7

Practical Design of Control and Monitoring Circuitry

CONTROL CIRCUITRY

Other chapters in this book have dealt with aspects of the power stage of a power supply. This chapter deals with the rest of the power supply, that is, the small-signal circuitry used to control the power stage and the circuitry used for monitoring the operation of the supply. (Except for the design of the error amplifier, which was covered in Chapter 6.) We start with some practical control circuitry.

Start-Up

It's always a problem: there's no power until the converter is running, but you need power to get the converter to run. The usual solution is to use a resistor and capacitor directly from the input power line to get things started, and then when the converter is running, provide power with a bootstrap winding from the main transformer (see Figure 7.1).

The way this works depends on the PWM having a UVLO (undervoltage lockout) with hysteresis. When V_{in} is applied, the capacitor is charged up through the resistor. When the UVLO threshold of the IC is reached, the circuit begins switching. It draws power out of the capacitor to run both itself and the switching transistor until the winding can supply enough power to run the IC. Figure 7.1 also shows a zener on the line, to prevent the voltage from rising so high that it damages the IC; this might be typically a 12–18V diode.

This circuit can require a quite large capacitor to store enough energy to keep things going until the converter is running. Consider a typical example: the PWM is a UC3825, which can require up to 33mA of supply current when running. Let's throw in an additional 10mA for gate drive, and a few milliamps for everything else, and we'll say it requires 50mA. Suppose the converter takes 10ms to come up to speed. (The winding on the transformer is voltage-limited by the other windings, and so typically won't provide

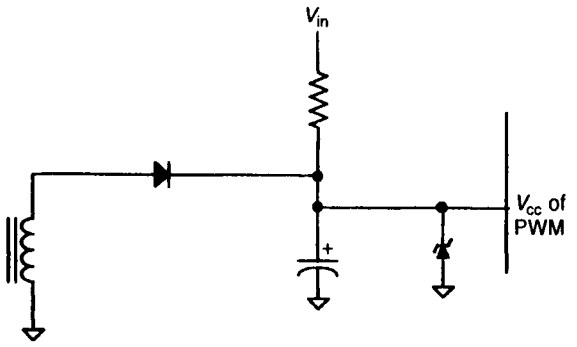

Figure 7.1 A bootstrap winding provides power after start-up.

any power until the main output is approximately in regulation.) The guaranteed hysteresis on the UC3825 is only 400mV, which means that if the voltage on the capacitor droops more than this, the PWM will go back into undervoltage lockout, cyclically hiccoughing. So we need to supply 50mA × 10ms = 500μC, with a voltage drop of only 400mV, needing a capacitance of $C = 500\mu C/400mV = 1.25mF$ (=1250μF)!

This particular problem can be solved by creating a larger hysteresis band for the PWM, as illustrated in Figure 7.2. This circuit has added in a PNP transistor that passes the current to the PWM (and any other loads). It works because the MOSFET isn't turned on until the capacitor voltage has reached a level set by its gate zener; once turned on, however, it stays on. The MOSFET in turn turns on the PNP, which passes current. By selecting, for example, a 12V zener, you can get approximately 5V of hysteresis (12V zener + 2V gate threshold for the MOSFET = 14V, and 14V − 9V UVLO = 5V hysteresis), so that the capacitor is reduced in size by a factor of 5V/400mV = 12.5, from 1250μF down to 100μF, a gigantic saving in size.

Of course in both these schemes the resistor continues to dissipate power as long as V_{in} is present: the winding provides the power to run things, but there is still voltage applied across the resistor. This problem is mitigated, however, by the consideration that

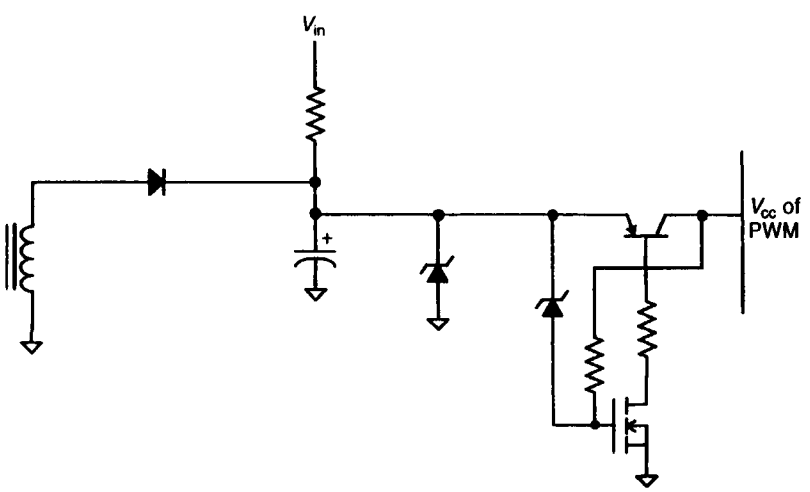

Figure 7.2 Increasing the UVLO hysteresis makes the start-up capacitor smaller.

