# 9

# Practical EMI Control

## AN OVERVIEW

One of the nightmares of a power supply engineer is being assigned to design a supply to a specification that calls out some EMI limits. This is because a common timeline in such case is:

Go ahead and design a supply to meet all the other specs.

After the design has been breadboarded, measure the noise and find that it's way outside limits.

Throw some inductors and caps in the front end of it, and find this doesn't help.

Seek advice from other engineers, whose suggestions also don't pan out.

Hire a consultant, who tells you to totally re-layout the board and the mechanical design that has already been through CAD.

Sound familiar? But whereas some books offer only generalities; this chapter gives some practical advice that will make compliance *really* easier.

The first step in the field of EMI is to straighten out some terminology. **EMI** (electromagnetic interference) refers to electrical noise from one device or system that causes malfunction in another device or system; but this term is now used generically to refer to noise regardless of whether it causes problems. The exact meaning of "noise" is discussed below. Other terms of related or overlapping significance are **EMC** (electromagnetic compatibility), which describes conditions under which two or more systems can work simultaneously in the presence of each other's noise; **susceptibility**, which is a measure of how much noise is required to upset a given system; and **EMV** (electromagnetic vulnerability), which is a new term that means the same as susceptibility.

There is no end of books about EMI, and it wouldn't be reasonable to try to cover the whole subject in one chapter. Instead, we're going to concentrate on EMI as it is generated by power supplies, and more particularly, on the practical aspects: where it comes from, how to measure it, and how to fix it—or, better, avoid it. Two topics we won't touch on except in passing are electromagnetic susceptibility (since the power supply is often the major source of system noise) and the response and protection of supplies to transients (since the same actions that are taken to prevent power supply noise from contaminating the environment to some extent protect the supply from damage by the environment). Where transients are larger, transient protection may well end up being a separate box from the power supply. While this chapter makes no claim to being comprehensive, if you follow the practical rules laid out here, you'll be far along on the road to meeting even the most stringent EMI requirements. Noise control need not be a black art!

## Radiated and Conducted

Perhaps the most fundamental distinction in types of EMI is that between conducted and radiated noise: respectively, noise that is carried by conductors and noise that does not rely on conductors. Note that "radiated" noise is something of a misnomer. Measurements are usually taken at one or a few meters' distance from the supply, and at the lower frequencies this is actually near-field signal, meaning that you are measuring components of the field that do not propagate energy to infinity (propagating energy to infinity being the definition of radiation).

Not a lot can be done about radiated noise, and in fact you can do nothing at all about it once it's outside the system. So your goals are first to avoid generating it, and then to ensure that any unavoidable noise doesn't get outside. You avoid generating radiated noise by means of the same types of strategy we'll discuss in detail below to prevent the generation of excess conducted noise: attaching switching devices to grounded conductors, pairing cables that leave the box with their returns, and so on. The two tasks are related because radiated noise has to be radiated from an antenna (read: cables coming into or going out of the supply); so if there's no signal along the antenna (no conducted noise), there won't be any radiation either.

## What to Do About Radiated Noise

What to do about radiated noise once you have it is an interesting problem. The very first thing to look for, both because it's the most common problem and because it's cheapest to fix, is whether each wire that comes out of the box (both power and signal) is matched with its return. "Matched" here means that the wire and its return are physically close together both as they exit the box and as they are routed to their destination or source out of the EMI chamber. Matching is important because the signal level (noise, in this case) is directly dependent on the loop area formed by the signal wires; putting them close together, or better, twisting them together, minimizes this area and thus the noise. What you *absolutely don't want* is a single signal wire that routes out to an instrument somewhere. Rather, give the wire its own ground return, even if the line has no high frequency signal on it or carries only DC: noise can be picked up on this wire while it's still in the box, and then you have a beautiful antenna.

The next cheapest thing to look at for fixing radiated noise is to see whether the box is adequately sealed. The supply should certainly have some type of metallic container around the box, if only to have a place to attach ground. No, plastic has no effect whatsoever (and won't serve as a ground, either). Remember that frequency (in hertz) and wavelength (in meters) are related by the speed of light:

$$\lambda = \frac{300,000,000}{f}$$

and considering quarter-wavelength antennas, it's clear that a 1cm hole will allow signals of frequency greater than about 600MHz free passage, and will probably allow some signal out at a tenth of that. The 1cm hole doesn't have to be round for this to be true, however; a 1cm slit (see Figure 9.1) can radiate pretty much the same frequencies as a 1cm diameter hole. The only holes in the box ought to be those where lines are entering or exiting.

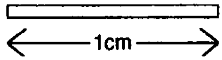


**Figure 9.1** A 1cm slit will allow passage of signals as low as 60MHz.

Once you have achieved control of the radiation from the system by enclosing it in an EMI-tight box, the only source of radiation will be the signal and power lines entering and exiting the box. Since you're going to be controlling the conducted noise on the power lines anyway, this design feature will control their radiation of noise. This leaves only the signal lines. You may want to consider putting filter pins on the signal lines, starting with those that carry high speed signals, such as digital clocks. But even static lines may potentially cause a radiation problem because of pickup onto them inside the box: that is, as the static lines go through the box toward their exit (or entry) point, various devices irradiate them, so that they carry noise; then once they exit the box, they are antennas, and radiate the noise to the outside world. So in many cases, it is advisable to simply save the hassle by getting a complete filter pin connector.

## What Kind of Box Material?

From a practical standpoint, as long as the enclosure around the supply is metal, it doesn't matter too much what the material is; because of cost, then, it will almost certainly be aluminum. When people get into trouble with EMI, sometimes they try something like a mu metal enclosure. Mu metal shields low frequency magnetic fields; the material is very expensive and difficult to shape mechanically. Although this approach can be made to work (for best results, the enclosure should be sandwiched between layers of grounded aluminum), it shouldn't be necessary if you pay attention to the signal and power lines.

> **Practical Note** Get the conducted noise under control first, as this will solve 80% of your radiated noise problem. Pay attention to signal lines. If a consultant recommends mu metal, don't even think about it; get a new consultant.

## Common Mode versus Normal Mode

Concentrating from now on on conducted noise, there are two basic types, common mode and normal mode (also called differential mode). It's easy to explain the difference: Normal mode (see Figure 9.2) is noise that flows in on one power line and returns on the other (neutral), whereas common mode (see Figure 9.3) is noise that flows on both the power lines simultaneously and returns on the ground line.

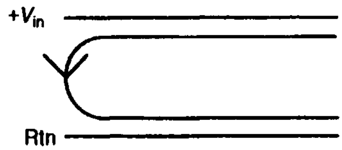


**Figure 9.2** Normal mode noise flows in on one power line and returns on the other.

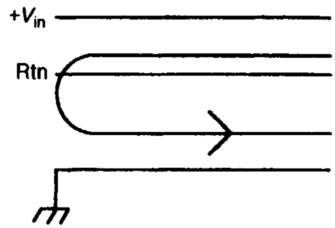


**Figure 9.3** Common mode noise flows in on both power lines and returns on ground.

## Return versus Ground

"But I thought return and ground were the same thing—the black wire going back to my power supply."

Even for your lab supply, return and ground aren't the same. Every good lab supply will have a third terminal which is the ground. The supply will have its output isolated from the AC line, as in Figure 9.4, and ground will be attached to the metallic box in which the supply is housed. You *may* then strap the ground and return together at the posts. This isn't necessary, though, and it's not always desirable.

