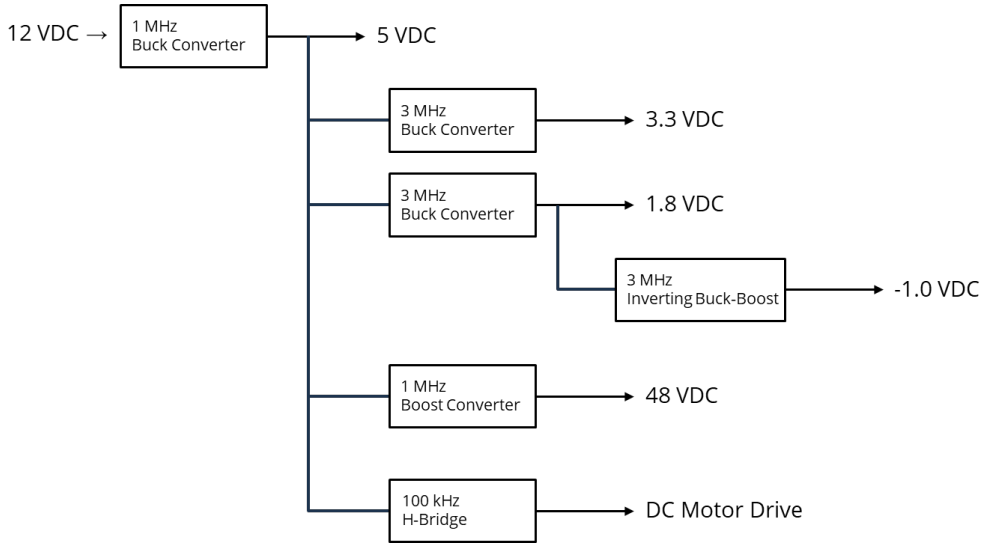

Figure 11.28. Six voltages derived from one 12-volt supply.

Since the layout depends on the converter topology (i.e., the fundamental components and how they are connected), it's important to recognize which topology a converter employs before the components are placed. There are many different converter topologies, and each has its own layout priorities for meeting conducted and radiated emission requirements. In this section, we will start by looking at the four basic topologies (buck, boost, H-bridge, and flyback) shown in Figure 11.29. Once we understand how to lay out these common converters, we will have the basic tools we need to make good layout decisions for virtually any other switching converter topology.

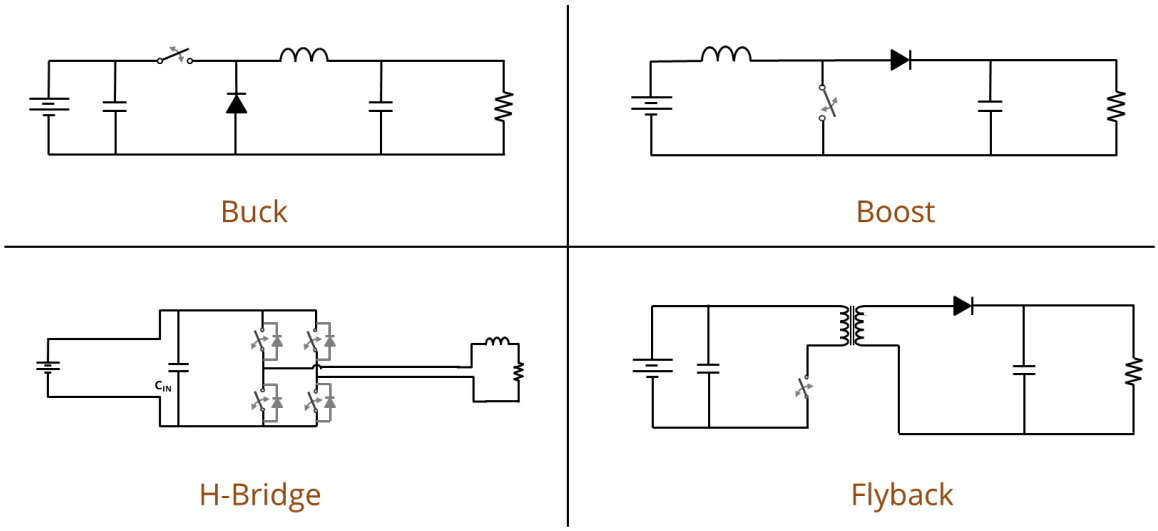

Figure 11.29. Four common power circuit topologies.

Buck Converters

Buck converters are relatively simple and employ low-cost components. They are widely used in circuit boards to efficiently convert a DC voltage down to a lower one. They can also be used in situations where a constant-current source is required. The basic topology is illustrated in Figure 11.30.

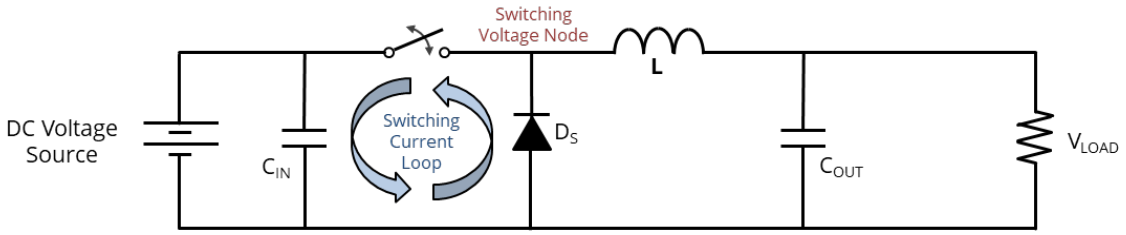

Figure 11.30. DC-to-DC buck converter topology.

The cycle starts when the switch closes, allowing current to flow from the voltage source to the load. The voltage across the load rises until the desired voltage (or current) is attained and the switch opens. With the switch open, the inductor keeps the current flowing through the load and returning to the inductor through the diode. However, as the energy stored in the inductor is depleted, the voltage across the load falls. Before it falls too far, the switch closes again, and the cycle is repeated.

