CHAPTER 16

Electronic Construction Techniques

Experience is a hard teacher because she gives the test first, the lesson after.

—Vernon Law[†]

16.1 INTRODUCTION

Advanced mixed-signal electronic hardware can be designed on a computer and taken straight to multilayer surface-mount circuit cards without prototyping. Then—unless you're a Certified Wizard—it will fail to work and have to be cut apart with a Dremel grinder, laboriously debugged, partly redesigned, and laid out again.

This chapter is mainly about electronic prototyping and the reduction of a working prototype to a working PC board, with specific application to high frequency, high dynamic range, mixed-signal circuitry. We'll talk about laying circuits out intelligently, calculating the circuit board strays, assembling breadboards, and tuning *LC* filters. By the end, you should be equipped to build what you've designed, without suffering from avoidable layout and assembly bugs such as ground loops. You'll need to be comfortable with test equipment, know how to solder (preferably blindfolded), and have basic prototyping skills such as stuffing boards, wiring connectors, and drilling boxes.

16.2 CIRCUIT STRAYS

Simulations often don't match up with experimental reality, due to modeling errors and circuit strays. Strays arise from Maxwell's equations and the known properties of well-behaved conductors and dielectrics. We may not know the exact boundary conditions well enough to produce a perfect model of the self- and mutual inductance and capacitance of everything in the circuit, but we can usually take a good first cut at it; since we hope that their effect is not dominant, that's often enough. It doesn't replace prototyping, but it helps a lot in reducing the prototype to a circuit board, where the strays are usually much worse.[‡]

[†] *This Week*, August 14, 1960.

[‡] It is surprising to many people that a small circuit card can have worse strays than some big ugly prototype with wires in mid-air, but it usually has.

Building Electro-Optical Systems, Making It All Work, Second Edition, By Philip C. D. Hobbs

As a general philosophy, if you're using circuit simulators, for any critical analog functions such as front end amplifiers, you should check the simulation against the prototype, and adjust both until they match and you know why they match, and you can explain it to a sceptical colleague. Routine stuff can be designed, simulated, and laid out without always needing a prototyping step, especially with the rise of rapid-turn PC board websites, which often have their own CAD software you can use.

16.2.1 Circuit Boards

Most PC boards are made of epoxy-impregnated glass cloth, type G-10 or FR-4 (a flame-retardant G-10 variant). It's not a terribly well-controlled material—ϵ_r varies between 4.3 and 4.6, but it's cheap, strong, and reasonably stable. In North America, the copper thickness is conventionally specified in ounces per square foot and comes in 0.5, 1, and 2 oz weights (thicknesses of 0.017 mm, 0.035 mm, 0.07 mm, respectively).

If you're wise, you'll make a habit of building production circuits on ground plane PC boards exclusively; the extra fabrication cost is more than repaid by stable, predictable performance. Mixed-signal boards (those with both analog and digital circuits) must always be built on ground planes, and if you don't use a ground plane with a board containing ADCs and DACs, it will quite definitely fail. We'll specialize our discussion of circuit cards to the ground plane case, where all traces are really microstrip transmission lines.

16.2.2 Microstrip Line

We've already encountered microstrip line, and since we know how to calculate the impedance of a piece of mismatched transmission line, we have a completely general solution that can be used with traces of any reasonable length.

As shown in Figure 16.1, the impedance of microstrip line of width W and negligible thickness, with dielectric of thickness h and dielectric constant ϵ, is approximately[†]

$$Z_0 \approx \frac{377}{\sqrt{\epsilon}} \frac{h}{W} \frac{1}{1 + 1.735\epsilon^{-0.0724}(W/h)^{-0.836}}, \tag{16.1}$$

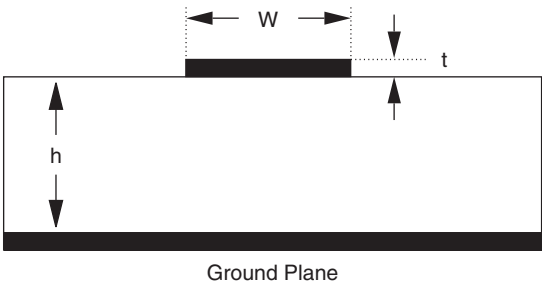

Figure 16.1. Microstrip line, that is, a metal trace over a ground plane.

[†]Ralph S. Carson, *High Frequency Amplifiers*. Wiley, Hoboken, NJ, 1982, pp. 78–79.

and the propagation speed $c' = c\lambda/\lambda_0$ is

$$\begin{aligned} c' = c\frac{\lambda}{\lambda_0} &\approx \frac{1}{\sqrt{1 + 0.63(\epsilon - 1)(W/h)^{0.1255}}}, \qquad \frac{W}{h} > 0.6 \\ &\approx \frac{1}{\sqrt{1 + 0.6(\epsilon - 1)(W/h)^{0.0297}}}, \qquad \frac{W}{h} < 0.6. \end{aligned} \tag{16.2}$$

If the trace thickness t is nonzero, we can adjust the width to take this into account, replacing W with W_{eff}, where

$$W_{\text{eff}} = W + \frac{t}{\pi}\left[\ln\frac{2h}{t} + 1\right]. \tag{16.3}$$

These formulas look ugly, but they're pretty accurate; anyway, since you've probably only got three different widths of trace on your board, calculating Z_0 and c' isn't onerous. You can ask your board vendor the width of a 50-Ω line in your geometry, and what their guaranteed impedance tolerance is.

Because of the asymmetry of microstrip line, its impedance and propagation speed are both frequency dependent, even in the absence of material dispersion. The low frequency limit resembles a TEM mode with $\epsilon_{\text{eff}} \approx (\epsilon_r + 1)/2$, whereas in the high frequency limit almost all the field is inside the dielectric, and $\epsilon_{\text{eff}} \approx \epsilon_r$.

Aside: Hook. Ordinary FR-4 or G-10 board is a reasonable high frequency dielectric; however, its dielectric properties are too dispersive for building really good wideband amplifiers. This effect is known as *hook*, due to the shape of the ϵ versus f curve. If you need really wide bandwidth, say, DC-1 GHz with good time-domain response, choose a different material such as alumina or Teflon-glass. Microstrip is slightly dispersive above 1 GHz anyway, but is quite acceptable on a good substrate. Boards are available with ϵ from 2.2 to 10.

16.2.3 Inductance and Capacitance of Traces

To calculate the inductance and capacitance of traces, we use the equation for the impedance of a mismatched transmission line. For lines whose electrical length $\theta \ll 2\pi$, we can take the low frequency limit of (14.13) using a shorted line to get the inductance and an open-circuited one for the capacitance. We find that a length x of line of propagation velocity c' and characteristic impedance Z_0 has low frequency inductance and capacitance of

$$C = \frac{x}{Z_0 c'} \quad \text{and} \quad L = \frac{xZ_0}{c'}. \tag{16.4}$$

Note that this assumes that x is much longer than W and h; otherwise the strays at the ends of the line become dominant. Thus we can't use this simple idea for pads, only traces. Also, the temperature coefficient of the dielectric constant of FR-4 board is almost 1000 ppm/K, which is really inconvenient sometimes.

Example 16.1: Inductance and Capacitance of PC Traces. Consider the PC board trace of Figure 16.2, which is built on 0.33 mm thick FR-4 board with a ground plane. The line

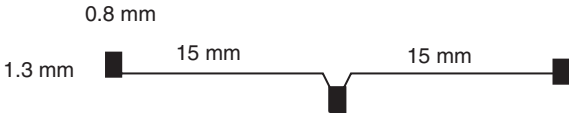

Figure 16.2. PC trace.

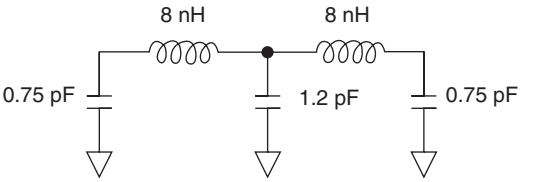

Figure 16.3. Circuit model.

is a 7 mil (0.18 mm) wide trace of 1 oz copper, which has $Z_0 = 91.4\Omega$ and $c' = 0.57c$. Its total length is 3 cm, and it has a pad measuring 0.8×1.3 mm at each end and another one in the middle. The parallel-plate capacitor formula isn't adequate for the pads, so we have to do a few measurements of the actual capacitance, and we find that it's 0.2 pF per pad. By (16.4), the two bits of trace look like 8.0 nH of inductance and 0.95 pF of capacitance. This model is only valid far below resonance, so to leading order it doesn't matter where we put them in the model; since the circuit is symmetrical, we'll add them symmetrically, so our final low frequency model is as shown in Figure 16.3.

16.2.4 Stray Inductance

We saw in Section 14.3.7 that a straight b centimeter length of round wire of diameter d centimeters in free space has an inductance of about

$$L = (2 \text{ nH/cm}) \cdot b \ln\left(\frac{2b}{d}\right). \tag{16.5}$$

This is a special problem in low impedance circuits (e.g., 1 inch of #22 wire (22 nH) will have a reactance of 50 Ω at 360 MHz) and in power supply wiring, where the supply lead inductance often resonates with the bypass capacitors, making a mockery of our supposedly well-decoupled supply. Inductance is reduced by shortening the wire and by running it closer to the ground plane.

A single-turn circular ring of diameter D made of the same wire has an inductance of[†]

$$L = (6.3 \text{ nH/cm}) \cdot D\left(\ln\frac{D}{d} + 0.08\right). \tag{16.6}$$

Thus a 1 meter diameter loop made of the shield of RG-58A/U coax (5 mm diameter) has an inductance of only 3.3 μH, an interesting fact that we'll come back to in Section 16.5.2 in connection with ground loops.

[†]Bureau of Standards Circular No. C74, US Government Printing Office, Washington, DC 20402 (1924).

16.2.5 Stray Capacitance

Stray capacitance comes from two places: the self-capacitance of objects and the mutual capacitance between the object and other objects, mostly ground. For small objects, the self-capacitance becomes important, because it goes as the radius whereas mutual capacitance goes as the area.

A large pad over a ground plane with a thin layer of dielectric constant ϵ has roughly the parallel-plate capacitance,

$$C_{\text{pp}} = (1.12\ \text{pF/cm})\frac{\epsilon_r A}{4\pi d}. \tag{16.7}$$

This formula considers only the region between the plates, and neglects fringing fields. It thus gives an underestimate, which becomes progressively more serious as the pad gets smaller; for a 1 mm square pad and 0.3 mm G-10 dielectric (typical of four-layer surface mount boards), it's too low by a factor of 2, which is somewhat surprising initially.