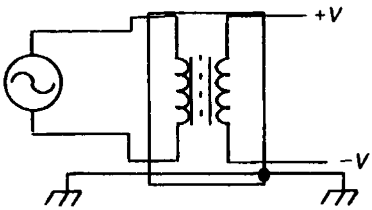


**Figure 9.4** A lab supply should have its return separate from ground.

In an AC system, ground and return are the same only at DC: the national wiring code requires that they be attached together where power enters the building, which can be a long way away from your system. When this is the case, ground and return are effectively isolated from each other at AC frequencies of concern for EMI. Thus it makes sense to talk about common mode noise as flowing back from the power and neutral to ground. Let's be explicit again:

Normal mode current is current flowing from $+V$ to $-V$ in Figure 9.4; it is what is normally thought of as delivering power.

Common mode current flows simultaneously through both $+V$ and $-V$ and returns on chassis ground; it does not normally deliver power.

### Military versus Commercial Measurements

The final major distinction to address before getting down to practice is what EMI tests you're trying to meet, which is to say, what kinds of measurements you're required to make.

The goal of all measurements of EMI is to ensure that the noise generated by the supply doesn't cause malfunction in anybody else's equipment (radio, garage door, helicopter navigation system). But military and commercial regulations have ended up with measurement techniques, and thus noise specifications, that are diametrically opposite, and this to some extent influences the corrective actions open to you if you are a noise generator.

To start with military measurements, the controlling document, MIL-STD-461, requires that the power and return lines be (AC) coupled together with a very high quality (meaning low ESR) 10μF capacitor for the measurement. You then measure the current flowing in each line, and the limits are set by the amount of current allowed at each frequency.

In a completely opposite way, commercial measurements [domestic (FCC), European (VDE), and others] require that the power and return lines be independently isolated from the power source with a 50Ω impedance that is more or less constant with frequency. (The impedance box is called a line impedance stabilization network, or LISN.) You then measure the voltage on each line, and the limits are set by the amount of voltage allowed at each frequency.

Even though these standards seem completely different, with military requirements measuring current and commercial requirements measuring voltage, they are related, as they must be: the current the military will measure flows through the impedance of the commercial system to generate the voltage the commercial system will measure. However, the techniques used for meeting these requirements will also, in a certain sense, be mirror images, as explained below.

## HOW CAN I SEPARATE CM FROM NM?

No doubt you noticed that neither the commercial nor the military measurements said anything about measuring the ground wire. This is because it is assumed that currents in the ground are of no practical concern to any system. But since you measure only one line at a time, the common mode and normal mode noise get mixed together: referring to Figures 9.2 and 9.3 and considering the power line for a moment, you can see that a measurement of the noise on this line includes both common mode, which is returning via ground, and normal mode, which is returning via the return. It is this "partial measurement" that is responsible for EMI's annoying habit of disappearing at one frequency but then re-appearing at another. Adding a normal mode filter to reduce normal mode noise can seemingly increase common mode noise, and vice versa. (A technician once told the author that EMI resembles a balloon—when you push down in one place it pops up in another.) Of course, in reality common mode and normal mode are independent, and to meet specifications, you must be able to control them both.

What's needed, then, is a way to separately measure common mode and normal mode, so that adequate filtering can be done for each, independently. This is fortunately very easy, especially in a military measurement. To measure common mode current as in Figure 9.5, you want to measure the current that is flowing simultaneously (in phase) through both wires. Thus all you need do is place the current probe around both wires with no twist. To measure normal mode current, you want to measure the current flowing in the power wire that is flowing out of phase with that flowing in the return, as in Figure 9.6. Thus you put a foldback in the return wire and measure the current while it's pointing in the same direction as the power. [The broken line in Figure 9.6 signifies that that part of the return line isn't inside the current probe.]

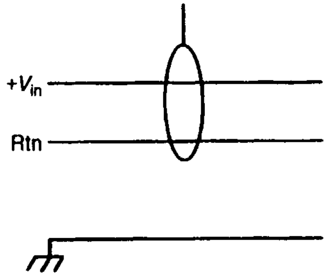


**Figure 9.5** Measuring the common mode current.

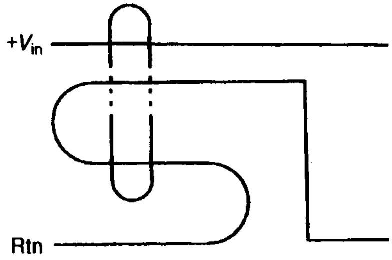


**Figure 9.6** Measuring the normal mode current is done with a loop in the wire to eliminate the common mode component.

When you're doing commercial measurements, there's no such convenient method, although recently Lee et al. [1] proposed doing something equivalent with transformers that couple the wires together appropriately. But as indicated above, the current and the voltage are certainly related to each other, so you can use a method like the one shown for military limits. Measure the normal and common mode currents with a current probe, just as presented above. The ratio of common mode to normal mode currents will reflect the ratio of common mode to normal mode voltage; so by measuring the noise voltages the commercial way, you will then be able to partition the noise between common and normal mode, and design filters appropriately.

---

**SIMPLISTIC EXAMPLE**

The current for the common mode is measured to be 300μA at 100kHz, and the current for the normal mode is 3mA at the same frequency. This ratio of normal to common mode is 10 : 1. The total noise voltage is measured on the power line at 100kHz to be 101dBμV = 110,000μV = 110mV. Presumably, then, 100mV of this comes from normal mode noise, and 10mV from common mode, because 100mV/10mV = 10 : 1, and 100mV + 10mV = 110mV, the total.

---

## WHERE DOES THE NOISE COME FROM?

The first step in controlling noise emissions is to understand where the noise comes from: that is, what generates it and how it gets to the lines being measured. Given knowledge of the origins of the noise, the first and best control technique will be arranging the circuitry to prevent noise from escaping to the measurement lines at all; filtering the noise that does get there is decidedly second best.

## Switching Waveforms

The major source of conducted (and radiated) noise from a switching power supply is, not surprisingly, the switching. This is not surprising because the switching involves the highest power in the circuit (and thus highest currents) and highest $dV/dt$, and also has the highest frequency components in the supply: for example, a MOSFET being switched from on to off in 50ns has a fundamental at something like 1/50ns = 20MHz, and also odd harmonics (at 60MHz, 100MHz, etc.). A diode has similar types of spectrum because we want it to turn on and off as fast as possible also, and for the same reason: fast switching minimizes power loss.

And in fact, one need only trace through the power path of a converter to see which elements are likely to be serious noise offenders: the switching transistors and the rectifying diodes (or synchronous rectifiers). If there is an inductance on the secondary, the high frequency, high power components of the spectrum won't conduct through it (though it will still radiate), so everything after the diodes is less noisy. Moreover, if the power transformer is designed well, its core material will form a partial shield and so it won't generate too much noise either.

## Capacitive Coupling

Having identified the major noise sources (see Figure 9.7), let's think about what can be done to reduce their generation of noise. We said that the high speed switching is desirable because it keeps down losses, and we don't want to do anything that might hurt the converter's efficiency.