An important component that is not shown in Figure 11.30 is a microcontroller that monitors the voltage across the load and operates the switch. Typically, the switch turns on at a set time interval and the turn-off time is varied to control the amount of power delivered to the load. So, the switching voltage and current are pulse-width-modulated (PWM) signals with a specific fundamental frequency. Today's low to medium power DC-to-DC converters typically switch at frequencies between 1 and 3 MHz, with PWM modulations designed to range from about 10% to 90% of the maximum pulse width.

The design shown (with the diode) represents an *asynchronous* buck converter. In *synchronous* buck converters, the diode is replaced with another switch. The optimum

layout for a buck converter doesn't depend on whether it is synchronous or asynchronous, and neither has an inherent advantage in terms of its conducted or radiated emissions.

The key to a good DC-to-DC converter layout is recognizing that today's converters are typically way too small to radiate significant amounts of power directly. That means that to cause a conducted or radiated emissions failure, noise from the converter MUST couple to other nearby circuits or objects. The first priority for a successful layout is to identify the sources of the switching noise and prevent them from coupling to anything that can carry that noise away from the immediate area.

There are three ways for EM noise to couple from one place to another when the source is too small to radiate. Noise can be conducted, E-field coupled, or H-field coupled. We typically deal with conducted noise by providing filtering. We want to allow DC current to flow into and out of the converter, but we can filter everything else. Filter design and power bus decoupling were both discussed in other sections of this book. Beyond providing adequate filtering, the main focus of our converter layout is controlling the field coupling.

Electric-field coupling is proportional to the voltage rate of change (dV/dt). Referring to the circuit in Figure 11.30, we see that there is only one circuit node that exhibits a relatively uncontrolled dV/dt relative to the reference node at the bottom. In the figure, that node is labeled "Switching Voltage Node." All buck converters have a switching voltage node, and that node is the primary source of electric-field coupling from the circuit. Since electric-field coupling falls off significantly with distance at distances greater than the size of the source, it's very important to minimize the size (surface area) of this node (trace) on the board. It's also important to keep this trace away from any other traces that could carry the coupled noise to other parts of the board or system.

In most cases, buck converters should be laid out on the top (or bottom) surface of the board with a solid return plane on the layer directly below (or above) it. Fields that couple to the return plane are not a problem, because they immediately return the current to the source. Fields that couple to anything else potentially provide unwanted coupling of noise to the external world. The basic guideline to apply to the switching voltage node is *minimize the surface area and keep it away from any conductor that could carry the coupled current away from the converter footprint*.

Magnetic-field coupling can be just as important and is often overlooked. Magnetic-field coupling is proportional to the current rate of change (dI/dt). Referring to the circuit in Figure 11.30, there is only one circuit loop that carries a virtually uncontrolled dI/dt. That is the loop consisting of C_{IN}, the high-side switch, and the low-side diode or switch. In many converters, the switches, diodes, and the controller are in the same IC package. In this case, the priority is reducing the loop area of the connection between the IC and C_{IN}. C_{IN} should be located as close as possible to the voltage input of the IC. It should be connected to the pin on one side and the return plane on the other side. The loop area must be minimized and kept away from other circuit loops that might carry a coupled voltage to other parts of the board or system.

Consider the two synchronous buck converter layouts illustrated in Figure 11.31 The converter IC contains the high and low-side switches, so the relevant I/O pins are VIN, SW, and RTN (or GND). SW is the switching voltage node connected to the external inductor. The components are all on the top side of the board and the layer right below

them is a solid RTN plane. To review this layout, we can start by evaluating the coupling from the switching voltage node. The surface area of the SW trace is reasonably small in both layouts (slightly larger than the endcap of the inductor). When we look to see what it might couple its electric field to, we see that the SW trace in the left layout comes very close to VOUT and the C_{OUT} capacitor. Capacitive coupling from SW to VOUT could potentially carry noise away from the converter's footprint. Another problem with the left layout is that the C_{IN} capacitor is mounted right over the top of the SW trace. Coupling from SW to this component can carry noise to VIN. The left layout does a poor job of containing the electric-field coupling from the voltage switching node. The right layout, which uses the same SW trace as the left layout, does not provide a good opportunity for the switching noise to escape the immediate area.

It's also important to evaluate potential magnetic-field coupling from the switching current loop. In the layout on the right, that loop is very small. The switching current flows out of the IC on VIN through C_{IN} to the RTN via. Once the current is on the RTN plane, it returns under the body of C_{IN} to the GND pin of the IC. The layout on the left has a similarly small switching current loop, however the RTN via of C_{IN} is shared with C_{OUT}. This allows the magnetic flux in the switching current loop to couple a voltage through the C_{OUT} capacitor directly to VOUT.

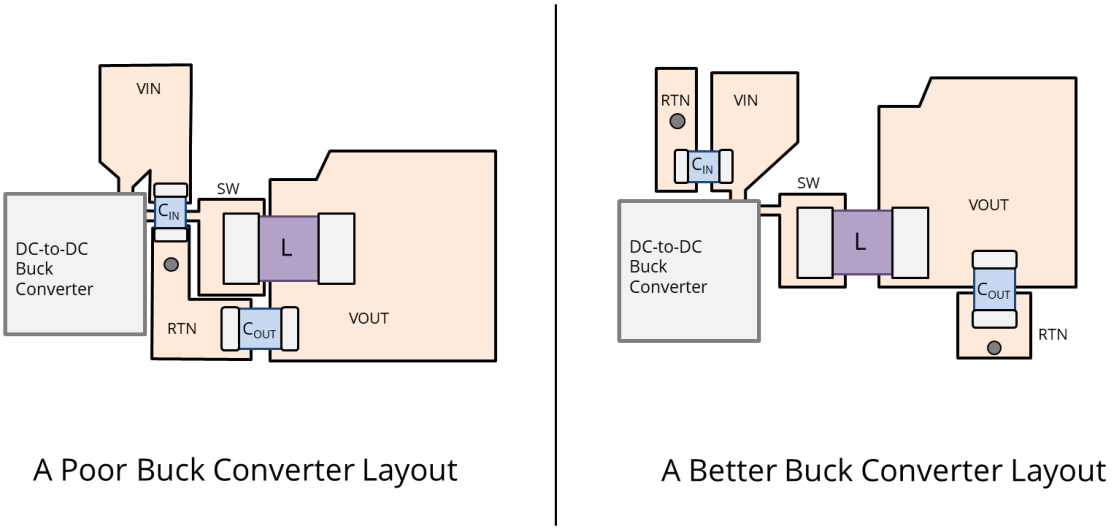

Figure 11.31. Two possible board layouts for a buck converter.