The explanation is the image effect from undergraduate E&M: a charge q above a conducting plane effectively generates an image charge $-q$ the same distance below the ground plane (Figure 16.4). Thus that 1 mm pad and ground plane have the same field configuration as two 1 mm pads, separated by 0.6 mm (now the pad looks a lot less infinite), and a major fraction of the fields are outside the pad area. The fringing capacitance will tend to go as the perimeter of the pad, whereas the parallel-plate capacitance goes as the area.

As a corollary, the area of ground plane contributing significantly to the pad capacitance extends about half a dielectric thickness past the edge of the pad, so that if you need to chop a hole in the ground plane to accommodate a high impedance, high frequency node, you need to make the hole a bit bigger than the pad.

You can make a guess that, on a surface mount board, the stray capacitance is about 0.15 pF times the number of components attached to that node. It's somewhat worse in through-hole boards. The other thing to remember is the self-capacitance of large objects; a 1 cm radius sphere has a self-capacitance of 1.12 pF, besides mutual capacitance with everything else. Fast, high impedance circuits have to be kept physically very small.

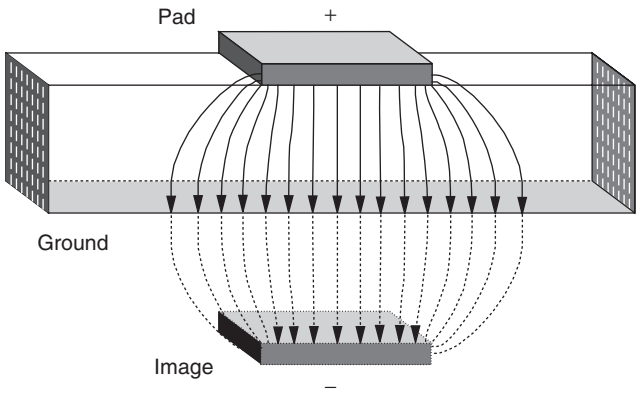

Figure 16.4. Electrostatic field lines for a pad over a ground plane.

16.2.6 Measuring Capacitance

It is possible to measure the capacitance of a 1 mm pad, but it isn't too easy. A better method is to make a 100× scale model, whose capacitance will be 100× larger, and measure that instead. At 10 cm on a side, it's easy to handle, and its capacitance will be a convenient 20 pF or so, rather than 0.2 pF. You can build it out of G-10 board or heavy-gauge aluminum foil and cardboard. Use air for the dielectric and multiply the measured capacitance by ϵ, or for more accurate results, dunk the model in oil of known ϵ to represent the dielectric.

This method is accurate and very general—you can use it to find the capacitance between some horrible stack of pads and vias and the ground plane, which is difficult with any other method. It takes a bit of work, but you only have to do it once for each pad stack design, and usually a judicious use of scissors can turn one model into another pretty fast. Alternatively, there are free programs available on the Web to automate this, but it's much more work than a quick measurement. Whichever way you do it, you'll rapidly discover that the parallel-plate capacitance is a wild underestimate if the separation is anything like the pad dimensions.

16.3 STRAY COUPLING

16.3.1 Capacitive Coupling

As we saw in Section 14.4.5, the fields of a balanced line fall off more rapidly laterally than those of a single wire, and the scale of the falloff is the distance between the conductors. Because of the image effect, wires near a ground plane have the same lateral confinement of the fields, and bringing them closer to the ground plane increases the confinement.

Thus two traces above a ground plane have much less mutual capacitance than the same traces alone. (That's one of the ways that ground planes help.)

16.3.2 Inductive Coupling

The mutual inductance of two identical microstrip lines whose self-inductance is L, separated by a distance s, is roughly[†]

$$M = L\left(\frac{1}{1 + s^2/h^2}\right). \tag{16.8}$$

16.3.3 Transmission Line Coupling

Two transmission lines in close proximity will couple together due to the overlap of their field patterns, in just the way we saw in Section 8.7.1. Transmission line coupling is only distinguishable from the other two when you've got a long run of line, but if you have, terminating the line on each end with Z_0 will send the forward and reverse wave coupling to the front and back end termination (this is how directional couplers work).

[†]See, for example, Johnson and Graham.

16.3.4 Telling Them Apart

Inductive coupling depends on the direction of the current flow, whereas capacitive coupling depends only on the voltage swing. Weak stray capacitive coupling is high-Z and moves the whole trace up and down together; weak inductive is low-Z and moves one end of the trace with respect to the other. Thus one way of telling whether your coupling is capacitive or inductive is to open-circuit the far end of the lead; inductive coupling will go away, and capacitive coupling won't. Similarly, short-circuiting one end to ground will get rid of the capacitive coupling but preserve the inductive.

If you have a combination of inductive and capacitive coupling, you can sort them out by changing the direction of the source current; the inductive coupling will change sign whereas the capacitive won't.

16.4 GROUND PLANE CONSTRUCTION

16.4.1 Ground Currents

Kirchhoff's current law (aka charge conservation) says that the net current flowing into any node of the circuit is always 0, so currents flow in loops; for every current flowing in a wire, there's an equal and opposite current flowing back again somewhere else. Though elementary, this is often overlooked.

Most of the currents we think about flow in wires, so that their routes are predefined. Where there are multiple routes (e.g., parallelled wires or a ground plane), however, we need to be a bit more careful. Whenever there is more than one path, the current divides itself according to the admittance of each path; at DC, this means that the current through each path is proportional to its DC conductance G. At AC, the resistance of the path rises slowly with frequency (as $\sqrt{f}$) due to skin effect, but the reactance rises linearly, since $X_L = j\omega L$. Accordingly, inductive effects become dominant at higher frequencies.

The inductance goes up as the loop gets bigger, so as f rises, the current tends more and more to return directly underneath the trace. In building circuits over ground planes, we must be careful not to obstruct these ground currents by cutting holes in the ground plane for other conductors. Even a narrow gap forces the current to go around, which adds significant inductance to the ground path, as well as leading to EMC vulnerability (see below).

> I once turned a beautiful 2 GHz amplifier into a 400 MHz bookend by deleting the ground plane and thereby effectively placing large inductors in the circuit.
>
> —Steve Roach[†]

16.4.2 Ground Planes

A ground plane provides very low impedance paths from everywhere to everywhere else, so that different parts of the circuit agree pretty well on just what ground is, and (through field confinement) provides isolation between different parts of the circuit.

Its main drawback is increased stray capacitance to ground. In most places in a circuit, this isn't too horrible if you plan for it.

[†]Steve Roach, Signal conditioning in oscilloscopes and the spirit of invention, in Jim Williams, ed., *The Art and Science of Analog Circuit Design*. Butterworth-Heinemann, Boston, 1995.

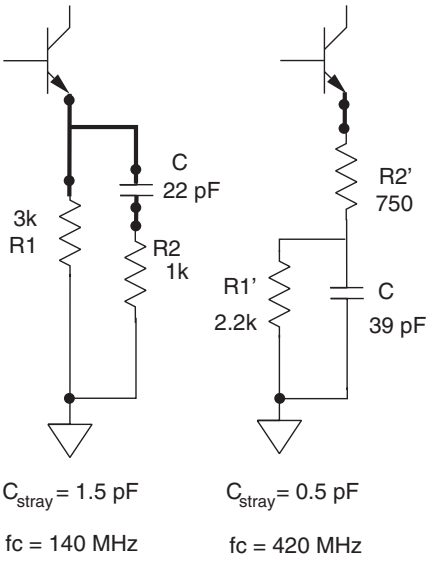

Figure 16.5. Minimizing capacitive loading by minor circuit changes (bandwidth-limiting meshes shown in bold lines).

Example 16.2: Stray Capacitance in an Emitter Load. The common-emitter amplifier of Figure 16.5 has a series *RC* circuit from emitter to ground, to set the AC gain of the stage. With the capacitor on top, in parallel with the bias resistor, there are five pads and three traces to provide capacitance from *E* to ground; with 2.5 mm of 7 mil width per trace (0.33 mm dielectric) that's about 0.16 pF per trace, for a total of 1.5 pF. If we split the bias resistor with the bypass instead, most of this appears harmlessly in parallel with *C*, so there are only two pads and one very short trace, about 0.5 pF. The common-emitter stage will exhibit a rising gain with a 3 dB point at 140 MHz in the first case, but 420 MHz in the second.

16.4.3 Relieving the Ground Plane

Sometimes, though, you can't stand that extra 0.5 pF (e.g., at the collector of a cascoded photodiode amplifier; see Section 18.4.4), and you have to chop small holes in the ground plane under a few pads. This is okay as long as you don't make ground currents detour even slightly.[†] You have to make the hole about one dielectric thickness bigger than the pad because of fringing. If the ground plane is near one side of the board, you can often do additional relieving with a milling machine or Dremel tool; failing that, lift the pads and run the offending circuitry in mid-air, dead bug fashion.

16.4.4 Skin Depth

Down at DC, current in a long piece of ordinary imperfect conductor flows uniformly throughout its cross section. At AC, the changing current produces an AC *B* field, which

[†]In extreme cases, you can bridge the hole with jumpers, for example, 0 ohm resistors.

by Faraday's equation,[†]

$$\nabla \times \mathbf{E} = -\frac{1}{c}\frac{\partial \mathbf{B}}{\partial t}, \tag{16.9}$$

produces an AC E field inside the conductor. Since it *is* inside a conductor, this E gets mostly shorted out, preventing B from changing much, which in turn suppresses the AC current in the bulk of the conductor. The result of all this is that the current flow is progressively confined to the surface of the conductor as f increases, and distributed so that $B = 0$ deep inside (for a round wire, that means that the current density J is independent of angle). The finite conductivity of real wire means that these effects occur over a finite thickness of the conductor; at lower frequencies, $\partial B/\partial t$ is smaller, so the induced E is smaller. In circuit terms, this means that the source impedance of the induced E is smaller at low frequencies, and so it takes more conductor thickness to short it out effectively. Since these effects are linear and homogeneous, we expect some sort of exponential falloff of J with depth, and that s in fact what we get.

An infinite half-space ($z > 0$) of imperfect conductor has a current density J that falls off with depth as $\exp(-z/\delta)$, where δ is the skin depth,

$$\delta = \sqrt{\frac{\rho}{4 \times 10^{-9}\pi^2 f \mu}} \tag{16.10}$$

(in centimeters) and ρ and μ are as usual the electrical resistivity ($\Omega\cdot$ cm) and relative magnetic permeability. This applies both to conduction currents from applied voltages and to eddy currents caused by ambient magnetic fields.

"Sufficiently highly conducting," in the section just above, means that the skin depth in the ground plane is much less than the thickness of the ground plane or of the dielectric. That points to where ground plane construction breaks down—low frequency magnetic coupling. For copper ($\mu = 1$, $\rho = 1.7\mu\Omega\cdot$ cm), δ is 7 μm at 100 MHz but 0.85 cm at 60 Hz.