The realization that the noise mechanism is the high speed switching of power perhaps brings to mind the idea of resonant conversion, since the switching in such a supply is by definition done at low power. (The current or voltage across either the FET or diode or both is zero when the switching occurs.) The possibility may seem tempting, but overall the disadvantages of resonant converters outlined in Chapter 2 outweigh the noise reduction benefit gained. Many types of resonant converter also change switching frequency with line and load, causing changes in the noise spectrum. This can make it harder to filter these converters than it would have been if you had stuck to a fixed-

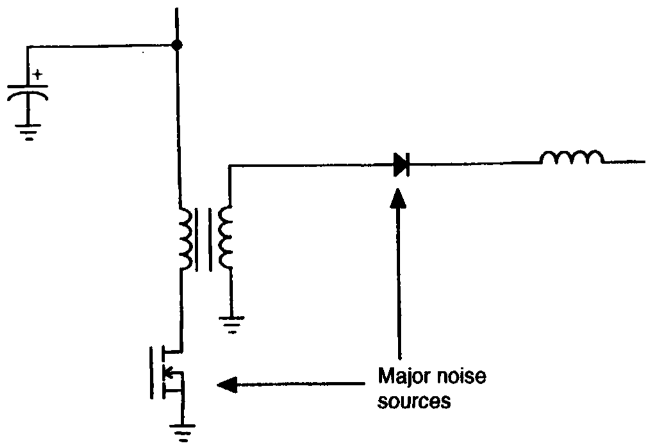


**Figure 9.7** Major noise sources in a switching power supply.

frequency, hard-switching converter. The practical aspect remains the same: avoid resonant converters. However, the same arguments show that the best of both worlds may be attainable by using a quasi-resonant converter, which retains zero switching, but with a fixed harmonic spectrum.

Nonetheless, acceptable noise performance can be achieved even with very high speed hard switching by considering the mechanisms whereby the switching noise gets out to the world and its measurement lines. The most obvious is that the current is pulled from the input line by the converter at the switching frequency. Aside from choosing a topology that has continuous conduction rather than discontinuous (which decreases edge amplitudes), filtering is the only option, as discussed shortly. A less obvious, but still potentially very serious mechanism for noise conduction is capacitive coupling of the switching waveforms onto ground. The conduction paths are shown in Figure 9.8.

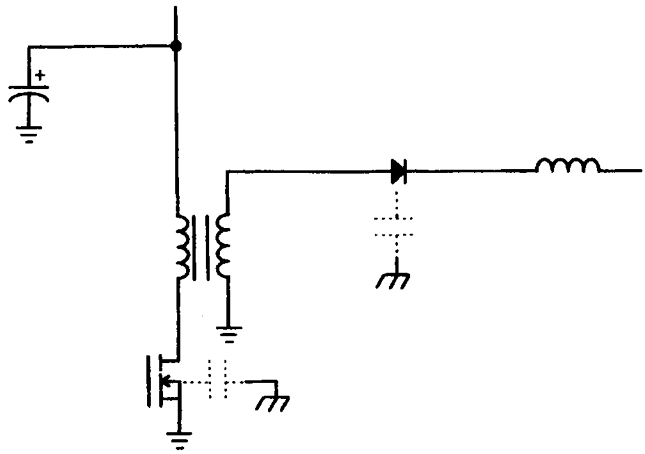


**Figure 9.8** Transistors and diodes can couple common mode noise through their body capacitance.

Power switching devices are typically mounted to the case of the power supply for heat sinking; this case is grounded. Since there is a small distance (e.g., the thickness of the MOSFETs case) between the actual die and the case, and a large area, there is a significant capacitance between the two, which will conduct high frequency signals to ground. The signals then return through the power and return lines, which is to say it is common mode noise.

Instead of filtering this signal, a better tactic is to reduce the coupling—that is, reduce the capacitance to ground. The area of this capacitance is fixed by package size, but the distance can be increased. The trick is to use a thermally conductive insulator between the case of the device and the power supply case, preferably one with a low dielectric constant. Typical choices are silicone-based plastics and beryllium oxide. This reduction in capacitance can mean very substantial benefits in reduced common mode noise filtering. In addition, for isolated supplies, it can cut down on noise transmitted between secondary and primary through the two series capacitors, the path in question being diode to case, then case to FET.

# CONCEPTS OF LAYOUT

Having been discussing the origin of conducted noise on the power lines, let's turn to another aspect of noise generation control, namely, the placement and routing of components and traces to prevent noise from upsetting the operation of the circuitry in the power supply. Such an upset can be a very serious problem; in the worst case, supplies won't work at all because of noise. We'll get to filtering shortly.

## Signal Ground versus Power Ground

A signal ground is, by definition, a ground trace that carries low currents; a power ground is a trace that carries high currents. This isn't exactly what you'd call a quantitative pair of definitions, but in practice the concepts usually are clear enough: the ground from the resistor that generates the timing signal from a PWM IC is a signal, the source of a power MOSFET is attached to power ground, and so on.

Maintaining separate signal and power grounds is essential to good operation of a power supply at all stages of the design. It will save endless trouble in your breadboards, and will make the difference between a good PC board and one that requires endless troubleshooting with noise filtering of signals. The reason becomes clear from a consideration of Figure 9.9.

Any trace (or wire, or even ground plane) has some resistance and inductance.

**Practical Note** The resistance of a trace is approximately given by the formula:

$$R = 0.5\text{m}\Omega\frac{\text{length}}{\text{width}} \quad \text{1oz. copper}$$

at room temperature. Two-ounce copper is half this, etc.

If a high current passes through the trace, there is a voltage drop across it because of its resistance; if the current is high frequency, there is additional voltage due to the inductance. If this high current passes through the same trace that is being used to ground a signal component, the signal component doesn't see the proper ground; rather, it is

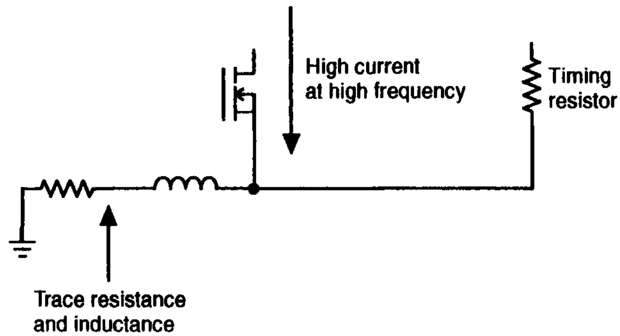


**Figure 9.9** Power paths can upset signal grounds because of trace impedance.

elevated off ground by $IR + L(dI/dt)$. Potentially even worse, the high frequency component elevates the signal component's ground periodically, and quite possibly even *synchronously* with the signal that the component is supposed to be processing! This is a prescription for disaster. The solution is to create separate grounds for signals and power and attach them together at a single point, preferably at a bypass capacitor at the power entry point (see Figure 9.10). This configuration is known as a star ground.