You may be thinking that the layout on the left is so obviously bad that nobody would ever attempt something like this. But similar layouts appear in many vendor application notes and datasheets. In general, you shouldn't rely on EMC layout recommendations from power converter vendors. Many of them are based on rules that were developed many years ago when switching frequencies were lower and components were much larger.

In most cases, simply controlling the surface area of the switching voltage node and the loop area of the switching current loop is sufficient to keep high-frequency switching currents from spreading to other parts of the board or system. However, there are additional design options that can be useful in situations where the layout alone doesn't limit the coupling sufficiently. For example, most converter vendors offer low-noise options such as spread spectrum clocking and building the C_{IN} capacitor into the IC. Low-noise

converters may also offer a more thoughtful placement of their I/O pins to facilitate an optimum layout.

Boost Converters

Boost converters are used in situations where it is necessary for the output DC voltage to be higher than the input DC voltage. The basic boost converter topology is illustrated in Figure 11.32. Like the buck converter, the boost converter employs a switch and a diode (or two switches for a synchronous boost converter). The inductor is on the input side of the converter. Closing the switch in Figure 11.32 causes current to flow through the inductor. Once a sufficient current has been reached, the switch opens causing the voltage on the switching node to jump to the value of the input voltage plus L dI/dt. With the switch open this voltage starts to drop as current flows to the load. Closing the switch sends the voltage on the switching node to zero, but the diode and capacitor prevent the voltage across the load from dropping rapidly. Meanwhile, with the switch closed the inductor builds up stored energy and the cycle is repeated.

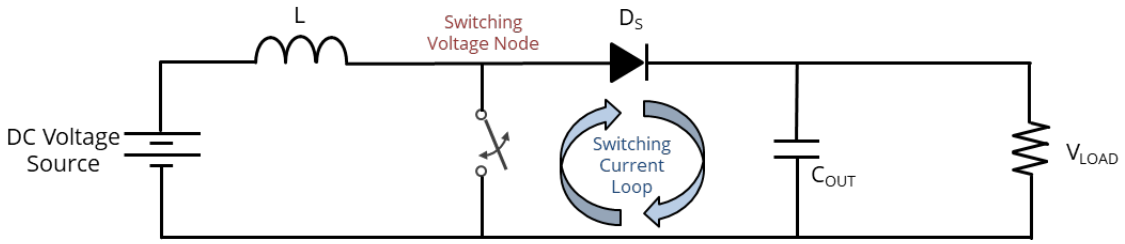

Figure 11.32. DC-to-DC boost converter topology.

Like buck converters, boost converters usually operate at a fixed frequency and use pulse-width modulation to control voltage at the output. Like buck converters, they have a switching voltage node and a switching current loop. The key to a good boost converter layout is controlling the electric-field coupling from the switching voltage node and the magnetic-field coupling from the switching current loop.

Figure 11.33 shows two possible boost converter layouts. The layout on the right resembles the buck converter layout in Figure 11.31 with VIN and VOUT swapped. The switch node is reasonably small and isolated. The switching current, which flows from the IC through C_{OUT} and back on the RTN plane, is a relatively small, isolated loop.

The layout on the left uses a PGND trace to carry the switching current back to the IC, where it makes a single-point connection to IC GND pin and the RTN plane. Isolated power current returns like this are reminiscent of the days when power converters switched at kHz frequencies. Back then, the switching currents were a source of common-impedance coupling to other circuits sharing the same “ground.” With switching frequencies now routinely in excess of 100 kHz, common-impedance coupling is no longer a problem on boards with a solid return plane. The high-frequency switching current stays largely confined to the small region of the plane directly below the components. It minimizes its own loop area. Providing an isolated PGND plane as shown in the figure is not only unnecessary, it also increases the inductance of the loop.

Another mistake in the layout on the left is that the C_{IN} capacitor connects to the PGND. This creates a significant mutual inductance between the switching current loop and VIN.

With this layout, switching noise will be strongly coupled to the input voltage and eventually back to the board's power supply.

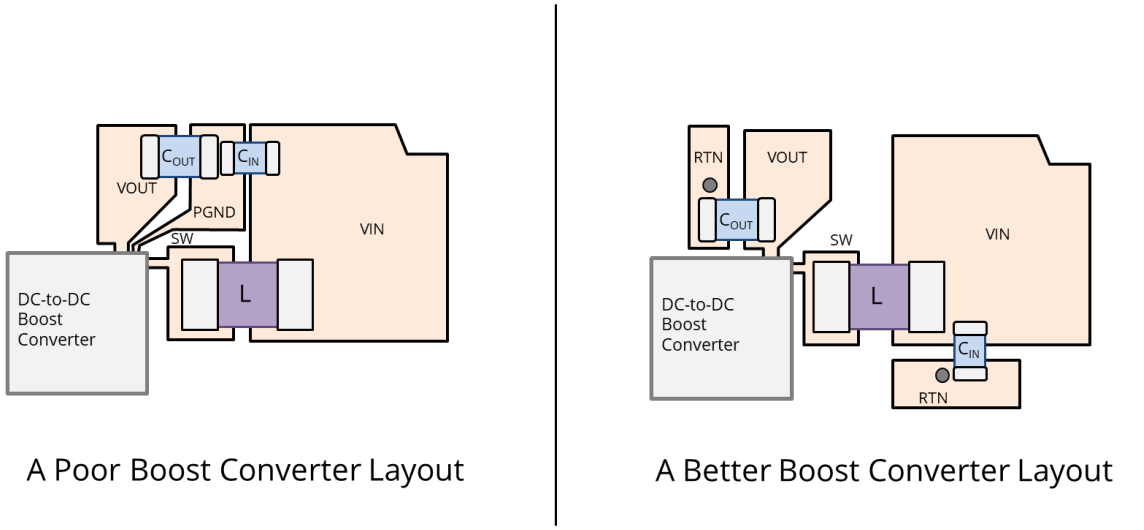

Figure 11.33. Two possible board layouts for a boost converter.