The presence of a shield will strongly suppress the B field nearby, but even neglecting that, a shield 7δ thick will give >60 dB isolation, and 10δ will give >85 dB. Thus a double layer of 2 oz copperclad board (0.14 mm total copper) will give >60 dB isolation anywhere above 10 MHz, but progressively less below there; that isn't too much of a problem, though, since below 10 MHz we're usually using magnetically shielded coils anyway, to get high enough inductance without massive copper losses, and that together with the lower frequency reduces $\partial \mathbf{B}/\partial t$ substantially in the surrounding circuitry. As a practical matter, you'll probably never see an inductive coupling problem in a circuit above 1 MHz if you use tall, continuous copper-clad shields; if you're worried about it, use ordinary sheet steel instead (note the factor of μ in (16.10)).

Aside: Split Grounds and Electromagnetic Compatibility. If you physically split your ground plane into "analog" and "digital" regions, you've made what is called a "patch antenna." It's really amazing how well structures like that couple electromagnetic interference (EMI) in and out. Commercial products have to pass stringent electromagnetic compatibility (EMC) tests, and split ground planes are one excellent way to make yours flunk. Do the layout carefully on a single featureless ground plane—vias and a few very small cutouts to relieve a pad are OK, but no big holes and *no slots*.

[†]Okay, we do need one Maxwell equation in this book, but we got pretty far without it.

16.5 TECHNICAL NOISE AND INTERFERENCE

16.5.1 What *Is* Ground, Anyway?

The first *koan* of prototyping is: When is ground not ground? (Appropriate Zen answer: always and never.) Ground is a very useful approximation, but that's all. All conductors have resistance and inductance, so as soon as your circuit gets turned on and current starts to flow, all your grounds will be at different voltages. This isn't necessarily obnoxious, but sometimes it is; you have to watch, or those voltage drops in your grounds will bite you. Typical problems include ground loops at low frequency and ground pigtails at high frequency.

16.5.2 Ground Loops

Ambient magnetic fields produce voltage swings in conductors, and multiply-connected grounds (*ground loops*) will allow them to produce large ground currents. Remember that 1 m loop of RG-58A/U shield (3.3 μH) in Section 16.2.4 that we were going to come back to? With a 1 milligauss field at 60 Hz, it produces 30 μV of signal with a reactance of 0.0012 Ω; the shield has a resistance of around 0.02 Ω. This impedance is lower than almost anything—it's the impedance of a big car battery. If the coax had been RG-214 (half-inch coax with two layers of braid), it would be a dozen car batteries in parallel.

Low voltage plus extremely low source impedance make ground loops do unintuitive things, and this is what makes them resemble gremlins—we don't usually expect to encounter a millivolt-level signal with a circulating current of hundreds of milliamps. Once you realize that the source impedance of the ground loop can be smaller than the contact resistance of the coax jack, it becomes obvious why wiggling the cable sometimes causes the hum to change by 10 dB.

This low impedance can easily cause a voltage drop across a chassis or ground plane, which is one source of the trouble. Even more important, magnetic coupling will induce the same voltage in the shield and the center conductor; shorting out the induced voltage in the shield while leaving the center conductor alone turns the common-mode swing into a differential swing that the amplifier sees as signal. Fortunately, it isn't hard to fix once you know how it occurs; just putting a 1 ohm resistor in series with the shield will attenuate ground currents by 30–60 dB, and adding a differential receiver will help even more.

Still, we shouldn't create ground loops where we don't have to. Plug all your instruments into the same wall circuit where possible, and use tape to insulate coaxial connectors that touch each other or the optical table (the author has seen cases where a file card slid between a patch cord and the optical table made a 60 dB improvement). Grounds in house wiring and water pipes can easily have 10 m loops.

Example 16.3: Ground Loop in an AC Amplifier. Consider the circuit of Figure 16.6, a 15 cm long by 5 cm wide, 1-ounce dead bug prototype with a 1 m loop of RG-58A/U attached. A weak stray field of 1 mG at 60 Hz will induce a voltage of 30 μV, with a $20 + j1.2$ mΩ source impedance. The ground plane has an end-to-end resistance of 3.0 mΩ, so 13% of the open-circuit voltage is dropped across the length of the ground plane. The center conductor will have the same open-circuit induced voltage, of course, but will not be shorted out by the ground plane and will not exhibit the large gradient

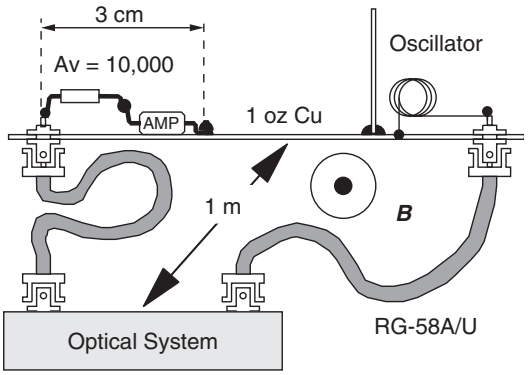

Figure 16.6. Ground loop example.

across the circuit. If the AC amplifier has a gain of 10^4, and its ground connection is 3 cm away from the input connector, it will see a differential signal of $30\mu V \cdot (1 - 0.13 \cdot 12cm/15cm) = 27\mu V$, and produce 270 mV of hum at its output.

16.5.3 Floating Transducers

It's a really bad idea to put transducers (e.g., photodiodes) on cables, but sometimes there's no alternative. If one end is going to be grounded, all is easy; connect that end to the shield, and ground it at the receiver end of the cable. Sometimes, though, both ends have to be floating, and that's more difficult. The key trick in this case is to use shielded twisted pair and ground the cable shield *at the transducer*.

The reason is that, if the shield is bouncing up and down with respect to the transducer's idea of ground, capacitive currents will flow between the shield and the wires. The cable won't be exactly perfect, so they won't be quite balanced, and a differential mode signal will result. If you ground the shield at the transducer end, there is no capacitive current and hence no interference. This remains true if the transducer has a differential line driver attached to it instead of really being a floating device.

16.5.4 Mixed-Signal Boards

Mixed-signal and RF boards require special treatment, especially when high speed, high resolution A/D converters are involved. Keep the analog and digital traces far apart; fast logic slews at 1 V/ns or faster, which will couple spikes into everything, especially when several outputs are switching at once. Metal shields, continuously connected to the ground plane by solder or finger stock, can be a big help. (See Figure 16.7.)

It's often tempting to split the ground plane into analog and digital sections, the way many ADC application notes suggest. While this is one way to keep digital junk out of your analog section, it's such an EMC nightmare that you're way better off not doing it. Ground currents drop off fairly steeply from the lowest-inductance return path, so just keeping the two sections physically separate is the better way to go. If you have something really really horrible to deal with, such as the A/D clock, which can hardly be kept entirely out of the signal processing path, you can run that in coax with a separate ground.

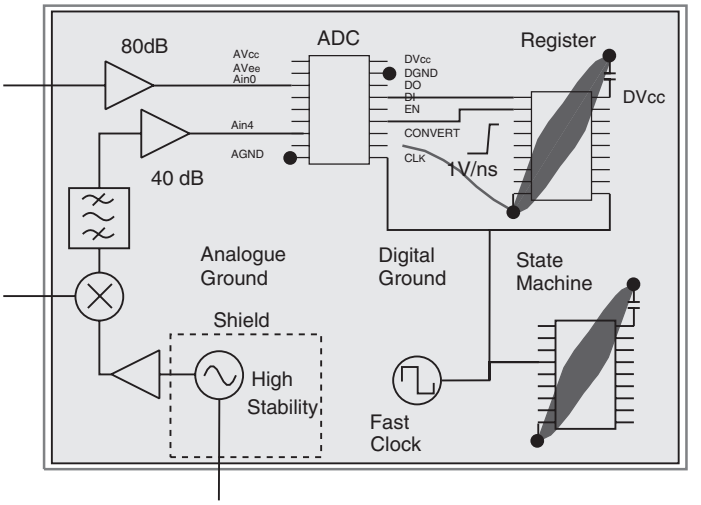

Figure 16.7. Connecting analog and digital grounds: use a single plane and separate the analog and digital sections so their ground currents don't interact.

Digital systems usually have separate power and ground planes, with densely sprinkled bypasses and the interplane capacitance to short them together at AC. Don't do this with your analog sections, unless the circuitry can be limited to the side nearest the ground plane; the supply plane always bounces a bit, and coupling between it and the traces running over it will cause you headaches (see Example 19.1). You can split the power plane, but it's usually best to have only ground planes in analog circuitry and run the power in traces.

A few things that can help quite a lot if you have analog and digital sections on the same card: put them on opposite sides of the board with two ground planes (analog and digital) in between and a shield over top, and don't use an external memory bus on your microcontroller. Serial I/O peripherals work very well for most things, and you don't have 30 wires jiggling up and down on every instruction cycle. This need not be a serious constraint, either; you can even get 64 megabit flash memories with serial interfaces.

16.5.5 High Impedance Nodes and Layout

Getting the layout of high-Z circuitry right is extremely important. Pickup is so insidious that you'll spend a long time finding it if you're not very thoughtful about layout. Life has gotten easier in one way, because SMT components pack more tightly and hence have less trace length and loop area to pick up junk, but at the same time, more circuitry is packed into a smaller space, so that the sources of pickup are that much closer. SMT isn't always a win; see Example 19.1.

16.5.6 Connecting Coaxial Cables

The ground loop example shows that connecting cables isn't as simple as it looks. One good way to eliminate ground loop problems is shown in Figure 16.8: ground the shield through a 1 Ω resistor, and use a differential amplifier or transformer to restore a local ground reference.

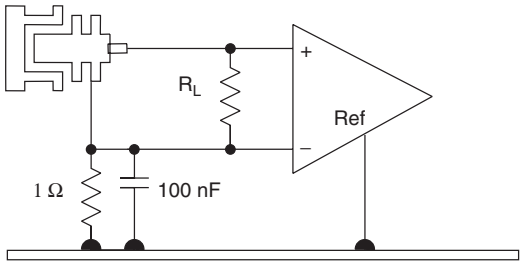

Figure 16.8. Isolating ground loops.

A common-mode choke, made by wrapping the coax many times around a low frequency toroid core (even a link of heavy steel chain works), is another possibility in bench setups, but expensive and inconvenient in commercial instruments. You need at least 5 mH or so to help very much.

16.5.7 Bypassing and Ground/Supply Inductance

Power supplies have to be considered as parallel resonant circuits.[†] The resonances are of several kinds, but the ones that cause problems are the parallel ones, because the supply impedance becomes high at resonance. These have the common property that what causes trouble is overkill—too much bypass capacitance and not enough resistance for effective damping of the resonance, resulting in a high-Q circuit and excessive ringing. At high frequency, supply resistance is your friend, not your enemy.

