Phil Perkins

22. My Approach to Feedback Loop Design

I like designing feedback loops. I have been designing and building feedback controlled systems for audio and low frequency control since high school. My interest in high fidelity audio started in the late 1950s. Transistors were scarce and not very good, so I worked with vacuum tube circuits. I learned that negative feedback would improve just about every characteristic of an audio amplifier. I built Heathkits and modified them to suit my own preferences. I experienced oscillation when there was too much feedback.

For a freshman project at M.I.T., I learned how negative feedback could transform an unstable device into a stable one. I built a device to suspend a steel ball a fraction of an inch below an electromagnet driven by a tube amplifier. The position of the ball was sensed by a light shining on a photocell. The ball would partially interrupt the light as it was pulled higher. The photocell output was fed back to the amplifier input to control the magnet. After I got the hookup right, the first thing the circuit did was oscillate. I tried out my newly acquired capacitor substitution box and discovered a network that would tame the oscillation. I later learned that it was called a lead-lag network. I was developing an intuitive feel for what to do to make feedback stable.

During my studies at M.I.T., I learned about circuit theory, circuit analysis, and feedback circuit analysis. M.I.T. taught methods for analyzing a circuit "by inspection" as well as the usual loop and node equations and mathematical analysis. I learned the theory of analyzing circuits and transforming between the time domain and the frequency domain. Then I could relate my early experiences to the theory. Along with learning the theory, I really appreciated learning methods of analyzing a circuit by inspection to get approximate results.

Much of the feedback loop design work I do is satisfied during the design phase with only rough approximations of performance. Actually, there are so many variables and effects you cannot consider during the design of a circuit, it is often useless to analyze with great precision. (You don't need a micrometer to measure your feet for a new pair of shoes.)

Since graduating from M.I.T., I have worked in the semiconductor Automatic Test Equipment (ATE) field, designing instrumentation and other parts of ATE systems. First I worked for Teradyne and now I work for LTX. At Teradyne and LTX I have designed several programmable power sources. These programmable sources make heavy use of feedback loops. I have developed a method for design and analysis, which I would like to describe here. I work and communicate better when I use a specific example to illustrate what I am designing or describing. The example I will use here is a programmable voltage/current source I designed for LTX. The drawings are based on sketches I made in my notebook during the development of that product.

My Approach to Design

First, I need a specification of the instrument I am going to design. Then I make a block diagram of the circuit. Also, I will draw a "front panel" of the instrument to show its functions and how they are controlled. This "front panel" has knobs and switches and dials, even though the finished product may be software controlled. The "front panel" helps to evaluate the functions that were specified and to investigate interactions between functions. In other words, does it do what you wanted, the way you want it to?

After I have a block diagram, I like to start the circuit design with a greatly simplified schematic made with a few basic building blocks of ideal characteristics. These blocks are simplified models of the real circuit elements I have to work with. I prefer to design the final circuit with blocks that work similarly to these basic blocks. The basic circuit blocks I like to use include:

- Amplifier with flat frequency response
- Ideal op amp
- Ideal diode
- Ideal zener (voltage clamp)
- Voltage output DAC

To analyze a specific aspect of a design, I make a new schematic that eliminates anything that won't significantly affect the results. I want to be able to analyze by inspection and sketches. After reaching a conclusion, I might check my assumptions on a more complete schematic, or I might build a prototype and make measurements.

I use a notebook to record, develop, and analyze my designs. I draw block diagrams, schematics (from simple to detailed), and sketches of wave forms and frequency responses. I keep the notebook in chronological order and draw new drawings or sketches when there is a significant change to consider. During the preliminary design phase of a project, I might draw dozens of similar sketches as I develop my ideas and the design. I draw with pencil so I can make small changes or corrections without having to redraw the whole thing. I date most pages in my notebook, and I usually redraw a diagram if I invent a change on a later day. I also record the results of experiments and other measurements.

I have my notebooks back to the beginning of LTX. They have been a great source of history and ideas. Sometimes a question comes up that can be answered by going back to my notebooks rather than by making new calculations or experiments. There is real value in having diagrams, sketches, notes, and test results all in that one place. Recently, though, I have been using various CAD and CAE systems to record some of my design developments. Sometimes, the precision that the computer insists upon has been helpful, and other times it's a hindrance. With CAE results, I now have two or three places where parts of the design process are documented. I need to develop a new system for keeping all the historical data in one place. Even with CAE, I don't expect to ever give up hand writing a substantial part of my design development notes.