The layout on the left appears to be designed to maximize conducted and radiated emissions. It is a terrible layout. Nevertheless, it is based on design advice that can be found in a number of application notes and datasheets.

H-Bridge Inverters

An H-bridge can produce voltages with a positive or negative polarity. They are often used to drive DC motors both forward and backwards. They are also capable of taking a DC input voltage and producing an AC output, so they can be used to drive AC motors as well.

An H-bridge inverter employs four switches as illustrated in Figure 11.34. When the switches on the upper-left and lower-right are closed, current flows to the load on the upper wire and returns on the lower wire. The amount of current can be regulated by pulse width modulation. Typically, the upper switch is modulated while the lower switch remains closed.

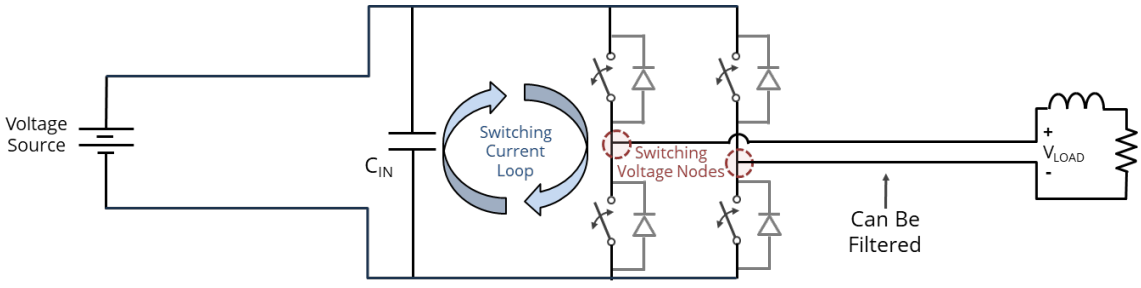

Figure 11.34. H-bridge inverter topology.

To achieve a negative voltage across the load, the switches on the upper-right and lower-left are closed. Now current flows to the load on the lower wire and returns on the upper wire. This current can also be regulated using pulse width modulation. Figure 11.35 illustrates how pulse width modulation of the positive and negative voltages can be used to produce a sinusoidal current flowing to the load.

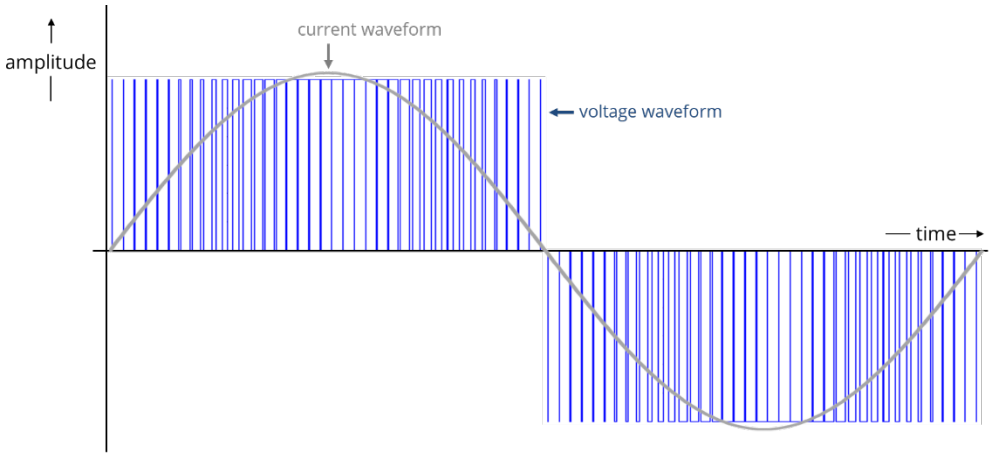

Figure 11.35. PWM representation of a sine wave.

An H-bridge converter has two switching voltage nodes, the connections between the upper and lower switches. Unlike buck and boost converters, the switching voltage nodes in an H-bridge converter carry current away from the inverter. We can't limit the noise by simply keeping their area small and far from other circuits. However, since the intentional current flowing to the load is relatively low-frequency, it is easy to filter the conductors connected to these nodes. In many cases, it is also possible to control the transition time of the switches to limit the higher harmonic content of the switched power. And in situations where filtering is not practical, shielding of the inverter, cable and load can be used to limit the electric-field coupling.

There are effectively several switching current loops in an H-bridge. All the switches have a capacitance when they are open, so high-frequency current can flow from the high-side voltage through any pair of high-side and low-side switches to the low-side voltage, then return to the high-side voltage through any other pair of switches or through the input capacitor (C_{IN} in Figure 11.34). To reduce magnetic-field coupling from the inverter circuitry, it is important to keep these loop areas small and isolated from other circuits that could carry the coupled voltage away from the inverter.

Figure 11.36 illustrates two possible board layouts for an H-bridge motor driver. This design uses switches that require mounting along the board edge to facilitate thermal heatsinking. Both layouts include a filter on the motor output. Both layouts keep the motor-driving circuitry away from the digital control circuitry. However, the layout on the left has connectors on opposite sides of the board. This allows the voltage that develops across the return planes to drive one cable relative to the other. To make matters worse, the reference for the cable connecting to the motor is different from the reference for the cable on the right because the layout uses an isolated motor "ground."