the resistance value can be almost arbitrarily large. All a large resistance does is give a delay period between the time power is applied and the time the converter starts—it has no effect on the slow-start time of the converter. For example, with the second scheme, suppose that the resistor is 10kΩ, the input voltage V_{in} is 28VDC, and the normal output of the winding is 15V. Then the capacitor charges up to the required 14V, and the converter starts, in

$$14 = 28\left[1 - \exp\left(\frac{-t}{10\text{k}\Omega \times 100\mu\text{F}}\right)\right]$$

or $t = 700$ms. The steady-state power dissipated in this resistor will be only

$$P = \frac{(28\text{V} - 15\text{V})^2}{10\text{k}\Omega} = 17\text{mW}$$

If the converter has an output short, the winding will produce no power, and so the resistor might have to provide power continuously for extended times. But even in this case, it is dissipating only 78mW, which can be handled by a 100mW resistor.

There are other, more complex, schemes for start-up power—for example, using a MOSFET in series with the resistor to turn it off completely once the converter is running. This supposedly allows use of a smaller resistor value without impacting component size (relying on the pulse power rating of a wirewound) and allowing less turn-on delay. During an output short, however, there is still going to be dissipation in the resistor, and of course the smaller value resistor will dissipate more power, requiring a large component anyway. Thus, it seems that nothing is to be gained from more complex schemes, unless it is necessary to minimize turn-on delay (not slow-start speed, remember).

Soft Start

Soft start (or slow start) has been mentioned, but without specifying what was intended. The idea is straightforward. When a control IC first receives power, the output voltage the feedback senses is of course zero (or at any rate lower than it should be, in the case of a nonisolated flyback). This causes the duty cycle of the converter to want to go to its maximum value. Were it allowed to do so, very high and potentially destructive currents would be drawn from the input and through the power devices in an attempt to charge up the output capacitors. Instead, the duty cycle of the converter is limited to a maximum value that increases linearly with time, usually controlled by the charging of a capacitor. Once the capacitor is fully charged, the duty cycle will be whatever it needs to be to regulate the output voltage.

> **Practical Note** Always use soft start to protect both the line and the converter.

Soft start is also frequently used for recovery from a fault such as overcurrent: when an overcurrent condition occurs, the soft-start cap is discharged, causing the duty cycle to come up slowly again while recovering from the fault. When this recovery is cyclical, it is known as hiccough mode. Since the soft-start cap is charged by a constant current source (whence the linear increase in voltage), this suggests a way of making the hiccough period different from the soft-start time: when an overcurrent is detected, switch on a transistor

that pulls some current from the soft-start through a resistor, decreasing the speed at which the soft-start cap is charged.

On some older ICs, a soft-start pin may not be available. In this case, the same effect can be had by attaching an *RC* from the reference to the noninverting pin of the error amplifier: this causes the voltage to which the converter is attempting to regulate to slowly increase.

Sequencing

Related to start-up is a requirement, occasionally seen, that certain voltage(s) be up and stabilized before others come up, or that one output always have a higher voltage than another. For example, if a +5V supply runs some TTL that controls some +12V relays, the TTL may need to be operating before the relays are powered to ensure that the relays don't go into unwanted states.

Using a flyback converter can be a good choice to meet this requirement, because the output voltage on every winding is clamped by the other outputs. Therefore each output is proportional to its final output voltage: that is, if a +5V output is at 2.5V, a +12V output will be at 6V, etc.

Use of a converter with an inductor doesn't provide this scaling. Instead, output voltage depends on both the output capacitor and the load. Thus for this case, relative voltages for the various outputs during start-up can be controlled to some extent by how much capacitance is placed on the output: placing a large capacitance on the +12V can ensure that it comes up last.

Finally, if an output has to be completely up and stabilized before some other output is allowed up, there may be no choice but to use a switch such as a p-channel MOSFET. The MOSFET could be controlled by a comparator that detects that the first voltage is above minimum regulation.

Rather less commonly, there may be a requirement for turnoff sequencing when the converter is turned off: again with the example of the relays, it may be required that the +12V be entirely removed before the TTL goes down. In this case, neither the flyback (since the converter is not delivering power, the windings don't clamp each other) nor the amount of output capacitance (because of load current ranges) is really enough to guarantee sequencing; it is in practice mostly dependent on the loads. A switch seems to be the only way to perform this function.