Example 16.4: Supply Bus. Consider a 150 mm long supply bus made up of a 6 mm wide copper trace, 1.5 mm above the ground plane, with 0.1 μF capacitors to ground every 25 mm. The impedance is about 30 Ω, and the velocity factor is 0.52. Figure 16.9 is a SPICE simulation of the voltage at the middle and one end of the trace, for unit AC current at the midpoint. The bypassing rapidly gets worse as ESR drops below 0.3 Ω—that is, as the capacitors get better.

The moral of the story is to be content with good enough bypassing, and to remember that killing the Q of a resonance requires putting in some loss: a series resistor, a lossy ferrite bead, or a lossy capacitor such as an aluminum electrolytic. Use a 0.1 μF monolithic ceramic capacitor on each supply pin of each IC, and put in a 1 μF solid tantalum on each supply for every 8 or 10 chips using it. When the layout comes back from the draughtsman, *then* you can eliminate the unnecessary ones. If you have ground and supply planes, the board itself will provide some nanofarads of very quick capacitance, so the need for individual bypasses is reduced.

We often need to use the same supplies for analog and digital circuits, or for all sections of a high gain or high isolation circuit. Use series resistors and ferrite beads in the supply traces between sections for high frequency isolation, and capacitance multipliers for low frequency (Example 14.1). Use a low frequency and high frequency bead together (Ferroxcube 3E2A and 4B ferrites work well), and be ready to use two or three sections if necessary.

[†]Paul Brokaw, *An IC Amplifier User's Guide to Decoupling, Grounding and Making Things Go Right for a Change*, Application Note AN-202, Analog Devices.

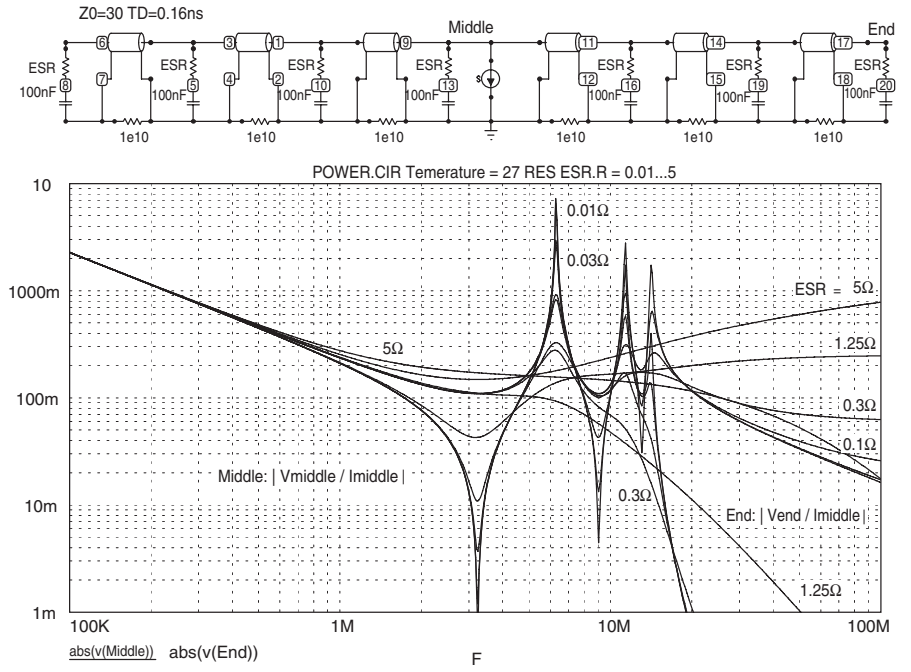

Figure 16.9. Resonances in a well-bypassed supply bus. Vertical scale is impedance in ohms.

16.5.8 Bypass Capacitor Self-Resonances

Bypass capacitors have a series-type self-resonance, with f_{SR} somewhere in the VHF or UHF range. From what we've seen of supply resonances, that sounds bad, but (surprisingly) it is entirely benign unless you make the lead inductance excessive. The reason for this is that a well-bypassed supply is one with very low source impedance at all frequencies, but just what the phase angle of that impedance is matters very little. The ESL and wiring inductance will eventually dominate the impedance of the bypass, but remember that parallel and series LC impedances are symmetrical functions of $\log(f)$—if your bypass is effective at $f_{SR}/10$, its impedance will have just the same magnitude at 10 f_{SR}, so it will be just as good a bypass there. It's parallel-type resonances you have to worry about, and you can fix those by spoiling the Q with a series resistor or lossy capacitor.

16.6 PRODUCT CONSTRUCTION

We'll only sketch the most important things to think about in constructing products, because by the time your prototype instrument works, you're over the hump.

16.6.1 Cost Versus Performance

The cost of many purely electronic systems is dominated by boards, cables, and connectors. A careful trade-off is required between flexibility, reusability, manufacturing yield, and per-piece cost.

16.6.2 Chassis Grounds

Ground the power supply common lead to the chassis only at the backplane or system board. When grounding something to a chassis, case, or safety ground lead, always use a star lock washer or a lug with a lock washer pattern. Assuming the plating of the washer and the chassis are compatible, this makes a dozen or so gas-tight contacts between the screw, chassis, and lug. This doesn't work as well with PC boards, because the board flows slightly at elevated temperatures (the glass transition temperature of G-10 board is about 125 °C), which loosens the connections. A combination of a star washer with a spring washer such as a Belleville may be best for this, if your boards will be used or stored at high temperatures.

16.6.3 Magnetic Shielding

We saw that copper shields become ineffective at low frequencies, so that we'd need several centimeters of copper to shield 60 Hz really effectively. Equation (16.10) for the skin depth contains a factor of μ in the denominator, so a magnetic conductor is a better shield than a nonmagnetic one; from the skin depth alone, a material such as super-permalloy, with μ up near 10^6, reduces the skin depth by a factor of 10^3 for the same conductivity, and mu metal ($\mu \approx 4 \times 10^4$) can manage a factor of 200.

High-μ materials can shield magnetic fields just the way a Faraday cage shields electrostatic ones. If the ambient fields are above an oersted or two, your high-μ shield is liable to saturate and become transparent; a mu-metal enclosure in a high field environment will need a steel shield around it. This is particularly true when the shield is held together with fasteners; the field concentrates in the regions of good contact and causes early saturation. High-μ materials are vulnerable to shock and bending, which pin the magnetic domain boundaries; although mu-metal comes in sheets, as soon as you bend it into a box and spot-weld the edges, it isn't mu metal anymore and has to be carefully annealed once again.

All in all, low frequency magnetic shielding is a lot harder than just soldering together some chunks of FR-4 board. Most of the time, you're much better off moving the offending transformer or switcher away from the sensitive circuitry and reducing its strays, either by using a wall wart power supply or by going to toroidal power transformers, using twisted pair wiring (which cancels the magnetically induced voltage on every twist), breaking ground loops, and doing everything you can to reduce the area of the loops collecting the stray flux. Small loop area, differential amplifiers, and wide separation of source and receiver are your best defenses, with steel and mu metal as the strategic reserve.

16.6.4 PC Boards

Circuit boards have to be debuggable, so you have to be able to get at the wiring. In ground plane boards, running signals on the top and bottom facilitates ECs[†]—you just cut the trace with a scalpel, instead of having to drill through it. One thing to remember is that the placement of the ground plane matters. Sensitive circuits can easily be driven nuts by a single pad lying over a noisy supply plane, for instance (see Example 19.1),

[†]Engineering changes (*i.e.* hacks), usually involving adding *roach wires*.

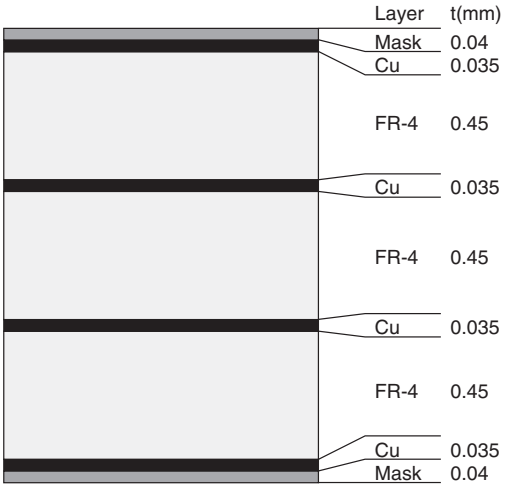

Figure 16.10. Cross section of a $\frac{1}{16}$ inch (1.57 mm) four-layer board.

whereas ground planes are really quiet. Figure 16.10 shows a typical four-layer board; you can get almost 1 mm thickness if you put signal wiring on the top, nothing on layer 2, ground on layer 3, and noncritical traces on layer 4. PC board houses can make the layers asymmetric in thickness, which is worth considering in some cases. If you have more than one ground plane, stitch them together with lots of vias. (There is some difference of opinion about this; if your signals are all at high frequency, the capacitance between the planes may be enough that you don't need so many vias.)

There are a number of quick-turnaround PC board houses with decent Web interfaces and free design tools, so you can lay your board out, hit the button, and have boards in a few days for a very reasonable price.

16.6.5 Design for Test

Before your board goes to layout, use the knowledge you gained in debugging the prototype to add test points and places to break the circuit (e.g., jumpers on power supply leads). They should be pads with via holes where possible, so that you can solder a pin to grab with your oscilloscope probe. Vias will stand a lot more soldering and messing about than an ordinary pad, which matters for test points, and your scope probe will put a lot of torque on the pin. Each test point should have a grounded pad and via within half an inch or so, because that probe needs a good ground if you want it to tell you the truth.

16.6.6 Connectors and Switches

Connectors, cables, pots, and switches are the main sources of flakiness in electronic systems. Choose them carefully; they're expensive to buy but even more expensive when they fail in the field. Don't run signal lines through panel mounted switches; to avoid pickup, use diodes, relays, or analog switches to switch signals, and run only control currents through the switches. Since most switches are especially unreliable at very low currents, this helps circuit reliability too. As for a panel pot, at least one of its leads

will have a high impedance to ground, and that invites capacitive pickup. In both cases, control lines can be heavily bypassed, but signal lines can't.

Use rugged BNC and SMA connectors for panel mount applications, but stick to the cheaper and smaller SMB or SMC connectors inside the instrument.

16.6.7 Multicard Systems

Multiboard systems have the advantage of flexibility, but are at an increasingly severe cost disadvantage, and are nowadays less necessary since circuit density has increased, logic parts are now programmable, and most of the display and control tasks are done on a computer. A 50,000 gate FPGA chip costs the same as one 200-position backplane connector. Consider using a single-board design, perhaps with a daughter card if necessary, or *population options*: parts of the circuit that can be omitted by simply not stuffing that part of the board. If you really need multiple cards, check out the Eurocard system. It's a good all-round way to assemble a system using multiple boards, because it combines widely sourced hardware with a wide choice of single contacts, high current points, and coax connectors. Card edge connectors are now nearly obsolete, except in computer plug-in cards such as PCI Express, because they lack this flexibility.