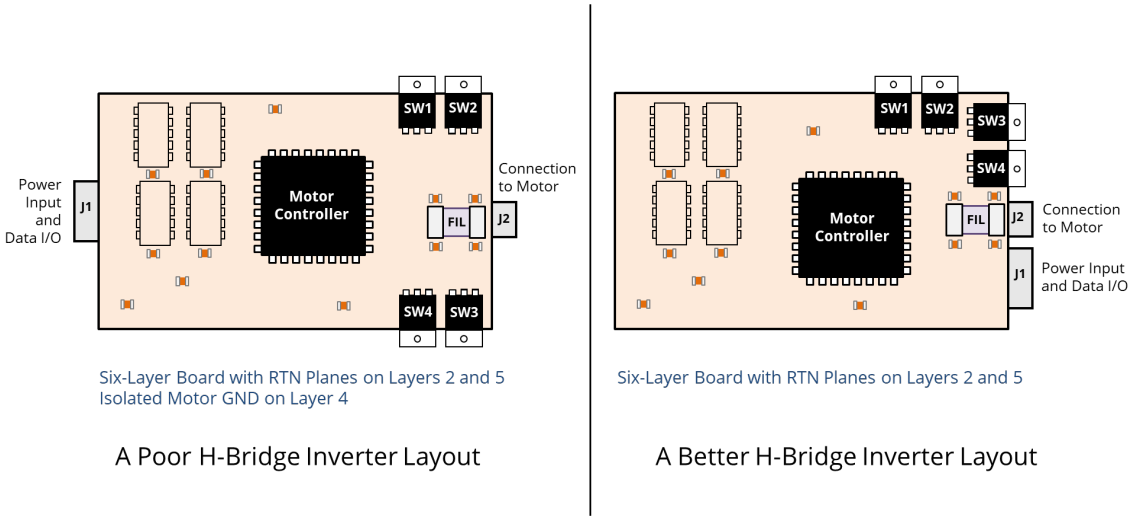

Figure 11.36. Two possible board layouts for an H-bridge inverter.

The layout on the right puts both I/O connectors near each other, while still maintaining reasonable isolation between the digital circuits and the motor-driving circuits. It does not have an isolated motor ground, which would rarely be required in applications like this. Since this board is mounted to a large heatsinking structure, the board RTN planes should be connected to that structure along the edge with the external connections.

Flyback Converters

Flyback converters are typically used in situations where it is necessary to isolate the current return paths on each side. A common example of this is a power supply that plugs into a wall outlet. In this case, the input voltage is high enough that none of the power current is allowed to flow on the chassis. On the other hand, it is usually advantageous to connect the current return on the low-voltage side to the chassis to meet EMC requirements.

The basic topology of a flyback converter with its secondary grounded is illustrated in Figure 11.37(a). Galvanic isolation between the input and the output is provided through a transformer. The primary side of the transformer is connected to the voltage source through a switch. As the primary side current is switched on and off, a voltage appears on the secondary. This voltage can be higher or lower than the primary side voltage depending on the transformer's turns ratio. The power on the secondary side is rectified to provide a DC voltage across the load.

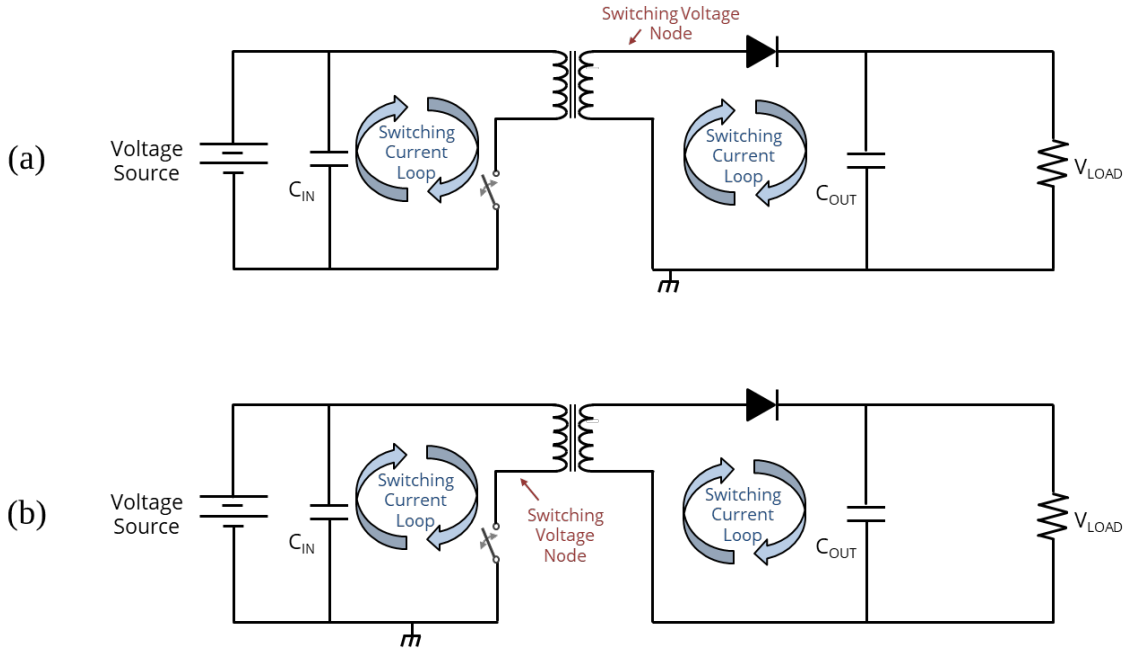

Figure 11.37. Flyback converter topology.

The topology in Figure 11.37(b) is identical to that in 11.37(a) except the primary side is grounded and the secondary side is not. This would be common in situations where a relatively low voltage was stepped up to a very high voltage. From a layout standpoint, one difference between the two topologies is the location of the switching voltage node. From an electric-field coupling standpoint, we are most concerned with voltages that rapidly change relative to the ground structure. So even though the switch is on the primary, the switching voltage node on the secondary in Figure 11.37(a) should not be overlooked.

The voltage across the switch in Figure 11.37(a) can also be a source of electric-field coupling, so it is important to pay attention to the layout on the floating (primary) side of the converter as well. In most cases, since neither the power nor the power return conductors are ground, it is best to have them mirror each other with neither conductor being substantially larger than the other. In other words, the power distribution on the floating side should be balanced.

In situations where the floating side is unbalanced (e.g., the power is routed on a trace and the power-return is routed on a plane), then the plane conductor should make a good high-frequency connection to the ground. In this case, both sides are considered to have the same high-frequency ground and the node on the other side of the switch is another switching voltage node.