Feedback

Chapter 6 on control theory discussed in great detail the design of a compensation network for the control loop, including selection of feedback resistors for the voltage being controlled. Frequently, though, the secondary voltage you're trying to regulate must be galvanically isolated from the primary where the error amplifier is. That is, no DC connection is allowed between the two. In such a case, a method of transferring the DC information across the boundary is required before the resistor feedback can be implemented.

There is no end of methods for accomplishing this isolated feedback. We mention a few that are popular ones:

1. Some people use an optocoupler, and perhaps attempt to linearize it with a second opto as feedback. (This approach has problems with optical gain affecting converter bandwidth if a single opto is used; it has problems if the two optos are not in the same package; and temperature and aging bring up additional problems.)
2. Other designers do a voltage-to-frequency conversion, send the frequency- (or pulsewidth-) modulated signal across the barrier with an opto or a transformer or just a capacitor, and then convert frequency back to voltage. (This approach is quite parts intensive.)
3. Still others use an instrumentation amplifier. (This is OK until you get a request for a 500VDC hi-pot test!)

The author's favorite method is shown schematically in Figure 7.3. It provides true high voltage isolation, uses few parts, has wide constant bandwidth, and can be made almost insensitive to temperature variation.

This method for isolating feedback works as a forward converter running from the (output) voltage to be measured. A BJT is switched by a secondary winding of the main power transformer. (If the freewheeling diode is not present, as in a flyback, the BJT is driven directly from the transformer, and it may be necessary to add a series base diode to prevent the BJT from emitter–base zenering during the on-time of the power FET.) When the BJT turns on, the output voltage is applied across the primary of a very small transformer. In a typical example, the voltage being sensed might be 5V. Then the transformer might be a step up of 5 : 1, so that the voltage applied on the schottky on the secondary side of the small transformer, which is grounded on the primary side of the converter, is 25V. The schottky and the capacitor then form a peak detector, and this

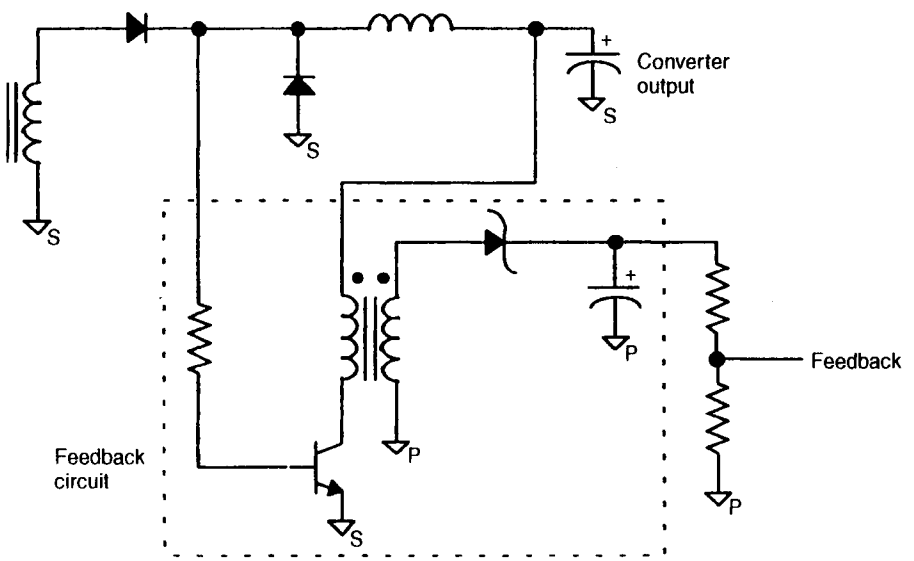

Figure 7.3 Isolated feedback using a small forward converter, driven by the main transformer.

voltage is divided down to get to a level suitable for the error amplifier. Note that the two grounds are denoted by "S" (for secondary) and "P" (for primary).

The inaccuracies of this method are due to the collector–emitter voltage of the BJT, which can be very low at low currents; the winding resistance of the transformer's primary, which can be very small because the current is only a couple milliamps; and the forward drop of the schottky, which is only a few tenths of a volt out of 25V. It is straightforward to achieve 2% accuracy with this method, surpassing all the other methods substantially. The bandwidth can also be made quite high—it basically depends on the time constant set by the peak detecting capacitor and the resistor divider network.

Current Limiting

A frequent requirement of power supplies is that they be current-limited; that is, they must be self-limiting in some fashion with respect to the amount of output current they can source. Such a requirement can be a little more complex than it seems at first blush, because "shorts", the generic term for output faults, can have differing impedances, ranging from levels that draw just slightly more current than they should to basically 0Ω. A short that has negligible impedance is usually called a "hard" short, and all others are referred to as "soft" shorts.

