1

Introduction

This book is targeted at designers of power supplies in the low to medium power range, roughly defined as 0W to 10kW. If you are in this group, you probably already have some experience with converters, at least to the extent of realizing that there are many different kinds. (Chapter 2 on topology talks about the various types.) There is an excellent reason for having many different kinds, rather than having all power supplies be mere variations of parameter values on a single type. This reason is twofold: the wide variety of sources from which converters are expected to run, and the similarly wide variety of loads converters are expected to provide power to. Unless you have spent many years designing converters, you probably don't realize how truly diverse these two groups are—this is one of the things that makes power supply engineering far more challenging than, say, digital design. To start off this book, then, examples of sources and loads, both common and less so, are presented in some detail, to give you a feel for the sorts of thing you may encounter. The samples here are of course not exhaustive, but rather represent some of the (occasionally not so nice) experiences of the author in the power range under consideration; you can start your own collection. Don't take these discussions to be comprehensive, as some of the sources and loads have a large literature attached to them; rather, this material is intended to give the flavor of the sorts of environment in which power supplies often need to operate.

This introductory chapter also makes a few comments on lab safety. This is a subject that seems to be always ignored, both in the lab and in texts, or at least pushed to the side when time becomes short; and yet it is of critical importance both to you and the people who work for you, and for visitors to areas in which your lab work is taking place. Make sure to take the time to read it!

SOURCES

Lab Supplies

Everybody uses lab supplies to begin development work on a new converter. Still, there are a few surprises to be had, even with these supposedly ideal sources.

The most obvious difference between lab supply types is between old lab supplies and new: you can easily tell which is which because the old ones seem to weigh a million pounds. The reason for their great weight is that there is a big hunk of steel inside, acting as a 60Hz (or 50Hz) transformer. The old lab supplies then act by linear-regulating the voltage down, and end up with a really big capacitor [10s of millifarads (sic) and more]. The rest of the volume of the supply houses a fan to keep the linear regulator from burning up. This reliance on steel guarantees that some supplies that are old now will still be around when you're ready to retire! They just never seem to break; the author is personally aware of labs with quite a few old lab supplies dating from WW II.

New supplies are (almost) invariably based on switching regulators. (The caveat is for the arena of lab supplies that are required to be *extremely* free of electrical noise—these still tend to be linears, but usually only low power ones.) The switching regulator design makes them of course vastly lighter than their older counterparts, but does leave them prone to the ills that afflict switchers, a subject that will be recurring throughout this book. For one thing, although they also usually sport large output capacitors, any switching regulator can be made to oscillate if you attach enough capacitance to its output. The manufacturer's intent, of course, is to have the output capacitance of the lab supply dominate anything you may reasonably hang on the lab supply's output, and the internal control loop is compensated to provide stability with this capacitance; but if you put *enough* extra capacitance there, eventually the loop will break up and the supply will oscillate. The 60Hz transformer style supplies seem to be immune to oscillations in practice, although it ought to be possible to make a linear regulator oscillate too; maybe it's just really difficult to hang *that* much capacitance there.

Another problem for lab supplies is 60Hz (line) feedthrough. Here, switchers are much better, because they have gain at 60Hz, which radically reduces the amount of 60Hz signal that appears on their outputs. Also, newer switchers are better than older ones, as some of the older ones were actually thyristors whose control was based on line phase angle; these tend to have much more 60Hz noise than modern switchers using MOSFETs with high bandwidth control loops. The old linear supplies relied on their gigantic output caps to filter the ripple, but it could still be quite a nuisance.

One other area where sometimes the old linears are better exemplified is when you need to parallel two or more supplies to get up to the current level your application requires. (This of course requires remote sense to work at all, regardless of the type of converter.) With the old converters, the noise is all at 60Hz, and each converter produces the noise in phase with all the others, since it's just feedthrough from the line. Though additive, the noise at least has a well-defined spectrum. With switchers, there never seems to be a synchronization pin when you need it, so that each switcher runs at its own frequency. If you're unlucky, these frequencies may be fairly close together, and then you may get beats between the frequencies of the various paralleled converters at a low frequency. This is definitely undesirable for attempting to debug a new converter design.

AC Mains

When you plug the toaster in, the bread gets brown. If only it were so simple for electronics! The AC mains is actually a wide assortment of power types, with numerous types of problems. Being able to really guarantee that a power supply is going to work successfully and have a long life running off the mains requires a lot of work, and lots of research into relevant (national and international) standards. Perhaps this shouldn't be too surprising, when you remember that there are engineers who spend their entire lives working on it!

