6

Practical Feedback Design

INTRODUCTION

Design of the feedback network of a converter is one of the most important aspects of a converter's design, and yet (with the exception of magnetics design) it is also the one most often misunderstood. Power supply designers can be classified into two groups: those who really understand what compensation (=feedback design) is about and those who don't. The former can compensate a converter within a few hours, and have it perform precisely as designed: this chapter will promote you into this group.

Feedback design is slightly more sophisticated mathematically than other aspects of power supply design, but this shouldn't frighten you off: the math is not used in day-to-day work, but is important only to be able to deal with new questions that sometimes arise. So to remind you of the necessary background, and unlike the other chapters, this chapter discusses some topics that aren't practical in the sense of being immediately visible in the lab; we start with a refresher on complex numbers, and ramp up to transfer functions, concepts that are essential to understanding what compensation is about. After this beginning, we show the practical part of feedback design in great detail: it really is possible to compensate a converter in a couple of hours, and have it perform exactly as desired on the first try!

REFRESHER

Logarithms and dB

Decibels are the standard unit for measuring the gain of a transfer function (which we'll define below), such as the open loop or closed loop gain of a converter. We start our refreshment with dB, and with logarithms, which are almost the same thing.

In engineering, the logarithm base 10 is almost always used, rather than the natural logarithm, which is base e. Remember what this means: if something gets 10 times bigger, its logarithm is increased by adding 1; multiplication is replaced by addition, and division is replaced by subtraction, making the system very convenient.

Decibels (dB) refer to logarithms that have been multiplied by 10 or 20 to make nice-looking numbers (i.e., approximately integers). For our purposes, the multiplying factor is always 20. Thus 10V is expressed in dB as

$$20 \times \log(10\text{V}) = 20\text{dBV}$$

100mV is

$$20 \times \log(0.1\text{V}) = -20\text{dBV}$$

Amplitudes less than 1 correspond to dB less than 0.

Decibels, being logarithms, convert multiplication into addition. So if you have a gain of 8 [$=20 \times \log(8) = 18\text{dB}$] and you want to reduce it to 4, you must multiply it by a factor of 0.5, so that in dB this is adding $20 \times \log(0.5) = -6\text{dB}$: multiplying corresponds to adding. Of course multiplying by 0.5 is the same as dividing by 2, so that in dB this is subtracting $20 \times \log(2) = 6\text{dB}$: dividing corresponds to subtracting. Either way, either multiplying by 0.5 or dividing by 2, the end result is the same, $18\text{dB} - 6\text{dB} = 12\text{dB} = 4$, as it must be.

Just as an aside, you sometimes hear people jokingly talk about "bels" which are taken to be 10 decibels each; no such unit is ever used in practice.

Complex Numbers

Complex numbers are constantly used during feedback design and measurement, although the practice is not always obvious. The most transparent case of use is in making a Nyquist plot, which is basically a graph of the imaginary part of the loop's gain versus its real part. Nyquist plots are discussed later in connection with Figure 6.34. For now, suffice it to say that complex numbers and functions are essential.

A complex number is an object that consists of two numbers, a real part and an imaginary part, written as:

$$(\text{Re}) + i(\text{Im})$$

where Re is the real part and Im is the imaginary part. (Engineers call the imaginary unit j instead of i to avoid confusion with the symbol for current). For our purposes, we don't care about i being the square root of -1; all we need to know is that there are two numbers, Re and Im.

Since a complex number consists of two numbers, we can draw it on a graph with the amount of real on the x axis, and the amount of imaginary on the y axis, as in Figure 6.1. Redrawing this picture, we have Figure 6.2, from which it is obvious that a complex number can equally well be determined by two other numbers, one the distance from the origin (0,0), the second the angle counterclockwise from the positive real axis, which is to say a magnitude and a phase.

The upshot of this is that in electronics, whenever we need to represent something with both a magnitude and a phase (and this happens all the time), it is natural to use a complex number, which encodes both pieces of information together.

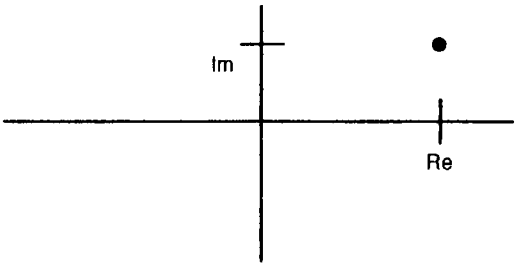

Figure 6.1 A complex number consists of two numbers, so it can be shown on a graph.

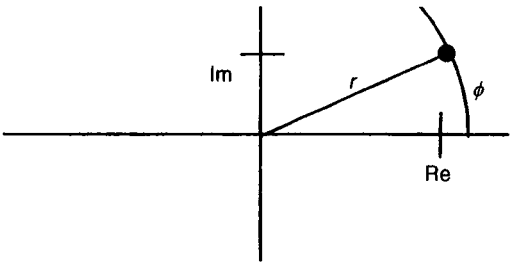

Figure 6.2 A complex number also can be shown as a distance and an angle.

Complex Functions

Since we can use a complex number to represent simultaneously a magnitude and a phase, we can have this complex number depend on a parameter (in particular, magnitude and phase depend on frequency). For the functions of interest to us in stability design, there are two interesting aspects of these functions: At what parameter value (read frequency) are they equal to zero, and when are they equal to infinity? These two conditions are referred to as zeros and poles of the function, respectively.

EXAMPLE

$$\frac{x-2}{x-3}$$

Clearly, this function is zero when $x = 2$, so it has a zero at the complex number of magnitude 2 and phase = 0, (Re) = 2 and (Im) = 0. (Think about Figure 6.1.) It becomes infinite when $x = 3$, so it has a pole at the complex number of magnitude 3 and phase = 0.

As another example, we'll soon see that a capacitor has a complex number associated with it (its impedance) that is a function of frequency, $1/sC$ (here s is related to frequency); and an inductor has an impedance sL. Figure 6.3 shows the output filter of a switching converter (without ESR of the capacitor, or winding resistance of the inductor).

So, this forms a voltage divider whose attenuation depends on frequency, the output is equal to the input times

$$\frac{1/sC}{sL + 1/sC} = \frac{1}{s^2LC + 1}$$

If it is conceptually easier, think of this as if it were a resistor divider with one resistor having a value sL and the other a value $1/sC$. This function has no zeros, but it will have a

pole if $s^2LC = -1$, or $s = \pm i/\sqrt{LC}$. So here, two poles occur at a magnitude equal to the resonant tank frequency, and at angles of 90° and 270° (being pure imaginary, no real part), as indicated in Figure 6.4.

The physical significance of this, of course, is that the LC tank has a resonance at this frequency, so that the output is indefinitely amplified from an input at this frequency. Of course, a real circuit always has resistance, and so the amplification is not infinite—that is, it doesn't really have two poles on the imaginary axis; it has some real part also.

For those already in the know, observe that there is no factor of 2π here, because s actually corresponds to $i\omega = 2\pi if$, not frequency.

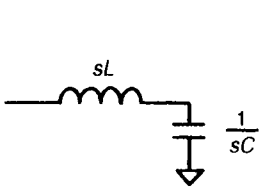

Figure 6.3 Complex impedance of a capacitor and an inductor.

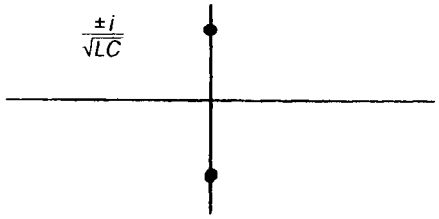

Figure 6.4 An LC tank has poles at its resonant frequency.

What Is a Transform?

Now that we know about complex functions, we can turn to transforms, which we use, for example, any time we look at a frequency spectrum. Most generally, a transform is a connection between two different ways of representing the same data. For our purposes, the two representations are

As a function of time. (*Example*: What was the voltage during the first second?)
As a function of frequency. (*Example*: There was 15V at 60Hz, plus 3V at 180Hz.)

Any function of time can be represented as a function of frequency, and vice versa, so if we talk about a spectrum (which is the transform of some function of time), we have exactly the same information contained in the original (time) function, just encoded differently, and perhaps more conveniently for some purposes.

Two Transforms

For engineers' purposes, there are two types of transform, Fourier and Laplace. Suppose then that we have a voltage that depends on time, $V(t)$. The Fourier transform is

$$V(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{i\omega t} V(t) dt$$

while the Laplace transform is

$$V(s) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{st} V(t) dt$$

What's the Difference?

The difference between the two types of transform is that the Fourier transform considers all time; that is, it assumes that you've been measuring the (in this case) voltage forever. A Laplace transform considers only voltages that start at time = 0: that is, the voltage was turned on at time 0 and was zero before that time.

Practically, this means that a Fourier transform is used for something like power supply noise emission, which can be measured as slowly as you please. A Laplace transform is used to describe how a waveform changes with time: for example, what does the voltage on this capacitor do when the power supply is turned on?

In this chapter, we only use Laplace transforms, because we are interested in how a power supply responds when something is done to it at time 0, (e.g., a load step is applied at that time). Fourier transforms are used implicitly in the chapter on EMI.

Transform of *C* and *L*

It is straightforward to determine the Laplace transform corresponding to the impedance of a capacitor. Of course, this impedance, and that of an inductor, is used all the time in designing filters, and in particular in filters used inside the loop of a converter, as mentioned above. We know that

$$C\frac{dV}{dt} = I$$

If we let $V = e^{st}$, so that the voltage is a sine wave [this is the same because $e^{i\omega t} = \cos(\omega t) + i\,\sin(\omega t)$], we find that

$$I = Cse^{st} = CsV$$

or that impedance

$$Z \equiv \frac{V}{I} = \frac{1}{sC}$$

as promised before. We didn't have to actually evaluate the integral in the definition of the Laplace transform because the integration is implicit in the solution of the differential equation—but this is not at all a practical concern. Similarly, we can find the impedance of an inductor by knowing that

$$V = L\frac{dI}{dt}$$

and substituting e^{st} for V again and integrating,

$$LI = \frac{e^{st}}{s} = \frac{V}{s}$$

or $Z = sL$. This then justifies the voltage divider example above.

TRANSFER FUNCTIONS

What Is a Transfer Function and Who Cares?

A transfer function is a compact method for expressing what a linear system will do at its output when you do something at its input. For example, the box in Figure 6.5 might be a complete power supply, and so the transfer function represents audio susceptibility: that is, how much input noise is attenuated (or amplified!) into the supply upon injection.

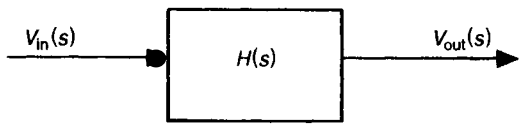

Figure 6.5 The transfer function describes what the system does to change the input into the output.

We already know that any function of time can be broken up into its constituent sine waves—that's what it means to take its transform. If I apply a sine wave of frequency ω to this system, the output will also be a sine wave of frequency ω, but with a magnitude and phase different from the input. The ratio of the output to the input is $H(s)$, the transfer function of the system. We use s, where s is complex, because it encodes both the amplitude change and the phase shift due to the system.

We can thus tell the response of the system to any input by breaking up that input into its constituent sine waves, applying $H(s)$ to each of them, and reassembling them at the output.

With power supplies, we are concerned when there is a small disturbance in the loop, or an error: instead of growing bigger, the disturbance should die away with time. This is what the transfer function can tell us. If the transfer function shows that a loop is unstable, as in Figure 6.6, then a disturbance will be amplified, with the result that the circuit oscillates.

Composition Law for Transfer Functions

When two systems are in series, there is an easy rule for deriving the composition of their transfer functions, as shown in Figure 6.7. This rule enables us to figure out the transfer function for a closed loop system like that of Figure 6.8, which could represent a power supply. The box labeled $H(s)$ is the power stage, since it transforms the input to the output, and the box labeled $G(s)$ is the feedback control, since it takes the output and uses it to control how much of the input the power stage sees. We know right away that $G(s)$ applied to Out is Out $\times$ $G(s)$, and we can then subtract this from In as shown in Figure 6.9.

But $H(s)$ times this has to equal Out, so we can solve:

$$H[\text{In} - (\text{Out} \times G)] = \text{Out}$$
$$(H \times \text{In}) - (H \times G \times \text{Out}) = \text{Out}$$
$$H \times \text{In} = \text{Out} \times [1 + (H \times G)]$$
$$T(s) \equiv \frac{\text{Out}}{\text{In}} = \frac{H(s)}{1 + H(s)G(s)}$$

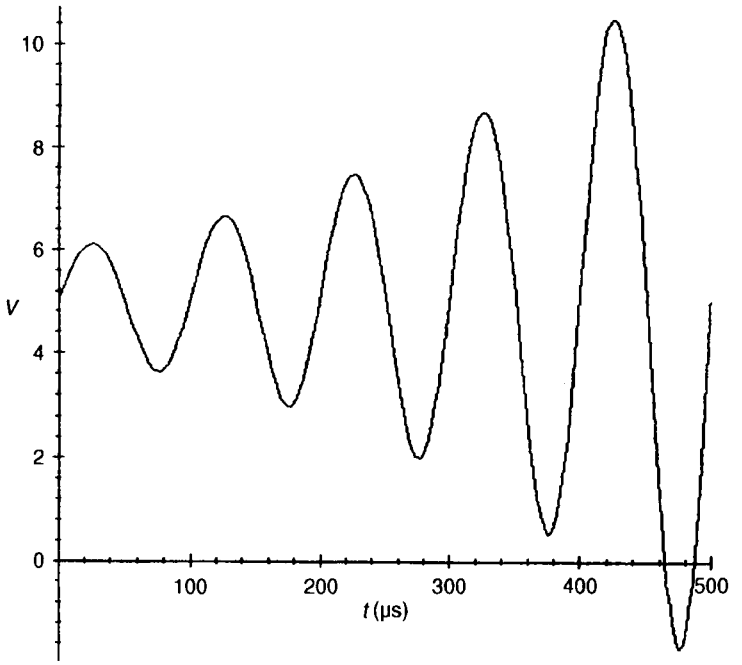

Figure 6.6 Response of an unstable system to a disturbance.

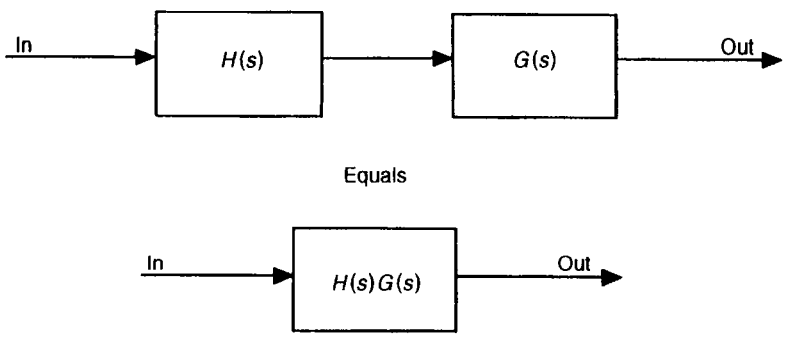

Figure 6.7 Composition law for transfer functions: two in series are multiplied.