The same reasoning applies to Figure 11.37(b), where the primary side is connected to ground. In this case, the secondary side should normally be balanced. However, when the secondary side is not balanced, its current-return should have a good high-frequency connection to the ground structure.

Flyback converters have two high-frequency switching current loops as indicated in Figure 11.37. When laying out flyback converters, it is important to minimize the area of these loops to prevent magnetic-field coupling to other circuits.

Figure 11.38 illustrates an example of a good flyback converter layout. In this example, power from a wall outlet is converted to a low DC voltage that powers the rest of the board. There is a solid return plane on Layer 2. The primary power does not connect directly to the plane, but it is filtered using X and Y capacitors as well as a common-mode choke. The input switching current loop (connector to switch to transformer and back to the connector) is kept reasonably small and away from other circuits. The output current loop (transformer to diode to capacitor to ground and back) is also minimized and doesn't couple readily to any other circuit. The switching voltage node (the connection between the transformer and diode) is very small.

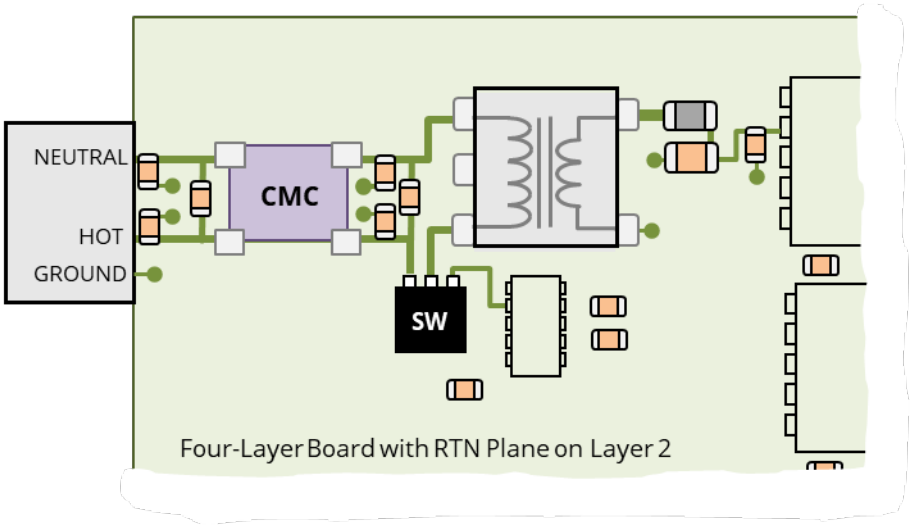

Figure 11.38. A good flyback converter layout.

Other applications may call for very different layouts, but the emphasis should always be to minimize coupling from the switching voltage nodes and switching current loops that could carry noise away from the immediate area. The layout in Figure 11.38 has all the components on the top side of the board. In general, both sides of the board can be utilized in power converter designs as long as the switching current loops do not pull current through board. In this design, the transformer, the switch, C_{IN}, C_{OUT}, and the diode should all be on the same side. On the other hand, the filter components and control circuitry could be mounted on the top or the bottom. The location of these components should be chosen to optimize their performance while minimizing any unwanted noise coupling.

Inverting Buck-Boost Converters

The four power circuit topologies described in the previous sections (buck, boost, H-bridge, and flyback) have the essential building blocks found in many other common topologies. Once you are comfortable identifying the switching voltage nodes and switching current loops in these power circuits, it is relatively straight forward to identify the same features in other power circuit topologies.

For example, consider the inverting buck-boost converter whose basic topology is illustrated in Figure 11.39. This circuit is commonly used to convert a positive DC input voltage to a negative voltage across the load. In order to identify the switching voltage node, we look for the connection to the switch that is shared with the inductor (just as it is in a buck or boost converter). To identify the switching current loop, we start with the

switch and find the lowest-impedance high-frequency path that completes the loop. The switching current path always involves at least one switch, but it never flows through an inductor. In Figure 11.39, we see that the switching current loop flows through the switch, diode, C_{OUT} and C_{IN}.

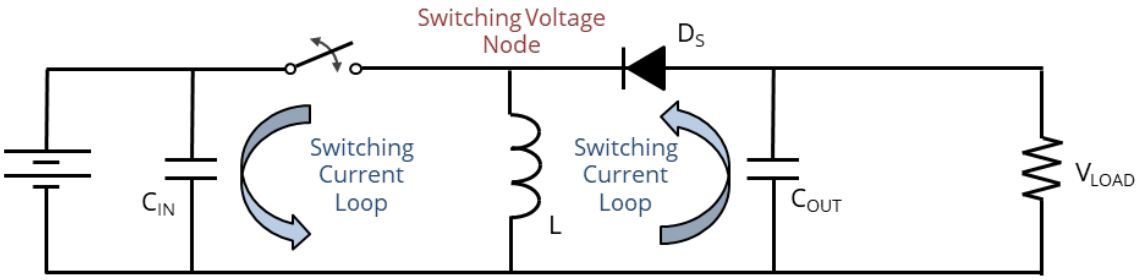

Figure 11.39. Inverting buck-boost converter topology.

Like other power circuit topologies, a good layout minimizes the electric-field coupling from the switching voltage node and the magnetic-field coupling from the switching current loop. Figure 11.40 shows two possible inverting buck-boost converter layouts. In this case, the switch and diode are located inside the IC with the control circuitry. As in the previous examples, Layer 2 is a solid return plane.