The most obvious difference between various AC mains is the different frequencies. In the United States the line is at 60Hz; in Europe it's 50Hz. Actually there's also a tolerance to these numbers, so if you're designing a supply for international use, it needs to work down to 47Hz, and in the United States up to 63Hz. This tolerance is necessary because the electric companies won't guarantee that their big turbines will always run at exactly the same speed. The requirement to work at 50Hz (really 47Hz) translates into considerably bigger capacitors than would be needed if the design only has to work at 60Hz.

There are *lots* of different mains voltages around. In this country, 110VAC (sometimes 120VAC) is the usual for wall outlets, but there's also 208VAC (for your washer), 480VAC three-phase (for industrial sites), and 277VAC (for fluorescent lighting, though it also runs off 120VAC, depending on the building). Then in Europe there's 230VAC... and in Australia it's 240VAC! Let's not forget the cable TV coax, which distributes what's *called* 60VAC but is really a quasi–square wave of pretty high impedance, with a peak-to-peak value of about 120V.

These are just the *nominal* values; each mains has all sorts of tolerances as well. Taking the 110VAC as an example, anything from 95V to 135VAC may be considered to be within normal range for a power supply to operate in without degradation of performance. Then there are sags and brownouts—basically the power supply has to be able to avoid damaging itself for any voltage from 0V up to nominal (use an undervoltage lockout to accomplish this). Also, it may be required to provide uninterrupted power even when its input disappears for several line cycles of 60Hz. (The only way to do this is with lots of capacitance or a battery; if the supply is power factor corrected, the capacitance has to be on the output, making it even bigger.)

Then there are overvoltage conditions. There are lightning strikes, which may be 6000V at an impedance of 2Ω, both line to line and common mode (see Chapter 9 on EMI for these concepts). The types of lightning come in two flavors, a short one (1.2μs rise time and 50μs decay time), and a much higher energy one that decays in 1ms. There are also transients: the line can go to $750V_{pk}$ for a half-line cycle! (This is a requirement for certain telecommunications supplies; the regulatory agency is anticipating that a high voltage wire will fall across the mains during a storm, and it will take some time for the circuit breaker to act.)

This short account doesn't even scratch the surface of the numerous problems a power supply faces when attached to the line. Altogether, the AC mains is an extremely nasty environment, and it can easily happen that as much time is required to make a supply robust and able to pass all the national safety requirements (and which of course are different in each country) as is needed to design the whole rest of the supply.

Batteries

Batteries represent something utterly outside the ken of most power engineers, since they involve chemical reactions and metallurgy. Indeed, when you talk with experts in electrochemistry, it turns out there's plenty *they* don't understand either. Compared with batteries, the AC mains is understandable, if nasty. Let's try to collect some of the basics here, to let you know some of the questions to ask when faced with designing a supply that will run off a battery.

First off, the author's pet peeve: batteries are NOT gigantic capacitors. Although you can put energy into a battery and get it back out, application of a sine wave will reveal that there is no phase shift between voltage and current. Batteries also are not much good as filters, as we'll discuss in a moment. So let's talk about what batteries are, at least as sources; what they look like when sinking current is discussed below in the section on loads.

A battery consists of a number of cells, usually, though not always, connected in series. It's useful to bear in mind the terminology: a cell is the basic unit of the battery, while a battery consists of one or more cells connected together.

Warning! Don't attempt to hook up cells in parallel to form your own battery, as this can be dangerous. Have the battery manufacturer configure the cells into the battery voltage and capacity you need. This caution is unnecessary if you have ORing diodes, and for series cell connections.

A single cell is basically a chemical reactor of sorts. It typically consists of two metallic plates, with some sort of conductive path between them, which can be either liquid or solid depending on the particular chemistry used. The key aspect of this arrangement is that it has a reversible chemical reaction dependent on electricity (for rechargeable, or "secondary" cells; the nonrechargeable ones are called "primary"). When you put electricity in by attaching a source to the two metal plates (which are the cathode and anode), there is a chemical reaction that causes some of the material to change chemical state; this stores energy. When you attach a load, the chemical reaction goes backward toward its original state, releasing electricity again.

What batteries look like as sources depends on what frequency is being considered. Let's start with the highest frequencies and work our way down. At typical converter switching frequencies, 20kHz or more, batteries look like open circuits, because they have some small amount of inductance associated with their terminals, internal plates, and so on; also, chemical reactions take a finite amount of time to occur, and so present the equivalent of some impedance. For example, a NiH (nickel–hydrogen) cell may have something like 200nH inductance; a battery of five of these cells in series (to get the voltage up to 6V) would have about 1μH. At a switching frequency of 200kHz, this is about 1Ω. Thus, you can't assume that a battery is going to sop up all the switching ripple your converter is generating; it is actually usually necessary to put some capacitors in parallel with the battery!