Despite users' frequent claims that the supply need be protected against hard shorts only, it is the author's firm belief that supplies should be protected against soft shorts as well. The goal here is not only to protect the user from experiencing currents that might damage the units, but also that the power supply should be able to protect *itself* from damage: not all shorts are 0Ω. The usual method of accomplishing this is routinely implemented inside typical PWMs with a two-stage current limit. The output of a current sense resistor (or a current sense transformer) in series with the switch is fed into a current limit pin of the PWM (this pin is frequently the same pin used for current feedback for current mode control). If the voltage on this pin exceeds a certain level, the PWM shuts off the current pulse going to the switch and doesn't restart it until the next cycle: this is called pulse-by-pulse current limiting. If the voltage on the current limit pin goes higher, and reaches a second level, the PWM terminates the pulse and re-soft-starts. This latter method is commonly called hiccough mode, burp mode, or various other names.

There is a major problem with relying on these methods if there is more than one output of the converter. The current limit sense on the primary has to be set up to prevent current limit from being reached under normal power (i.e., with all the power from every winding at maximum load summed up together). But now suppose that only one output gets a short. Then, if the other windings are at minimum power, almost the entire power rating of the converter has to go through that one output before current limit trips. The usual result of this situation is either a blown rectifier diode on that one output or an opened wire on the transformer (or inductor if the topology has one). In any case, the converter has failed to protect itself.

There doesn't seem to be a cheap solution hence. Assuming that you don't want to put current-limited postregulators on every output, the best that can be done is to sense each output current individually (with a sense resistor between the return of the output and secondary ground, so you don't need common mode rejection), let each one go into its own open collector comparator, and OR them all together. The ORed signal can then be used to control the current limit or shutdown pin of the PWM (if the converter is nonisolated), or it can control this pin through an optoisolator.

Switching Frequency

The switching frequency of a converter is defined as the number of times the convertor goes through the same set of states each second. Thus, for example, a flyback converter has a switching frequency of 200kHz if the switching transistor turns on and then turns off again 200,000 times per second.

It is necessary to be slightly wary when selecting the timing components for an IC to run at a certain switching frequency. Some ICs run an oscillator at one frequency and then use the first pulse of the oscillator to drive one output and a second pulse to drive a second output, or as a blanking signal to prevent duty cycles greater than 50%; the net result is that the actual converter frequency is half the oscillator frequency. Thus an IC that claims to be able to run at 1MHz may actually run a converter at only 500kHz.

There is also a practical limitation on maximum switching frequency. The problem is not with the controller ICs, some of which currently can run at 2MHz; the problem is with the gate charge of the MOSFETs. Gate current is proportional to frequency, so that as frequency goes up, so do losses in driving the gate; and of course, switching losses are also dependent on frequency. In very recent times, manufacturers of MOSFETs have started to come out with devices that have substantially reduced gate charge. For very high switching frequencies, this sort of MOSFET is a must.

Synchronization

A final topic of control circuitry is synchronization. It is sometimes required that the power supply switching frequency be synchronized with a master clock in a digital system, often with the idea that the noise spikes will have less effect on the digital parts' noise margin if the spikes occur at the exact time that the parts are undergoing state transitions. (The concept may well be flawed, since it ignores propagation delays inside the power supply.)

Anyway, data sheets never seem to explain what is required to synchronize a PWM, they just show an example. To rectify this oversight, here's a brief description. PWMs work by comparing the output of the error amplifier (approximately a constant compared with the switching frequency) with a ramp. As shown in Figure 7.4, they turn on a switch at a regular interval, and terminate the pulse when the ramp and the error amp output are equal. The next pulse then starts (a short time after) when the ramp reaches a certain level, internally set by the PWM. The idea of synchronization is to force the ramp to terminate prematurely by injecting a signal on top of it (Figure 7.5). The little pulse added "puts it over the top," and the next pulse begins sooner than it would have otherwise.

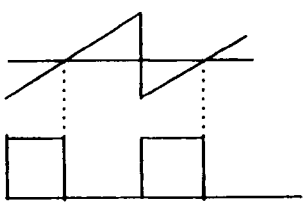

Figure 7.4 The gate drive turns on at a regular interval, and shuts off when the ramp is equal to the error amplifier's output.

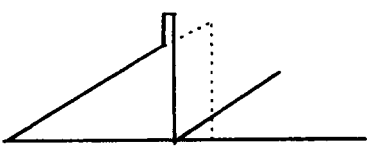

Figure 7.5 Synchronization occurs by adding a pulse to the ramp, starting the next cycle.