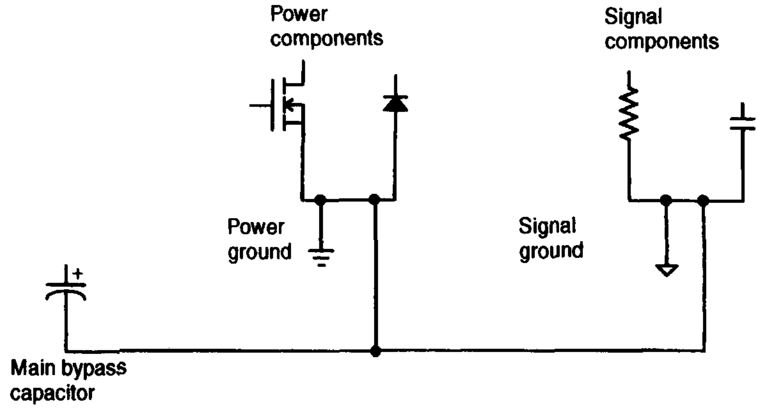


**Figure 9.10** Power ground must be separate from signal ground; only at the power entryway can the two grounds be tied together.

> **Practical Note** Create these separate grounds on both breadboards and PC boards. Do it religiously! As a practical matter, anything above about 100mA can be considered to be power. It is re-emphasized that the two grounds are to be attached together only at a *single point*. Otherwise there can be ground loops, which defeat the whole purpose.

Figure 9.11 is a picture of how not to lay out physical traces (i.e., with multiple paths for the ground currents).

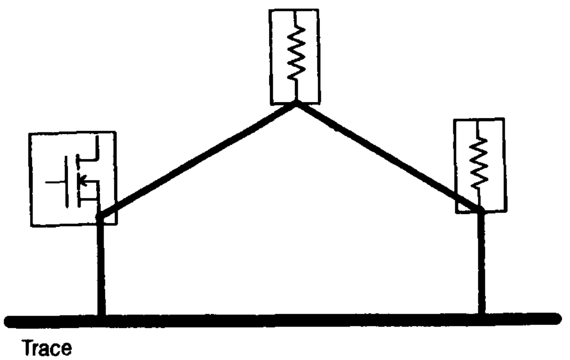


**Figure 9.11** DON'T lay out a PCB like this. Return currents have more than one way of returning to the point on the left. In this physical picture, bold lines are traces.

If you feel the need for additional copper for the power ground, thicken the traces instead of trying to run multiple traces all over the place. The correct way corresponding to our bad example is shown in Figure 9.12, where the bypass capacitors are located at the junction at the far left.

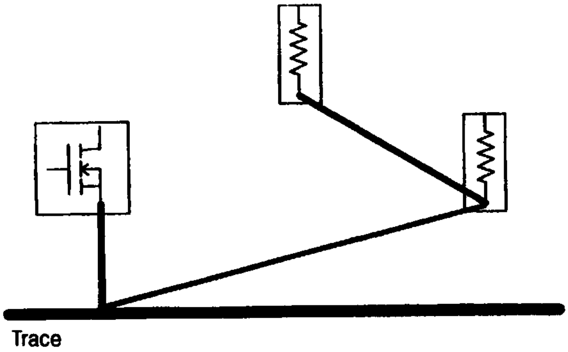


**Figure 9.12** The right way to lay out a PCB: connect all the power returns together before attaching them to the main return trace. Return currents have only a single path to return to the point on the left.

In the (presently unusual) case in which there is substantial high speed digital circuitry on the power supply, it would certainly be worthwhile to consider a third, digital, ground system separate from analog signal and power ground, again attached to the others only at one point.

## Grounding a High Current Driver; Ground Islands

A special case of needing a separate ground is nevertheless so common that it warrants mention here: the MOSFET gate driver. A gate driver works by pulling from its bypass caps current that it then delivers to the gate–source capacitance of a MOSFET. When the MOSFET is turned off, that gate capacitance is discharged, and the charge returned to ground. In the full cycle, there are two very fast pulses of high current (as high as 6A on some devices). The object of the layout is to ensure that these fast pulses are not seen by the rest of the board, which should only see the (much lower) average current. Figure 9.13 shows a highly recommended layout for this type of device. You can think of this as creating a little "island" of ground just for this current path. The current from the FET's gate is returned preferentially to the capacitors, preventing the high fast currents from flowing into the rest of the ground plane, which would cause upsets. Of course the drain current passes through to the source and then to the main ground trace. This arrangement

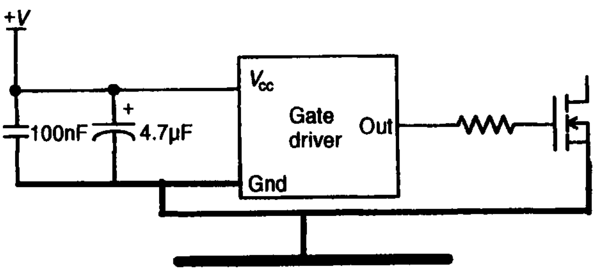


**Figure 9.13** Creating a ground island for a gate driver ensures that gate currents won't upset the ground.

is called an "island" because on a PC board these grounds are all large traces connected by a thin neck to the rest of the ground plane: (see Figure 9.14).

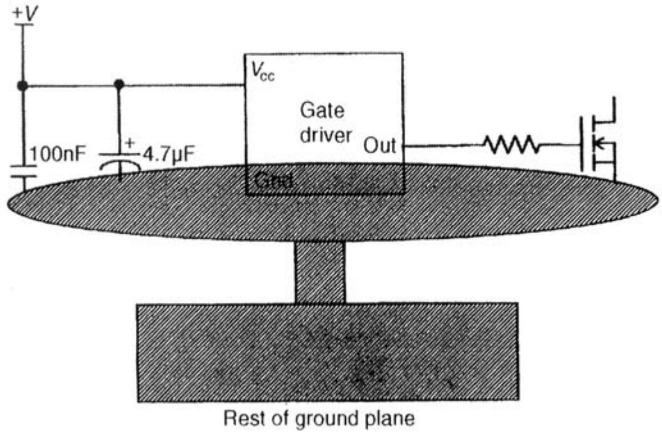


**Figure 9.14** How to lay out a ground island on a PCB ground plane.

## What If the Device Has a Signal Input But No Signal Ground?

Some of the best gate drivers have multiple pins for power ground (which is good) but neglect to have a separate pin for the ground of the signal that is driving the device (which is bad). In this case, it is still desirable to use both ground pins for the power ground, and let the signal take its chances with ground bounce (it usually has TTL noise margin). As long as the connection to the rest of the ground plane isn't too far away, this arrangement seems to work. If ground bounce becomes a real problem, you may have to select a different gate driver. (The qualifier "best" above is to be understood in the sense of gate drivers having the highest drive current.)

## Where to Put the Current Transformer

As long as we're talking about gate currents, it is useful to point out that the high current also may affect the current transformer if positioned improperly. Based on the analogy with a current sense resistor, the natural place to put a current sense transformer is in the source of a MOSFET (see Figure 9.15). However, this means that the gate turn-on current (which, remember, can be as much as 6A) also passes through the current sense transformer. Even for high power converters, this can be a significant fraction of the switching current you are trying to measure; for lower power converters, it can be the largest component of the signal. As a result, either the signal is corrupted by this irrelevant current pulse from the gate or it has to be so heavily filtered that the signal you want to measure gets filtered out as well. Whether the signal is then used for current mode control, or just a pulse-by-pulse current limit, the result is bad.