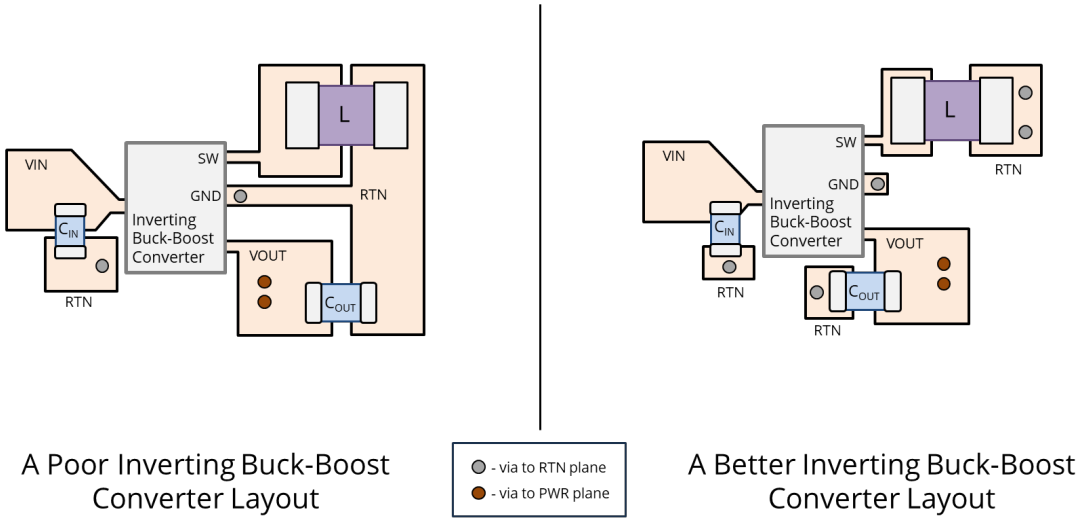

Figure 11.40. Two possible board layouts for an inverting buck-boost converter.

In the layout on the left side of Figure 11.40, the switching voltage node is larger than it needs to be. It couples to a return trace on the top layer that only makes a single-via connection to the plane on Layer 2. Depending on what is located nearby, this could allow electric field coupling to carry noise away from the converter.

A bigger problem with the layout on the left in Figure 11.40 is that the switching current loop is much too large. If we trace it from VIN through the IC to VOUT and through C_{OUT}, it must flow across the return trace on Layer 1 to the via, then to the plane to get to C_{IN} and back to VIN. This large loop facilitates magnetic-field coupling to nearby circuits.

The layout on the right of Figure 11.40 minimizes the surface area of the switching voltage node and minimizes the loop area of the switching current loop. Note that, for this

converter, the switching current between C_{IN} and C_{OUT} stays on the plane. It does not need to return to the GND pin of the IC. As a rule, for any of the converter layouts we've discussed, if there is a solid return (or GND) plane on Layer 2, return (or GND) traces are not required or helpful on Layer 1. All RTN or GND connections should be made directly to the plane.

Other Power Circuit Topologies

As previously indicated, other common power circuit topologies have the same elements found in the buck, boost, H-bridge, and flyback topologies. Figure 11.41 illustrates a SEPIC (Single-Ended Primary-Inductor Converter) topology, which is basically a non-inverting buck-boost converter. These converters can step a DC voltage up or down. This topology has two switching voltage nodes and one switching current loop.

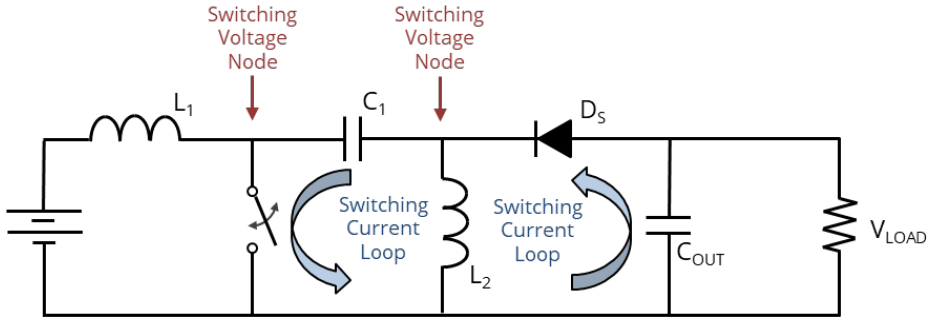

Figure 11.41. Inverting buck-boost converter topology.

Figure 11.42 shows the basic topology of an active-clamp forward converter. It utilizes elements found in the flyback and buck converters. Active-clamp forward converters can step DC voltages up or down and are often used in high-power applications. These converters are isolated like flyback converters and the location of the switching voltage nodes depends on which side(s) of the inverter is (are) connected to ground. The locations in the figure assume the lower node on both sides is ground. This topology has three switching current loop areas to minimize.

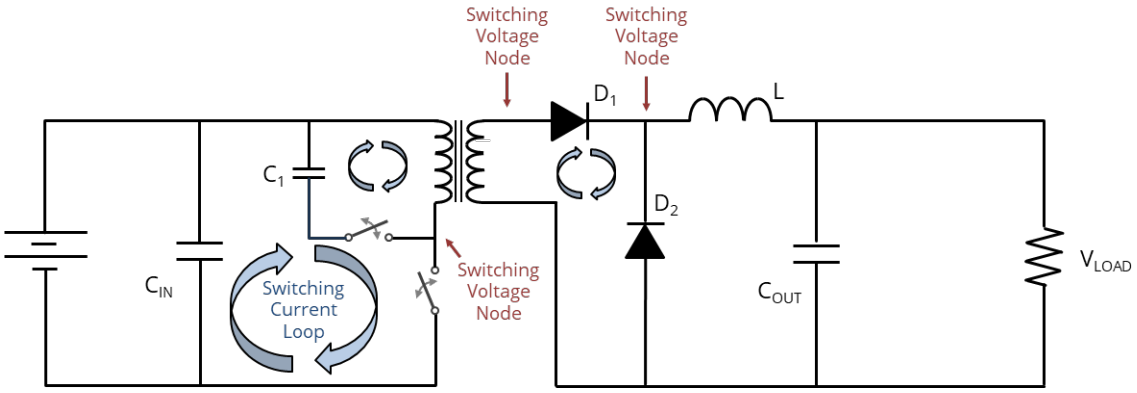

Figure 11.42. Active-clamp forward converter topology.

Figure 11.43 shows the basic topology of a three-phase inverter. It's basically an H-bridge inverter with an extra phase output. Three-phase inverters are widely used for driving medium-to-high-power variable-speed AC motors. Like an H-bridge inverter, each of the outputs is a switching voltage node and should generally be filtered or shielded if it extends

very far from the switches. There are also multiple switching current loops whose area should be minimized.