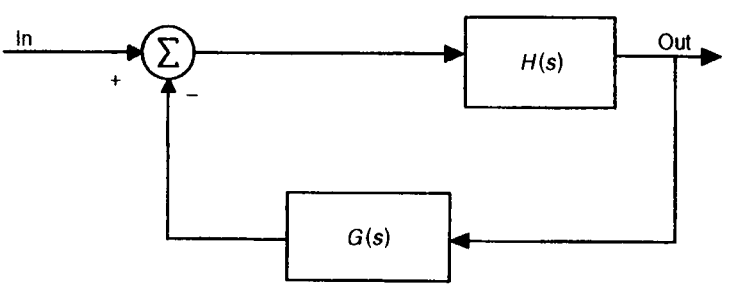

Figure 6.8 A closed loop system.

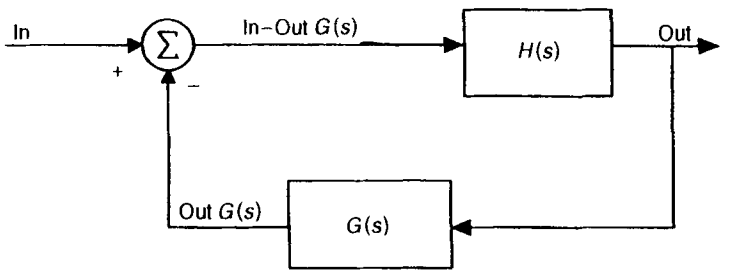

Figure 6.9 Calculating the transfer function of a closed loop system.

This transfer function, $T(s)$, will prove fundamental in determining the stability of a closed loop power supply. As an overview, you can say that if $T(s)$ becomes infinite at any frequency—that is, if the transfer function has a pole—the converter is unstable.

There's No Such Thing as a (Useful) Transform of a Nonlinear System!

Before we get carried away about how wonderful transforms are, let's note that they are really useful only for linear systems, which in terms of power supplies usually means small-signal (just what this means will be made more precise later). The typical problem with transforms usually occurs when part of the system depends on a parameter nonlinearly. Then the operating point of the system changes as the parameter changes, and the transform gives information about the system only at that one particular operating point, telling nothing about the new operating point.

EXAMPLE

A converter's input power source is a solar array. The input voltage changes as the logarithm of the current, $I = ae^{bV}$. If you do a transform on this system, it tells you only how the power supply operates at that one input voltage—**NOT** what happens at any other operating point! The transfer function does not provide useful information in this case. (But you might think about this: Could you devise a function that depended *both* on frequency *and* on the operating point?)

The usual solution to getting useful information from a nonlinear system is to time average it, so that the nonlinearities are averaged out—this is what state-space averaging does for power supplies. (We won't talk about state-space averaging in this book.) But for truly nonlinear systems, such as the solar array example, transform methods cannot be usefully used.

BASIC CONTROL THEORY

Remember that a closed loop system has the transfer function $H/(1 + HG)$. In most systems, and for power supplies in particular, H is given: it's the system you're trying to control (in our case, the power stage of a converter).

The goal of control theory is to select G, the feedback, so that the transfer function is always finite.

The transfer function can become infinite only if the denominator is zero, which happens only if $HG = -1$. So the quantity we measure is $H(s)$, and then we select a $G(s)$ in such a way that when the magnitude of $|HG| = 1(= 0\text{dB})$ the phase is not 180° (i.e., −1). The amount by which it differs from 180° is called the phase margin of the system.

Bode Plots

Rather than work directly with the complex numbers that constitute a transfer function, engineers find it easier to deal separately with the real and imaginary parts, or magnitude and phase. There are several ways of doing this, but the most popular calls for the use of Bode plots. (There's been argument over how to pronounce this name, boh-dee or boh-day? The first is by far the more common.)

A Bode plot consists of two graphs: one is the magnitude of the transfer function (in dB) versus logarithmic frequency; the other is the phase in degrees versus logarithmic frequency. Since for the magnitude plot, both axes are logarithmic, a plot of the transform of a single pole (like a capacitor, its impedance is infinite at 0 frequency) or a single zero (like an inductor, its impedance is 0 at zero frequency) is a straight line.

The Bode plot in Figure 6.10 shows the impedance of a 1μF capacitor. You can verify that this is the impedance of a 1μF cap by observing that the phase is flat at 90°, corresponding to capacitance; and the magnitude of the capacitance can be verified by

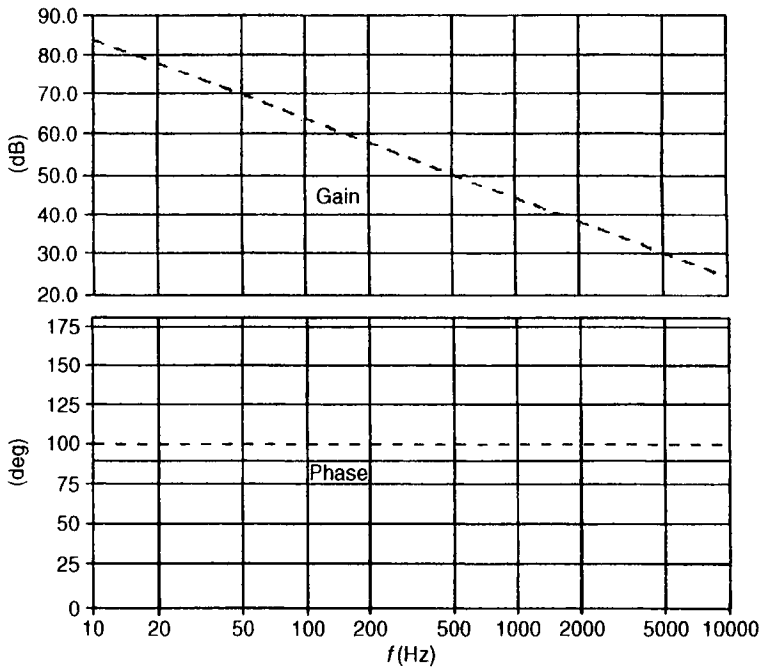

Figure 6.10 Impedance of a 1μF capacitor.

noting that, for example, at 10Hz, the magnitude is 84dBΩ = 15.9kΩ, which gives 1μF at 10 Hz ($10^{84/20}$ = 15,900Ω; 1/(2π 10Hz 15,900Ω) = 10^{-6} F).

The next Bode plot (Figure 6.11) shows a 1kΩ resistor in series with the capacitor. This could represent a capacitor and its ESR, except of course that 1kΩ is much larger than the real ESRs of most caps. You can see that the magnitude levels off at 60dBΩ = 1000Ω, and the phase returns to 180°. Thus at low frequencies, the series combination looks like a capacitor, and as the frequency increases beyond $1/(2\pi RC)$, the combination looks just like the resistor, because the capacitor's impedance goes toward zero.

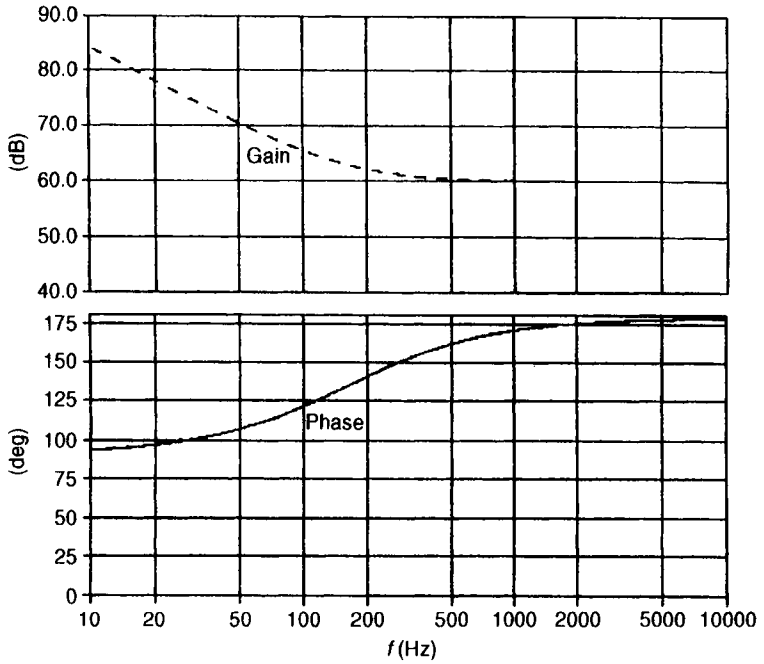

Figure 6.11 Impedance of a 1kΩ resistor in series with a 1μF capacitor.

Note: The phase goes to 180° at high frequency, not 0° as you might expect, because a 180° rotation on the complex plane corresponds to multiplication by −1 (start at the real number 1 and go counterclockwise half a rotation), and so doesn't affect the "mix" of real and imaginary parts. Indeed, many older machines for making Bode plots get confused (to anthropomorphize a design problem) about this 180° degree business, so you might as well get used to seeing such things.

Requirement for Stability

The requirement that the system be stable, stated above for the product $H(s) \times G(s)$, translates into a requirement on the Bode plot:

There must be positive phase when the gain = 1 (= 0dB)

To repeat, the amount of phase at 0dB gain is called the phase margin. Reworded again, the fundamental goal of the design of a feedback loop is to ensure that there is phase margin.

Figure 6.12 shows a loop that is unstable: we see that there is 0dB at a frequency of 10kHz, and that at that frequency the phase is −30°; therefore, the loop oscillates, as shown in Figure 6.13.

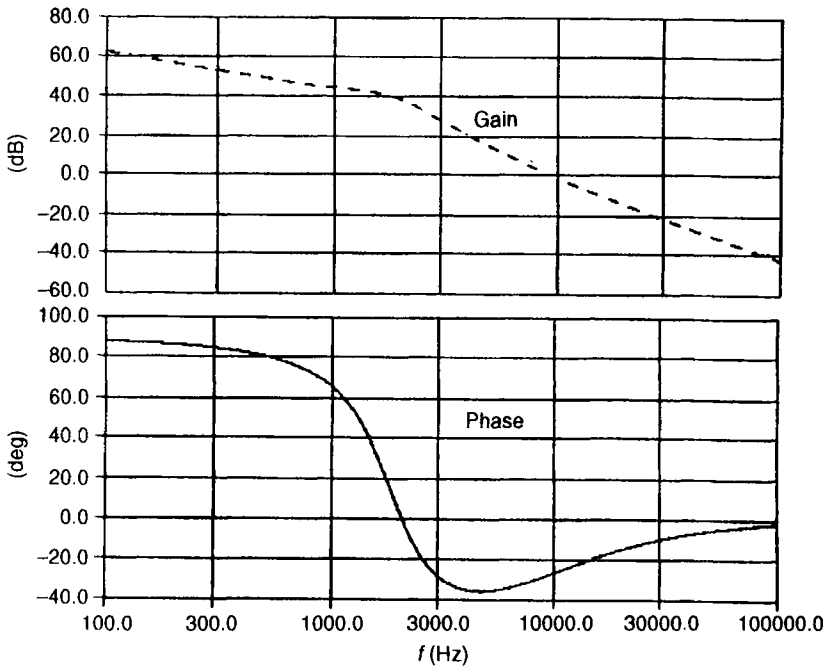

Figure 6.12 A Bode plot of an unstable loop.

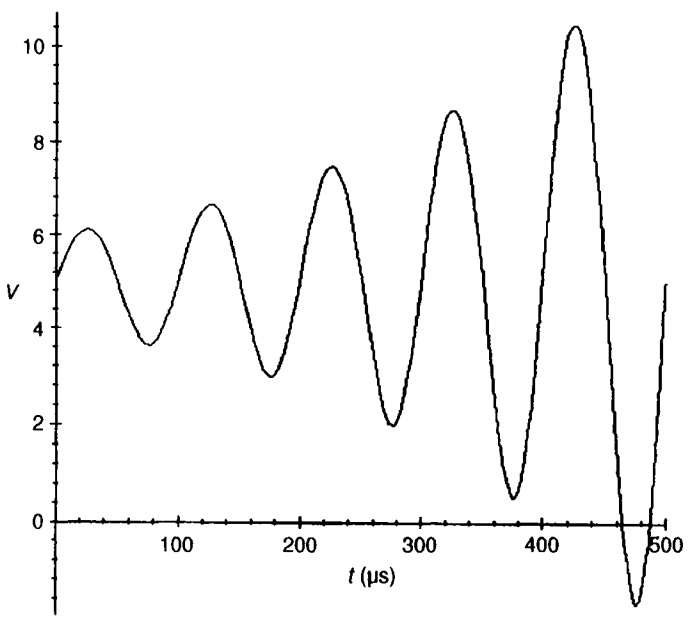

Figure 6.13 The response of an unstable loop to a disturbance.

In a stable loop (Figure 6.14), we see that there is again 0dB at a frequency of 10kHz, and that at that frequency the phase is +30°; therefore the loop damps oscillations, as shown in Figure 6.15. (The Bode plot shown is technically "metastable": see below.)

How Much Phase Margin Is Enough?

You sometimes hear arguments over how much phase margin is enough, and you frequently see differing requirements from spec to spec: 30° minimum, 45° minimum, 45° typical, an occasional 60° are all common.

To understand these disagreements, you need to realize that phase margin serves two separate purposes: it relates to the damping of output transients due to a step load on the converter; and it guarantees stability regardless of variations of component values (e.g., initial tolerance and temperature).

In the author's opinion, the first issue is not really relevant, because load steps invariably take the supply into the regime of "large-signal stability," whereas phase margin determines only "small-signal stability." (More on this soon.) Of course, you wouldn't want a supply to have, say, only 5° of phase margin, because it would ring for a very long time after a transient. But tiny phase margins like this are unrealistic anyway, because of the second purpose of having adequate phase margin.