The way around this set of problems is to put the current sense transformer's primary in the drain of the FET, where it sees only the MOSFET's drain–source current, not the

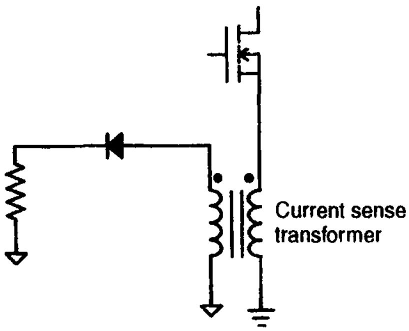


**Figure 9.15** It seems natural to put a current sense transformer in the source of the power MOSFET.

gate–source current associated with the gate capacitance (see Figure 9.16). This design has no deleterious effects, either on the current signal (since it is after all transformer-isolated), or on the operation of the converter (since the primary inductance, which is usually single turn, is negligible). The current sense transformer could even go between the main transformer primary and the power bus, as long as it comes after the main input cap.

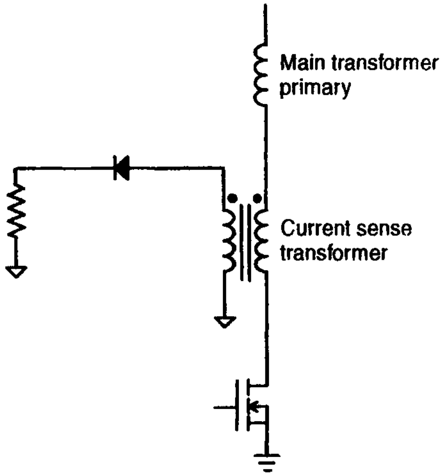


**Figure 9.16** To avoid measuring gate current, place a current sense transformer in the MOSFET's drain.

## Feedback Lines

While we're on the subject of positioning components, a few practical tips on feedback layout are in order. The two usual feedbacks are those for voltage and current, and the practical tips here apply to both.

When you build a breadboard, there are usually components every which way, and a mess of wires all over, quite possibly mixing power and signal lines. And while the converter has been compensated to have 45° of phase margin, if the noise pickup is too great onto either the current feedback or the voltage feedback, there may still be, if not instability, at least considerable duty cycle jitter.

**Practical Note** When one is building a breadboard, it can be helpful to use a twisted pair of wires for the feedback lines, to reduce the noise pickup onto these critical lines. Shielding the twisted pair is rarely necessary, but if shielding is required, ground the shield on the signal end only; leave the power end floating. (The power end for a voltage feedback line refers to the output voltage node. For a current feedback line, it refers to the secondary of the current sense transformer.) It is also desirable to have any small-signal components close to the PWM rather than located at the output. For example, if there is a voltage divider for feeding back the output voltage, locate these resistors close to the PWM, not the output, and run the twisted pair from the output voltage; don't try to put the divider close to the output, and *then* use the twisted pair.

The reason for the injunction with respect to voltage divider placement is that a low impedance source, such as the converter output, is more noise resistant than a high impedance source such as a $10k\Omega$ resistor.

When you go to a PC board, of course you can't use twisted pair anymore, but you can still try to run the trace containing the feedback signal in parallel (i.e., on top or bottom) with a ground trace, or even better, with a ground trace both on top and on bottom (in a multilayer PCB).

One more trick works for both breadboards and PCBs. Try terminating the twisted pair or trace for a voltage sense line with a 100nF capacitor. Yes, from a schematic viewpoint the capacitor is just in parallel with the output caps already present. From a noise viewpoint though, those output caps don't help: they're in the wrong place. A cap right at the termination filters noise very nicely, on the other hand, and obviously has no effect on the loop.

### Further Layout Tips

All the foregoing layout tips have been variations on the same theme: power and signal must be kept separate! In line with this restriction, a few additional specifics may be in order. When laying out the power stage of a converter, it is important to keep all the power components physically as close together as possible. Not only is this good for efficiency (by reducing trace resistance), it also minimizes loop area for radiating noise onto signal lines.

This rule is of the greatest importance in designing the conection between the gate driver and the MOSFET gate. Specifically, make the connection as short as possible. It is probably worthwhile to orient the IC so that its output pin faces the transistor's gate pin. And DO NOT connect the two together through any vias—this shortcut risks contamination of other planes!

## LOW FREQUENCY FILTERING

### The Basics

Now we're ready to find out what to do with the noise that's still there after you've done your best with layout and the physical aspect of design. The subject of filtering divides itself rather naturally into two parts, low frequency filtering and high frequency filtering.

Low frequency filtering is what can be done with discrete, lumped components, such as discrete capacitors and inductors; high frequency filtering is everything else, and into which we lump filter beads, feedthrough caps, etc.

Filtering consists basically of presenting a high impedance to the signal in the path you want to keep free of noise, and a low impedance on the path you do want the noise to travel.

## Normal Mode Filters

Low frequency filtering comes in two parts, normal mode filtering and common mode filtering. Following the discussion above, normal mode filtering tries to reduce noise on the power line that returns on the return line. Remember that this means noise on the power line *that exits the box* and returns on the return line. So the tactic for filtering is to shunt the noise on the power line to the return line *before it leaves the box*, thus ensuring that it will return without having been measured. This amounts to putting an inductance in line with the power line, to block it from getting out, and at the same time providing a capacitor from the power to the return line to provide a low impedance path for the noise to go through instead.

## Commercial versus Military

Now although the preceding discussion on commercial versus military filtering indicated that the two conventions are closely related, their difference shows up here, when a low frequency normal mode filter has to be designed. The question is whether (looking from the power supply out toward the power source) to design a filter that first has a capacitor and then an inductor, or one that has first an inductor and then a capacitor. The commercial measurement has a (relatively) high impedance source (50Ω) and measures voltage; thus we can use this source as a block for the noise, and it is advantageous to use first an inductor and then a cap, as illustrated in Figure 9.17.

In some situations, the noise may be so low that the inductor is unnecessary; the capacitor forms a voltage divider with the 50Ω that is small enough to divert away most of the noise. Remember, though, that for this circuit to work, the ESR of the capacitor is critical. Either a multilayer ceramic or a metallized plastic capacitor might be tried for this application.

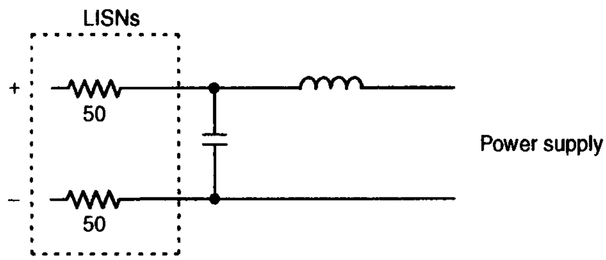


**Figure 9.17** A commercial supply should be filtered with a two-pole filter's capacitor facing the LISNs.

For a military measurement, conversely, the source is a low impedance source (10μF) and measures current; we thus need to block current from getting to this low impedance, and the filter should be first a cap and then an inductor (see Figure 9.18).

In this case (unlike the commercial case) the capacitor is doubtless already there in the form of a large electrolytic. With this bulk cap, however, it is advisable to parallel a ceramic 1μF or 100nF cap (or both—a 1μF capacitor loses its effectiveness around 1MHz, whereas a 100nF device continues working up to about 10MHz). This measure is advised to counter the effects of the poor high frequency characteristics of the large cap.