16.6.8 Computer Plug-in Cards

Be very careful if you're building hardware to plug into a PC. The external interfaces (serial and parallel ports, USB, Ethernet, GPIB, and so on) are pretty robust, but if you plug a wrongly wired board into a PCI slot, you can blow up your motherboard. Be very sure that your circuit is correctly wired, and that the bidirectional bus is working properly, before you plug it in, and don't test it on your development machine. If you can possibly avoid a proprietary plug-in card, avoid it.

16.7 GETTING READY

16.7.1 Buy a Stock of Parts

You should have a stock of at least 50 pieces each of 1% metal film resistors, in values spaced approximately 10% apart. Get the $\frac{1}{4}$ watt type with leads. The mil-spec goes only up to 301 kΩ, but most manufacturers make them up to 1 MΩ in the commercial grade, and you can get 10 MΩ in the half-watt size. Nowadays, your production boards will be surface mount, but leaded components are indispensable in prototyping. Make sure you get some values such that $\frac{1}{2}$, $\frac{1}{3}$, $\frac{1}{4}$, $\frac{1}{5}$, $\frac{1}{9}$, and $\frac{1}{10}$ of them are in the set too, for example 10.0k, 4.99k, 3.32k, 2.49k, 2.00k, 1.11k, and 1.00k, because that helps a lot in making amplifiers of integer-valued gains. Murata and other manufacturers make very nice collections of surface mount resistors, capacitors, inductors, and trimmers (both pots and capacitors). For RF work, and for stuffing your first few engineering boards, you should have a set of these kits too. It isn't a huge investment, and it'll make you fast.

16.7.2 Get the Right Equipment

You need at least a good fast analog scope, access to a digital one, plus some good way of measuring noise. The absolute minimum acceptable bandwidth is three times your maximum operating frequency; since a scope is a time-domain instrument, you need to

be able to see the shape of your waveform, and that means reproducing the fundamental and enough of its harmonics. The first, second, and third harmonic will allow you to see asymmetry and peaking or flattening of the tops of the waveform. You need at least the first 10 harmonics to be really sure you know what your signal looks like, for example, a 1 GHz vertical bandwidth for a 100 MHz amplifier.

The best way to measure noise is with a really well-calibrated spectrum analyzer[†] and low noise input amp, but setups like that don't grow on trees; for work up to 10 MHz, a filter whose noise bandwidth is accurately known[‡] plus a HP 3400A RMS voltmeter is a good substitute. (You can also use peak-reading meters such as the HP 400EL, but remember to add 1 dB to the RMS reading to get the noise power accurately—see Section 13.6.5.) For high frequency work, you also need access to a spectrum analyzer capable of detecting spurious oscillations up to twice the maximum f_T of your fastest component. Everything in the world resonates between 1 and 5 GHz, so if your circuit might have gain that far out, look for oscillations there, and look *hard*.

16.7.3 Soldering

Most soldering problems are caused by overcooking or undercooking. Surprisingly, both are caused by an iron that's too cold or too dirty. It may be a bit counterintuitive that a hot iron keeps components from overheating, but consider soldering with an iron at $T_{MP} + 1^{\circ}$: the parts would stay very hot for a long time before the solder melted. The amount of heat dumped into the package through the lead goes as the time integral of the temperature gradient; somewhat hotter irons used for a much shorter time win.

Accordingly, get a decent 60 W, temperature-controlled iron. The right temperature is about 750 °F (400 °C). Much cooler and you cook parts; much hotter and you can't keep the tip clean. Use a fairly fat conical tip (about 2 mm diameter, e.g., Weller PTA-7) for general purposes, and a skinnier one (1 mm, Weller PTK-7) for SMT prototyping. Stick with the conical ones and avoid the ones that neck down to a narrow cylinder; they look convenient but actually cook parts because they can't heat the leads fast enough. For RF boards, a large (100–150 W) soldering iron is very helpful for attaching the shields to the ground plane. You'll need a soldering station with a wet sponge and an iron holder. Get a good quality solder sucker (e.g., a Soldapullt). (Don't worry about the conductive plastic kind, they're not nearly as good as the plain Teflon tips.) Learn to use it by touch and with your less-skillful hand (cock it on the bench or on your knee). Get some solder wick for surface mount parts, and if you're going to be doing a lot of work with SMT ICs, get a SMT desoldering station, which will help keep the pads from being torn or cooked off. For those with bigger budgets, metals are very nice.

16.8 PROTOTYPING

16.8.1 Dead Bug Method

Building prototypes of high performance circuits needs a low capacitance, low inductance, flexible prototyping system: the *dead bug method*. You fit a piece of copperclad board into the lid of a suitable die-cast aluminum box, as shown in Figure 16.11, and use

[†]Emphasis on *well calibrated*. Be suspicious.
[‡]For example, from swept measurements or a first-principles calibration, as in Section 2.5.4

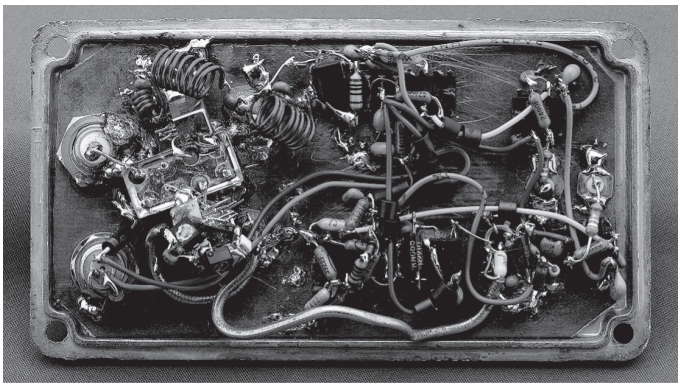

Figure 16.11. Dead bug prototype of a dual channel, linear phase bandpass filter and detector for modulation-generated carrier interferometry. A surface mount prototype appears in Figure 16.12. (Photo by Frank Giordano; a full-resolution color version is available at http://electrooptical.net/www/beos2e/photos.html.)

cyanoacrylate glue to attach ICs and transistors to the ground plane. These work like terminal strips, so you can hang your passive components right on the pins.

It's often ugly, but it's fast, it's easily modified, and it works up to 500 MHz or higher if you're careful.

Die-cast boxes are easy to drill and make good shields when you finish the prototype and attach the bottom. BNC bulkhead mount connectors hold the board into the box, and with the aluminum, they provide a very low impedance ground to minimize ground loop voltage drops (see Section 16.5.2).

Some people put the ICs upside down, some bend the leads and mount them right side up, depending on whether they have more trouble with breaking leads or forgetting which IC is which. If you go upside down, mark the package on the notched end, so you know which is pin 1. Try to keep all the chips going the same way, so you don't get confused.

Dead bug prototyping is easiest by far with DIP packages, and you should use those when possible. Unfortunately, they're starting to disappear, so we have to get good at prototyping with SMTs. SMT transistors are easy enough to get wires on, but their leads are so delicate that you have to be careful. As for SOICs, you can do the 0.050 inch ones by staggering the leads like saw teeth, but the 0.025 inch ones you can just forget. Buy or make an adapter to spread the leads to some useful spacing.[†]

It is possible to use SMT chip resistors and capacitors in dead bug prototypes, but it's fiddly, so do it only where really necessary; leaded components are so much easier to use that you'll be a factor of 10 faster with them.[‡]

16.8.2 Laying Out the Prototype

A common error is to run out of room. If this happens at one edge of the board, it isn't too serious; bolt another piece of board on that side, run some copper tape all along the seam, and solder it along both sides so that the ground plane is continuous.

[†]One good source of adapters is the Solder-Mount System, from Wainwright, Andechs, Germany.

[‡]There's not enough space here to go into all the details of dead bug prototyping high speed circuits. For more, have a look at the classic (and very large) Linear Technology Application Note AN47 by Jim Williams.

If it happens in the middle, where you perhaps were a bit optimistic about how much gain you could get out of one stage and now need two, it's much more awkward. So the best way to proceed is to build each functional block (an IF strip, say) on a small piece of board, and stick them together on a much larger sheet with screws and copper tape. Alternatively, you can just leave extra space in places where you're not positive it will work perfectly. Building prototypes that can be modified easily is vital—all prototypes eventually get to the point where they're too flaky to use, but good ones don't get there until you're finished with them.

If you're building something big that needs to be one piece (e.g., an RF back end), start with a 6×10 inch piece of copperclad board for a ground plane. You can attach smaller chunks of board to it as your prototype grows. Don't make the mistake of grounding the wrong side of the board, or you'll make very leaky shields.

Prototype your signal flow starting at one end and going to the other in a line. This keeps the input and output from coupling together and making the whole string oscillate. If your VHF circuit needs good isolation (≥ 60 dB) between any two places, keep them apart and use copperclad shield walls in between them. Use these shields to separate VHF or UHF gain stages into 20 dB chunks, and run the signal and power leads through small notches cut in the bottoms of the shields. Minimize loop sizes; a big loop will pick up RF on one side of the shield and reradiate it on the other side.

Think about ground currents; run especially noisy grounds right to the appropriate supply bypass cap if you can. If necessary, run two parallel wires a couple of diameters apart, which is nearly as good as foil because the **B** field between them tends to cancel. Remember that the inductance of a loop depends on its area; a big U made of wide foil is not necessarily lower in inductance than a shorter run of fine wire. Bridge any big loops.

16.8.3 Adding Components

The major skill in dead bug prototyping is making connections incrementally, in such a way that they don't come apart again when you remelt the solder to add the next part. You must avoid bending stresses in the component leads to prevent this.

To add a component (say, a resistor), gauge by eye how long the leads need to be. Cut and tin one lead. Using the other lead as a handle (bend it if that's more convenient), solder it to the desired node, keeping it as far as possible in its final position, to avoid bending stresses in the leads. Solder the other side. Now—and here's the vital step—*go back and reflow the joints, in the same order*. This releases the spring tension in the leads, and prevents their springing apart when you add another component. Now wiggle both ends, to make sure that you have everything properly connected, rather than two groups very near each other. Whatever you do, *don't* make the mistake of grouping all the leads to be joined, and then soldering them all at once. This will make a haywire prototype that will be hard to fix and hard to get working.

Dressing the leads short and close to the ground plane improves performance, reduces coupling, avoids shorts (since the component bodies wedge the leads apart), and generally makes it easy to build and modify. Most leaded components are amazingly strong—you can really crank on their leads if you need to, but really, really don't leave residual spring tension in them.

Wire the power supplies first, then the bypass caps on the ICs, then their supply connections. This puts the supply wiring on the bottom, which is good since it rarely needs changing, and gives the bypasses the shortest available path.