From this description, it is clear what needs to be done to synchronize a converter: (1) the free-running frequency of the PWM has to be lower than the synchronization frequency; (2) a short pulse, having the following characteristics, has to be added to the ramp: amplitude great enough to trip the PWM's comparator, and falling edge corresponding to the new period of the PWM. The pulse has to be short because the PWM is forced off while the pulse is present (i.e., the pulse acts like dead time).

Thus, a pulse can be generated by a TTL device, for example, and capacitively coupled into the timing capacitor (see Figure 7.6). The resistors R_1 and R_2 form a divider that can be used to scale the pulse. R_1 should be a low value (some tens of ohms maximum) to avoid disturbing the ramp, which will be integrated by the RC formed by R_1 and the timing cap. Additionally, the end of the synch circuit coming from the TTL signal should have a relatively high impedance to ground. One method of doing this is shown in Figure 7.6. This requirement exists because the coupling capacitor is quite large and is in parallel with the timing cap, and so would affect the free-running frequency, were it to be grounded.

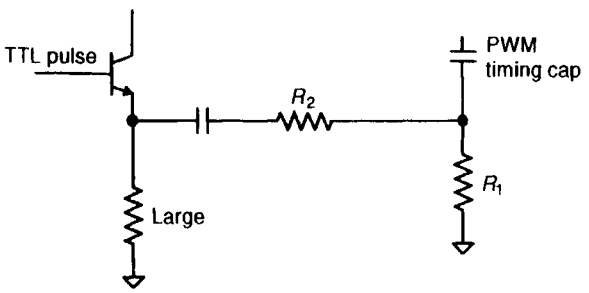

Figure 7.6 How to synchronize a PWM.

Let's mention one more little fact about synchronization. From our description of how synchronization works, it is clear that the peak-to-peak amplitude of the ramp is reduced by this process of premature termination. As we know from Chapter 6 on loop stability, the ramp amplitude is one of the factors directly determining the gain of the loop; by synchronizing a converter, therefore, you directly influence its bandwidth and phase margin. You should always set limits to the range of allowable frequencies for the synchronization, and check the power supply's loop at the maximum frequency (=minimum ramp amplitude) as well as when it is free-running.

Practical Note A practical limit is to not let the synch frequency exceed 1.5 times the free-running frequency of the converter, if it can be avoided.

MONITORING CIRCUITRY

How to Monitor Voltage

One of the most frequent monitoring requirements is to produce a signal indicating when the output voltage(s) are valid. This means at least that the voltage is above a minimum value, and usually it must be below a maximum value as well. Both monitoring requirements can of course be met with a single comparator with hysteresis. Choosing

the component values that accomplish the monitoring can be quite laborious, however; the author has found the use of a numerical or symbolic algebra computer program to be quite a time-saver in this regard. One just writes down the Kirchhoff equations and lets the computer find the values. (Assuming that a solution exists—it is possible to select a reference voltage for the comparison that forces some resistors to be less than 0Ω!) Generally, the reference voltage needs to be between the minimum and maximum trip points.

Voltage References

There is a need for a certain degree of caution when accepting a spec that calls for voltage regulation tighter than ±5%. Typical references on PWMs have several percent tolerance, and then the 5% may include monitoring tolerance, which means that a lot of your tolerance is eaten up by using 1% resistors—and since you probably can't get the exact resistor value you want, you always have to round off, adding another 0.5% or so. Finally, requesting better than 5% in an *isolated* feedback is really pushing it, because there are additional errors in crossing the isolation barrier. As indicated above, even quite good schemes have something like 1% error. If a spec calls for much better than 5% output regulation with isolation, it's best to plan on a postregulator on the secondary side.

As an example of what can be expected from a PWM reference, consider the UC3825, whose "features" section proudly announces a "trimmed bandgap reference (5.1V ± 1%)." The first thing to notice is that the 1% is for industrial and military grade parts only (the UC2825 and UC1825); the commercial part is 2%. This 2%, however, is at nominal conditions only; over line, load, and temperature, the commercial part is 3% (and the others 2%). Additionally, there is long-term drift: after 1000 hours, the parts may have changed an additional 0.5%. Thus, a typical off-the-shelf PWM claiming 1% reference voltage really gives 3.5%!

Furthermore, if you need a divider to bring the monitored voltage down to the reference voltage, there is an additional 1% error in the output voltage (assuming 1% resistors), and now you're up to 4.5%.