Thus, phase margin really has to guarantee only that the loop will remain stable with component tolerances, load variations, and temperature variations. For this reason, I recommend the following design practice:

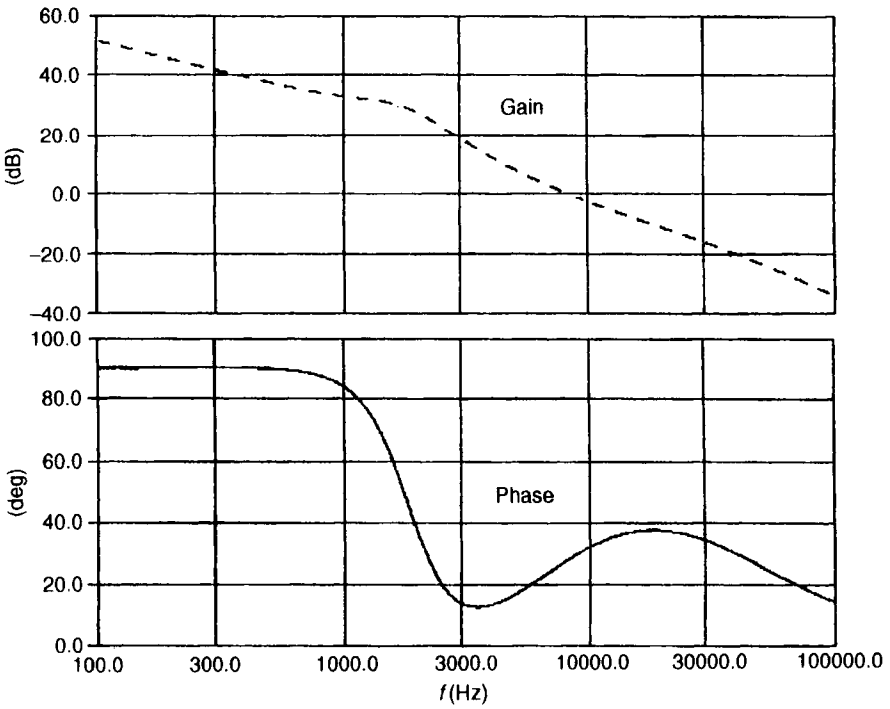

Figure 6.14 Bode plot of a stable loop.

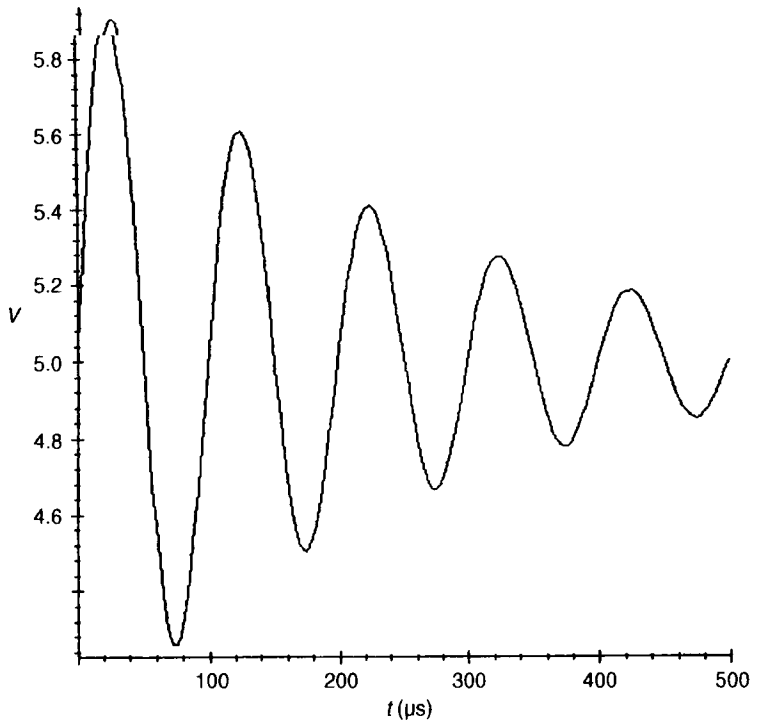

Figure 6.15 The response of a stable loop to a disturbance.

> **Practical Note** Loops should have 45° of phase margin typical at room temperature, with nominal values and nominal load. This usually is enough to guarantee stability under all variations and tolerances. If the load swing or the input voltage range is exceptionally large, you might instead consider a specification for a loop having 30° minimum at all loads and all input voltages.

Gain Margin?

Sometimes you also see gain margin specified. Gain margin is sort of the inverse of phase margin; it measures how much gain (actually attenuation) the system has when the phase reaches 0°. The most important circumstance when this is relevant is in a system with a right-half-plane zero, such as a boost converter. (We'll be discussing RHP zeros below.) The magnitude part of the Bode plot in this circumstance may have an appearance similar to the *RC* impedance plot of Figure 6.11, with flat gain at higher frequencies; the phase however, continues to decrease. Although you may compensate this system to have adequate phase margin, changing the load may change the crossover frequency a great deal because of the flatness in the gain, resulting in a 0dB frequency at which there is no phase margin left. By specifying minimum gain margin (typically 12dB), you can avoid this situation. The author recommends that gain margin be specified only if a flyback converter is being designed; in other converters, naming a gain margin serves no practical purpose.

About Conditional Stability

Normally, in a closed loop system, both phase and gain decrease with increasing frequency (ignoring for now intentionally introduced phase boost, such as in the feedback systems we will be designing later in the chapter). Thus in a normal system, when gain is reduced, phase margin typically increases. But in a conditionally stable system (also called "metastable"), phase goes up again before going down (see Figure 6.16). So if gain were reduced for some reason, the phase might go below 0° while there is still some gain, proceeding to oscillate at this frequency, even though it is "stable" and "has adequate phase margin." Conditionally stable systems should thus be avoided when possible.

As an exercise, you might think about how you could word a specification to avoid this sort of problem.

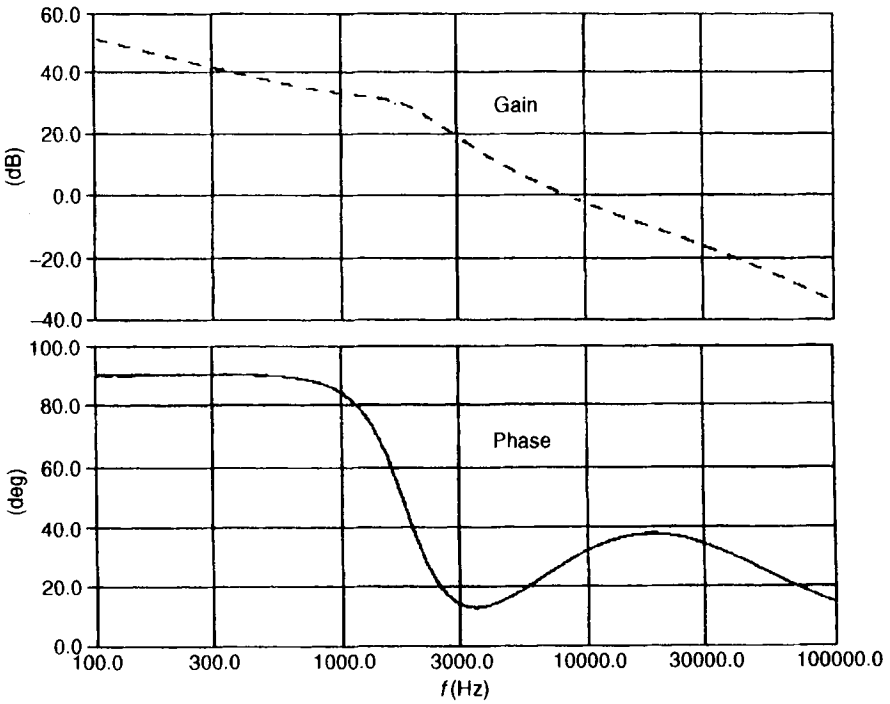

Figure 6.16 A stable loop exhibiting metastability.

Small- versus Large-Signal Stability

Bode plots relate to the response of a system to a small (ideally, infinitesimal) disturbance. But if a disturbance is large, the system response probably will not be determined not by the linear aspects of the feedback but by nonlinear aspects, such as op amp slew rate or rail voltages, or maximum and minimum achievable duty cycles, etc. When factors such as these enter into a system's response, the system is exhibiting a nonlinearity, and so the transform method doesn't apply. Thus, while small-signal stability is necessary, it is not sufficient to guarantee stable operation.

Practical Note If possible, a loop should be designed to avoid large-signal operation.

In any case, large-signal bandwidth of a converter is always less than or equal to small-signal bandwidth, because before the converter's loop can run into some nonlinearity, it first has to respond, and the response is set by the small-signal bandwidth. One way around this, which has been tried on occasion, is to have a completely separate nonlinear loop for large-signal response: but then, it must be decided when each loop should be active, how interference can be avoided between them, and so on.

EXAMPLE

A true story A 1.2V power supply was designed to have very wide bandwidth, and measurement of its closed loop response showed 45° of phase margin. Unfortunately, when a load step was applied, the system began to oscillate: the op amp, which had insufficient gain bandwidth and slew rate, spent all its time trying to catch up with itself, first hitting its positive rail, then swinging down to ground, then back up again. This oscillation was eliminated by swapping the op amp for a pin-compatible one with higher gain bandwidth (and higher slew rate).

HOW TO STABILIZE A VOLTAGE MODE BUCK CONVERTER

With all these preliminaries out of the way, we can turn to the task of the practical design of a feedback compensation. We start by demonstrating the method for measuring and stabilizing a very simple converter, a voltage mode buck. The same methods demonstrated here are directly applicable to stabilizing other types of converter as well.

Initially, we have just the power stage (Figure 6.17). We have a power MOSFET controlled by a PWM switching at 250kHz; there is a freewheeling diode and an inductor and an output cap; and the PWM has an error amplifier but is otherwise just a block for the

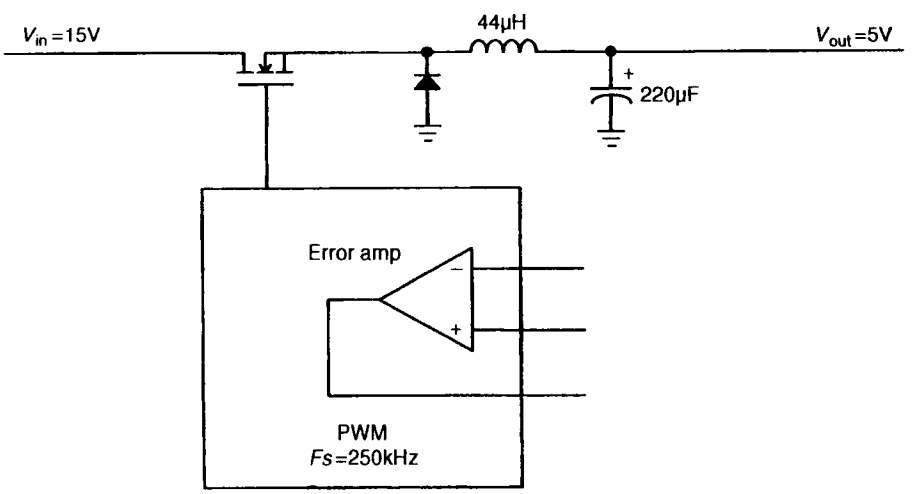

Figure 6.17 Power stage of the buck converter to be stabilized.

moment. The inductor is 44μH (at 0A—it's in fact just the DC inductor we designed in Chapter 5, the one that swings down to 35μH at 2A). The output capacitor is 220μF. The author has taught lab sessions on loop compensation with exactly this setup, so if you like, you can assemble the components as you go along and duplicate the measurements shown in the book.

Our overall strategy is going to be this. First we will measure the open loop response of the system, which can be done without using the error amp. From the data we collect, will be able to design a compensation that ensures stability. Finally, we measure the closed loop response, to verify that we did it right. It sounds easy, and it actually is; once you've done it a couple of times, the entire process of stabilizing a converter can be done in a morning—honest. No more two weeks wasted swapping components in and out. Stabilization can be successful on the first try!

How to Measure Open Loop Response

The open loop response of the converter is by definition the response of the power stage—that is, what does the output look like when you jiggle the duty cycle? To measure this open loop response, you need to have the converter operating more or less at its correct operational parameters: since the output is supposed to be 5V when the converter is working, you don't want to measure the open loop response when the output is 3V or 13V!

So the plan of action is going to be to find the DC operating point, and then add a small AC on top of it, as illustrated in Figure 6.18.

The error amplifier of the PWM is made unity gain; don't forget and leave anything attached to it, or you'll get crazy-looking results!

To make this measurement, add together (using a mixer, described below) an adjustable DC voltage and a swept sine wave from a network analyzer. Starting the adjustable DC voltage at zero, increase it very slowly until you get approximately 5V at the

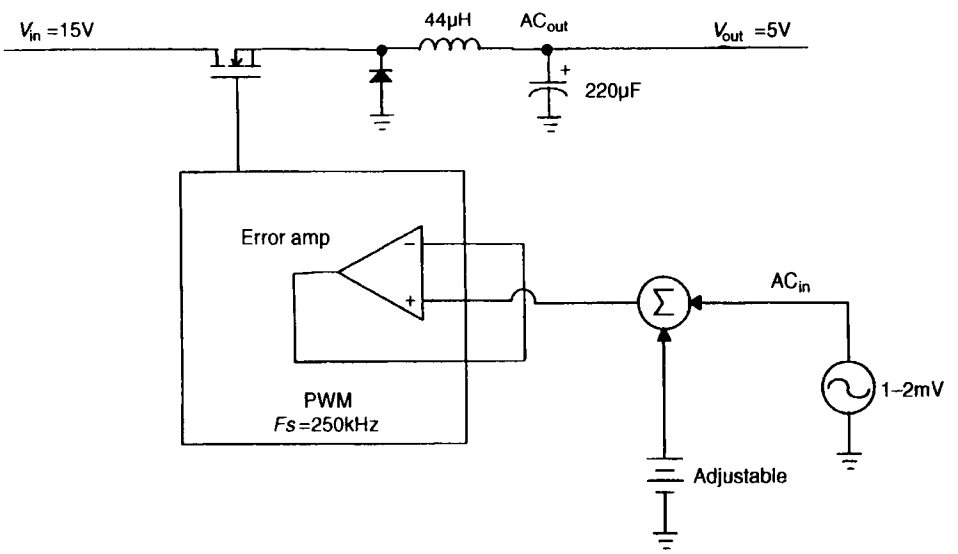

Figure 6.18 An open loop measurement will be taken by setting the DC operating point and adding an AC component.

output of the converter. Make sure that a resistive load (maximum specified load is preferable) is already attached to the output before you bring the adjustable DC voltage up from 0V. This is important because changing the load in this open loop mode can cause the output voltage to fluctuate wildly, potentially blowing up your output cap.

Practical Note Some PWMs have an offset voltage, so the duty cycle remains at zero until you reach a volt or so.

Practical Note Power stages with high gain will clearly be very sensitive to the adjustable DC voltage. With an ordinary lab supply, it may be frustratingly difficult to adjust the control voltage to get exactly the output voltage you want. In this case, you can probably settle for anything within 5% of the actual output. An alternative is to buy or build a precision voltage source that can be adjusted in millivolt steps. Better yet, see below for a method that avoids this measurment altogether.

Once you have the adjustable DC voltage set so that the output voltage is correct, measure the open loop transfer function by dividing AC_{out}/AC_{in}, at the points shown in the schematic (Figure 6.18), with the network analyzer which is generating the swept sine. (See Chapter 4, on instrumentation, for detailed operational information on how to run a network analyzer.)