16.8.4 Hookup Wire

Most of the time, you'll use the clipped-off component leads littering the bench for hookup wire, and strips of copperclad board glued to the edges of the card as power supply bus bars. When you do need insulation, use #26 stranded, PVC jacketed hookup wire,[†] and use a consistent color code. Get the stuff with the really skinny jacket (OD is about 33.5 mils or 0.85 mm).

Strip only about one jacket diameter back from the end, and always twist the wire and tin it before bending the stripped part—otherwise strands will get loose. This is especially important if you like stripping the wire in several places to run daisy chain power connections. Always twist the wire at each stripped place and tin it.

Most people try using wire wrap wire for solder prototyping—usually only once. The trouble is that it's so fragile and the insulation melts too easily. The insulation is very useful by itself, though, for insulating things like thermocouples and thermistor leads.

16.8.5 Wire It Correctly and Check It

This is another wrinkle on the build-what-you-designed principle. Your gizmo has never worked yet, so there may be design errors in it, but wiring it wrongly will have you chasing your tail trying to figure out what's wrong. Orient the components so that you can read their markings, and always check the value or type number of each part as you install it; errors of factors of 10, or simple misfilings, are common, and lead to mystifying problems that take a long time to solve. SMT components are marked poorly or not at all, so pay special attention to getting them right, for example, by installing all the 100 kΩ resistors first, then all the 10 kΩ ones, and so on. Capacitors are usually unmarked.

Pay special attention to the wiring of power, ground, chip enables, and resets. On each chip, make sure you account for every last pin; if you wire only the ones you've thought about, there's sure to be a chip enable or latch input dangling somewhere. Use a yellow highlighter on a copy of the schematic to mark off the parts you've wired, so that if you're interrupted you'll know where to begin.

Building the whole prototype before turning it on is a mistake. You'll have to do the smoke test eventually,[‡] but test the circuit in small bits first: the voltage regulators, the first stage, the power output stage, and so on. Testing each one individually is very helpful in localizing faults—it's probably in the last thing you did. Keeping the subunits small limits the damage if something's been wired wrongly (and it probably has). As a corollary, always wire the power first, and test it thoroughly before connecting it up to the other bits.

16.8.6 Cobbling Copperclad Board

Ordinary G-10 or FR-4 copperclad board is a wonderful construction material, both for small structures and for boxes and shields. Use the $\frac{1}{16}$ inch, 2 ounce, two-sided kind. You cut it using either a big sheet metal shear if you have one, or by scoring and breaking it, and then cleaning it up with hand-held compound-action sheet metal shears. (Note: fiberglass board dulls shears quickly.)

[†]Stranded wire has a bad rep for prototyping, on account of shorts due to stray strands. It is much easier to work with, however, and that is more important.

[‡]That is, turning it on and seeing if smoke comes out.

Copperclad also makes good boxes and shields to separate gain stages and filter sections. Single isolated walls are good enough most of the time. Put sides on the sections when you get above 100 MHz, and bad cases such as power amplifiers and high-Q resonators will probably need a lid too.

Cut the board to size, tack-solder the sides in place, then use a high wattage (at least 60 W, preferably 100 W) iron with a broad screwdriver tip to run a continuous solder bead along all the seams. Remember that the inside and outside foils are insulated from each other, so make sure that at least one of them is continuous all the way around.

Sometimes people solder the box together using the inside foil, then attach it to the ground plane by the outside foil, and wonder why the shield is ineffective. Ideally, you should solder both, because it makes the whole thing stronger. Use copper tape, sold for PC board repair and for leaded glass hobbyists, to bridge the outside edges, and solder it down both sides (the adhesive is insulating).

If you do need a lid, you can make the hinge from copper tape, soldered all along both edges, and just tack-solder the lid here and there to the other sides of the box when you flip it down. Solder copper tape along all the edges of the lid to short the top and bottom foils together firmly.

16.8.7 Perforated Board

You can use plain 0.100 inch grid perf board (without pads) for small prototypes using DIPs and discretes, but make sure you use copper tape for power and ground; the ICs need some sort of solid mechanical connection in order for the prototype to be solid and reliable. There also exists copper-plated perf board, where you cut out pads yourself with a pad cutting tool, which is a bit like a miniature hole saw with a guide pin in the middle. The problem is the shreds of metal they leave behind, and the great big holes in the ground plane that result from cutting pads on adjacent holes (e.g., for an IC). If you put the ground plane on the component side, you can touch each hole with a drill to avoid shorts, which avoids pad cutters but isn't as mechanically robust.

16.8.8 Perf Board with Pads

Vero makes ground plane boards with pads and supply busses, which are good for more complicated digital and mixed-signal stuff where you want to use sockets. The ground plane is relieved around each hole, and solder mask is applied, which avoid shorts but make the ground plane more of a mesh (which works fine). When using this kind of board, make sure that you run the wires (longer than 0.3 inch) on top of the board, away from the solder pads. That way you can avoid solder bridges by running a small scraper (the edge of a thin metal ruler or the back of a broken X-Acto blade works well) between the pads before applying power. This is a sensitive tactile test for solder bridges if you pay attention. Using something sharper runs the risk of slicing off copper shavings and causing shorts instead of curing them. Make sure the ground plane on the top is stitched to the ground traces on the bottom at least every 2 cm, and more often if you have a high current, high frequency (or sharp edged) waveform coming from some chip.

The cheaper breadboards with the interdigitated power and ground areas, intended for digital ICs, are fine for slowish logic if you use enough bypass capacitors, but are not good enough for mixed-signal or RF work.

16.8.9 White Solderless Breadboards

These superficially attractive chunks of white nylon allow you to shove component leads right into an array of socket pins, which are interconnected in rows. A simple circuit can be wired up without using any solder at all. That's the good news. The bad news is leakage due to the low grade hygroscopic plastic, flakiness due to poor contacts and metal debris scraped off the leads, high stray capacitance and inductance, and poor bypassing due to ground inductance. Nobody can do good work on one of these. If you're a hopeless case, at least resist using white protoboards above 50 kHz, 100 mA, 50 V, or when a randomly sprinkled 0.1 Ω or 100 nH in series or 10 pF and 10 MΩ in shunt will screw you up.

16.8.10 Prototype Printed Circuit Boards

Your circuit will almost certainly be built on a PC board eventually, so why not do one for the prototype? You can if you like—there are quick-turn PCB houses such as PCB Express that will turn around small boards in a few days inexpensively. The bad news is that you then have to use CAD for your initial design, which will probably slow you down. The one exception to this rule is microwave circuits or very high impedance RF circuits, where to a large degree the strays are the circuit, and you expect to have a couple of board iterations anyway. Such boards are usually small—if they aren't, you've put too much on them, so consider putting the special stuff on its own card.

16.8.11 Blowing Up Prototypes

Prototypes and development hardware get a lot of abuse, and you don't want to spend all your time fixing things that you broke; thus you should spend a little on bulletproofing. For example, supply leads are always falling off, which can result in reverse-biasing some of your components, and—confidentially, now—sometimes we all put the banana plugs in the wrong binding posts on the power supply, or drop a screw in the wrong place. Protect yourself against power problems by using big beefy Schottky diodes in inverse-parallel with your supplies, adjusting the voltage and current limits on the supply so they will contain the damage, and turning power off at the supply each time you pick up a soldering iron.

16.9 SURFACE MOUNT PROTOTYPES

You can do dead bug prototypes with SMT devices, but it requires a delicate touch; the best bet is to use a mixture of leaded components and surface mount, which isn't as hard. The only thing to worry about with SMTs and dead bug is that the SMTs won't take much torque, so you can't just solder one end of an RN55C resistor to an SOT-23 transistor and then twist it to where you want it; you have to take the strain with long nose pliers instead, or the transistor lead will break off. Chip resistors and capacitors are worse still. The 1206 size is best, with the 0805 size[†] being barely tolerable, and anything smaller, you can forget.

[†]Short for footprints of 0.120 in. × 0.060 in. (3.05 × 1.52 mm) or 0.080 in. × 0.050 in. (2.03 × 1.27 mm).

16.9.1 Stuffing Surface Mount PC Boards

Stuffing surface mount boards by hand is a fiddly job, because of the small size of the components, their close spacings (inviting solder bridge shorts), and especially their unforgiving character. It takes about 5 or 10 times longer than when you're using leaded parts, on account of their small size and their lack of leads to allow sequential assembly and to use for hookup wire. You have to get rid of solder bridges with solder wick, because the usual trick of touching it with a clean iron usually doesn't work. You can't usually get a solder sucker in there, either. Solder wick is very hard to use in mid-air, because you have to grind it in with the soldering iron to get it to work. It is very important to rest your hand on a stable surface while working on SMT boards, so get a good selection of high quality, sharp-pointed, curved-jaw tweezers, which make that a lot easier. You have to get every lead in good contact with its corresponding pad, and well aligned, before applying any solder. You can use small amounts of glue on IC packages, but that makes them hard to remove, which is bad when you're debugging. A better way is to use a rare-earth magnet under the board plus a nut on top of the IC to keep it flat, so that all the leads match up and stay in contact with the board. You can hold capacitors and resistors with tweezers.

Wear a 3× eye loupe or some drugstore reading glasses,[†] and angle the soldering iron so that the handle tips toward the IC package, which helps the solder jump across the gap. Make sure you lean your elbow on the table, and the heel of your hand on the vice—it's steadier, you can keep it up longer, and by making you wrap your arm around, it keeps the soldering iron away from your cheek.

Alternatively, but more slowly, you can do it under a low power stereomicroscope (see Section 9.3.8). Using a single loupe lets you switch magnification by switching eyes, which saves time but will give you a headache if you do it all day—don't stuff a whole big board at once this way. Use solder paste and an SMT oven instead.

A chair of easily adjustable height is very important for this, because otherwise SMT prototyping is an ergonomic nightmare—if you use one of those tall lab stools, run and don't walk to get a pneumatically adjustable chair. You'll get a kink in your neck if the microscope isn't set up well, and severe eye strain if your magnifiers are of poor quality. Go with the good stuff—with an aching head and a sore neck, you're not at peak debugging efficiency. The author has a 1960s-vintage Zeiss surgical stereomicroscope, which is perfect.

16.9.2 Debugging SMT Boards

Debugging surface mount boards is a true pain. Make sure the signal leads are on the outside! One of the main troubles is that it is so hard to stuff an SMT board incrementally (see above). Another way, if you have an SMT oven for production, is to use solder paste and the oven to stuff a couple of boards completely, and hack them up. They don't survive very much hacking, so you'll need two or three just for test purposes.