Looking now at lower frequencies, say 1kHz down to a few hertz, there are a lot of nonlinearities due to the chemical processes. As you draw increasing current out of the battery, the voltage drops (the relationship is approximately a hyperbolic sine). Figure 1.1 shows a nominal current–voltage curve of a 12V NiH battery. Nominal voltage is 12V, and

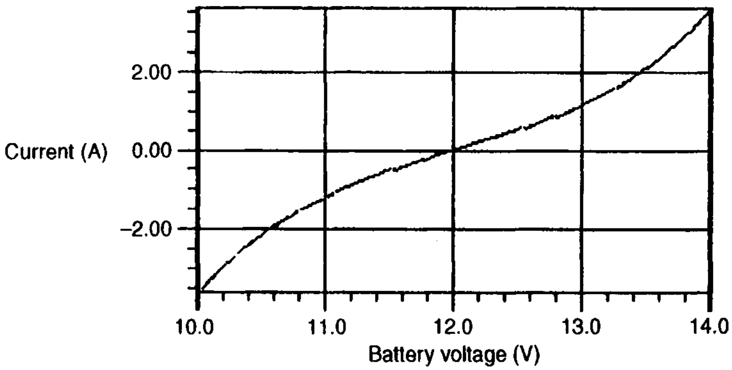

Figure 1.1 Typical *I-V* characteristics of a 12V battery.

current coming out of the battery is defined to be negative. At small currents drawn from the battery, there is a voltage drop at the battery's terminals that is approximately like a resistor: doubling the current doubles the voltage drop. At higher currents, the voltage drop becomes relatively less, until (not shown in the figure) you can pull quite gigantic currents out of the battery before the terminal voltage reaches 0V. (DON'T TRY THIS—if you short a battery, it may explode!) The author saw test results on a NiH cell that produced 1500A short circuit!

The relationship between current pulled from the battery and output voltage is dependent on temperature, and also on how much charge the battery has left. And you can damage a battery if you try to pull too much current out of it. You can damage almost any battery if you try to pull current out of it below its rated operating temperature; for example, sealed lead–acid batteries don't work very well much below −10°C, which is why your car doesn't want to start when it's cold.

Continuing down in frequency, on the time scale of minutes to hours, the capacity of a battery is measured by manufacturers by how many "amp-hours" of charge it has (current × time = charge). Confusingly for power supply designers, this has no simple relationship to how much energy you can get out, which is *not* equal to the capacity times the output voltage; unfortunately the output voltage depends on the current being pulled! The behaviour of every one of these parameters is described in curves from the manufacturer, but the curves never seem to cover the operating point at which you're actually operating. Lots of interpolation and hope are required; it's almost always impractical to do your own tests on batteries. And it should be borne in mind that each manufacturer makes batteries a bit differently, so you can't assume that just because two batteries have the same chemistry and amp-hour rating, they're going to have the same run time in the field.

Another phenomenon in this approximate frequency range is self-discharge. If you leave a charged battery sitting around, it will gradually lose its energy all by itself, without any load attached. The time required to lose a substantial portion of the stored energy varies widely depending on chemistry, from 24 hours for NiH to years for some lithium batteries.

Finally, on a scale of years, after many charge/discharge cycles, the battery will no longer store its rated capacity. This time, which may be considered end of life for the battery, depends on how it is operated: how many charge/discharge cycles it's undergone, how deep the discharges were, and so on. Even a battery used only for backup, and so being "float-charged" (always held fully charged) will need to be replaced in 5–10 years.

Every type of battery chemistry—lead–acid, NiCd (nickel–cadmium), zinc–air, whatever—has its own set of characteristics. So you get the idea: you could spend a lifetime studying batteries. The best plan is to find a manufacturer who is willing to work closely with you and lean heavily on that person's technical expertise.

Solar Cells

Yet another entertaining power source is solar cells. A solar cell is a diode that produces a current when exposed to visible light. Actually all ICs respond to light (this is why an EEPROM can be erased by UV), but solar cells are optimized for producing a maximum output of electricity per unit light exposure. The current is produced at a voltage with characteristics pretty much like that of a regular rectifier, if you imagine it putting out energy rather than dissipating it. It thus has an *I-V* curve that is logarithmic in current, as idealized in a typical curve shown as Figure 1.2. Note that contrary to the way you at first expect, this curve shows current versus voltage, not the other way around. This is standard for solar cells. If you use this curve to determine power output as a function of current (power = current × voltage), you find (see Figure 1.3) that there is some curent at which output power is maximum; of course it is not at open circuit, because then current is zero, nor at short circuit, since then voltage is zero. A converter always needs to operate on the

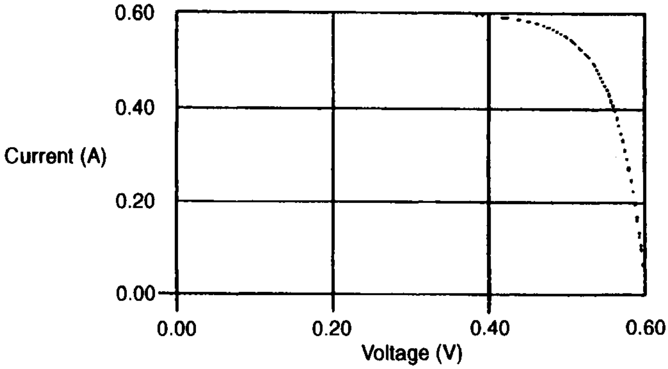

Figure 1.2 Typical *I-V* curve of a solar cell.