What Is a V/I Source?

Integrated circuits need to be tested at several stages of their manufacture. Electrical testing is done with automatic test equipment (ATE). One of the instruments of an ATE system is the programmable voltage source. It is used to apply power or bias

voltage to a pin on the device under test (DUT) or to a point in the DUT's test circuit. Programmable voltage sources usually can measure the current drawn at the output, and sometimes include the capability of forcing current instead of voltage. In that case, the instrument is called a V/I source.

A V/I source I designed at LTX is called the DPS (device power source). It is part of the Synchromaster line of linear and mixed-signal test systems. The DPS can force voltage or current and measure voltage or current. It's output capability is ±16 V at 1 A to ± 64 V at 0.25 A. The current measure ranges are 62 μA full-scale to 1 A full-scale. There is great opportunity for analog design sophistication and tricks in designing a V/I source. Some typical performance requirements are:

- 0.1% forcing and measuring accuracy (at the end of 20 ft of cable)
- 100 μsec settling time after value programming or load change
- Minimum overshoot and ringing (Sometimes 1 V overshoot is acceptable, other times overshoot must be zero.)

There are many interesting aspects to the design of a V/I source. Perhaps the most challenging is to design a V/I source that can force voltage with acceptable performance into a wide range of capacitive loads. Why capacitive loads? In many (perhaps most) of the situations in which a voltage source is connected to the DUT or its test circuit, there needs to be a bypass capacitor to ground. This is the same bypass capacitor you need when actually using the DUT in a real end-use circuit. Occasionally, bypass capacitors need to be gigantic, as in the case of an audio power amplifier IC that requires 1000 μF.

At the same time the V/I source is forcing voltage, the test may require measuring the current drawn by the DUT. The current measure function usually requires a selection of resistors in series with the source output. The current is measured as a voltage across the resistor.

An Ideal V/I Source

It would be ideal if the capacitive loading on the V/I source output had no affect on the force voltage response to programming changes. However, one response effect we should accept is a reduction of the output voltage slew rate as the capacitive load increases, due to the limited amount of current available from the source. Since the V/I source must have some slew rate limit even at no load, the ideal voltage wave form would look like the one in Figure 22-1.

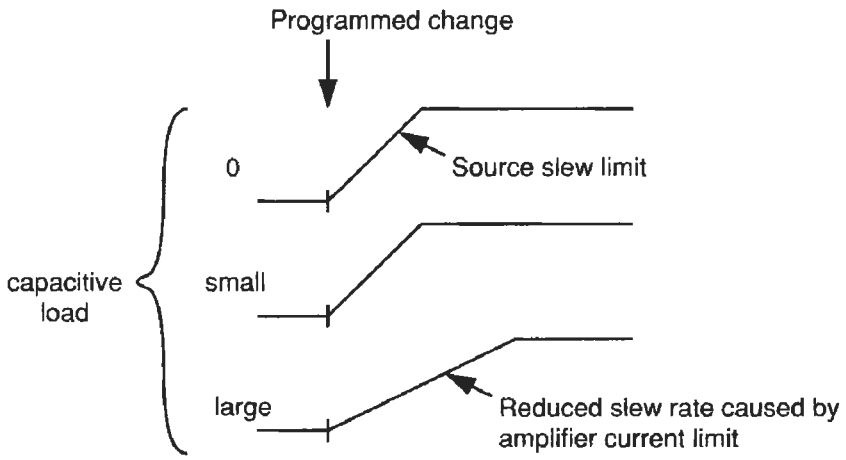

Figure 22-1. Ideal voltage wave form.