Let's take a careful look at the actual measurement of the open loop response of the converter (Figure 6.19). At very low frequencies (below 10Hz nothing is happening in this circuit), the gain is 11.8dB, and the phase is 0°. Zero phase is to be expected at very low frequencies, since when you increase the duty cycle on a buck, the output voltage increases (transfer function = Out/In).

As for the low frequency gain, consider a hypothetical increase in the DC control voltage of, say, 100mV. The oscillator ramp for the PWM used (a UC2825) is $1.8V_{pp}$, so the 100mV causes a change in duty cycle of 100mV/1.8V = 5.6%. (Note that the ramp amplitude directly affects converter gain.) Now the PWM actually has two outputs, of which only one is being used. So increasing the duty cycle increases only the one output—the other is zero (as seen by the MOSFET) no matter what: that is, the MOSFET has a maximum duty cycle of 50%. So the real increase in duty cycle is only 2.8%. Now when the duty cycle increases, the effect is to increase the average voltage applied to the inductor, because the MOSFET is on longer; that is, the 15V input is applied longer. The output voltage thus increases by 15V × 2.8% = 420mV. Since this was caused by a control voltage increase of 100mV, the gain at low frequency is 420mV/100mV = 4.2 = 12.4dB, which is quite close to what we actually measure. (The 0.6dB error = 7%, and is probably caused by finite rise and fall times of the MOSFET switching; it is unimportant for our loop measurements and can be eliminated by methods discussed further on.)

Looking again at the actual measurement, we see that as the frequency increases, the gain rises and the phase falls, which is caused by the *LC* tank resonance. As confirmation we may check that the resonant frequency should be

$$f = \frac{1}{2\pi\sqrt{LC}} = 1618\text{Hz}$$

in close agreement with the measurement.

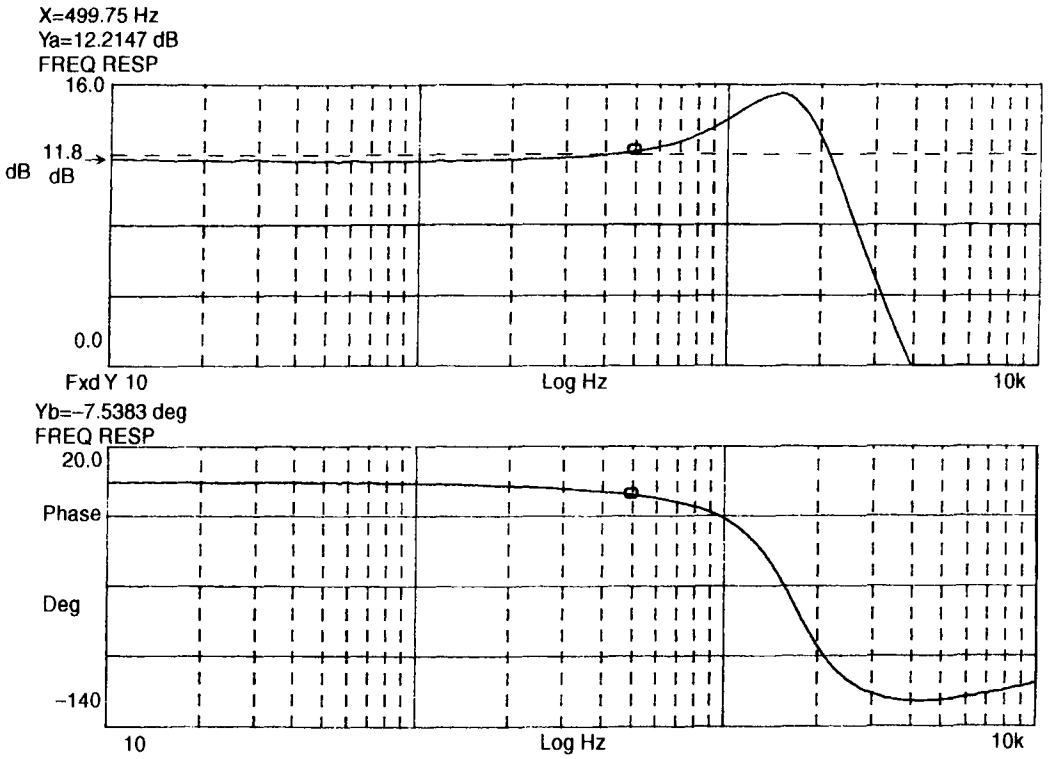

Figure 6.19 Buck converter open loop measurement.

Finally, at high frequencies, the gain rolls off at 12dB per octave (i.e., the gain at 4kHz is one-quarter the gain at 2kHz; an octave is a factor of 2 in frequency), corresponding to the two poles, the inductor and the capacitor. The phase at the topmost frequency in the measurement is starting to rise again, as a result of the ESR of the output capacitor. The measured ESR of the cap was 120mΩ, so the *RC* frequency is

$$f = \frac{1}{2\pi RC} = 6\text{kHz}$$

(The phase doesn't go all the way to −180° because of this zero.) If the gain measurement had been continued to even higher frequencies, you would have seen the gain decreasing at only 6dB/octave, which is a single pole, because the capacitor had become resistive at these frequencies owing to its ESR, just as in the example earlier in the chapter of a series *R* and *C*.

Venable's *K*-Factor Paper

Now that we've measured the open loop, we can design a compensation network to make the converter stable in closed loop. We're going to choose a bandwidth of 500Hz, which is well below the resonant frequency of the output tank, because we are in voltage control mode. Although it is possible to compensate the loop so that it is stable with a bandwidth above the resonance, it is a better idea to go to current mode control for this. (Alternatively,

it is possible to stabilize a voltage mode converter well above the *RC* frequency, because the phase is back to only −90°; we won't be demonstrating that here.)

The method described here, first demonstrated by Venable [1], amounts to computation of the amount of phase the error amplifier needs to give the desired phase margin ("phase boost"), followed by selection of one of three types of compensation based on this calculation. The actual component values are then computed based on the idea of placing zeros and poles symmetrically around the desired bandwidth frequency: zeros to cause the phase to rise below the bandwidth, poles to cause the gain to decrease above the bandwidth. The three amplifier types are shown in Figures 6.20, 6.21, and 6.22.

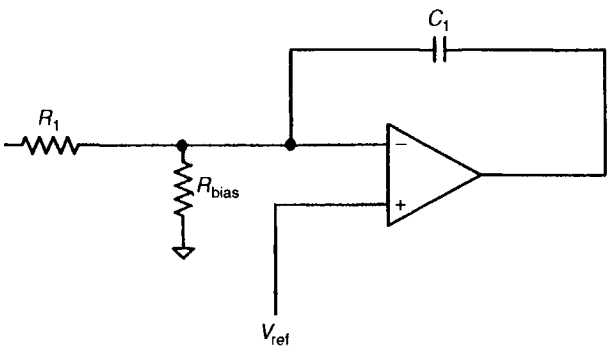

Figure 6.20 Type I amplifier.

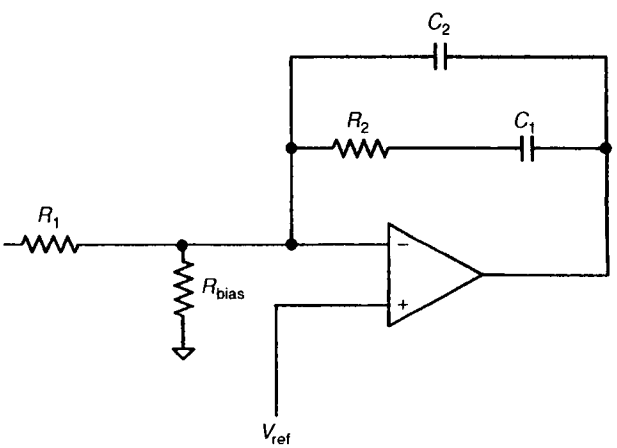

Figure 6.21 Type II amplifier.

Now the cursor of the network analyzer in our open loop measurement of the buck converter was set at our intended bandwidth of 500Hz (a convenient practice if you need to make presentations), so that we know the open loop gain there is 12dB, and the phase is −7°.

We've already done the first two steps of the *K* method, namely, making the measurement and choosing the cross over frequency (another name for the bandwidth.) The third step is to choose a phase margin; based on our earlier discussions, we will choose 45°.

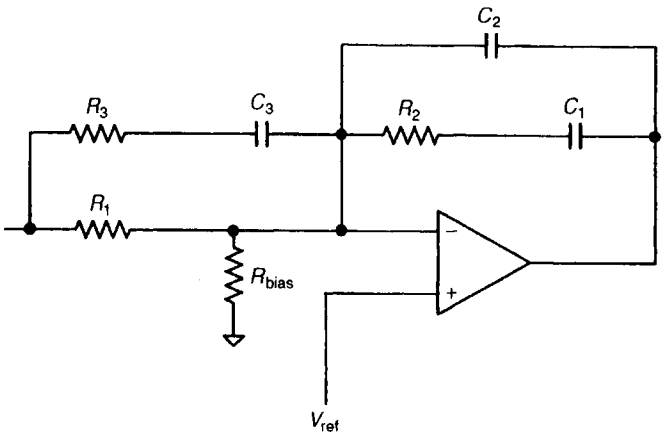

Figure 6.22 Type III amplifier.

The fourth step is to compute the gain necessary for the error amp. Since the open loop gain is 12dB, we have to reduce the gain by 12dB to get 0dB at 500Hz. So the error amplifier gain should be $G = -12\text{dB} = 0.25$ at 500Hz. Be careful in this step, don't leave the gain in decibels!

As a fifth preparatory step, we compute how much phase boost we need from the compensation:

$$\text{boost} = M - P - 90$$

where M is the desired phase margin and P is the measured open loop phase shift. For the case we are considering, we have boost $= 45 - (-7) - 90 = -38°$; since this is less than 0°, no phase boost is required, and we can use a type I amplifier (see the discussion below).

One final preliminary step is to select R_1. Since this is a 5V output, and the UC2825 PWM IC has a 5V reference, no R_{bias} is required, and a good starting value for R_1 is 10kΩ. If the output voltage had been higher (or the reference voltage lower) an R_{bias} would have been used to divide down the output voltage. However, R_{bias} has no effect on the gain or phase of the error amplifier; as long as you leave R_1 alone, the output voltage can be adjusted by adjusting R_{bias} without affecting the error amp! (Note, however, that adjusting V_{out} does affect the open loop, and thus potentially loop stability. This is why the open loop measurement is done close to the actual operating point.)

Now, the compensation for our buck converter only requires C_1 to be computed. For the type I this is done as follows

$$C_1 = \frac{1}{2\pi f G R_1}$$

[Type I Amplifier]

For us, $C_1 = 1/(2\pi \times 500\text{Hz} \times 0.25 \times 10\text{k}\Omega) = 127\text{nF}$. Rather than trying to parallel capacitors, simply round this down to the nearest common value, 120nF—remember that the purpose of phase margin was to make it unnecessary to worry about component tolerances! In the lab, the value we actually measured on our 120nF capacitor was 135nF. (This was a CKR06 style cap, with a tolerance of 20%. Generally, it is better to use 10% NPO or C0G style capacitors in the feedback loop.)

Before getting back to our buck converter to verify performance, let's consider the other types of error amp. As long as the necessary boost computed is less than 0°, as in our buck converter, a type I amplifier is all that's required. However, if more than zero degrees of boost is required, one of the other error amps must be used. The type II amplifier theoretically can provide up to 90° of phase boost; practically, trying to get more than about 75° out of it results in component values that are either too large or too small to be practical. Above 75°, and less than about 160°, a type III can be used. If you need more than 160°, chances are you're measuring wrong!

Practical Note Limit the boost of a type II amplifier to 75°, and a type III to 160°.

For reference, the design equations for the other two types of amplifier are:

$$K = \tan\left(\frac{\text{boost}}{2} + 45\right)$$

$$C_2 = \frac{1}{2\pi f G K R_1}$$

$$C_1 = C_2(K^2 - 1)$$

$$R_2 = \frac{K}{2\pi f C_1}$$

[Type II Amplifier]

$$K = \left[\tan\left(\frac{\text{boost}}{2} + 45\right)\right]^2$$

$$C_2 = \frac{1}{2\pi f G R_1}$$

$$C_1 = C_2(K - 1)$$

$$R_2 = \frac{\sqrt{K}}{2\pi f C_1}$$

$$R_3 = \frac{R_1}{K - 1}$$

$$C_3 = \frac{1}{2\pi f \sqrt{K} R_3}$$

[Type III Amplifier]

Remember to use degrees in these equations, not radians!

Practical Considerations

Much of the time you will end up using a type II compensation. There should be something unusual about the power stage if you determine you need a type III—otherwise, either your measurement or your calculation may be wrong! A little bit of caution here will save quite a bit of aggravation later.

In any case, as stated above, you should never need more than 160° of phase boost. If >160° seems to be necessary, something's wrong for sure.

What happens if the values you calculate turn out to be very large (say, 10MΩ) or very small (say 7pF)?

Practical Note Avoid using resistors larger than 1MΩ, or capacitors smaller than 22pF. Strays will make them too inaccurate for loop compensation. If your calculations show a need for parts that are outside these bounds, a good plan will be to start over, using a different value for R_1, such as 1kΩ instead of 10kΩ.

Finally, what do you do when you calculate an oddball value, such as 900pF? Do you need to parallel caps?

Practical Note As long as you're not at the top end of the possible phase boost from the type of error amp compensation you're using, just round off the capacitor values to their nearest standard values; the compensation will still be good enough. For the question just asked, 1nF will be good enough instead of 900pF if you're trying to get 60° phase boost. If you're right at 80° phase boost with a type II, try going to a type III.

Other Comments

The comments above suggest that there is a practical limit on the maximum bandwidth you should try to give a converter. Of course, a voltage mode converter shouldn't be stabilized above its resonant tank frequency unless you are going to stabilize it above its output cap—ESR zero frequency; and of course no normal switch mode converter can be compensated to have a bandwidth close to its switching frequency. Beyond these limits, however, even current mode converters have limits, set by the compensation values that can be reliably attained. To get bandwidths of 100kHz or more not only will require very high switching frequencies, but probably, in addition, some sort of integrated or hybridized circuit, to control stray impedances. The typical practice is to give substantially more phase margin than is used with lower bandwidth converters, because the strays give increased variability to the values used.

One final comment on Venable's method. Symmetrical distribution of poles and zeros is enough to uniquely determine a type II (and, a fortiori, a type I) amplifier, but it doesn't uniquely determine the compensation values for a type III. There may be ways to adjust the position of the poles and zeros of a type III error amp for special purposes.

How to Measure Closed Loop Response

So now the buck converter circuit looks like Figure 6.23. In addition to adding the compensation on the schematic, a summer has been added (the circle with the Σ in it) showing injection of the swept sine wave, and the correct points to measure the closed loop response of the converter, AC_{out}/AC_{in}. Observe that the loop remains closed during this measurement, so that it automatically controls the output voltage to the 5V level; no external DC control is required.