Of course, there are things that can be done to shave this percentage a bit. The most obvious (and the least costly) is to go to 0.1% resistors—2 cents a piece is all they cost. Then, you could eliminate line and load variations to the chip; but it's still going to be more than 2%. If you start checking through available zeners and other 2- and 3-pin devices, the story is much the same. The bottom line is, if a spec calls for much better than 5% tolerance on the regulated output (never mind the tertiaries), you're going to end up with a moderately expensive IC for a reference. (The REF01 is an excellent choice.) Perhaps a good plan is to carefully inquire *why* the user thinks such a tight tolerance is necessary. Is it just for the sake of a safety margin?

How to Monitor a Negative Supply Without a Negative Rail

Sometimes you'll have a negative output voltage in a nonisolated converter, and it sure would be nice not to have to stick in a transformer to monitor that rail. But since you are running all the control circuitry from +12V and ground, you can't just feed the negative supply into an IC somewhere. The solution, once seen, is obvious: the negative supply can be inverted (and divided, if necessary) by using the virtual ground of an op amp. It is clear from the circuit shown in Figure 7.7 that essentially any negative voltage can be monitored

this way, even hundreds of volts below ground. There is only one caution: to keep the negative voltage from being applied directly to the IC pin, the power to the opamp should be present before the negative voltage is applied. If this condition can't be guaranteed, the IC can still be protected from damage by ensuring that the resistor from the inverting pin to the negative voltage is large enough to ensure that the current is limited to, say, 1mA or so. Then the diode (which needs to be a schottky) will prevent the IC from seeing voltages more than 0.3V below its negative rail.

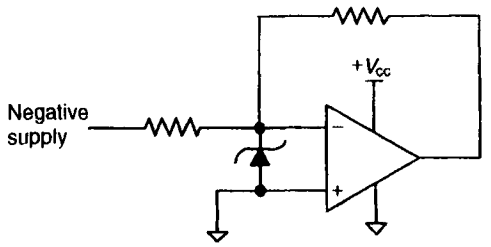

Figure 7.7 Monitoring a negative rail with a single supply: use a virtually grounded op amp. Also shown is a protection diode.

Why You Should Always Use Hysteresis on Comparators

A little hysteresis costs only one resistor, so don't think about saving 2 cents by dropping this component. Since a comparator has only finite gain, there will be a small range of voltages in which the comparator will try to operate in some sort of linear mode. The result is that the comparator may oscillate, slewing back and forth against its rails at its slew rate. This may be OK if it's driving a one-shot latch, but most of the time such oscillation gives oscillating failure signals, great for swamping a microcontroller's interrupt line. Worse yet (and the author has seen this), if the comparator is driving some sort of shutdown, you may have a closed loop system that sits at precisely the point where the comparator is oscillating and the system is just on the verge of shutdown, but doesn't quite get there. Surely one resistor isn't worth all this?

Although it is obvious post facto, occasionally people don't realize that hysteresis can be used on a comparator regardless of whether it is inverting or noninverting. Just for reference, then, if the signal you want to monitor goes into the noninverting terminal of a comparator, do it as illustrated in Figure 7.8. On the other hand, if the signal goes into the inverting terminal, use the circuit arrangement shown in Figure 7.9.

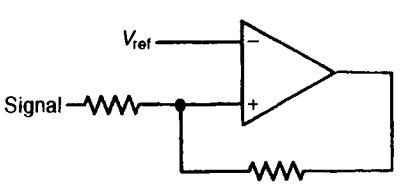

Figure 7.8 Comparator hysteresis when the signal is applied to the noninverting input.

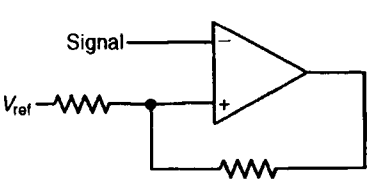

Figure 7.9 Comparator hysteresis when the signal is applied to the inverting input.

The important thing to remember is that the hysteresis always goes back to the noninverting terminal, regardless of where the signal is. It should be noted in Figure 7.9 that some amount of current will be fed back from the output of the comparator to the reference voltage, so the reference needs to be stiff enough to prevent feedback from affecting its value.

Resistors and Shunts

Using a resistor to measure current is obvious, but a little thought shows that the resistor value has to be pretty small to avoid excessive power loss at moderate currents. Additionally, since resistors have inductance, if the current being measured has an AC component, the voltage across the resistor may be dominated by the AC current times the impedance of the inductance, rather than by the DC current times the resistance. (Don't even think of using a wirewound for such an application, unless it is noninductively wound.)