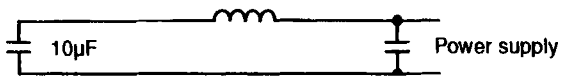


**Figure 9.18** A military supply should be filtered with the two pole filter's inductor facing the source cap.

## Selecting the Values

Selecting the values for the $L$ and $C$ is relatively straightforward. You already know from measurement the unfiltered spectrum, and we know that the two-pole filter we are designing is going to roll off the noise at 40dB/decade. So here is the procedure for deciding the position for the filter to start rolling off.

> **Practical Note** Find the lowest frequency component that is out of spec (preferably of the normal mode measurement, rather than the specification measurement that combines normal and common mode, as discussed above). Suppose, for example, that it is out of spec by 20dB = 10 at 100kHz. The filter that would bring this into spec would start at a freuqency of $100\text{kHz}/\sqrt{10} = 30\text{kHz}$ (square root because it's two poles). Now, on top of (a copy of) the noise measurement, draw a straight line, starting at this frequency (30kHz here) and rolling off 40dB/decade. If none of the other peaks are above this line (and frequently they won't be), you're done: you need a filter with an $LC$ resonant frequency of 30kHz. If one or more of the other peaks *are* above the line, repeat the whole calculation with the lowest frequency peak until you have a frequency that guarantees that all the noise peaks are below it. [It is interesting to observe that the noise specifications typically roll off at 30dB/decade, which is between one and two poles.]

With the frequency for the $LC$ filter determined, one more parameter still needs to be determined to calculate the magnitudes of the components. In a general sort of way, inductors are more expensive than capacitors, both in cost and in power loss (since they have series resistance). So the preference is to use more capacitance and less inductance. A practical recommendation is given in the section below entitled Optimal Filtering.

## Common Mode Filters

Common mode filters are easier to design than the normal mode variety because the former have fewer available choices. A common mode filter consists of common mode capacitors (called "Y caps" in the commercial world; "X caps" are normal mode) and a balun (the word comes from "balanced–unbalanced transformer") or a common mode inductor: see Figure 9.19. The common mode caps shunt current on both lines to ground, and the balun (note the polarity dots) presents a balanced impedance—that is, the same impedance on both the power and return lines, which constitutes a high impedance path for common mode noise. (The Z shape shown on the balun is standard, as is the set of two dots, although you will generally see just one or the other.)

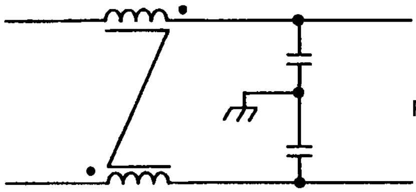


**Figure 9.19** A common mode filter consists of a balun and capacitors to ground.

## Selecting the Values

The common mode case is quite the opposite of the normal mode case with respect to selecting components: here, the capacitors are more expensive than the inductor. The reason is twofold. First, the capacitors go to ground and must be rated to take a transient that may be 3kV or even 6kV, making them quite large. Furthermore, there are strict safety limits on the amount of current allowed to flow to ground, which limits the maximum size of common mode capacitance that can be used, typically to a few nanofarads. Thus, you select the maximum allowable capacitance and then use the same technique sketched above to determine where the *LC* needs to be; this procedure uniquely determines the balun's inductance value.

> **Practical Note** Don't panic if you calculate that a large value is required for the common mode inductance. Baluns carry the same current in both windings, and so, ideally, see no net current. They can thus have a large number of turns without saturating the core.

Furthermore, when calculating the inductance required, remember that the two capacitors are in parallel (doubling the net value) for purposes of common mode noise; and the individual windings of the balun are in series. Thus you have double the number of turns, which is four times the inductance. This gives you an extra factor of 8 noise suppression for free (see Figure 9.20). The capacitance in Figure 9.20 equals 9.4nF, and the inductance is 4mH, giving a cutoff frequency of 26kHz!

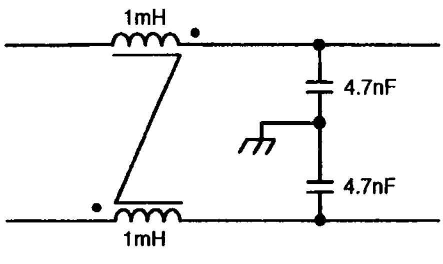


**Figure 9.20** The windings of the balun are in series and the capacitances are in parallel for the common mode noise.

## Caps and Inductors and Their Limits

Capacitors have limitations in their frequency response, and this translates into limitations on their usefulness for EMI filtering. Electrolytic capacitors have relatively large ESR, which means that above the *RC* frequency, they look resistive and no longer are a pole. For example, a 100μF cap with 100mΩ ESR becomes resistive at about 16kHz. It is thus of no use for EMI control.

In practice, ceramic or plastic capacitors are always used for EMI. Even these have their limitations, though, because of lead inductance.

> **Practical Note** 1μF caps are good only up to about 1MHz. Above this, use a 100nF cap, which in turn is good only up to about 10MHz. It may thus actually make sense to have a 1μF, a 100nF, and a 10nF all in parallel for noise suppression.

Inductors have limitations too. (Their winding resistance, though bad for power consumption, is too small to be significant at noise frequencies.) The most important limitation is distributed capacitance, which may be thought of as occurring in parallel with the inductor. Above a certain frequency, the capacitance is lower impedance that the inductance, and so the inductor no longer blocks noise above this frequency in the way that might be naively expected.

---

**EXAMPLE**

Suppose that a 1mH inductor has a capacitance of 100pF. Then its impedance will stop increasing above 500kHz, and will actually start decreasing. Of course, a smaller inductance also has a lower capacitance, and thus a higher frequency.

---

It might thus make sense to have two inductors in series, rather than a single larger inductor, which would put the inductances in series (increasing inductance) and the capacitances in series (decreasing capacitance). Inductors are so expensive, however, that this is never done in practice.

## MOVs Have Capacitance

Here's a freebie! Many, if not most, designs will require an MOV on the power lines to deal with transients. An MOV has a small capacitance and so can be used as part of the filter. This is also a caution to those who put MOVs from power and return to ground: the MOVs capacitance adds to the leakage current, so the size of the Y caps must be suitably reduced.

## Two for the Price of One

Now here's a good idea. You can build both a normal mode inductor and a balun into a single magnetic device, saving the price and space of an extra device. This is done in a way that is straightforward to describe but requires care in design. You wind up a balun, using, let's say, 47 turns on each side; then on the power line side, you *add one extra turn*. The device is still a balun, but now it also has a series inductance, equal in magnitude to

$(48^2 - 47^2)A_L$, far larger than a single turn on the same core. The catch is that now the core has to support a flux density, and you have to check to be sure that it won't saturate at maximum current.

### You Can't Get 100dB Attenuation!

You've no doubt noticed that the discussion of low frequency filters only talked about two-pole filters, not higher order filters with more than one capacitor or inductor or both. There are reasons for this. First, such filters are much harder to design; indeed, there are people who make a career of designing them. Another reason, at least for commercial designs, is that a higher order filter requires another inductor, and this often means unaffordability.

The most basic reason, though, is that except in special circumstances, such filters shouldn't be necessary. Four-pole roll-off is pretty fast already, and if you require six poles, something is probably wrong.

**Practical Note** If your calculations indicate a need for more than about 60–80dB of attenuation in the low frequency regime, you had better go back and work on your layout. Another possible way to reduce the filter requirements is to increase the switching frequency.