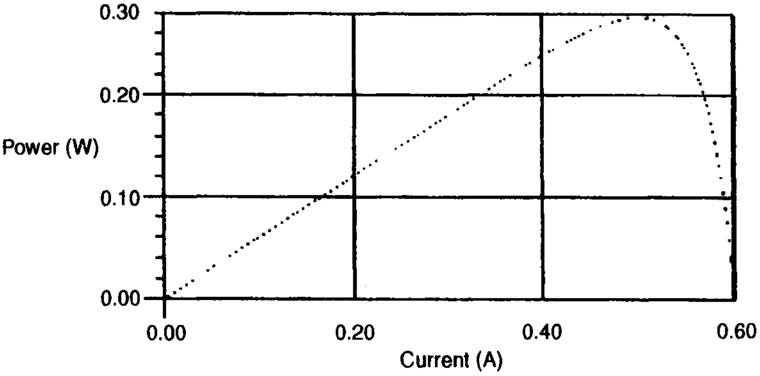

Figure 1.3 There is a peak in the delivered power from the solar cell.

left side of this peak power point, so that pulling increasing current will produce increased power. If you ever go beyond the peak power point (to the right), the system becomes unstable: if the converter wants to pull more power, it pulls more current, which reduces delivered power, causing it to pull more current, etc. That is to say, it falls down the curve. Finding the peak power point and ensuring that the system doesn't go past it is always a major challenge for designs utilizing solar cells.

Some types of converter utilizing solar cells act as a flyback; that is, the current is shorted to ground for some duty cycle and then released to power an inductor during the remainder of the period. This brings in another difficult aspect of solar cells: being semiconductors, they have a voltage- and temperature-dependent capacitance. This capacitance produces a momentary surge of current into the shunting device until it is discharged. Then, of course, the voltage doesn't come right back up, it ramps up while recharging the capacitance! This time to recharge is typically a few microseconds, and so sets a maximum switching frequency for the attached flyback converter.

LOADS

High Speed Requirements

Many people have by now heard about requirements on power supplies for microprocessors running at 3.3V: the data sheets are calling out a load step of up to 30A/μs. So referring to Figure 1.4, let's suppose the load changes from no load to 7A: this takes less than 1μs. If your switching supply has a bandwidth of 20kHz (no mean feat), it still takes something like 1/20kHz = 50μs to change to the new load level, and so you have a deficit of approximately (7A/2)50μs = 175μC to support. If you need 2% tolerance on the 3.3V line, which is 66mV, you need 175μC/66mV = 3mF (sic) of capacitance to hold up the output during the transient!

It is worth observing that you can't just stick a 3300μF capacitor in this job either, you have to parallel multiple smaller caps. This is because the initial voltage drop on the bus (see Figure 1.5) is going to be due to ESR of the caps, not bandwidth limitations of the converter: you need 66mV/7A = 9mΩ of ESR maximum. If each cap has about 100mΩ ESR, you need at least 11 caps in parallel to achieve this, so maybe 330μF tantalum chip caps would be a good choice here. Of course, this calculation assumes that the connection from the output of the converter to the load has no resistance and no inductance—if there's any trace length, you need even better power supply performance!

Another assumption in this calculation is that the large-signal response of the converter is adequate, also. This is discussed in detail in Chapter 6 on stability, but basically you have to make sure that the error amplifier has a slew rate adequate to track the small-signal response of the converter; this may not always be true. The large-signal bandwidth of the converter can't be greater than the small-signal bandwidth, and if there is inadequate slew rate, it may be considerably smaller.

This sample calculation makes it clear that having wider bandwidth converters and higher speed amplifiers is essential to keeping converter size down. In industry today, this is the dominant reason for continuing to push to higher switching frequency converters (since bandwidth can't exceed roughly half the switching frequency). Certain actually working converters now switch at 2MHz, and have bandwidth of over 100kHz.