As discussed in the chapter on components, a shunt is a current-measuring resistor typically consisting of a wide, thin strip of manganin. Typical values range from 50mV at 5A (=10mΩ) to 50mV at 500A (=100μΩ) and less. The next section discusses how these small voltages can be measured. Here, it is to be observed that even shunts have some inductance. It might be tempting to try to compensate out the inductance, thus obtaining a high frequency, high current sensor, by paralleling the shunt with a capacitor, but only until you realize the size of the putative cap. Suppose the smallest one, a 5A shunt, has 20nH of inductance. Then the shunt's time constant is $L/R = 20\text{nH}/10\text{m}\Omega = 2\mu\text{s}$. To compensate this requires $C = t/R = 2\mu\text{s}/10\text{m}\Omega = 200\mu\text{F}$! A more practical method is shown in the next section.

Differential Amplifiers

To sense output current on the high side, a resistor (or shunt) is placed in series with the line (see Figure 7.10). This requires a somewhat difficult measurement to be made: that is, a small voltage difference has to be sensed on top of a large common mode voltage. For this purpose, a differential amplifier is used. ("Instrumentation amplifier" is another name for the same thing.)

It is certainly easier to measure current on the return side of a line rather than on the high side, because all you need is an op amp amplifier. Sometimes, though, a return side measurement is undesirable (see Figure 7.11): for example, there can be problems with

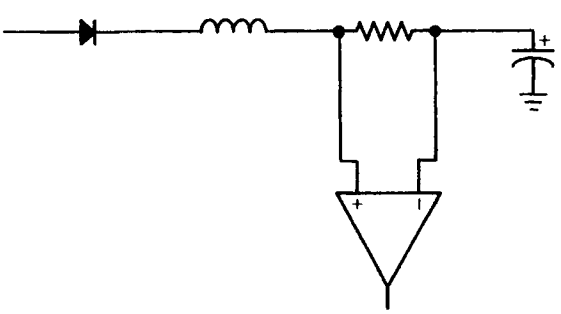

Figure 7.10 Measuring current on the high side requires a differential amplifier to eliminate the common mode signal.

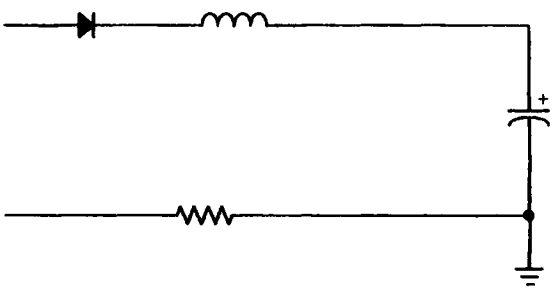

Figure 7.11 Monitoring current on the low side is not recommended.

having your load at some 50–100m VDC above "real" ground, and potentially (pun intended) much higher than that at AC. Furthermore, having a load at one ground and a power system at a different ground can have adverse effects on EMI control. So for various reasons, it may be necessary to monitor current on the high side.

> **Practical Note** Always do the current sensing in the power line, not the return line; don't use a ground elevated from the other grounds.

The simplest differential amplifier can be built with a single op amp as illustrated in Figure 7.12. In detail, this works as follows. Suppose a gain of 10, as in Figure 7.12. The more positive input gets divided by 10/11, which is the resistor divider value. Then the op amp works to force the inverting input to also be $(10/11)V_+$, where V_+ is the positive input voltage and V_- is the negative input voltage. The current through the 1kΩ to the inverting pin is thus

$$I = \frac{V_- - (10/11)V_+}{1\text{k}\Omega}$$

This same current flows through the 10kΩ, so that it has a voltage drop of

$$V = 10\left(V_- - \frac{10}{11}V_+\right)$$

The output voltage is then the voltage at the inverting pin minus this, or

$$V_{\text{out}} = \frac{10}{11}V_+ - 10\left(V_- - \frac{10}{11}V_+\right) = 10V_+ - 10V_-$$

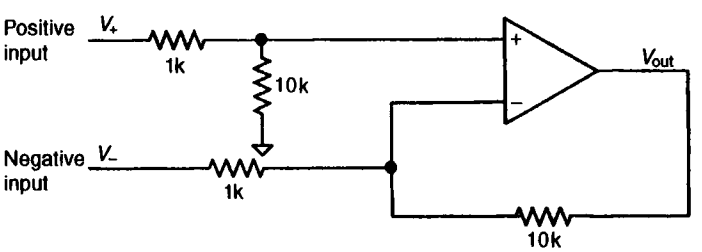

Figure 7.12 A single op amp differential amplifier.

which is to say, the voltage common to both inputs has been rejected, and only the difference between the two (with a gain of 10) is at the output.

As always, component tolerances produce limitations on this result. The tolerance of the resistors used is the dominant contributor to the error. Any error in the matching of the two ratios causes some portion of the common mode voltage to feed through to the output.