Here's another point along the same lines. We've already mentioned that components have imperfections which limit their performance. Additionally, layout on real boards entails leakage paths and cross talk between traces. The bottom line is the same as in the practical tip above: you can't get 100dB of attenuation. If you think you need it, try again.

It's possible to get somewhat more attenuation using a commercially available filter than building your own, primarily because the commercial device has paid very close attention to avoiding cross talk in the layout, and typically has enclosed the filter in a metallic box. Of course, you can achieve the same results by building your own with the same techniques, and at far less cost.

## HIGH FREQUENCY FILTERING

High frequency filtering comes into play at frequencies where lumped components start having imperfections that make them poor filters. A "high" frequency might start at 10MHz. From this frequency up to maybe a couple of hundred megahertz, there are still components that can be added to help the spectrum; above this frequency, all you can do is seal the power supply enclosure better.

### Where Should I Use Beads?

Ferrite beads have excellent high frequency characteristics, with impedance continuing to increase even above 100MHz. Unfortunately, they are also very easy to saturate with just a little DC current—a parameter some manufacturers don't even bother to specify. So for the most part, a bead is pretty useless for input noise filtering.

As a sidelight, you may see a design that uses a bead on the gate (or drain) of a MOSFET. On the gate, this is a really bad idea: the reduction in noise is caused by a reduction in the FET switching speed, which directly translates into power loss. On the drain, it is typically ineffective, again because it saturates at a low current. If the idea is to block current flow for a few tens of nanoseconds (such as in synchronous rectification) and then let it saturate so that it doesn't add inductance to the power path, however, a bead on the drain may have a use. Even then, though, the power stored in the bead has to be dissipated or otherwise dealt with each cycle, just like leakage inductance in the main power transformer.

### Feedthroughs

Feedthrough caps and filter pins are roughly the same, though pins are typically used for signals and feedthrough caps for power. They are very high quality capacitors, effective out to hundreds of megahertz, occasionally with tiny inductances that can take as much as 10A, depending on size. They begin providing filtering at about 10MHz, unless you get a large one, which may have some modest attenuation as low as 1MHz. Note, though, that their attenuation is typically rated in a 50Ω system, which gives almost no clue to what it might do in a military measurement.

In many cases you can get by without a great deal of feedthrough capacitance. This is because at the high frequencies at which the feedthroughs work, the power cable leading up to the power supply box has very significant impedance: about 1μH, which at 10MHz is 60Ω, for a meter of wire. Since all these numbers are quite dependent on details of physical layout, there are no rules for how to filter the high frequencies: you try a filter, and if it doesn't work, you try a bigger one. Sorry.

## SOME OTHER TOPICS

### Noise Estimation

The amount of noise a converter will generate can be estimated before the device is built. For example, a buck converter draws current in rectangular pulses. This pulse train can be decomposed into its spectral (frequency) components. Each such component can then be divided up between the input capacitors and the source impedance. The resultant currents in the source (for military; multiply by the impedance to get the voltages for commercial) can then be compared with the spec limits to design a filter as detailed above.

---

**EXAMPLE**

Suppose you have a buck converter with 5V input and 2.5V output, so that it conveniently has 50% duty cycle. Suppose the input current is conveniently 0.5A. Since the input is a square wave, the peak current is 1A. We can easily look up in a math book [2] that a square wave has spectral components at odd multiples of the fundamental, with amplitude inversely proportional to the multiple. For our case, let's suppose the switching frequency is 100kHz, so we have $4/\pi$ amps at 100kHz, $4/3\pi$ at 300kHz, $4/5\pi$ at 500kHz, etc. If we have an input cap of 1000μF, the 10μF source-impedance cap for the military measurement will source about 1/100 of the current, that is, $4/100\pi$ amps at 100kHz, $4/300\pi$ at 300kHz, etc. For the higher frequencies, the rise and fall times of the current waveform dominate. Suppose that the rise and fall times are both 100ns. These will generate spectral components at odd multiples of 1/100ns, (i.e., 10MHz, 30MHz, etc.) and the resultant currents can be estimated if some knowledge is available about trace inductances. A filter can now be designed using these various estimates.

---

The trouble with this procedure, despite its seeming utility, is that in practice it is common to add 20dB to the estimates to ensure that the filter produced will really work. This doesn't seem like much until you remember we are talking about *10 times the size of the inductor or the capacitor*! The trouble is that even for the low frequency components, there are numerous "sneak paths" that add in to the actually measured current. The high frequency components have so many capacitively coupled paths (and radiated paths, too) that your estimate is likely to be wildly wrong. Estimating using a simulator has, of course, exactly the same problems.

> **Practical Note** Estimating noise is unlikely to get you to a reasonable filter design. The only thing that works is a measurement on real hardware.

### Optimal Filtering

The question of how to select the $L$ and $C$ values of a low frequency filter was not fully answered above, in that the pole frequency was found in terms of the required attenuation but actual values were not determined. Subject to stability criteria, though (discussed in the next subsection), the best possible selection of component values can be based on optimizing the cost, or volume, or some other parameter of the filter. This is covered in an article by the author, which, as it is now hard to get, is reproduced here.

## OPTIMAL MILITARY EMI-FILTER DESIGN*

Military power supplies must comply with conducted emi limits specified in MIL-STD-461. This standard calls for measuring, across a wide band of frequencies, the noise current flowing into a 10$\mu$F capacitor. This is often done by installing a filter, measuring the current, installing another filter of similar design, remeasuring the current, installing still another filter and repeating the process until the emi limits are met. This method is not only inefficient, but can also result in a filter that is substantially larger, heavier, and more costly than it need be. By making suitable measurements of the noise source, however, it is possible to design a filter that works the first time and that occupies the minimum possible volume.

A DC input power supply with a single ground (nonisolated secondary) typically has two main sources of conducted noise: the switching transistor(s) and the output rectifiers. These two sources are, in turn, associated with two main frequencies: the basic switching frequency of the converter and the inverse of the transition times. The transition times for the switching transistors are the rise or fall times; for the diodes, the reverse recovery

* *This article appeared originally in Powertechnics Magazine* (April 1989, pp. 47–48).

times. There is also some noise generated as the result of ringing in various parasitic circuit elements, but the contributions from these sources is minimal.

A Thevenin source is characterized by its open-circuit voltage and short-circuit current, the two being divided to yield the Thevenin impedance. The Thevenin equivalent open-circuit voltage of the power supply [$V_{oc}(\omega)$] is a function of frequency. It can be measured using a high impedance connected in series with either the high or low power line, as shown in Figure 1. A spectrum analyzer is used to measure the frequency spectrum of the voltage existing betwen the line and chassis ground. Note that the voltage measurement is referenced to the chassis and not the return line. This is because MIL-STD-461 requires that the 10$\mu$F capacitor in which the noise current is measured be referenced to the chassis, as shown in Figure 2.