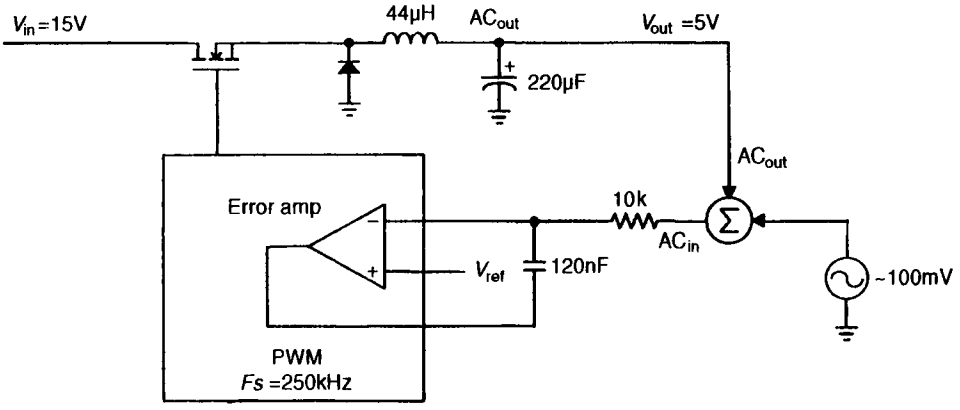

Figure 6.23 Loop with compensation added. Also shown is diagram for measuring closed loop response of converter.

For the astute reader, let's mention that what is being measured here is not the power supply transfer function, V_{out}/V_{in}. We're instead measuring a related quantity inside the loop. Relabeling our block diagram, we come up with Figure 6.24.

We are measuring the point labeled $GH(\text{In} - V')$ divided by the point labeled V'. Now the thing to notice is that the network analyzer measures only AC signals, using a bandpass filter centered at each frequency in its sweep. So the term ($GH \times \text{In}$) is filtered out before the division, because it is a constant, (i.e., DC). So the result from the analyzer is $-GHV'/V' = -GH$. So the design criterion that $GH \neq -1$ can be replaced by the criterion that what the network analyzer measures must not equal 1; that is, we have to avoid having gain = 0dB and phase = 0° at the same time.

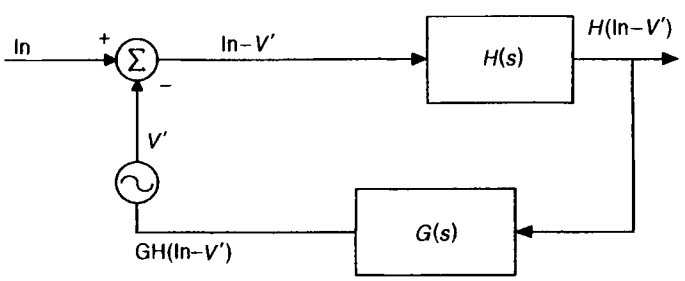

Figure 6.24 The closed loop measurement is really measuring GH.

How to Measure It : Transformer Method

Up to now the summing function used in measuring the loop has been shown as just that, a big Σ. In this section and the next, this omission will be remedied, showing the two popular methods for implementing this function, and pointing out practical aspects of their use. We start with the transformer method.

The transformer method (Figure 6.25) works just as you might suppose, by transforming the AC drive signal over into the loop. The 50Ω in parallel with the secondary is not strictly necessary, but accomplishes the following convenient function: the transformer can be soldered on top of the resistor, allowing you to measure the loop without having to desolder anything on the PC board. The 50Ω is small compared to the 10kΩ, and so does not affect the output voltage all that much.

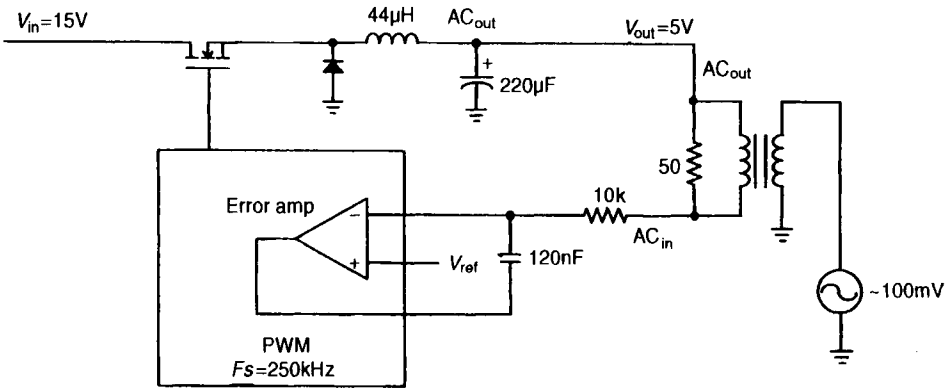

Figure 6.25 The transformer method of measuring the closed loop response of the converter.

The transformer method enjoys some popularity, and this ability to tack it in is its best feature. Of course, it also provides isolation between the network analyzer and the circuit, which can be important if you are measuring a high voltage output.

On the other hand, the transformer has to be carefully designed to be responsive over a very broad range of frequencies, both very low (where it mustn't saturate) and very high (where it mustn't have so much capacitance that the signal will be shorted out).

How to Measure It: The Mixer Method

A second method, and the one the author recommends, is using a mixer, which is just an op amp adder (see Figure 6.26).

The mixer (unlike the transformer) works at arbitrarily low frequencies, and with proper op amp selection can work at very high frequencies.

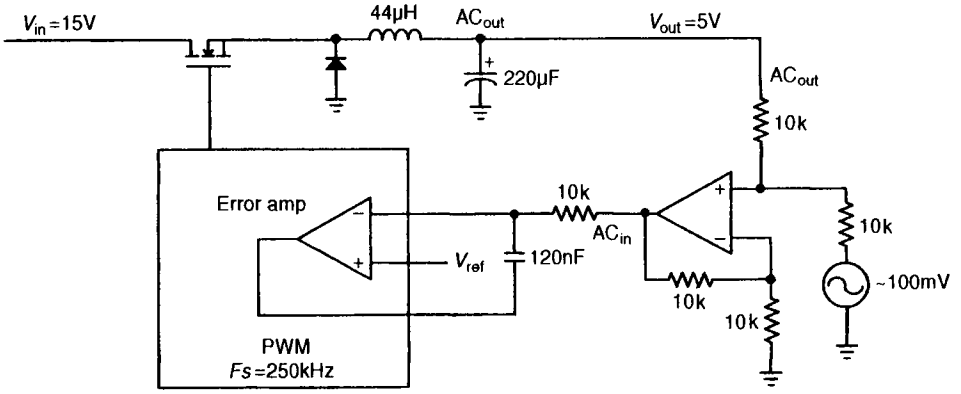

Figure 6.26 The mixer method of measuring closed loop converter response.

Practical Note Be careful! Before you use the mixer, measure its response, particularly its phase, with the network analyzer. Some high gain bandwidth op amps have lots of phase shift, and some don't; usually you can't tell from the data sheet. Don't forget 100nF bypass caps for the op amp's power lines. And for really high frequency measurements, you might want to make everything surface mount and use BNC connectors for inputs and output.

The major drawbacks of the mixer are as follows: (1) you have to desolder a component on the PCB to insert the mixer in the loop, and (2) you can't run it in a loop whose output is greater than the op amp's supply voltages. On the other hand, the only real requirement for the placement of the mixer is that it have low impedance inputs, and output to a high impedance node. So anywhere that is low impedance could go into the mixer—maybe even in the divider network, which will have to be present anyway for a high voltage output.

Converter Closed Loop

We used the mixer method to measure the response of the actual buck converter circuit in the lab (see Figure 6.27). Since we had targeted 500Hz as the crossover frequency, we expect to have 0dB at 500Hz, and a phase margin of $90 - 7 = 83°$. What we actually measure is -1.8dB and 83° exactly. The 1.8dB error ($= 23\%$) is explained in part by the 7% error in the capacitor value (135nF instead of 127nF), the remainder is due to a slight error in the output voltage when the open loop was measured. Below, we show how to avoid even this small error.

Let's examine the plotted measurement carefully. At low frequencies, the gain is higher as the frequency decreases, due to the pole at the origin (i.e., the capacitor in the feedback). At a very low frequency (not shown in this measurement), there is a maximum gain, set by the open loop gain of the error amp. The pole at the origin is desirable, and is standard in closed loop converters because it means that the DC output voltage will have

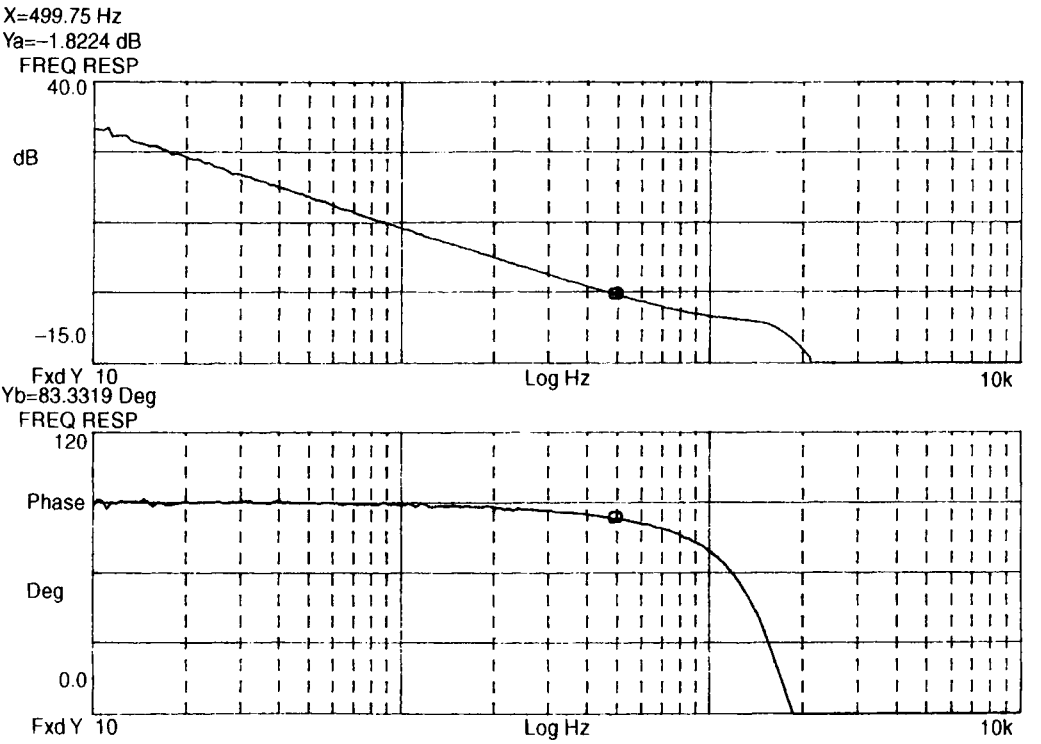

Figure 6.27 Closed loop response of the buck converter measured with the mixer method.

only a tiny error (with respect to the reference voltage); all three error amp configurations have a pole at the origin. For example, the gain of the error amp at very low frequencies might max out at 80dB. Since 80dB = 10,000, this forces the output voltage to match the reference voltage to 1 part in 10,000, 0.01%. This is of course much better than the accuracy of almost all references, and consequently the error in the output voltage is entirely due to the error in the reference.

At frequencies above the resonant tank frequency, the phase drops rapidly, but there is about 12dB of gain margin. Altogether, then, calculated performance and measurement match well.

> **Practical Note** When doing your closed loop measurement, always make sure to hit all four corners: maximum and minimum input line voltage, and maximum and minimum output load. The open loop changes with these parameters, and thus so does the closed loop response; you need to be stable for all four conditions.

How NOT to Measure a Loop

Naturally, there are any number of ways of doing something wrong. Over the years, the author has encountered at least four incorrect approaches to measuring a loop. Here they are, with the reasons for inadequacy briefly noted.

1. Don't measure the "In" signal at the AC source rather than at the output of the mixer. This fails to take into account the feedback loop. It's more like measuring $H(s)[G(s) - 1]$ rather than $H(s)G(s)$.
2. Don't try to inject the AC at the reference pin instead of inside the loop. (Where would you measure the loop response then?)
3. Don't put the "In" and the "Out" probes at the same place (because the transformer's just a piece of wire, right?).
4. Don't try to stick the output of the network analyzer right into the loop directly, without an amplifier or transformer. Although this can be fine for a 5V input converter, generally it can be a good way to blow up the equipment.

A Better Method of Measuring the Open Loop

As we've seen in our buck converter, there can be some difficulty in measuring the open loop, both in the required equipment, as well as in achieving a desirable accuracy. And when the noninverting pin of the error amplifier isn't available (as in 8-pin PWM ICs), measuring the open loop can be rather problematic, since there's no straightforward way of making the error amp unity gain. Fortunately, there is a better way, devised by the author some years back. I now use this method to exclusion of all others because of its substantial superiority.

At low frequency, a power stage is always a fixed gain with 0° phase shift: you put in some duty cycle and get some voltage out, and increasing the duty cycle increases the output voltage. This implies that a converter can always be stabilized by picking a low enough bandwidth, and in particular, by using a type I amplifier. Following this thought, the preferred method for determining the open loop response is to use a large capacitor for compensation in a closed loop configuration, measure the closed loop, and then subtract the effect of the large capacitor. The compensation to be used in the actual circuit can then be derived as usual.

To illustrate the method, we perform a closed loop measurement on our buck converter with a 1μF capacitor as compensation, to determine the open loop characteristics (see Figure 6.28).

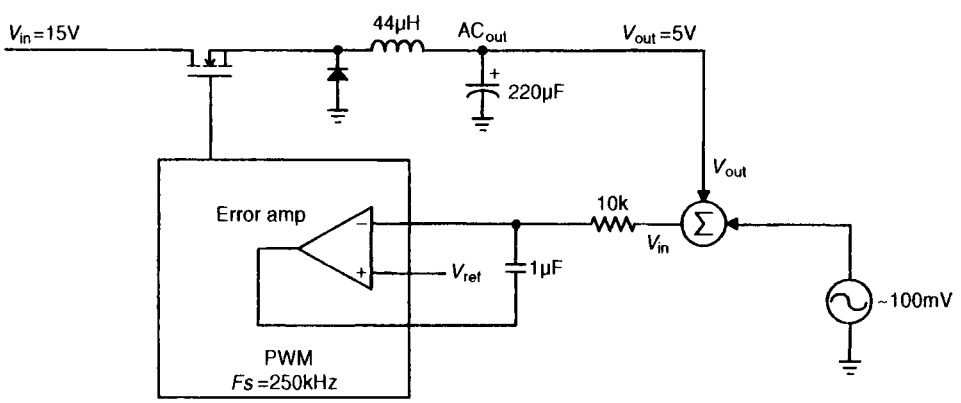

Figure 6.28 The best method of measuring the open loop characteristics of the converter is to compensate it at a low frequency.