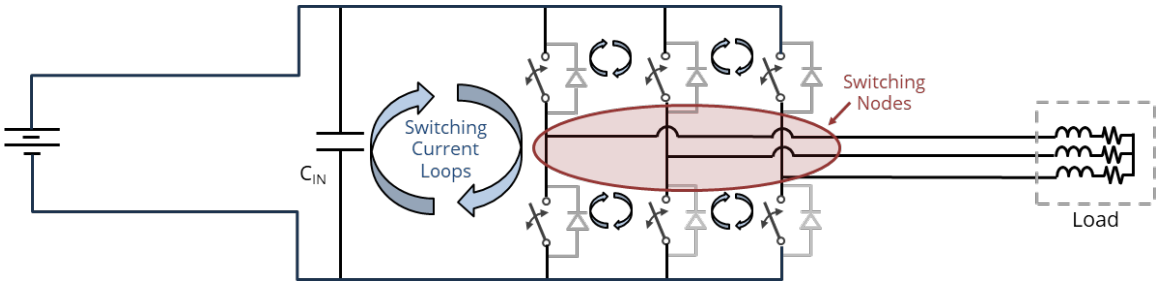

Figure 11.43. Three-phase inverter topology.

If shielding is required on the wires between the inverter and the load, all of the wires should be inside the same shield. Shields on individual wires are forced to convey relatively strong differential switching currents. Any small impedance in the shield's connection to the ground structure can then result in a voltage that drives the cable shields relative to the structure. On the other hand, a shield around all three wires only carries the relatively small current associated with any unintentional common-mode voltage driving the wires.

Snubber Circuits

The switching voltage node of a power circuit sees a very high dV/dt. In many cases, the voltage on this node overshoots its target and rings as depicted in Figure 11.44. This ringing is generally associated with the inductance of the switching current loop and a capacitance in the loop (usually an open switch capacitance). In the frequency domain, this ringing appears as a relatively narrow-band peak in the coupled noise.

The energy that drives this ringing typically comes from two sources. The first is the energy stored in the magnetic field associated with the switching current loop. This is proportional to the inductance of the loop and the square of the current flowing in the loop. The second source of ringing energy is the *reverse recovery current* injected into the loop when a conventional MOSFET or IGBT switch transitions from a closed to an open state. Reverse recovery current is due to a brief injection of charge that results when a diode transitions from a forward bias to a reverse bias. It can be the primary driver of ringing in power circuits employing conventional silicon switches. On the other hand, wide band gap switches (e.g., GaN or SiC devices) have virtually no reverse recovery current.

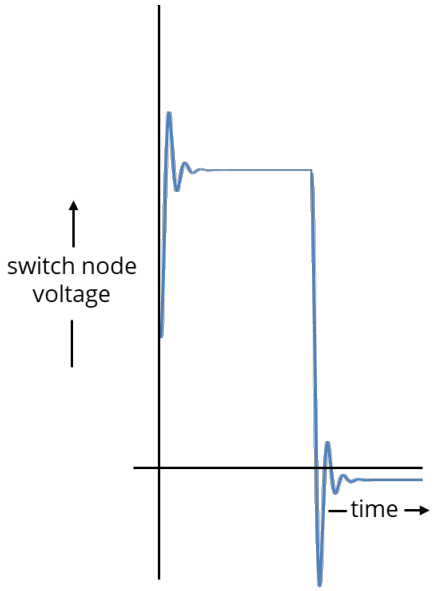

Figure 11.44. Ringing in a voltage waveform.

One method to reduce ringing in a switch node is to apply a *snubber circuit* across the switch (or between the switching node and circuit ground). In most cases, the snubber circuit is a small resistor in series with a capacitor. Snubber circuits applied to an H-bridge inverter are illustrated in Figure 11.45.

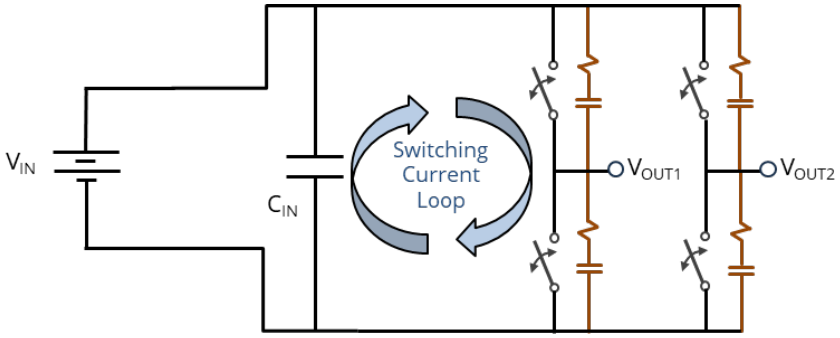

Figure 11.45. Snubber circuits applied to an H-bridge inverter.

To be effective, the snubber circuit must be the preferred path for the ringing frequency current. The series resistance damps the ringing. If its value is too small, it won't provide sufficient damping. If its value is too large, the snubber won't be part of the preferred current path. The optimal value is typically $R = \omega L$ at the ringing frequency. In many cases, this is a few ohms.

Decades ago, snubber circuits were an essential feature of nearly all switching power circuits, and they still appear in many design guidelines. However, snubber circuits are rarely necessary or appropriate in modern medium-to-low power converters that have low-inductance switching current loops and small switching voltage nodes.

Currently, snubber circuits can still be found on high-power converters and inverters. These devices typically require large components due to the relatively high currents and large spacing due to the high voltages.

Linear Voltage Regulators

Finally, no discussion of power circuit layout for EMC would be complete without mentioning linear voltage regulators. Linear voltage regulators convert one DC voltage to another (lower) DC voltage, by essentially forming a voltage divider that shunts excess current to ground. They are only used when the difference between the source voltage and output voltage are relatively small, because shunting a lot of current to ground wastes power and generates heat. On the other hand, compared to switching power converters, they are small and inexpensive. Most importantly, from an EMC perspective, they don't produce any switching noise.