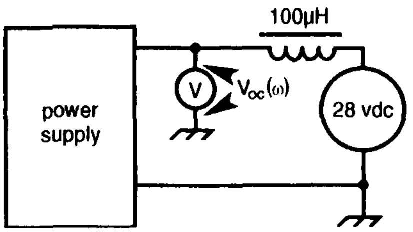


**Figure 1** Measurement of open-circuit noise voltage.

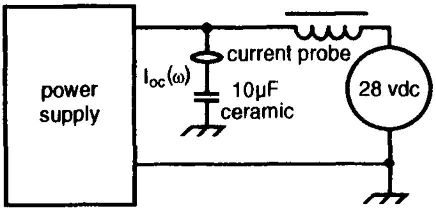


**Figure 2** Measurement of open-circuit noise current.

The capacitor used to measure short-circuit noise current [$I_{sc}(\omega)$] should exhibit a very low impedance at the noise frequencies of interest. A multilayer ceramic (MLC) capacitor is well suited to this application. This capacitor must be connected as close to the supply as possible, since even a few inches of wire presents significant inductive reactance at the frequencies being measured. With the capacitor in place, the short-circuit current can be measured with a spectrum analyzer using a current probe enveloping the wire lead to the capacitor. The Thevenin equivalent impedance at any frequency can then be calculated by dividing the equivalent open-circuit voltage by short-circuit current at the frequency of interest.

With the noise source characterized, the filter can be modeled. As shown in Figure 3, the basis of the filter is a discrete two-pole $LC$ ladder ($C$ and $U$), the values of which are such that the filter provides the required attenuation in the minimum volume. Also present in the filter model are an additional capacitance and inductance ($C1$ and $L$) representing a high frequency filter that attenuates noise in the spectrum at which the parasitic elements of the discrete filter give it poor performance. In addition, there is an inductance ($L1$) representing the one-meter wire required by MIL-STD-461 to connect the power supply to

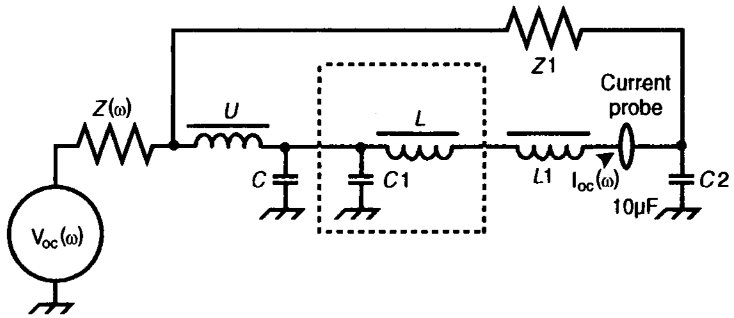


**Figure 3** Model of complete emi filter

the 10μF capacitor ($C2$). This wire has an inductance of about 1μH. In parallel with the filter is an impedance ($Z1$) representing all other paths for the noise.

***Calculating Attenuation*** The calculation of how much current will be measured by the probe [$(I_0(\omega)$] is straightforward. For example, if the shunt impedance is infinite and the 10μF MLC capacitor approximates a short circuit, the current will be:

$$I_0(\omega) = \frac{V_{oc}(\omega)}{a + b + Z + c + \omega U} \tag{1}$$

where $\omega = 2\pi f$

$a = \omega(L + L1)$

$b = Z\omega^2(L + L1)(C + C1)$

$c = \omega^2(L + L1)(C + C1)U$

A filter slug has already been chosen on the basis of current-carrying capability and high frequency rejection; thus $L$ and $C1$ are approximately known. The inductance of the one-meter wire represented by $L1$ can be measured directly. The value of $C2$ is set by MIL-STD-461 to be 10μF. MIL-STD-461 also sets the maximum value of $I_0(\omega)$ that can be allowed to flow at each frequency within emi spectrum. This means that the value of $I_0(\omega)$ is known at the switching frequency and its harmonics, the frequencies at which the greatest attenuation will be required. Similarly, $V(\omega)$ and $Z(\omega)$ are also known at these frequencies. Note that if $Z1$ is not infinite, it can be measured. Consequently, the equation can be solved for $U$ and $C$ at these frequencies.

***Minimizing Volume*** Clearly, there will be one frequency at which the equations relating $U$ and $C$ represent the worst-case condition; that is, the condition at which maximum filtering is required to bring the value of $I_0(\omega)$ within the limit set by MIL-STD-461. Depending on the topology, this is typically the fundamental or one of the first harmonics.

A second equation is required to determine the total volume of the filter. The inductors' size roughly tracks the maximum energy that can be stored. Typically the inductor in a DC bus filter is wound on a molypermalloy-powder (MPP) toroid, which permits high currents to be carried before the core saturates. Allowing inductance to fall by a maximum of 20%, the volume can be roughly estimated as:

$$\mathrm{Vol}_L \approx \frac{200\ \mathrm{in.}^3}{\mathrm{H} \cdot \mathrm{A}^2} \tag{2}$$

where the current is the maximum current the inductor must carry and still provide filtering.

Similarly, a ceramic capacitor's size also roughly tracks the maximum energy that can stored. The volume of, for example, a 1μF, 50V type CKR06 ceramic capacitor can be approximated as:

$$\mathrm{Vol}_C \approx \frac{3\ \mathrm{in.}^3}{\mathrm{F} \cdot \mathrm{V}^2} \tag{3}$$

where the voltage is the maximum voltage the capacitor will have to withstand. These calculations are, however, fairly rough because of the discrete nature of the components; one cannot get an arbitrary-voltage capacitor, not an arbitrary-size core. Because of this, the component values yielded by the calculations must be rounded up to the values of the real-world components.

The total volume occupied by the filter is essentially equal to the sum of the volumes of the capacitor $C$ and the inductor $U$. Equation 1 provides a relationship between $U$ and $C$ and, as a result, total volume can be expressed in terms of just one variable. Thus, total volume can be determined by taking the derivative of the volume with respect to this variable and setting it equal to zero. This yields a quadratic equation that can be solved for the variable selected with the result then being used to solve equation 1 for the other variable. Once the optimal filter component values have been determined, the inductor can be designed using standard practices with the capacitor value rounded up to the nearest available standard value.

A simulation of the filter is highly recommended to make sure there are no resonances at frequencies close to the switching frequency or its harmonics. This is important because the reactances involved have relatively high $Q$s, which can boost the noise at the resonance above the maximum allowable level. If resonances are close to critical frequencies, they can be shifted by increasing the value of the offending component(s). This not only moves the resonance to a spectral region of background noise where it does no harm, it also maintains the quality of the filter at other frequencies by increasing attentuation.

## Converter Stability versus EMI Filtering

In filter design there are limits (though usually not *practical* limits) to how large the inductor can be, and how small the capacitor. As follows from the discussion in Chapter 6, on stability, if the source impedance the converter sees is too high, the system may oscillate, and this is true whether the source is a filter or another converter. *Middlebrook's criterion* is the plan that the output impedance of the filter should be at least 20dB lower than the input impedance of the converter. Clearly this satisfies the rule of thumb given in Chapter 6 about stability; but it is not a necessary condition, only a sufficient one. The real criterion for system stability is the same as that given in Chapter 6:

**The phase margins of the system consisting of the filter and the converter must be positive.**

## REFERENCES

1. T. Guo, D. Y. Chen, and F. C. Lee, Separation of common-mode- and differential-mode-conducted EMI noise, *IEEE Transactions on Power Electronics*, PE-11(3), 480 (1996).
2. *CRC Standard Mathematical Tables*, 24th edn, W. Beyer, Ed. CRC Press, Cleveland, OH, 1976, p. 406.