> **Practical Note** A good estimate of the rejection of the differential amplifier is that the common mode voltage is rejected by the tolerance of the resistors. For example, if there is 5V of common mode, and the resistors are 1%, there will be approximately 50 mV of output even when there is no difference between the two inputs.

This in turn leads to a limitation on the signal size; assuming that 0.1% resistors are the best that can be obtained at a reasonable cost, the ratio of common mode to signal should be less than 1000 : 1.

You also want to think very carefully before placing capacitive filters on the dividers to clean up noisy signals. Any imbalance in the capacitors will show up as reduced common mode rejection of AC. Even if you don't add capacitors, there will still be imbalances in stray capacitances, and unless your resistances are small, these strays can cause lack of adequate common mode rejection at moderately high frequencies.

There are also differential amplifier arrangements with two and three op amps (as well as IC versions). They are used in some applications because of their essentially infinite input impedance, which doesn't load the signal source. (The single op amp differential amplifier shown in Figure 7.12 loads it with 11kΩ.) For monitoring current, this is irrelevant, as the output line presumably has very low impedance anyway, and a single op amp amplifier is usually all that is required.

Compensating Shunt Inductance

Using a differential amplifier, it is possible to compensate out the inductance of a shunt with reasonable value capacitors, as illustrated in Figure 7.13. Using the same example as above, to get a 2μs time constant, a capacitor is required $C = t/R = 2\mu s/1k\Omega = 2nF$. Conceptually, what is happening is that the unwanted sudden rise in voltage across the shunt due to its inductance is integrated away by the *RC* filter, without affecting the DC response.

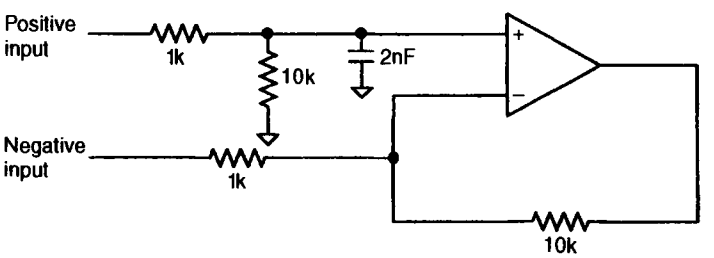

Figure 7.13 Compensating the inductance of a shunt.

Fail Should Be Low

Here's a practical tip that's obvious after reading—but surprisingly common to see violated. It is frequently required that a signal be provided indicating a problem with the converter in the event of failure. (Of course, this assumes that somehow whatever is monitoring this signal has a way of remaining powered—check that this assumption is warranted before acquiescing in a "converter fail" requirement!) Since the converter has failed, there is no guarantee that a signal can be pulled up to *indicate* failure, so:

> **Practical Note** Fail signals should always be low.

If sink capability is needed, a PNP with a base pull-down resistor provides a nice passive low signal; it has to be actively driven to be high.

Driving That Red LED

In multicard cages, there may be a requirement that a red LED be illuminated to indicate converter failure. This entails the same dilemma posed by a "converter fail" requirement: How can an LED be powered when the converter has failed? (This almost, but not quite, ranks up there with a requirement the author once saw that a signal be sent when the converter was about to fail!)

The solution diagrammed in Figure 7.14 may be acceptable. As long as the converter is active, the current to the LED is shunted away from the red LED by the BJT, so the red LED is off; and the other BJT pulls current through the green LED, and so the green LED is on. If the converter fails, the red LED turns on and the green LED turns off. If main input power fails, both the red LED and the green LED are off. Thus, the status of both the converter and the input power can be ascertained visually. This is probably the best that can be done to alert operators in the event of failure.

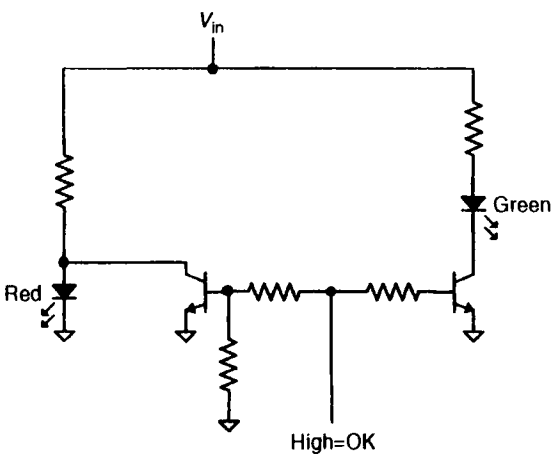

Figure 7.14 Driving signaling LEDs so that both converter failure and power input failure can be ascertained visually.