Once you've soldered wires to SMT pins on the board a few times, the connections between the pins and the board become unreliable, and flakiness sets in. Put in a few test pads, ideally with largish via holes, to allow attaching wires without damaging the board. Vias are pretty robust even on SMT boards, but plain pads will get cooked off

[†]Drugstore glasses are made with an interpupillary distance of 66 mm. If that's right for you, you're in good shape, but otherwise it will lead to severe eye strain. You can mail-order custom glasses for $10 or thereabouts.

after only a couple of solder/desolder cycles. Pads with vias also allow you to bodge in some leaded components dead bug style while debugging, or to tack on a 450 ohm resistor in series with an SMB jack to connect to a 50 Ω scope input—the best kind of ×10 probe for RF signals.

It is worth using a few 0 ohm resistors (SMT ones are available) that can be removed to isolate sections of the circuit for debugging purposes; in low power circuits, we often use 10 Ω resistors in series with the supply lead to each section for decoupling, which is a natural place to break the circuit. Put vias on those pads, too.

16.9.3 Probe Stations

One of the main problems with SMTs is that the pins are either inaccessible or too wimpy to clip a probe to. The solution is a poor man's probe station: a corner of an optical table with a few magnetic probe holders, and the board held firmly in a vice. That way you can make firm connections to the test point without needing to grab onto it (and risk tearing it off or shorting it out).

16.9.4 Hacking SMTs

Sometimes you won't have the right physical size resistor or capacitor—you'll need an 0805 but only have the larger 1206 ones. One good way to handle this (and the problem of needing to change values generally) is to mount your chip components *edgeways*—so that they stand up like miniature billboards. That way you can get your iron and tweezers in there much more easily, and can tilt the component to avoid shorts. Figure 16.12 shows an SMT board that was used for debugging. (It's the analog back end of a 12 Mpixels/s, hand-held, micropower, interferometric 3D surface scanner.)

16.9.5 Board Leakage

High impedance circuits such as pyroelectric detector front ends live or die by the quality of their insulation. Novolac epoxy, as used in IC packages, is a very good insulator, Teflon is better, and air is better still. Clean glass–epoxy circuit cards are pretty good, fingerprints are bad, solder flux is lousy, and damp solder flux is atrocious. A board that hasn't been cleaned is a failure waiting for the next rainy day. Cleaning surface mount

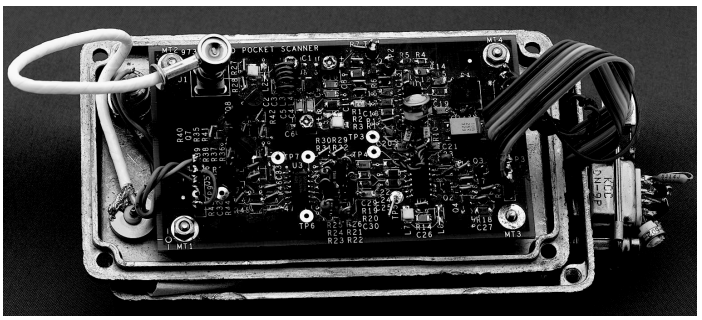

Figure 16.12. Hacked-up SMT board: this is a follow-on to the dead bug prototype of Figure 16.11, with only one channel plus some additional function. (Photo by Frank Giordano; higher resolution version at http://electrooptical.net/www/beos2e/photos.html.)

boards is especially difficult, because the cleaning solvent is held under the components by capillary action. The best method for getting rid of flux seems to be an ethanol or trichloroethylene ultrasonic rinse followed by a trip through an ordinary domestic dishwasher with ordinary dishwasher detergent. A good conformal coating applied in a low humidity environment will prevent condensation from causing problems. (The equilibrium thickness of the adsorbed water layer on a clean surface reaches one monolayer at about 40% relative humidity, and increases rapidly, to infinity at 100%—the board will leak even when it isn't dripping wet.) Guard rings around sensitive nodes are a good idea, but not a cure-all; if the leakage resistance gets low enough, you may run into funny noise peaks and instability caused by leakage from the guard ring becoming an important part of the feedback network.

At high temperatures, glass–epoxy board starts to leak through the bulk, so that wider guard rings (a few times the board thickness) and guarding on internal layers may become necessary.

16.10 PROTOTYPING FILTERS

Complicated filters are not that easy to prototype, even if you have network analyzers and such. Finding components is no problem; variable capacitors and inductors are easy to get. For small inductors (5–500 nH) you can easily wind your own, using (14.4), and squash or stretch them to adjust their values (that's why the coils in Figure 16.11 are irregular). The problem is tuning, that is, how to go about adjusting them to give the desired response. Filters with only a couple of lowpass sections aren't hard to tune, but complicated bandpass filters are more difficult; your starting point has to be close to the desired filter.

The best way to ensure this is to make sure the component values are very close to correct before you start, and that means using an accurate standard inductor or capacitor to resonate each one somewhere near the operating frequency, and then calculating the value. *LCR* meters aren't a good solution at high frequencies.

Bandpass networks are much more sensitive to the resonant frequencies of their individual sections than to the exact component values, so for phase-insensitive applications you may be able to get away with just measuring the resonant frequency of the combination (e.g., with a grid dip meter). See Section 16.11 for a more organized way to tune high-Q filters.

16.10.1 Standard Capacitors

Highly accurate picofarad capacitors are scarce, and so are accurate inductors of any sort. So how do we do this? A good way at HF and VHF (10–300 MHz) is to use shorted transmission line sections. You can measure the propagation velocity of your transmission line very accurately by making a shorted stub of a known length, and then finding where it is exactly a half-wave long (i.e., where it looks like a short again). Use the setup of Figure 16.13, and adjust the frequency until the signal goes away (verify that it's $\lambda/2$ you've found and not λ). If you want to use a fixed generator, you can adjust the stub's length using a thumbtack as a movable short, then measure the length of the cable. Because you don't know the exact impedance of the generator, scope, and tee connector (which will change the effective length of the stub), find the first *two* nulls, which are exactly $\lambda/2$ apart, and measure the distance between them.

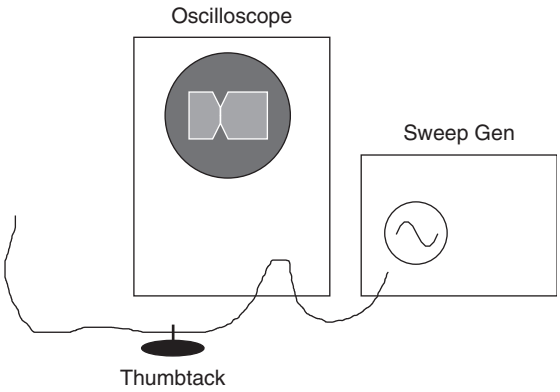

Figure 16.13. Calibrating test capacitors: generator, oscilloscope, open-circuited cable, and thumbtack.

Once you know the propagation speed, measure Z_0 with a long length of cable; sweep the frequency while looking at the generator end with a high impedance probe (i.e., a 2 kΩ chip resistor and the 50 Ω input setting of the scope, *not* a ×10 probe) and adjust the resistance at the far end until the ripples in the envelope go away completely. Use a very small pot, or (ideally) solder a 56 Ω chip resistor right across the coax at the end (no pigtail) and then scribble on the resistor body with a soft pencil to adjust its resistance to Z_0 (most 50 Ω coax is really 51 or 53 Ω, so a 49.9 Ω 1% resistor won't do). Alternatively, use an emery board to adjust the value of a 49.9 Ω metal film resistor. Either way, measure the resistance with an ohmmeter, and you know Z_0. Note: Don't skip this step, because there's a factor of Z_0 in all the formulas for backing out LC values from coax reflections.

16.10.2 Calibrating Inductors and Capacitors: A Hack

A quick and quite accurate calibration method uses the setup of Figure 16.13 minus the thumbtack. Cut a coax patch cord to a length that has its first open-circuit resonance (λ/4 electrical) in the frequency band of interest, and just short the inductor or capacitor across the open end (no pigtails—stick it right on the featureless chopped-off end). Then if the first OC resonance is at f_0, and the first resonance with the shunt element is f_1, the inductance is

$$L = -\frac{Z_0}{2\pi f_1} \tan\left(\frac{\pi f_1}{2 f_0}\right). \tag{16.11}$$

Similarly for a capacitor,

$$C = \frac{1}{2\pi f_1 Z_0} \cot\left(\frac{\pi f_1}{2 f_0}\right). \tag{16.12}$$

This is OK as long as (1) the frequency shift is much larger than f_0/Q, (2) you get the right null, and (3) you really know Z_0 accurately (see above).

We can use this to get a set of beautifully calibrated 1206-series C0G chip capacitors epoxied to popsicle sticks, that we can use for a long time to come for measuring inductors and so on. We can do the same with the inductors, but the caps are more

temperature stable and less susceptible to outside influences, for example, proximity to nearby magnets or conductors, which can change their inductance.

Aside: Variable Standards. It is sometimes possible to make *a priori* standards that don't need extra calibration. For example, consider using a trombone line as a low frequency capacitor. Provided that the shaft is long enough for end effects to have died away, adding a length Δl doesn't change the end contributions; it just lengthens the radial-field region in between. Thus the added capacitance $\Delta l/\Delta l$ is that of a section of an infinitely long coaxial line, which (in MKS units) is

$$\frac{\Delta C}{\Delta l} = \frac{2\pi\epsilon_0\epsilon_r}{\ln(b/a)}. \tag{16.13}$$

Finding the offset C_0 is easily done by putting an inductor in parallel, measuring the resonant frequency at various Δl values, and fitting the curve to extract L and C_0. Similar approaches include the cutback method for waveguide attenuation measurements, the crystal ring-down calibrator of Section 15.7.2, and the turbidity meter of Example 13.11. It's worth looking for this sort of idea anywhere your calibration requirements start looking awkward.

16.10.3 Filter Layout

There are two other difficulties with filter prototyping: the inductance of the leads and stray coupling between sections. Stray coupling is mostly a problem when you need really good rejection in the stopband. Orienting the coils at 90° to one another helps, and in bad cases, copper shields between sections make a big difference. Make sure you don't screw up your inductors by mounting them right next to a metal surface nearly perpendicular to the coil's axes: that looks like a shorted turn, which reduces the inductance and Q by as much as 25%.

16.10.4 Watch for Inductive Coupling

The main drawback of large air-core coils is that they tend to couple to other coils and even to component leads. If you have a mental picture of what the field lines of a solenoid look like, you can orient the coils to reduce the coupling fairly easily: tip them so that the *field lines* due to the two coils are roughly orthogonal inside the coils. The commonly heard advice to put the axes at right angles is wrong unless they intersect at the center of one of the coils. This approach is good for about a 20 dB reduction in coupling. If that isn't enough, use shielded coils or shield the air-core coils from each other with copperclad board.