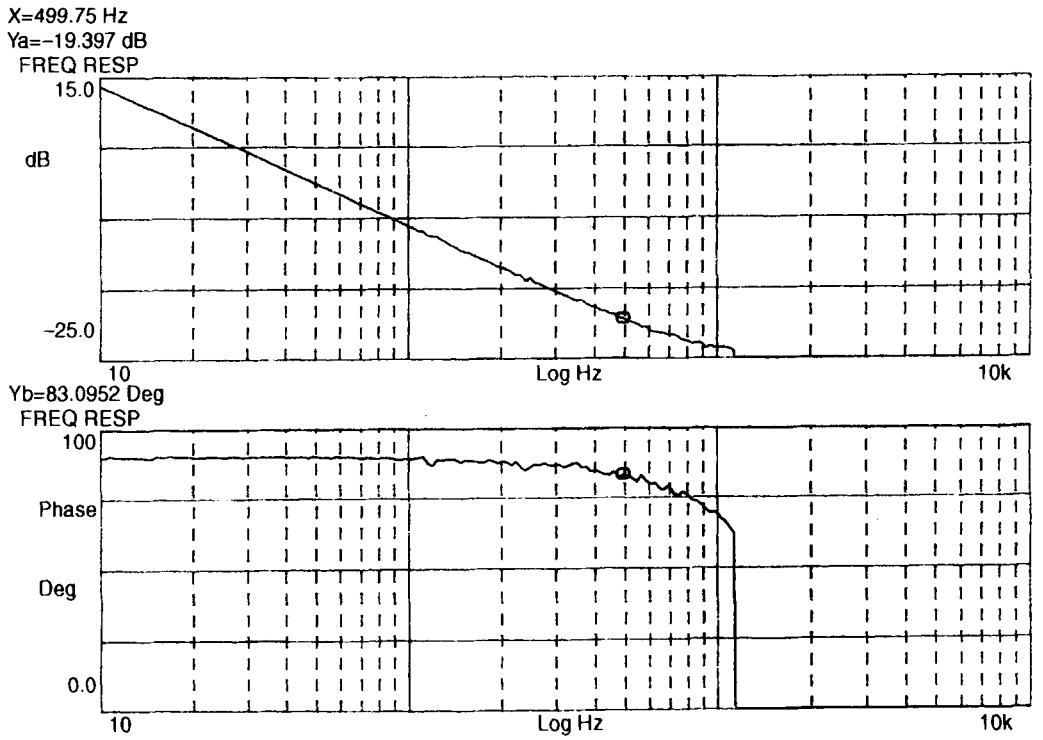

Figure 6.29 Loop characteristics measured with an initial low frequency compensation.

The measured response is shown in Figure 6.29. The bandwidth appears to be about 50Hz—but this is not what we're interested in. Closing the loop was just a stratagem for obtaining the open loop characteristics. At 500Hz, where we want to close the loop, the gain is measured to be −19.4dB, with a phase margin of 83°. This means we should select a capacitor smaller than the 1μF by −19.4dB = 0.107, which is to say 107nF, which will still give us 83° of phase margin, since it is still a type I amplifier. Using the 107nF instead of the 135nF actually selected would increase the measured gain by 135/107 = 1.26 = 2.0dB, which would bring the measured gain (−1.8dB) into agreement with calculation to within 0.2dB—dead on. Thus, this method can be expected to give almost perfect results, as seen here.

> **Practical Note** Sometimes the gain of the power stage is low, and if you are trying to compensate the loop to a high frequency, the gain measured by this method at that frequency may be low enough to be in the noise (e.g., −60dB). In such a case, you should probably try decreasing the 1μF cap to 100nF, thus increasing the gain by 20dB.

If the phase margin we measured using this method had been too low, say only 20°, we would have selected a type II compensation, with a phase boost of 25° (because 20° + 25° = 45°) at 500Hz and a gain of −19.4dB at 500Hz, thus bringing us back to 45°

at a crossover frequency of 500Hz. Compensation with this method is thus almost trivially easy; you can astound your boss with how quickly you can stabilize a converter.

What If the Noninverting Pin of the Error Amp Isn't Available?

It's not uncommon to use an 8-pin PWM IC for a supply, in which case it is typical for the noninverting pin of the error amplifier to not be pinned out, but rather to be referenced internally to a reference voltage. As mentioned above, in this case it is quite hard to figure out how to measure the open loop gain directly, since there's no convenient way of making the error amp unity gain. If, however, you use the closed loop method described, such measurements are straightforward, and the inability to access the noninverting pin is no drawback. Figure 6.30 shows the measurement setup for determining the open loop characteristics of a converter using such a PWM.

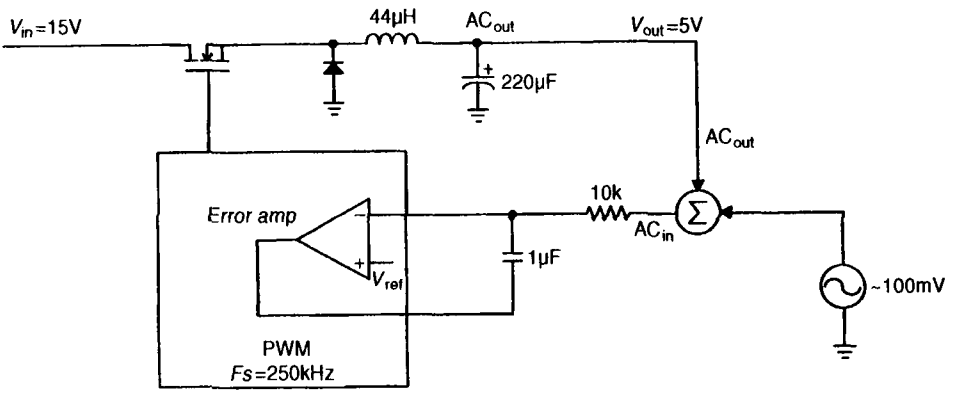

Figure 6.30 Using the low frequency compensation method to measure the open loop of a converter in which the noninverting pin of the error amplifier isn't accessible.

CURRENT MODE CONTROL

Theory

The difference between current mode control and voltage mode control is that current mode control has two feedback loops, one to control the inductor current, the other to control the capacitor (output) voltage.

The control theory for this system can be worked out the same as for the voltage mode control system, though we're not going to. Let's note, however, that there are some subtleties involved with the high frequency response of current mode control converters that have been worked out only in recent years (notably by Ray Ridley).

For practical purposes, the reason for having the second, inner loop (see Figure 6.31) is that controlling the inductor current acts to remove the effect of the inductor on the power stage's transfer function. This is because the transfer function of the power stage already includes the closed loop of the current loop, so that the effect of the inductor is

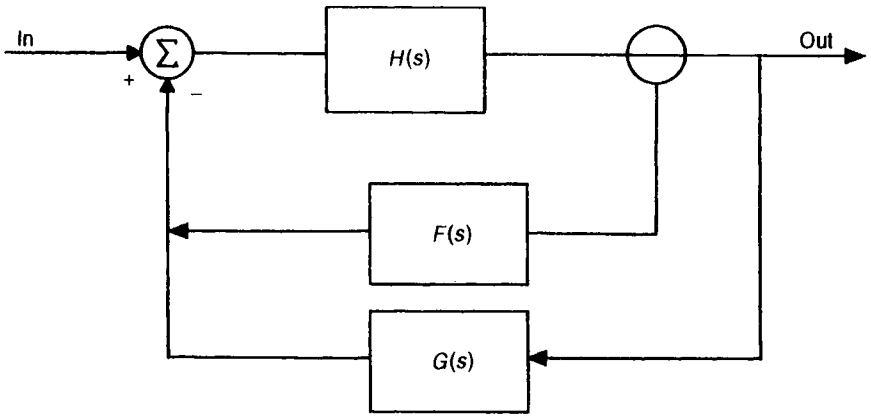

Figure 6.31 Control block diagram of a current mode converter.

entirely absorbed by the loop controlling it and doesn't appear in the response. Thus, there is no resonant tank to worry about, and at high frequencies there is only a single pole (the output cap), so that the phase shift goes to $-90°$ rather than $-180°$. These effects make it much easier to control current mode than voltage mode, and make it possible to give current mode controlled converters high bandwidth.

A Limitation of Current Mode Control

The usual implementation of current mode control is to sense a current with a resistor (or a current transformer going to a resistor, which is the same thing) and feed it into a PWM IC. However, as load current decreases, the magnitude of this signal must also naturally decrease. If the load is light enough, the current signal will be negligible, and the current feedback loop has no effect on the system; thus,

Current mode control becomes voltage mode control at light loads.

So if you give the converter a lot of bandwidth at maximum load, you need to carefully check that at minimum load the extra pole (the inductor) isn't cutting back in and causing instability. Counteracting this effect, though, is the fact that converters typically have less open loop gain in the power stage at light loads than at heavy loads.

> **Practical Note** From a practical standpoint, you won't have current mode control over the whole range of operation if your load range is more than 10:1.

Slope Compensation

When the duty cycle of a current-mode-controlled converter exceeds 50%, the converter will oscillate at a subharmonic of the switching frequency, actually half of the switching frequency, unless slope compensation is added. The origin of the problem can be understood thus. A current mode control loop works by turning off the switch when the

current reaches a certain level (set by the error amp output). If the duty cycle exceeds 50%, the inductor current is being ramped up for more than 50% of the period. This means, obviously, that the inductor current is being ramped down for less than 50% of the period. This smaller time means that the inductor current has not yet returned to its steady-state initial value by the time the next period is getting started, so the current for the next period starts off too high. During this next period, therefore, the inductor current reaches the turn off level too soon, causing the duty cycle to be terminated early; in fact, it's terminated at less than 50% of the duty cycle. But now the off-time is too long (> 50%) and so the current at the start of the next cycle is too low, causing the duty cycle to again exceed 50%, and so on, oscillating between under- and overshooting the current every other cycle. There is a clear demonstration of this subharmonic oscillation happening in the literature [2].

Slope compensation fixes this problem basically by adding a fixed ramp to the current signal. Since this ramp has a constant value, the effects of variations in the current signal are better damped. In fact, the real effect of slope compensation is to make the control loop somewhat more like a voltage mode control. This makes sense: if you think about it, a voltage mode controller works with a fixed ramp against which it compares the error amplifier's output. So adding more and more slope compensation brings the converter back closer and closer to voltage mode; if the ratio of slope compensation amplitude to current–signal amplitude went to infinity, you'd be back entirely at voltage mode. This explains the statement above that at low load power, your current mode control converter is back to voltage mode.

It is also to be observed that adding slope compensation puts the converter somewhere between voltage mode (with two poles) and current mode (with one pole); this means that when you take a loop measurement, and measure the slope of the Bode plot, you find that it is intermediate in value between one and two poles. Of course, it is the active circuitry that makes such a nonlinear transfer function possible.

As an aside, people sometimes tell me there's no such thing as a "subharmonic oscillation." I'm not sure what their reasoning is, but one need only try building a current mode power supply with duty cycle exceeding 50% and no slope compensation to see that there certainly is a subharmonic oscillation; the converter oscillates at exactly half of the switching frequency.

Adding slope compensation to a current-mode-controlled converter (see Figure 6.32) is as simple as adding in some amount of a fixed ramp to the current-sensing feedback. Without going into the technical details, it turns out that adding various amounts of fixed ramp accomplishes either perfect current mode control or perfect audio rejection of the

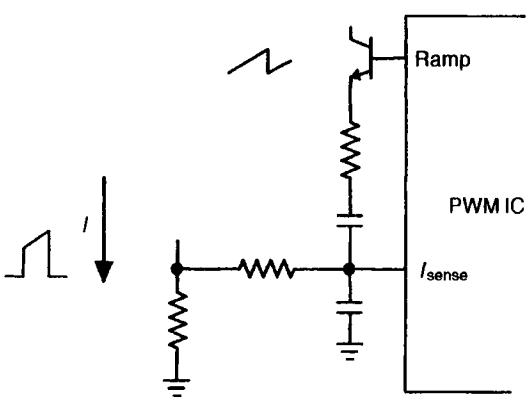

Figure 6.32 Adding slope compensation to a current mode control converter.

supply. But the reality is that tolerances in components and variations of the load make it impossible to attain either of these states.

> **Practical Note** If the duty cycle of a current-mode-controlled converter is going to exceed 50%, the converter needs to have slope compensation. The practical way to determine the right amount of slope compensation is to run the converter at maximum load current and add enough slope compensation to stabilize the converter against subharmonic oscillations. At lower loads the converter will then automatically be stable.

How to Compensate a Current Mode Controller

With one notable exception (discussed below), a current mode converter can be compensated in the same way as a voltage mode converter. Arrange the current-sense resistor to produce about the maximum signal the IC will take (typically, 1V) when the converter is at full load, and then forget about it. If you're going to have a duty cycle greater than 50%, remember to add in some slope compensation. Now you measure the open (voltage) loop exactly as you did for a voltage mode controller (10kΩ and 1μF), design the compensation, and you're through! Don't forget to check the four corners!

Can I Measure the Current Loop?

The foregoing methods have proven so successful and easy at measuring the voltage loop, that I'm sometimes asked about the possibility of measuring the current loop as well.

As a first part of an answer, let's comment on the desirability of making such a measurement. In a general sort of way, the current loop is always stable, with (it turns out) lots of phase margin, at least as long as you remember to add the slope compensation when necessary. So with the notable exception of average current mode control (discussed below), there isn't any need to measure the current loop: you just use it and forget it.

Making the measurement turns out to be quite difficult, and in fact can't be done with a normal network analyzer. You'll recall that a PWM works by having a comparator that changes logic levels when the ramp signal (here, the current signal plus any slope compensation) is equal to the error amplifier output signal. The theoretical aspect of this digital operation is that instead of using Laplace transforms, the system needs to be described with a z transform, or else at least via an analog approximation to the dynamics of the comparator (worked out in recent years by Ridley) involving two RHP zeros.

Practically, because of this digital portion, ordinary swept sine wave analyzers can't be used. Instead, some sophisticated digital modulators have been devised. In any case, this sort of thing is done only in universities, never by practicing engineers. As stated above, the current loop is basically always stable in normal converters.

Average Current Mode Control

As noted in several places throughout this book, *average* current mode control, which is mainly used in off-line power factor correction converters, constitutes an exception to the general rules about current mode control.

The idea of average current mode control is that instead of using a comparator to compare the current signal with the output of the error amplifier, a second amplifier is used to provide some gain for the difference between the current signal and the output of the error amp. Thus, while the standard current mode control current loop has a bandwidth equal to the switching frequency of the converter, the average current mode control current loop can have a reduced bandwidth. In the author's original invention of average current mode control, the current error amplifier could be arbitrarily compensated to achieve any desired bandwidth and phase margin (using the same techniques used in compensating the voltage error amplifier). In the systems in common use today, the current loop is heavily filtered down to around line frequency. With a loop such as this, it is straightforward to measure the bandwidth and phase margin by means of the closed loop techniques described above.