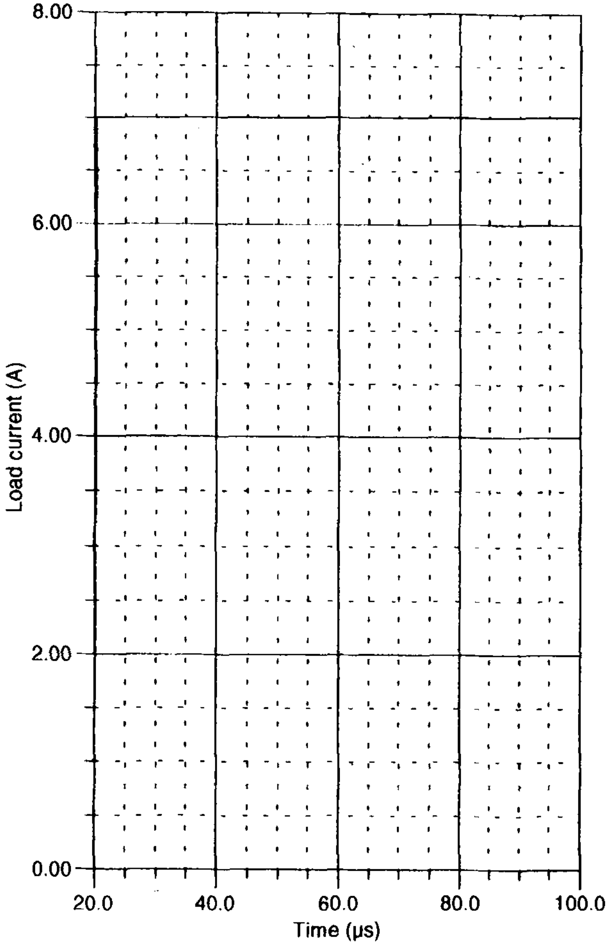

Figure 1.4 A 7A load step in less than 1μs at $t = 20\mu s$.

Low Noise Requirements

A load that again requires some thought is the low noise load. The power supply engineer should not be surprised to see a requirement for an output that requires, say, $1 mV_{pp}$ ripple noise. (This is in addition to the transient response mentioned in the preceding section).

Practical Note Somehow these ultralow noise requirements often seem to be on one of the outputs of the same converter that is also expected to provide 10A at 5VDC to a tankful of TTL ICs. Before spending weeks in the lab working on complex filters, it's a good idea to go talk with the users and make sure they really *need* that low noise: maybe they just didn't want to bother thinking about it and hadn't realized that their casual spec would make you go prematurely gray.

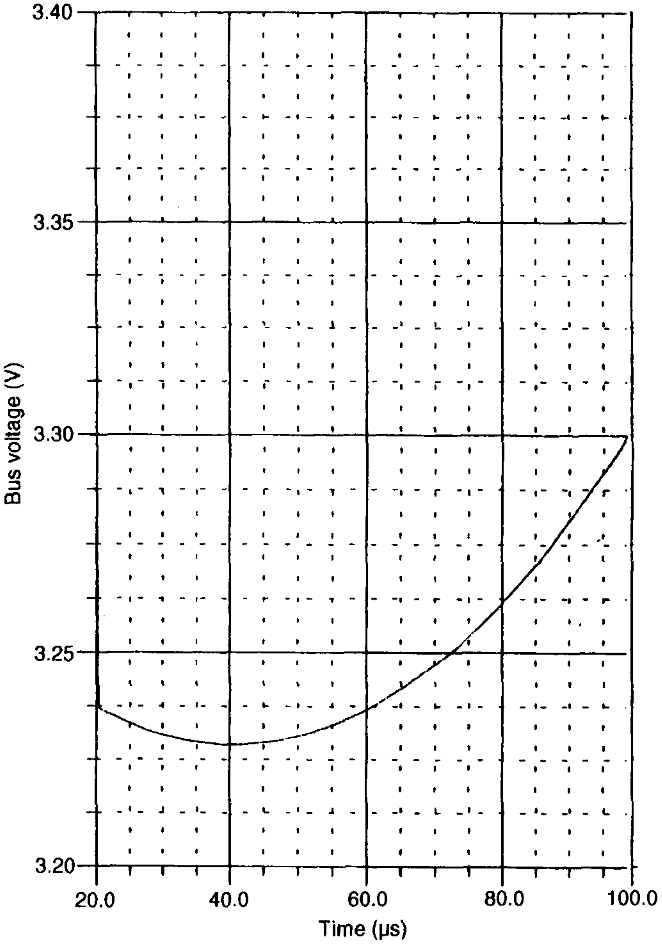

Figure 1.5 Response of bus voltage.

A common load that *does* require low noise is an rf (radio frequency) power amplifier—for example, as used in a cellular phone. The converter provides both gate and drain voltage to the amplifier (the amplifier is basically a FET). If there is ripple on these voltages at the switching frequency of the converter, the output of the amplifier will also have ripple, since the output power is determined by the gate and drain voltage—indeed, changing these voltages is the usual method of *controlling* the power level. Since the output of the amplifier is rf, the ripple shows up as sidebands on the carrier frequency. It's easy to see that you don't want any ripple (or harmonics), since this would create sidebands that would be demodulated as a signal by the receiver.