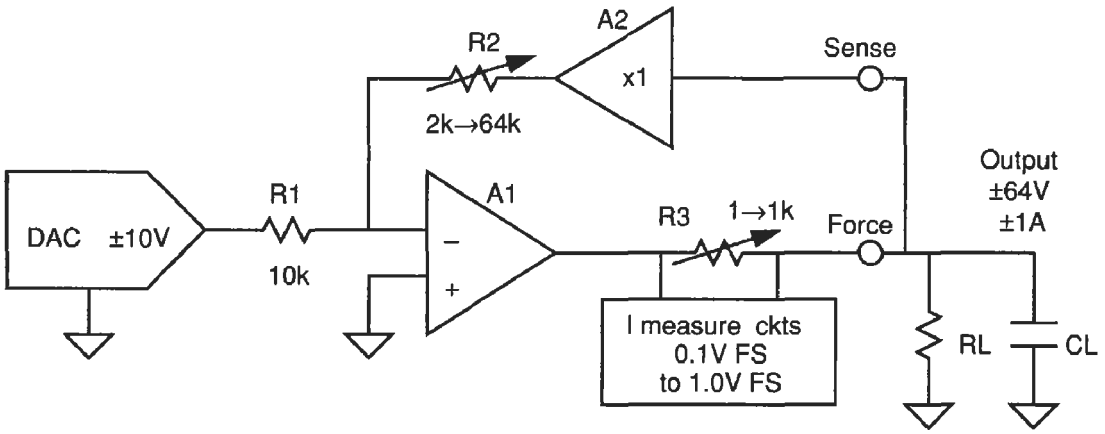

Figure 22-2. First model of the programmable voltage source.

In this ideal case, the voltage will settle to the programmed value just as soon as it is finished slewing. In practice, there is probably no way to achieve the ideal case. The design of a V/I source involves lots of compromises.

Designing a V/I Source

To illustrate my approach to feedback loop design, I will describe a circuit that I developed for the LTX V/I source called DPS. A feature of the DPS is compensation for capacitive loading. With capacitive load compensation, the DPS can drive loads of any reasonable amount of capacitance with good stability and without overshoot.

The figures in this sections are very close to the kind of drawings I do in my notebook. Here is a first model, shown in Figure 22-2, of the circuit to consider. I have simplified it to concentrate on the force voltage mode and the response when driving capacitive loads. This model meets the performance requirements listed above.

I picked 10 kΩ for R1 because that value works well in real circuits and calculations would be easy. I like to have a real value in mind just to simplify the analysis. R2 sets the output voltage full-scale ranges of 2 V to 64 V. R3 varies from 1 Ω at 1 A full-scale to 1 kΩ at 100 μA full-scale. The ×1 amplifier A2 eliminates the loading effect of R2 on the current measure function. With CL possibly as large as 1000 μF, I expect that the rolloff caused by R3 and CL will be a major source of stability problems.

To investigate the effects of capacitive load on this model, I will simplify it further (see Figure 22-3). One of the objectives of a model is to reduce the problem to the basics and make approximate performance estimates easy (or at least possible). Figure 22-3 shows the simplified version.

I have given the simplified model a single voltage and current range. I have left out the load resistor RL, retaining only the capacitor load. I have retained the ×1 amplifier to remind me that R2 does not load the output. If the compensation works in the simple case, then I expect it can be expanded to work on the complete DPS.

I like to design a feedback loop with just one integrator or dominant pole. The other stages are preferably flat frequency response. With this approach in the model, the analysis is easy. I want to design the real circuit the same way. Settling time is related to the lowest frequency pole in the loop. Therefore, extra poles

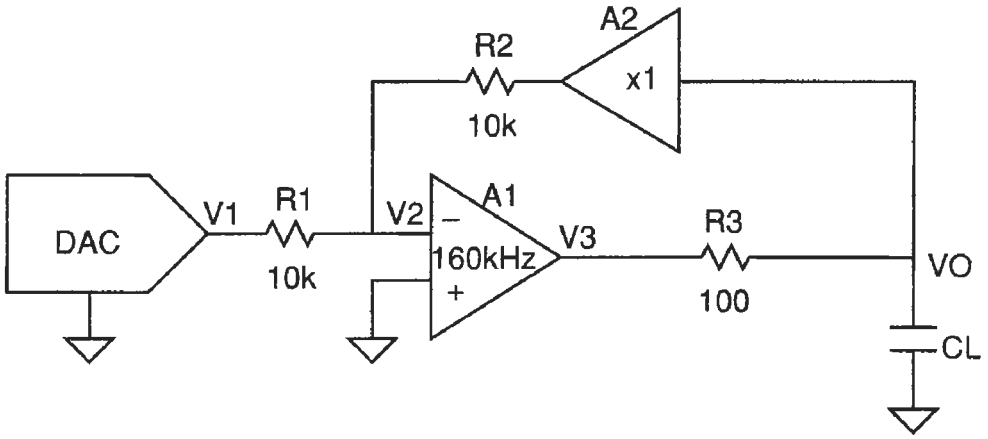

Figure 22-3. Simplified model to investigate capacitive loading effects.