A general requirement for stability is that the outer (voltage) loop have a smaller bandwidth than the inner (current) loop. Of course, this is a no-brainer for the usual current mode control, in as much as the current loop then has a bandwidth equal to the switching frequency.

NON-MINIMUM-PHASE SYSTEMS

Once in a while, you'll get a Bode plot that just doesn't make any sense, even though you're sure you've measured it correctly. For example, it will show the phase as being $-180°$ at low frequencies, crossing through zero degrees up to some maximum, and then coming back down again. This response is symptomatic of a non-minimum-phase system, for which a Bode plot is not sufficient to be able to determine stability.

A non-minimum-phase system is any system that has a right-half-plane zero in its open loop transfer function. What this means can be most easily understood by thinking about how a flyback converter works.

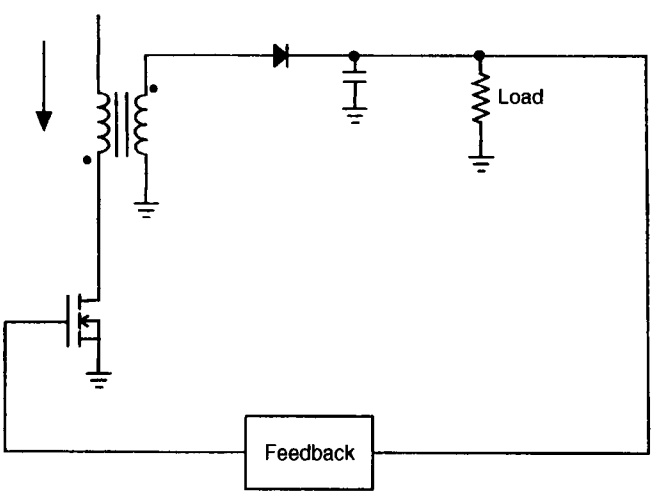

Figure 6.33 A flyback converter responds to a drop in voltage by turning on longer, causing the voltage to fall even further. This is what is meant by a right-half-plane zero.

Consider what happens when the load current increases (the load resistor value decreases) in the flyback converter shown in Figure 6.33. The output voltage begins instantaneously to drop. To supply more power, the feedback increases the duty cycle of the transistor, in order to store more energy in the primary inductance. Instantaneously, though, the transistor sees that it is turned on longer. This means that it doesn't deliver energy "on time" during that cycle, since it can't deliver energy until it's off again. But this causes the output voltage to drop even further; if the loop isn't designed to handle this, the voltage just keeps dropping. This, then, is a 180° phase shift, the essence of a RHP zero: *increasing* the duty cycle *decreases* the output voltage.

Remember from the earlier part of this chapter that an RHP zero causes the gain to go flat while the phase is still decreasing, making it difficult to guarantee stability. As a rule, the bandwidth of the converter will be designed to ensure that the RHP zeros occur at much higher frequencies than the bandwidth. But, caution! These zeros move with load. So, you need to check all four corners of the converter's operation to make sure you're not going to have problems with RHP zeros.

Nyquist Plots

Since the Bode plot of a non-minimum-phase system isn't enough to be able to determine stability, we need to use a different display of the information, called a Nyquist plot. Remember that the Bode plot consists of two graphs, one showing (the logarithm of) the root of the sum of the squares of the imaginary and real parts of the transfer function:

$$\text{magnitude} = \sqrt{\text{Im}^2 + \text{Re}^2}$$

and the other showing the phase:

$$\text{phase} = \arctan\left(\frac{\text{Re}}{\text{Im}}\right)$$

both as functions of frequency. Instead of this, the Nyquist plot plots the imaginary part on the y axis and the real part on the x axis, on a single graph (see Figure 6.34).

The important point on a Nyquist plot is "real part $= -1$, imaginary part $= 0$." Figure 6.34*A* shows the overall view of a Nyquist plot on a large scale (200 units per division). You can see that the graph goes around $(-1,0)$ once in the clockwise direction, but because of the scale of the axes, you can't make out details of what's happening close up to $(-1,0)$. The expanded picture, zoomed in to 1 unit per division (Figure 6.34*B*), reveals detail that isn't visible on the coarser graph: the graph goes once around $(-1,0)$ in the counterclockwise direction [it also contains another loop, but this doesn't matter because it doesn't enclose $(-1,0)$]. The net result is that $(-1,0)$ is encircled zero times: clockwise + counterclockwise $= -1 + 1 = 0$. This guarantees that the system is stable:

A Nyquist plot represents a stable system if there are zero net encirclements of $(-1,0)$.

As a practical matter, a measured Nyquist plot won't be closed the way this demonstration figure is because your measurement doesn't go down to 0Hz nor up to infinity hertz. Nevertheless, measuring the converter response over the normal frequency range, say 10Hz to 100kHz, is sufficient to determine the stability because the gain below

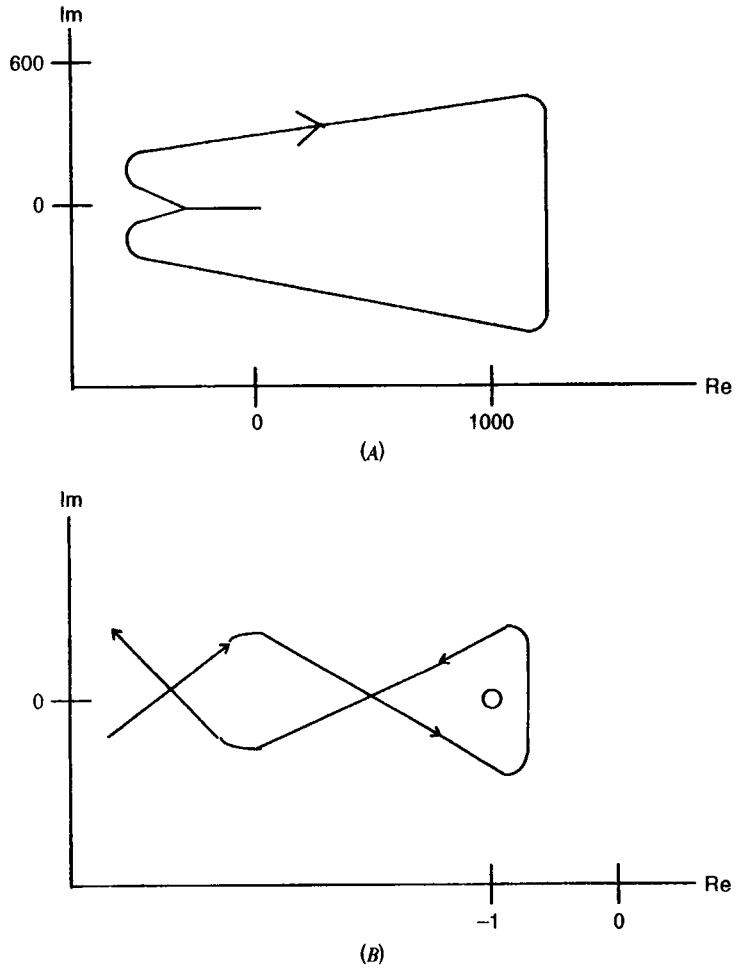

Figure 6.34 (*A*) Nyquist plot of a stable system. (*B*) Close-up of the area around (−1,0). (Modified from Ref. 3, p. 367.)

the low end is constant, and is less than 1 above the high end; encirclements of (−1,0) all occur in the same band of frequencies you would be using if you looked at the Bode plot instead.

SOME CONCEPTS OF SYSTEM STABILITY

Input and Output Impedance

As a final topic in practical feedback design, we're going to talk briefly about converter impedance and its relationship to system stability. Here "system stability" refers in particular to the stability of a group of converters interacting with each other. This is a situation frequently encountered in practical design work: for example, your 5V output converter is going to have a 3.3V output converter hung on its output; or your off-line

power factor correction converter produces 300VDC, and you're going to use a second converter to step it down to ±12V. The key question for stability, it turns out, is what input and output impedance the converters must have to ensure that the system as a whole is stable.

Converter input impedance is a measure of how much the input voltage changes when the input current changes. As such, it is closely related to the transfer function of the converter. Figure 6.35 illustrates a generic method of measurement.

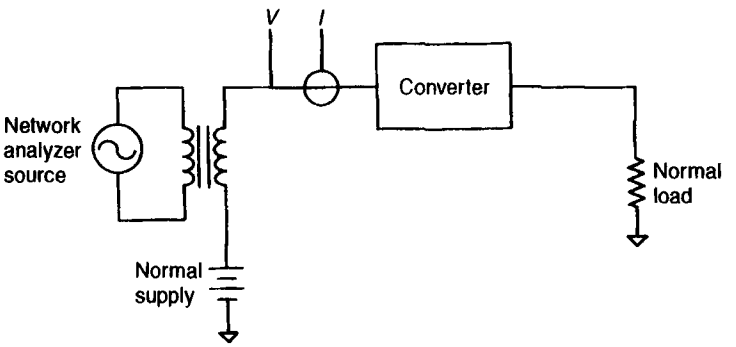

Figure 6.35 Block diagram of method of measuring converter input impedance: $Z_{in} = V/I$.

The idea is that the converter is run with a normal load and a normal input supply voltage. Superposed on top of the normal supply voltage, however, is a small AC signal, a swept sine from a network analyzer. As the frequency of this input voltage varies, the amplitude of the input current does also, and the ratio of these two is the input impedance, $Z_{in} = V/I$, as a function of frequency.

The only tricky part is remembering to use a scale factor: usually the output of a current probe is 10mV per division. So, for example, if you use a 1 : 1 probe for measuring the voltage, and the current probe is set to 1A/div, then $1\Omega = 1V/1A = 1V/(10mV/A) = 100 = 40dB$, that is, the 40dB mark on the network analyzer will be equal to 1Ω.

> **Practical Note** For high power inputs, it may be necessary to use an amplifier to drive the transformer, rather than driving it directly from the network analyzer. You'll find that a good audio amp works well here; actually, one of the old linear (vacuum tube) amps is best because of its low harmonic distortion.

Exactly this measurement was performed on the buck converter we've designed in this chapter. Figure 6.36 shows the plot that was generated.

Let's take a close look at the results of this measurement. At low frequencies, we see that the impedance is approximately flat. Indeed, we expect that the input impedance at low frequencies will be a constant: it should be just the input voltage divided by the input current. For this buck, that is $15V/0.78A = 19.2\Omega = 25dB\text{-}\Omega$, which should show up as $(25 + 40) = 65dB$ on the graph, which is right on the money. (Actually 780mA is the measured input current, not calculated. As an aside, you might notice that input power is

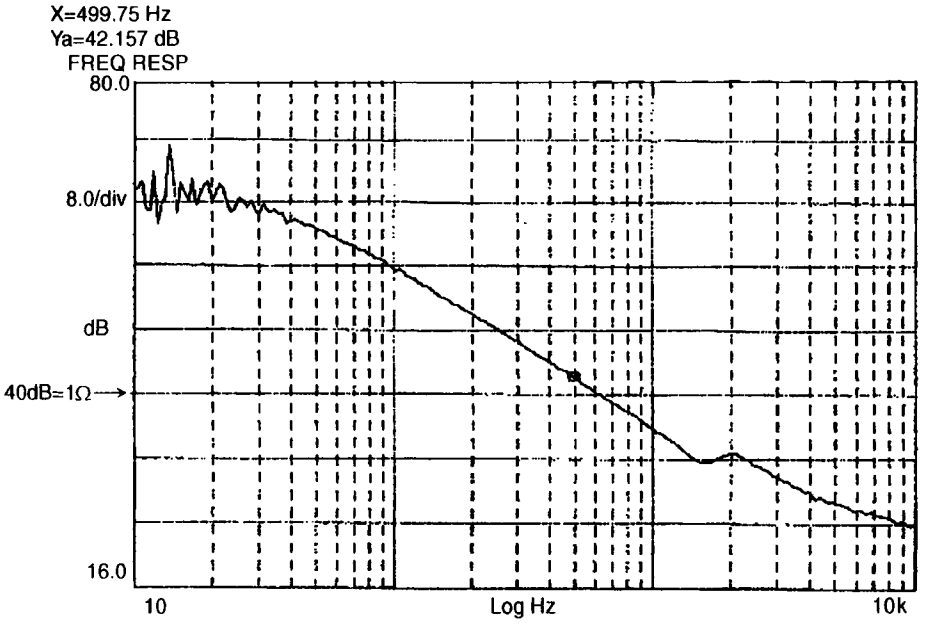

Figure 6.36 Input impedance of the buck converter.

15V × 780mA = 11.7W, and output power is 5V × 2A = 10W, so this converter has an efficiency of 85%.)

Remember that the phase is −180° at these low frequencies because a converter is a constant power load: if you increase the input voltage, the input current decreases! It is this effect that will cause problems when we consider system stability. (Phase is not shown in the measured plot. It's of no real importance unless you're cutting it very close when you get around to considering system stability; see below.)

As a further note, you hear people referring to the converter as a "negative impedance." In fact, they are referring to this 180° phase shift, and the statement is true only at low frequencies.

As the frequency starts increasing, the input capacitor becomes equal to the converter in magnitude of impedance:

$$f = \frac{1}{2\pi \times 19.2\Omega \times 220\mu F} = 38Hz$$

Above this frequency, the input is predominantly capacitive with a 90° phase shift. We can check this by looking at the cursor: 42dB scales to 2dBΩ = 1.26Ω at 500Hz, or

$$C = \frac{1}{2\pi \times 500Hz \times 1.26\Omega} = 253\mu F$$

which agrees reasonably well with the 220μF that was used.

At around the resonant tank frequency, we see a little bit of activity, but it is not nearly as pronounced as in the open loop case because the entire converter is in parallel

with the input capacitance. Then at the top frequencies being measured, we see the gain starting to flatten as the ESR of the input capacitor cuts in:

$$C = \frac{1}{2\pi\ 0.12\Omega \times 250\mu F} = 5.2\text{kHz}$$

with the measured ESR of 120mΩ.

The input impedance of the converter thus looks like a "negative resistance" at low frequencies, like a capacitor at intermediate frequencies, and like a positive resistance at high frequencies. Of course, if you go high enough in frequency, you'll start seeing inductance as well, but at these frequencies, it may become necessary to consider also the cabling used in the system. System cabling can be very important to system stability in certain cases.

Converter Output Impedance

Converter output impedance is very similar in concept to input impedance: When I jiggle the load current, how much does the output voltage change? Ideally, of course, we would like this change to be zero, because we want an output voltage that is independent of load.

A circuit for measuring output impedance is shown in Figure 6.37. In this circuit, the network analyzer provides both a DC offset and a swept sine. It drives an electronically controllable load, which pulls both DC and AC current from the converter. (Make sure the amplitude of the AC current is small enough to ensure that the load is always pulling current—it can't source current!) The output impedance is $Z_{out} = V/I$ as a function of frequency.