Both the ripple, which is due to the peak-to-peak inductor current times the ESR of the output caps, as well as the switching noise, which is due to transition times of diodes and transistors, have to be considered to meet low noise requirements. At these levels, it is no longer practical to try to get a big enough inductor and enough output caps in parallel: the only choices turn out to be a linear post-regulator or extra filter poles following the main converter.

The linear postregulator is never very desirable because it is inefficient. The extra filter poles mean specifically an extra *L* and *C* following the main output filter as shown in Figure 1.6. The only tricky part is deciding what to do with the feedback loop to the converter. The simplest solution is to continue to use feedback from the main converter output cap; the converter then sees only two poles and is easily stabilizable regardless of how large the additional filter is. However, the response of the additional filter is now uncontrolled, and it will probably ring when excited by a step load, defeating its purpose.

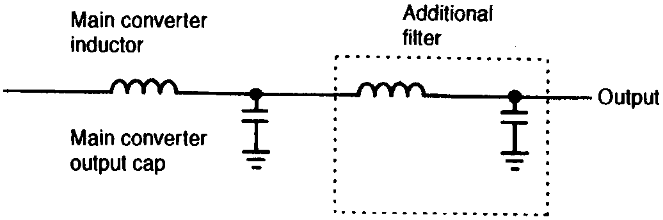

Figure 1.6 Low noise output can be achieved with an additional *LC* filter.

A better choice is to get the converter feedback at the output of the additional filter. This introduces two extra poles, making the converter unstabilizable if the poles are too low in frequency.

Practical Note A good choice is to make the resonant frequency of the additional filter approximately 10 times higher than the bandwidth of the converter. This then gives little phase shift for the compensation to deal with (see Chapter 6), and may still give adequate attenuation at the switching frequency. Generally, the inductance should be made small and the capacitance large, to decrease the converter's output impedance.

Practical Note You usually end up with quite small inductance required for this type of post-filtering, perhaps some hundreds of nanohenrys to a few microhenrys. Instead of trying to use a ferrite bead, which has trouble supporting DC current, try using one of the small MPP toroids as a bead, making it a single turn by just passing the output bus through it.

The worst load will occur when you need fast transient response and low noise together; then you will have to combine techniques from both these sections, and can expect to spend a lot of sweat on it.

Batteries, Again

Batteries were unpleasant sources, so, just as you would expect, they are unpleasant loads, too. The first thing about them is pretty obvious: when you need to charge a battery you can't just apply a voltage to it, because the amount of current the battery takes is exponential in the voltage. You need to have a way of controlling the current.

The way charging current is measured in databooks (discharge current is measured this way, too) is in terms of "*C*." This is best elucidated with an example: a 20A-h capacity battery is said to be being charged at *C*/5 (no one says 0.2*C* for some reason) if the current into it is 4A (20/5 = 4). That is, the 1*C* rate is a current that would nominally recharge the battery in 1 hour (20A × 1h = 20A-h). Notice that "*C*" is a current, measured in amps; the capacity of the battery is measured in amp-hours, making its units charge.

To get a battery fully recharged, (don't ever discharge a primary battery to zero, it damages the battery), you need to first recharge it at a relatively high rate. (How high? The higher the charge rate, the more inefficient the recharging will be, because the battery will actually warm up. If you don't expect to use the battery soon in the system—for example, if it's used in a standby application—*C*/20 is a good choice.) After recharging to maybe 80–90% of capacity, you can taper the current off, that is, reduce the charge rate, to get back to a fully charged state. You usually don't want to continue charging at the high rate for this last portion of charge because this can cause the battery to heat. Keeping up the high charge also makes it difficult to know the actual state of the battery, since the "resistance" of the battery means that the terminal voltage will be higher when a lot of current is being pumped in. Some types of battery can be severely damaged by being overcharged. Because of the inefficiency of the chemical reactions inside the battery, you typically have to put 5–10% extra charge in beyond what you took out. Actually, the battery self-discharges, so you have to keep on putting a little bit of current in forever, even when you don't use the battery.

This last state, called trickle charge (or float charge), is usually done with a voltage rather than a current, because the amount of current required can be so small as to be hard to measure. A typical float-charge regime for a 12V, sealed lead–acid battery might be $13.6\text{V} + 30\text{mV}\ (T - 25°\text{C}) \pm 0.2\text{V}$; in a fully charged battery, this may correspond to a rate of *C*/1000, a few milliamps.

So handling a battery properly requires fairly good measurement of current (down to milliamps and up to several amps), and voltage (200mV/13.6V = 1.5%), and long integration times (keeping track of current for a 20-hour recharge). This application cries out for a microcontroller in your converter. Are you ready for this?