below the loop unity gain frequency (UGF) make the settling time longer. That's bad in ATE applications, where fast settling is very important. Poles below the loop UGF may be needed, but they should be carefully considered.

I have given op-amp A1 a UGF of 160 kHz. Why 160 kHz? My past experience with programmable sources has shown 100 kHz to be the maximum practical loop bandwidth. More than 100 kHz leads to oscillation due to phase shift from the many amplifiers that will ultimately be in the loop and from the output cable. The output cable? The output cable may be 20 ft long; the cable inductance has a significant effect.

As with other parameters I choose, A1's UGF is a convenient approximation. 160 kHz corresponds to a time constant of 1 µsec. We often need to switch between time domain analysis and frequency domain analysis. I like to remember a simple conversion to eliminate the need to consult a table or calculator while designing. Just remember that 1 µsec corresponds to 160 kHz; you can extrapolate from there—2 µsec → 80 kHz, 100 µsec → 1.6 kHz, etc. Conversely, 1 MHz corresponds to 160 nsec.

Now we need to analyze the model. Bode plots for the individual stages and the complete loop will give us an idea of the loop stability and bandwidth. Let's start with CL = 0, as shown in Figure 22-4:

The loop gain is unity at 80 kHz and the phase shift is 90°, the "ideal" case for stability.

What happens when CL is increased? First, make new Bode plots as shown in Figure 22-5.

Figure 22-4. Bode plots for individual stages and complete loop.