16.11 TUNING, OR, YOU CAN'T OPTIMIZE WHAT YOU CAN'T SEE

Tuning a network is a knack, like aligning optical systems on the bench. You get good at it after a while, but it's slow starting out. The best way to learn is to sit at the elbow of an expert—as a 21-year-old junior designer, the author was taught it by a crusty old technician of about 26. You spend some time getting a good starting point, then go round the circuit, adjusting one thing after another until it looks the way the simulation says it

ought to. Like any iterative optimization process, the steps start out big and get smaller as you get closer. If there are only a couple of components (e.g., an L-network matching section between two RF amplifiers), tuning is easy. Much the same is true if the sections don't interact much, for example, a multistage tuned IF strip or an active filter using several op amps in cascade—it can be broken down into several easy problems.

On the other hand, tuning multiple-section LC filters with carefully optimized responses can be tough if you're new to it, because all the controls interact. What's more, the reason we've spent so much design effort on that filter is usually that its characteristics govern the performance of the whole instrument. A typical example is the last IF filter in an RF signal processor; its bandwidth sets the SNR, and the sidelobes of its impulse response determine how fast we can take data. That one just *has* to be right, and we often can't buy it off the shelf. (Custom filters can be had in a week for a hundred dollars or so from TTE.)

In numerical optimizers, each trial step can adjust all the L and C parameters, and every step produces a numerical value of the penalty function that can be compared accurately with all previous values. Real-world tuning isn't like that—we have to adjust the components individually, and we're susceptible to noisy readouts and our own mental inconsistency. Thus if we find ourselves down some long narrow valley in parameter space, fairly far from the optimum, we might never tune our way out of it. It is critically important to have stable, detailed, real-time indicators of network performance, and to start out near the optimum in the first place. After that, a few pointers will get you well on your way.

1. *Tune passive networks for return loss*. Figure 16.14 shows the right way to tune filters when we care only about amplitude response (e.g., wideband RF stage filters for image rejection) or when the desired signal looks like noise and not pulses.

You tune for return loss because the peaks of the transmission response correspond to $\Gamma = 0$; you're looking for easily spotted nulls rather than the very subtle peaks of the insertion loss. It really isn't possible to tune a high performance filter looking only at the insertion loss, since your cables, measurement setup, and generator variations with tuning will disguise small insertion loss changes. Do check the insertion loss when you're done, just to verify that you're not matching beautifully into the copper losses in your coil, and that the skirt selectivity is what you expect—the skirts are as invisible in reflection as the peaks are in transmission, and for the same reason.

If you have a fair amount of microwave work to do, it's really worth getting familiar with a magic labor-saving device: the Smith chart. It's beyond our scope here, but you can find it in microwave books and online tutorials: trust me, it's worth learning—it can make you 100 times faster at microwave design.

2. *Use Dishal's method for high-Q filters*. Sharp filters are made of high-Q sections (reactances large compared to Z_L), so when they're detuned, series sections look like open circuits and parallel sections look like shorts. This allows us to use a shortcut called *Dishal's method*.[†] Starting with the input section, use the return loss test setup to tune its resonance to $1/(2\pi\sqrt{L_1C_1})$. Go down the ladder and tune each section to its calculated resonance in turn, and when you're done, your filter is pretty nearly perfectly tuned. (For a lower-Q filter, you can do the opening and shorting with a soldering iron.) The reason

[†]M. Dishal, Alignment and adjustment of synchronously tuned, multiple-resonant-circuit filters. *Electrical Commun*., pp. 154–164 (1952)

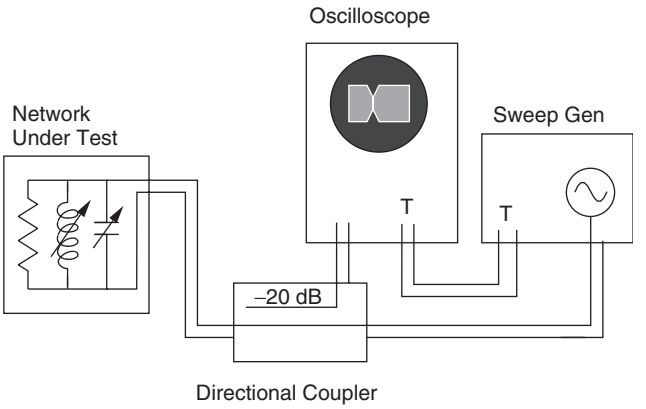

Figure 16.14. Filter-tuning setup with directional coupler.

this works is that (in the high-Q limit) series resonances don't move around when you put other reactances in parallel with them, and parallel resonances don't move when you put reactances in series with them. For low-Q filters this isn't terribly accurate, but it's usually close enough to get all the bumps showing (see item 4).

3. *Tune pulse filters for impulse response*. In pulse filters, the things that matter are the impulse response and the noise bandwidth , but what are usually specified are the amplitude response and phase nonlinearity limits, which is a bit indirect.

Whichever you choose, you have to have both things on the screen simultaneously if you want to optimize both. A vector network analyzer will do this, and is a must in a production environment, but there are other ways for the development lab. The easiest is to do a preliminary tune using the return loss method, and then use the setup of Figure 16.15 for the final tweak. Make sure the pulses are clean, and are at most 10% of the width of the desired impulse response (or 10% of one cycle of f_0 for BP filters), and that the duty cycle is 1% or so; since the pulse has a flat spectrum out well beyond the

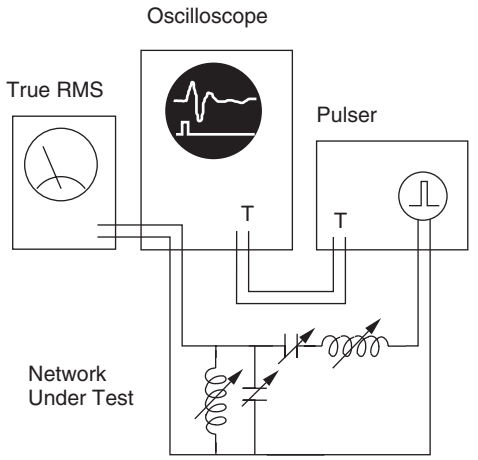

Figure 16.15. Setup for tuning pulse filters.

filter cutoff, the total AC power output represents the noise bandwidth, and the impulse response comes out directly. You'll need a true RMS power meter with really good crest factor performance, of course; a thermal detector is one good way to do it, or a calibrated square-law diode detector, or even a sufficiently fast digital scope plus some software. In any event, make sure you have a real-time, analog indication of the total RF power. That means a scope trace or a good quality electromechanical analog meter—the little bar graphs on DVMs are beneath contempt.

4. *Get the right number of bumps first, then make them look nice*. This advice applies even for filters that aren't supposed to have bumps (e.g., Butterworths); a (mildly mistuned) N-pole lowpass or highpass ladder filter should have N bumps in its transfer function, counting both positive and negative going ones, and bandpass ones should have $2N$. That way you're sure to have all the resonances in the right neighborhood.

5. *Get the center frequency and bandwidth right early*. In filter design, for a given number of sections, you have to trade off performance pretty rapidly to gain bandwidth. In tuning, it is tempting to get the filter looking nice at a slightly too-narrow bandwidth, and then expect to tweak it to get the bandwidth right. Don't fall for it—there are so many local minima when you do this that you might as well start over. The desired filter design may not even be nearby when you've done it this way. Get the right number of bumps, spread them over the right bandwidth, then optimize.

The reason for this is interesting: when we tune a capacitor or inductor, we're changing not only the resonance of the section, but its impedance as well; in moving the section resonances slowly around (as $(LC)^{-1/2}$), we change the coupling coefficient between sections rapidly (as C/L) in the process. If we had an adjustment method that changed f_0 and k separately,[†] this advice would not apply. The key is to make the adjustment ranges no wider than necessary to accommodate component tolerances; a circuit with tunable coils and trimmer capacitors as the only reactive elements is a bad circuit.

6. *Tune front ends for minimum noise figure*. RF front ends may exhibit minimum noise figure somewhere other than at best matching. Look at the SNR of a small modulation on the carrier to get the best tuning point.

7. *Don't be fooled by running out of range*. The capacitance of a trimmer has a value that depends on the area of overlap of the stator and rotor plates. Minimum and maximum capacitance are half a turn apart, and you can go either direction to get there. A network tuned with variable capacitors should thus exhibit *two* optima per turn, if the real optimum is inside its tuning range. If there's only one optimum per revolution, all that's happened is that you've run out of tuning range. This is a frequent problem.

Tunable inductors are a little less common, but will also exhibit multiple optima in general; the magnetic effects of the core are greatest when it is centered in the winding, so positions to one side are equivalent to those on the other. This is not true if there is coupling to other circuitry involved, however, since the two positions will exhibit different fringing fields and so different coupling. In any case, it's pretty easy to tell if you've run out of range on a coil—the slug falls out. If you have reason to know that the true optimum is somewhere in the adjustment range, don't worry about it; just keep tuning the other elements until the railed one comes back in range. Otherwise, try sticking in a fixed element to move the center of the range.

[†] We're talking about k the coefficient of coupling here, not k the relative permittivity, k the propagation constant, k the imaginary part of the complex refractive index, or k for Boltzmann's constant. Isn't interdisciplinary work fun?

Example 16.5: Tuning by Finite Differences. It is usually best to tune complicated networks by finite differences, especially when the readout is noisy. The author and his colleagues were once faced with a motorized π-network match box (one motorized tapped inductor and two motorized vacuum variable capacitors), coupling a 5 kW RF generator into a plasma etch chamber. Solid state RF power amps are much more vulnerable to mismatch than tubes; the big standing waves in V and I will blow up the rather delicate output devices. This one automatically reduced its gain to keep the reflected power below 100 W or so, producing nasty discontinuities in the reflected power with tuning. What's more, the RF impedance of a plasma chamber is wildly nonlinear, so we couldn't just turn the power down, and the onset of power limiting detuned the network severely, so that it had to be coaxed out of power limiting. Even worse, the forward and reverse power meters were digital, and the plasma was unstable, so that the digits kept bobbling up and down even at a fixed setting.

After a change in the cathode design, the network wouldn't tune at all; capacitor C_{fast} was very tweaky, and C_{slow} did almost nothing. Every adjustment of C_{fast} would send the amplifier into power limiting, so the usual iterative method of slowly trading one control off against the other while moving the system toward match failed—we were in one of those long narrow valleys, between cliffs, with lots of loose rocks. The solution was to tune by finite difference: tune the network for best VSWR with C_{fast} at one setting of C_{slow}, then move C_{slow} a fairly long way and tune it again using C_{fast}. If the minimum VSWR went up, that was the wrong way to move C_{slow}. By taking big enough leaps that the VSWR change was unambiguous, we got enough directional information to tune the network correctly. The final fix was to put another capacitor in series with C_{fast} to reduce its tuning sensitivity and confine it to the right neighborhood, but we couldn't tell what value to use until we'd tuned the network.