Telephones

Telephones, which have been around for 100 years, were designed with large pieces of steel and copper in mind, not semiconductors. They are powered by the phone line, not by the local mains, which is why your phone continues working when the lights go out. They are thus typically located hundreds of meters away from their power supply, which introduces substantial resistance and inductance between the supply and the phone.

A telephone can be modeled as having three different states: either it is not in use, or it is ringing, or else it is off-hook and in use. These three states have different characteristics, and the characteristics of each are (naturally) different in each country.

To appreciate how hard it is to drive a telephone in the ringing state, consider some sample numbers. In the ringing state, a phone looks like a resistor in series with a capacitance, and it has to be driven by a low frequency sine wave. This sine wave has to have a minimum voltage at the phone of 40V_{rms} (in the U.S.) or 35V_{rms} (in Germany); in reality, the voltages required from the supply are considerably higher because the output of the supply is divided down by the various line impedances before it reaches the phone. U.S. phones are approximately 7kΩ in series with 8μF, and are driven with a 20Hz sine wave. German phones look like 3.4kΩ in series with 850nF and are driven with 25Hz. Phones in France are required to be more than 2kΩ and less than 2.2μF, and can be driven at either 25 or 50Hz, depending on whether the driver is differential ("balanced") or not. Electronic phones can be almost any load whatsoever, from 6kΩ to 60kΩ or more! And yet the power supply has no way of knowing which of these telephones it will be powering, unless it is tailored for each country individually; indeed a supply running five phones needs to be able to power both conventional and electronic types simultaneously.

A quick calculation [$\chi_C = 1/(2\pi \times 20\text{Hz} \times 8\mu\text{F}) = 1\text{k}\Omega < 7\text{k}\Omega$] shows that U.S. phones, because of their large capacitance, are dominantly resistive, whereas German phones [$\chi_C = 1/(2\pi \times 25\text{Hz} \times 850\text{nF}) = 7.5\text{k}\Omega > 3.4\text{k}\Omega$] are dominantly capacitive. French phones [$\chi_C = 1/(2\pi \times 25\text{Hz} \times 2.2\mu\text{F}) = 2.9\text{k}\Omega > 2\text{k}\Omega$, while $\chi_C = 1/(2\pi \times 50\text{Hz} \times 2.2\mu\text{F}) = 1.45\text{k}\Omega < 2\text{k}\Omega$] can be either resistive or capacitive! Thus the power supply has to be able to produce this high voltage sine wave into a load that may have either 0° or 90° of phase shift. When you add in cabling inductance, it turns out the load could even be inductive and have a −90° phase!

As if this weren't nasty enough, when you are talking on the phone, it looks like a pure 200Ω resistance. So here you are driving a 120V_{pp} sine wave into a reactive load, and when the user picks it up, suddenly it becomes a 200Ω resistor! Of course, the supply must quickly change its drive—otherwise it would be supplying huge power (a single supply should be able to power five phones). But because of the differences in phones of different types, the same measurement technique can't be used even to determine when the phone has been picked up. In the United States, for example, they look for a certain level of current (since it is mostly resistive), but in Germany, with the big capacitor, there is no substantial change in current (although there is in power), so they look for a phase change.

Fluorescent Tubes

Fluorescent tubes are another unusual type of load, driven by a special type of power supply called a ballast. Tubes come in quite a variety of types, the ubiquitous 4-foot-long ones you never pay attention to overhead, 8-foot-long ones you see in supermarkets, circular ones, cold-cathode types you use on your desk, sodium lamps in parking lots, etc., etc. They all have different characteristics to contend with, but the fundamental distinction among them is whether or not they have heated filaments. Those that don't have heated filaments require only a single pair of wires; those that do work basically the same, but require in addition extra pairs of wires for the filaments. Since the two types are otherwise similar, this section concentrates on tubes with heated filaments.

A fluorescent tube can be thought of as similar to a vacuum tube, except it's not a very good vacuum. The glass tube has some gases in it (such as argon), and a drop of mercury liquid, which vaporizes when the tube is working. The glass in turn has some phosphors coating its inside (similar to a television tube). The tube works when a voltage is applied across the gas from one end to the other. (There is actually a cathode and an anode, but since fluorescent tubes are usually operated with AC, this is an unimportant distinction. AC is used rather than DC so that both ends have a chance to be the anode, reducing wear on the electrode.) The voltage is enough to cause the gas to ionize, which is to say, it forms a plasma. Getting a headache yet? The plasma gives off UV light, which the phosphor coating on the glass changes into visible light. Altogether, this is not a real efficient electrical process, but it is substantially more efficient than what normal incandescent bulbs do, which is to make a piece of metal so hot that it glows.