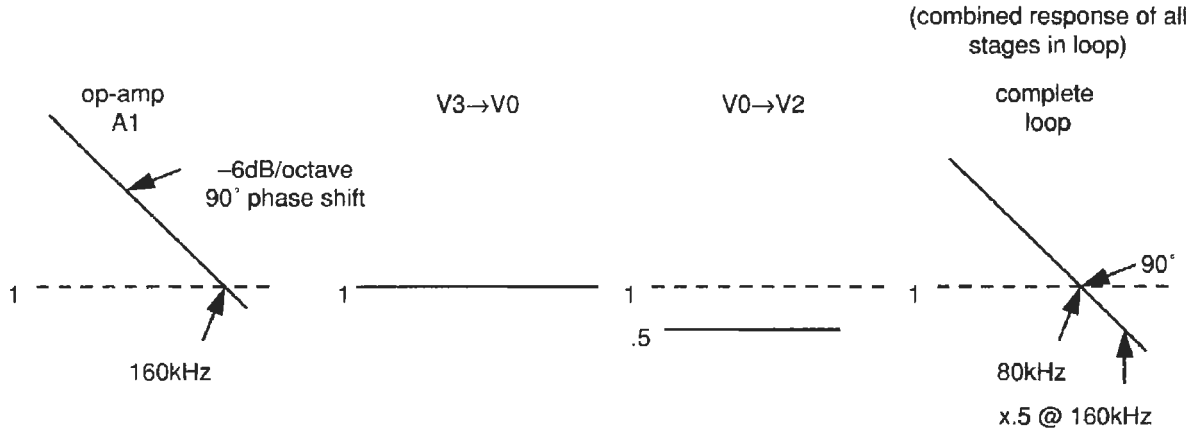

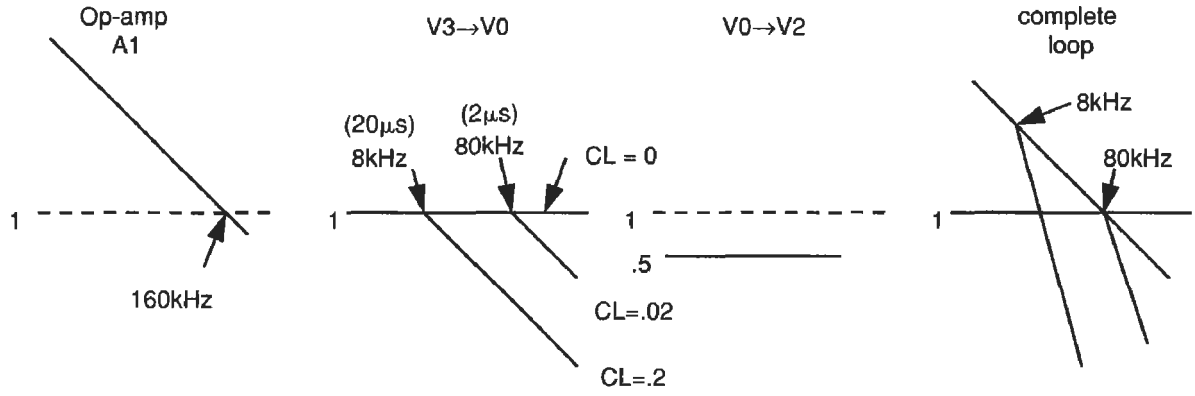

Figure 22-5.
Revised Bode plots.

For CL = 0.02 μF the R3CL rolloff starts at 80 kHz adding 45° phase shift at the loop UGF. For CL = 0.2 μF the added phase shift at 80 kHz is nearly 90°. From experience and by inspection, I would estimate the step response of this circuit to be as shown in Figure 22-6.

For CL = 0, the only time constant in the loop comes from the op amp A1. The step response will be a simple decaying exponential of time constant 2 μsec (based on loop UGF of 80 kHz). At CL = 0.02, there is 45° phase shift added to the loop at the UGF; I estimate a small amount of ringing, perhaps 1 cycle. At CL = 0.2, I know the added phase shift is higher (nearly 90°) and expect more ringing.

I don't need to estimate closer than ±30% or so because this circuit is a greatly simplified model. The main thing we need is to get within a factor of 2 or so of the real performance.

Capacitive Load Compensation

In a previous attempt at cap load compensation, I added an RC network as shown here in Figure 22-7. This technique is similar to adding a resistor and capacitor from force to sense on the output of a power supply. That is one way to stop oscillation

Figure 22-6.
Simplified step model response.