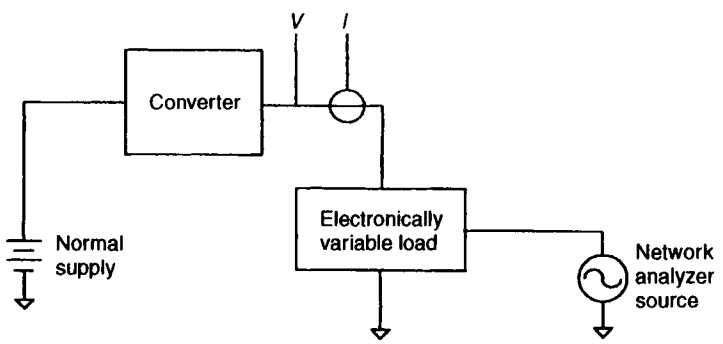

Figure 6.37 Block diagram of method of measuring converter output impedance: $Z_{out} = V/I$.

> **Practical Note** You don't want to put a resistive load in parallel with the electronic load, since changing the output voltage changes the current through the resistor.

Using the buck converter once more, a measurement was made of the output impedance (remember the scale factor, again). Figure 6.38 is a typical-looking plot for output impedance. Take note that the scale is offset from the one in Figure 6.36 for input

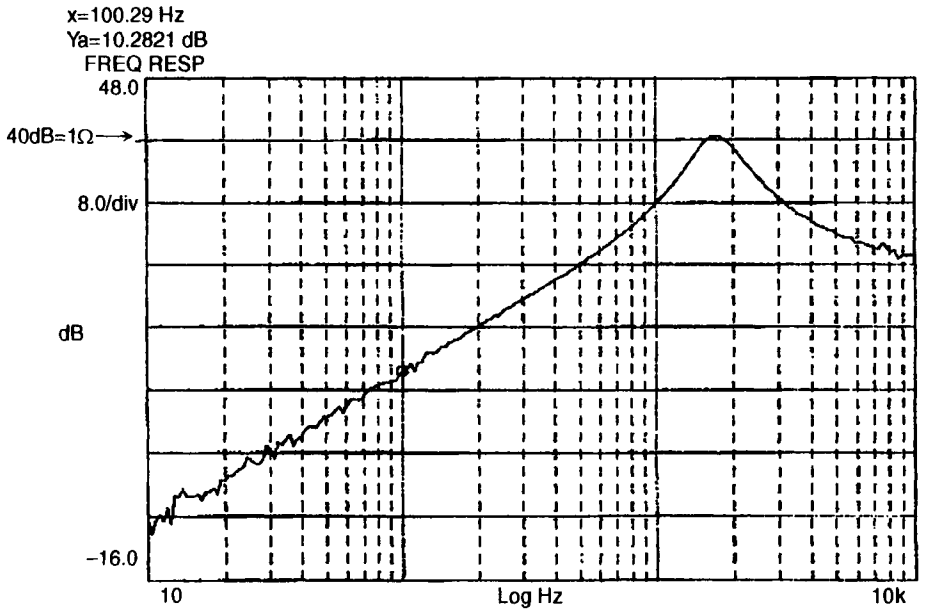

Figure 6.38 Output impedance of the buck converter.

impedance ($1\Omega = 40$dB is almost at the top of the graph). At low frequencies, the output is inductive. In fact, we have at 100Hz, from the cursor, that the impedance is 10.3dB, which is -29.7dB-$\Omega = 32.7\text{m}\Omega$, so the inductance is

$$L = \frac{1}{2\pi \times 100\text{Hz} \times 32.7\text{m}\Omega} = 49\mu\text{H}$$

agreeing well with the actual value of 35μH. At the output tank resonant frequency, the output impedance peaks and thereafter is controlled by the output capacitor (you can again see the ESR at the top frequencies).

Two Stable Converters Can Make an Unstable System!

This rather shocking idea brings us to the forefront of today's research in power systems. It is quite possible (and indeed frequently happens) that you have two converters, each of which is stable with plenty of phase margin; but when the one is connected as a load to the other, the system, meaning here both their output voltages, oscillates! Figure 6.39 diagrams the problem.

Here is a rule of thumb to ensure that attaching two stable converters together in series won't cause oscillations:

Practical Note If possible, ensure that

1. The output impedance of the first converter is less than the input impedance of the second converter at all frequencies.
2. The bandwidth of the first converter is greater than the bandwidth of the second converter.

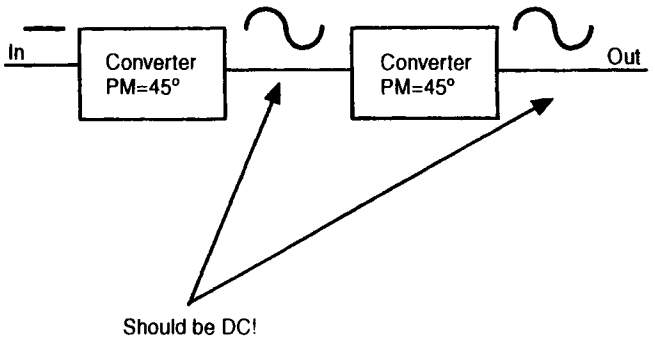

Figure 6.39 Two stable converters in series can form an unstable system.

This is merely a convenient way of ensuring stability; a system can be stable even if it doesn't meet this rule, although the question then becomes substantially more involved. When you are actually setting up such a system, of course, you need to verify that each individual converter is stable before attaching the set together!

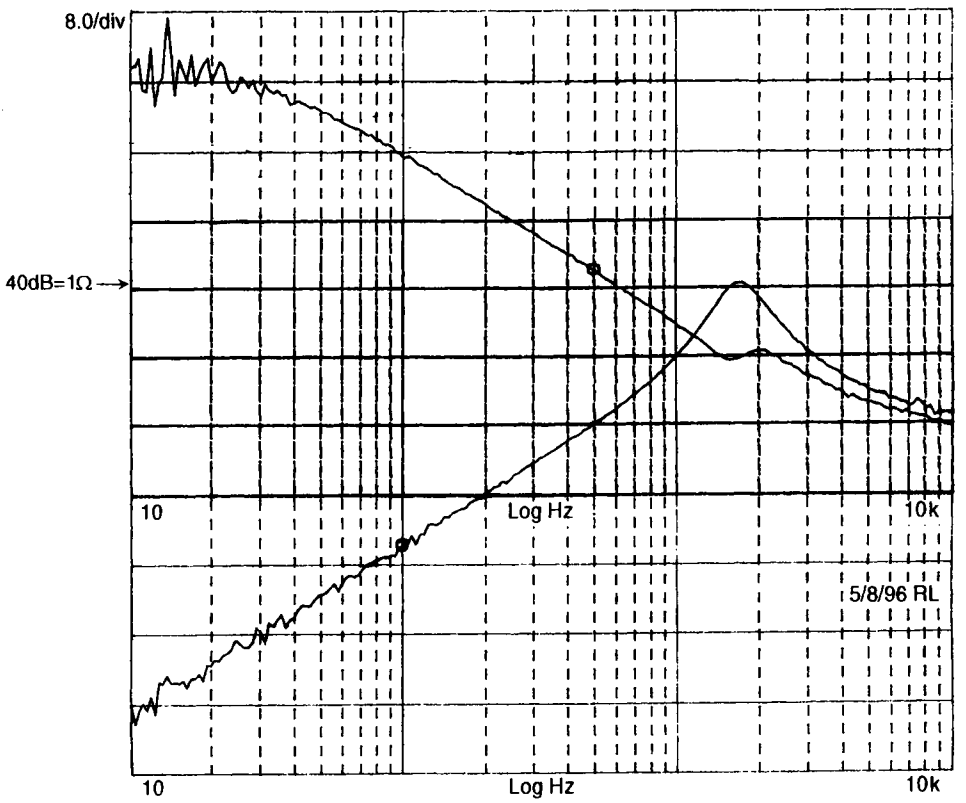

Figure 6.40 Input and output impedances superposed.

Example of an Unstable System

We have designed the compensation of our buck converter to produce a stable device. Let's suppose that for some (crazy) reason we connect two of these bucks in series. (That is, we would reset the first one to produce an output voltage of 15V so that it could run the second one. This could be accomplished by making the first converter run from $45V_{in}$, and tripling the switching frequency so that all the components retain their values; the loop and the impedances of the converter would then remain the same.)

How well does this system comply with the rules given above? Admittedly, rule 2 is marginally OK; at least the downstream converter doesn't have more bandwidth than the upstream one does. But rule 1 is definitely violated, as can be seen in Figure 6.40.

A look at the superposition of the input and output impedances shows that the source impedance (output impedance) is greater than the load impedance (input impedance) at frequencies greater than about 1.2kHz. This doesn't guarantee that the system will be unstable (you actually have to examine system, as opposed to converter, phase margin), but don't be surprised to see the 15V and the 5V outputs oscillating at about 1.2kHz!

SOME THOUGHTS ON THE ROLE OF SIMULATIONS

Although this book explicitly avoids talking about simulations, a few comments are in order here, because it is in the area of stabilizing converters that people most frequently turn to simulations.

For many years now, people (read managers) have been talking about the possibility of designing power supplies using just simulations, no breadboarding. Leaving aside the problematic existence of adequate computing power, one can see without going into the details that such a thing might be possible. To aid in this visualization, consider a simulation of the open loop of the buck converter we've been dealing with (Figure 6.41).

We're not going to discuss this simulation in any detail. Figure 6.41 shows a pseudodrawing of the blocks of an open loop buck, including a state-space-averaged switch model; the listing is part of a SABER listing, and the plots were generated using the SABER simulator. The names in circles are the node names, and the other names are the part names and values. Further information on simulations can be obtained from a number of books. At any rate, you can see that the simulation, shown in Figure 6.42, matches the measurements pretty well, certainly well enough to allow you to design the compensation, and that compensation would have the same values we came up with before. The trouble is, even assuming the data on ESR and so on are available from data books (which is only occasionally true), the model says nothing about switching. Thus you have to build either a much more complex model or a breadboard—in which case you've defeated the purpose of the simulation! Besides, unless you are *very* experienced (and even then), you need to measure a working unit to verify that you've built the simulation model correctly. It's so easy to type in incorrect numbers and assume that because the computer gave you an answer, this output corresponds to reality!

In the author's view, nothing is gained by using simulation to design a compensation, in part because doing it by real measurements is so easy. Simulation *is* useful in measuring worst case: it can be almost impossible in the lab to find components at their minimum or maximum values, whereas with a simulation it's easy. Simulation and breadboarding are

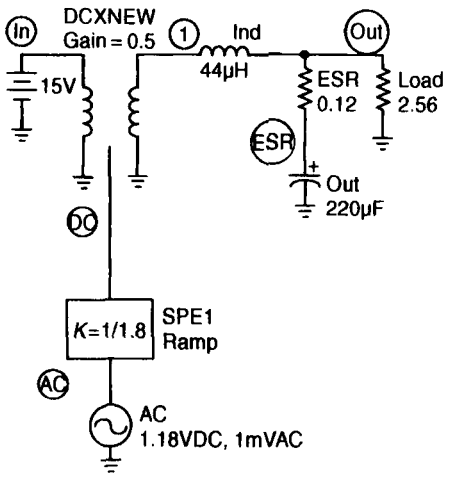

Simulation Listing (SABER)

r.esr	out esr	=120m
r.load	out 0	=2.56
c.out	esr 0	=220u
l.ind	1 out	=44u,r=18m
v.in	in 0	=dc=15
v.ac	ac 0	=dc=1.18,ac=(1m,0)
dcxnew.x1	1 in dc	=gain=0.5
spe1.ramp	dc 0 ac 0	=e=[0,1/1.8]

Figure 6.41 Simulation model of the buck converter open loop: circles enclose node names; see accompanying listing.

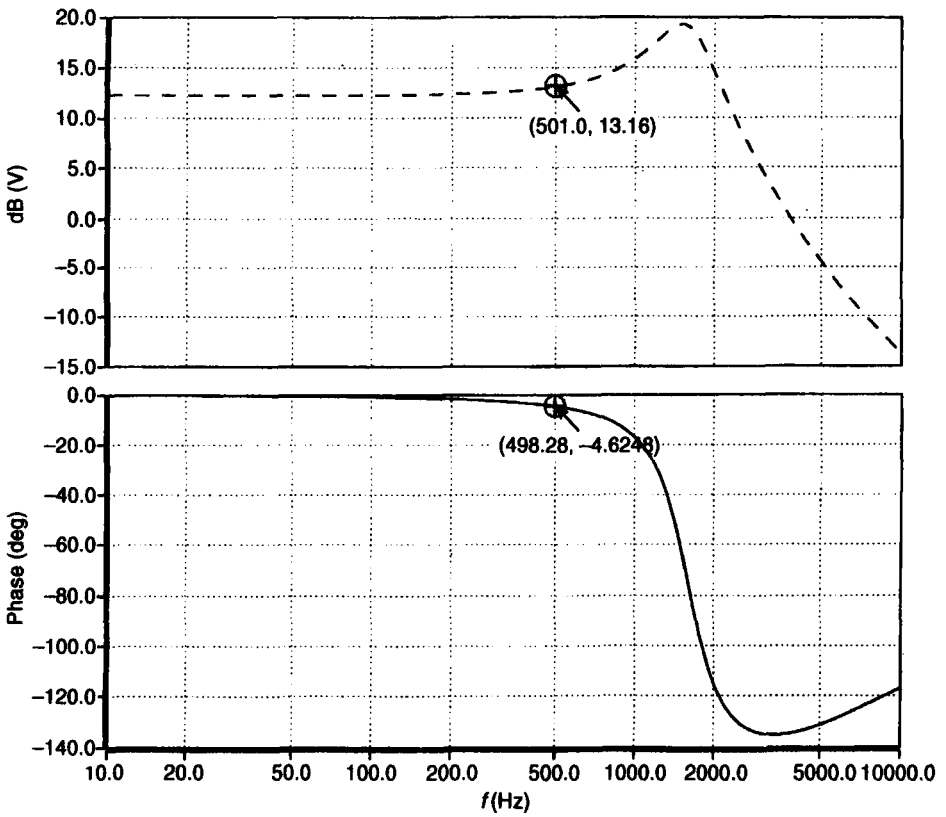

Figure 6.42 SABER simulation of open loop buck converter: both curves, OUT/AC.

thus complementary, not opposing: breadboarding is for doing the design, simulation is for worst-case analysis.

REFERENCES

1. H. Dean Venable, The *K* factor: A new mathematical tool for stability analysis and synthesis, *Powercon*, 10, H1–H12 (1983).
2. Thomas V. Papathomos, On the stability of peak current–controlled converters: Analysis, simulation, and experiments, *IEEE Transactions on Industrial Electronics*, IE-33(2), 176–184 (1986).
3. J. Kassakian, M. Schlecht, and G. Verghese, *Principles of Power Electronics*. Reading, MA: Addison Wesley, 1992.