Safety Tip Since fluorescent tubes contain mercury, which is highly dangerous, don't go smashing fluorescent tubes! Leave them intact to be handled by those who know where to dispose of them without contaminating either people or the environment.

When a fluorescent tube has been off for a while, it requires a high voltage to get it started (because the mercury is liquid). In this state the tube is a high impedance. Cold cathode types (i.e., those without heated filaments) just require this high voltage to be applied for a certain length of time, after which they turn on. Those with a filament require their filaments to be heated, preferably for several hundred milliseconds prior to application of the high voltage; failure to preheat seriously degrades the life of the tube. The whole electronic ballast industry got off to a bad start because early electronic ballast designers overlooked this fact.

After the filaments have been heated and the high voltage has been applied, the tube turns on. In this state, it is approximately like a zener: passing double the current through the tube changes the end-to-end voltage perhaps 10%. Of course, passing double the current the tube is rated for almost doubles the light output, but it also degrades the life of the tube.

In this on state, the filaments still have to be heated, but with considerably less power than during preheat. Since the filaments are basically just pieces of resistive wire, this can be accomplished by reducing the filament voltage.

Other Converters

The most common type of load for your converter is another switching converter. The troubles potentially associated with having two converters in series are discussed in some depth in Chapter 6 on stability. Here it is sufficient to state that two converters, each of which is individually stable, can both oscillate if attached in series! The reason has to do with the negative input impedance of a switching converter, that is, increasing the input voltage causes the input current to decrease, because the converter is a constant power load. It is well known that negative impedance loads are used in oscillators of many types (intentional or unintentional).

SAFETY

Power supplies often generate high voltages, or work off an AC mains. The author feels very strongly that he would be remiss not to at least touch on a subject that is often ignored in labs under the pressure of schedules: personal safety.

To start off with a true horror story, the author was once in a lab working on 277VAC when one of the engineers accidentally touched this voltage inside a circuit he was probing. The engineer fell backward over the chair he was sitting on, crashed to the ground, and lay there twitching uncontrollably for a minute or more. Falling over probably saved his life, since it disconnected his hands from the line.

The important thing to know is that (aphoristically) current is what kills, not voltage. If more than a few milliamps passes through your heart, it can fibrillate (stop beating).

How much current is required depends on all sorts of factors (How humid is it? Are your palms sweaty?), but here is the practical safety tip:

Safety Tip If you're working on anything higher than 5VDC, keep one hand behind your back (e.g., hold onto your belt). This prevents current from flowing into one hand, through your heart, and to the other hand to complete the circuit.

All the time people tell me that this rule is too conservative: It's only 12V! It's only 60V! It's only 120VAC!—(Someone actually told me this last one.) But you can actually get a nasty jolt from a *1.5V D-cell*, if you go about it right. It's better to be safe.

For the same reason, you don't want to have a good conducting path to ground:

Safety Tip Wear shoes with rubber soles in the lab. This prevents current from flowing into your hand, through your heart, and down your leg to ground, completing the circuit.

Did you know that the metallic case of an oscilloscope is attached to the ground of the BNC inputs? In many labs people look at signals that are not ground-referenced by floating the oscilloscope, that is, defeating the three-wire connection on the oscilloscope's power cord by using a "cheater" to convert it to a two-wire connection. Unfortunately, the oscilloscope's case thus sits at the potential at which the probe ground is attached, just waiting for someone to come along and get a shock by touching it. (The 10MΩ impedance of the probe is in the *signal* path; the *ground* connection is a short.) There's a good *reason* that the plug is three wire, and as power engineers you should be the most familiar with it.

Safety Tip Buy an isolator for the oscilloscope, and throw away all cheaters. The isolator allows each probe's ground to be at any potential. If you are trying to use cheaters because of perceived noise in the system, you'd better look at system grounding rather than covering it up. (Try connecting all the ground posts of the different instruments together, and attaching them to the converter's return at only a single point. Also try attaching all instruments' power cords to the same power strip.) You can now get isolators with bandwidths up to 50MHz, and isolation of 1500V, more than enough for most power work.

Buying isolators for each oscilloscope in the lab may seem expensive at first, but management will become very receptive to the idea as soon as they hear the words "wrongful death lawsuit".

Another practice to be careful about is leaving a power supply running while you go out of the room to do something else. It is a particularly bad idea to leave a power supply running overnight, as in a burn-in test. There's always a VP walking through at such a time who feels obliged to stick a finger in; or else, have you thought about the janitorial staff sweeping the floor and snagging a dangling line?

Safety Tip If you're not present, the power supplies you work with should either be off, inaccessible, or surrounded by a barrier. A plastic link chain would be good. Or how about a piece of magnet wire strung in front across two chairs? A warning sign would be a nice addition (*not* replacement), and maybe you can make it multilingual and graphically enhanced (with a skull and crossbones).