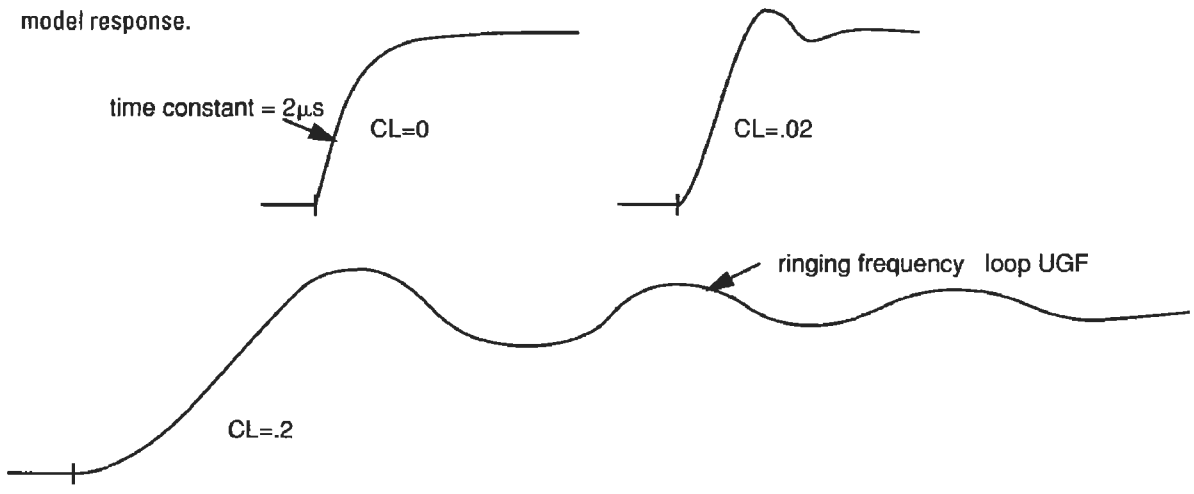

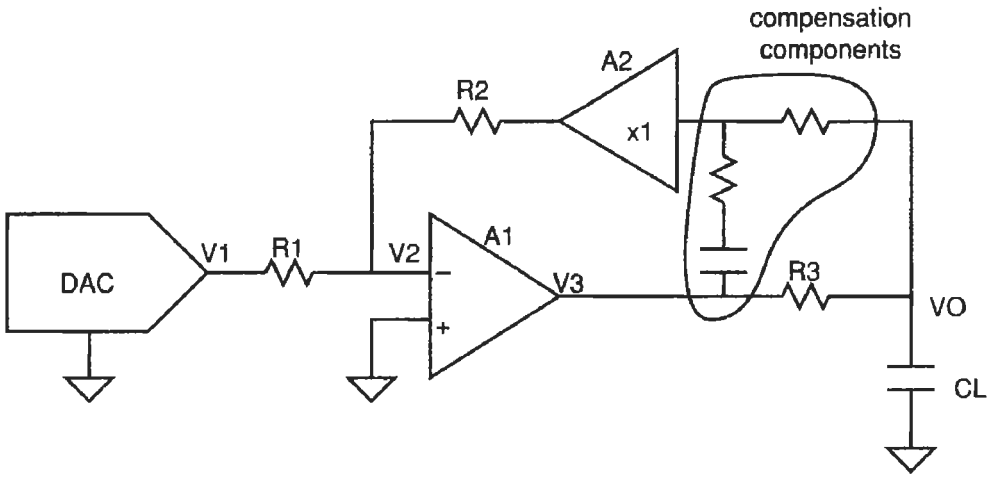

Figure 22-7. Previous attempt at capacitive loading compensation.

when there is a long cable to a capacitive load. This V/I source has a programmable choice of four RC networks for compensation. The RC network can be chosen to eliminate ringing for a range of capacitive loads, at the price of increasing settling time for no load. The overshoot is not reduced, however. Overshoot has become a serious issue in a few cases, so when I designed the DPS I wanted a better compensation technique that would reduce ringing and overshoot.

What Causes the Overshoot?

It will be easier to investigate the cause of overshoot with a better model. The above model is completely linear. It behaves the same for small or large signals. The amplifier in the model is an ideal op amp. A real amplifier will be limited to some maximum slew rate and will be limited in its current output. Here we can do a thought experiment. Assume instant settling time and perfect stability. Make a feedback loop from an amplifier with a slew rate limit and a current limit. Then the step response would look like the ideal, as shown in Figure 22-8.

These wave forms show the ideal effect of capacitive load on a voltage source: no effect until the capacitor current would exceed the amplifier current limit, then the slew rate is decreased according to the current limit. Let's see how close we can get to this ideal.

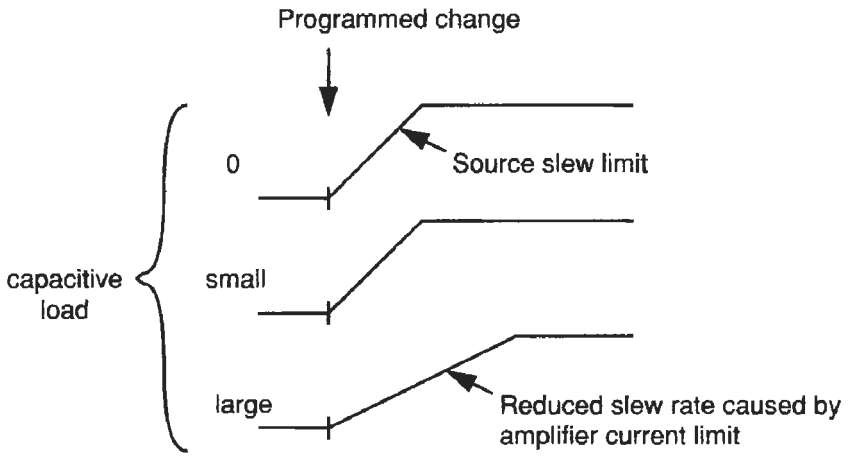

Figure 22-8. Ideal step response.

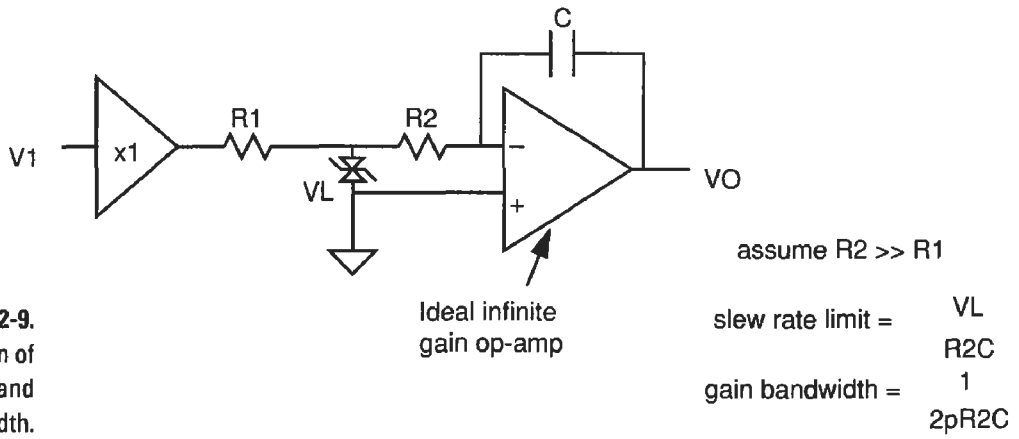

Figure 22-9. Interaction of slew rate and bandwidth.

Improving the Model

To add slew rate limiting to the circuit, we could simply redefine the characteristics of the op amp A1 to include a slew rate limit. However, that would add several parameters to remember when analyzing the circuit. I prefer to add a few ideal components to the circuit that would add the slew rate limit and make it easier to visualize the effect.

To model a slew rate limited op amp, I take an ideal op amp with very high bandwidth and put components around it. The components show the effect and interaction of slew rate and bandwidth, as shown in Figure 22-9.

The double anode zener represents an ideal bidirectional clamp that limits the voltage to VL. I have simplified the model to be an inverting amplifier only, since that is what our circuit needs. A "better" model with differential inputs would just make the analysis by inspection harder.

Note that slew rate and gain-bandwidth (GBW) can be set independently of each other by changing VL. Real op amps show this interaction of parameters. Compare a bipolar op amp and a FET op amp of similar GBW. The FET op amp has higher slew rate but needs a larger input voltage than the bipolar to get to its maximum slew rate. A feedback loop built from this model will be linear when V1 < VL and will be in slew rate limit when V1 > VL.

Model to Investigate Overshoot

I have made a new circuit model to include the slew rate of a real amplifier. I simplified the op amp model a bit to make the circuit easier to analyze, as can be seen in Figure 22-10. In this case the simplification should have no effect on the accuracy of the analysis. This op amp model is good here because it is easy to see at which point the circuit becomes nonlinear. When V2 < VL the circuit is linear. When V2 is clamped at VL, the op amp is at slew rate limit (see Figure 22-10).

With this model, I can estimate the time response. When CL = 0, the output will be at slew rate limit until V2 becomes smaller than VL. From that time on, the circuit is linear and the output is a simple time constant response.

This circuit *must* overshoot with capacitive load. That's because CL causes a delay between V3 and VO and there is also a delay from VO back to the amplifier output V3. When the output is programmed from zero to a positive value, V2 starts negative and stays negative until VO gets all the way to the programmed value. At that point, V3 is above the programmed value and positive current is still flowing into CL. In order to make V3 go down, V2 has to be positive. VO has to go above

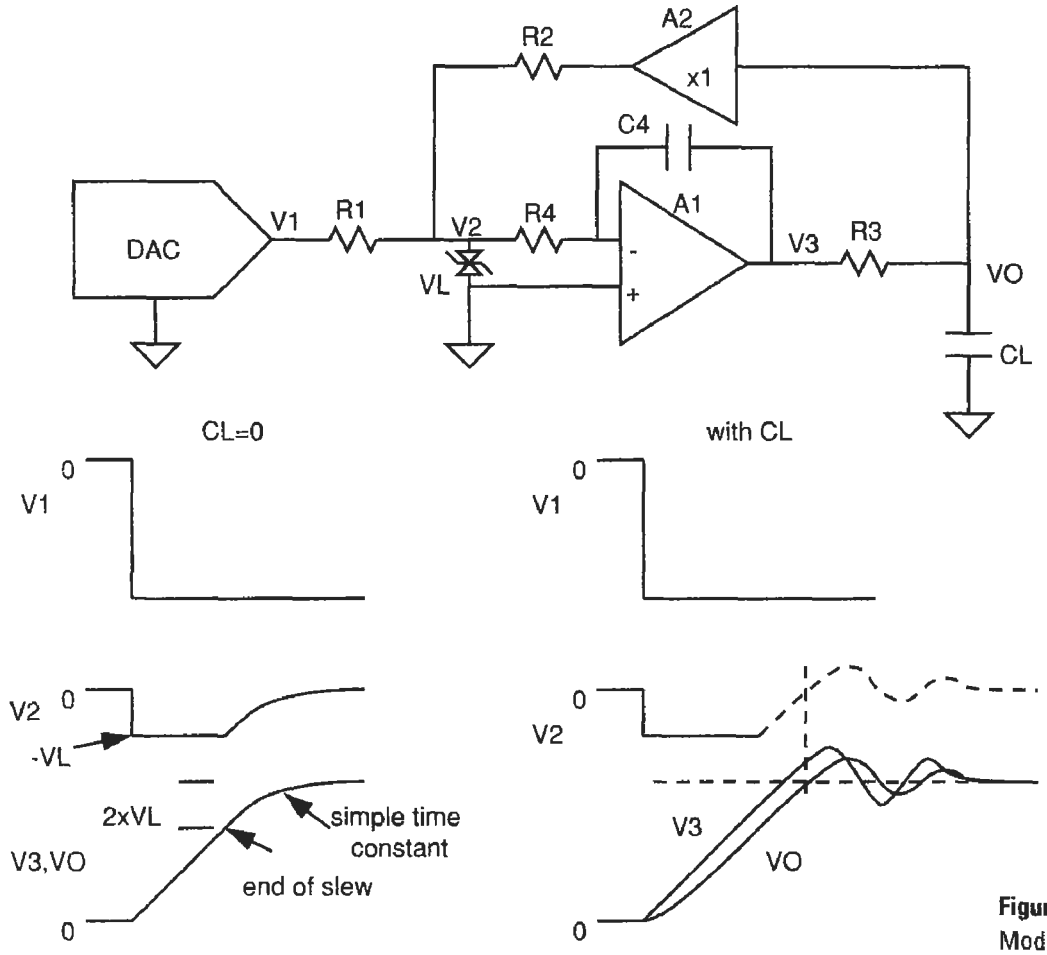

Figure 22-10. Model to investigate overshoot.

the programmed value in order for V2 to go positive, and that's overshoot. It's a little like trying to stop a boat the instant you get to the dock. There is no brake—all you can do is row backward, which takes time to take effect. Likewise, it's too late to stop VO from overshooting if you don't reverse the current in CL until after VO gets to the value you want to stop at.

How do you stop the boat at the dock? Start slowing down before you get there. In this model, we need an anticipator to reduce the slew rate of V3 before VO gets to the programmed value. One way to do this, as shown in Figure 22-11, is to add a capacitor C2 across the feedback resistor R2. The faster the output is slewing, the more current C2 adds to the null point (V2) and turns down the amplifier input in anticipation of VO getting to its programmed value (see Figure 22-11).

During the time that VO is slewing at a constant rate, the current in C2 is constant and provides a fixed offset current into the loop null point at V2. Without C2, VO started to slow down only when it was 2VL away from the final value. With C2, VO starts to slow down when it is an additional (slew rate × R2C2) away from the final value. That's the anticipation. What happens with different values for C2? If C2 is too small, then VO will overshoot. If C2 is too large, then VO will be too slow to settle. Since ATE applications need minimum settling time, C2 needs to be set just right. Therefore, the value of C2 would have to be programmable according to the load capacitance.

In addition to affecting the transient response, C2 changes the loop gain. At high frequencies where C2 is low impedance, the loop gain is increased by a factor of 2.

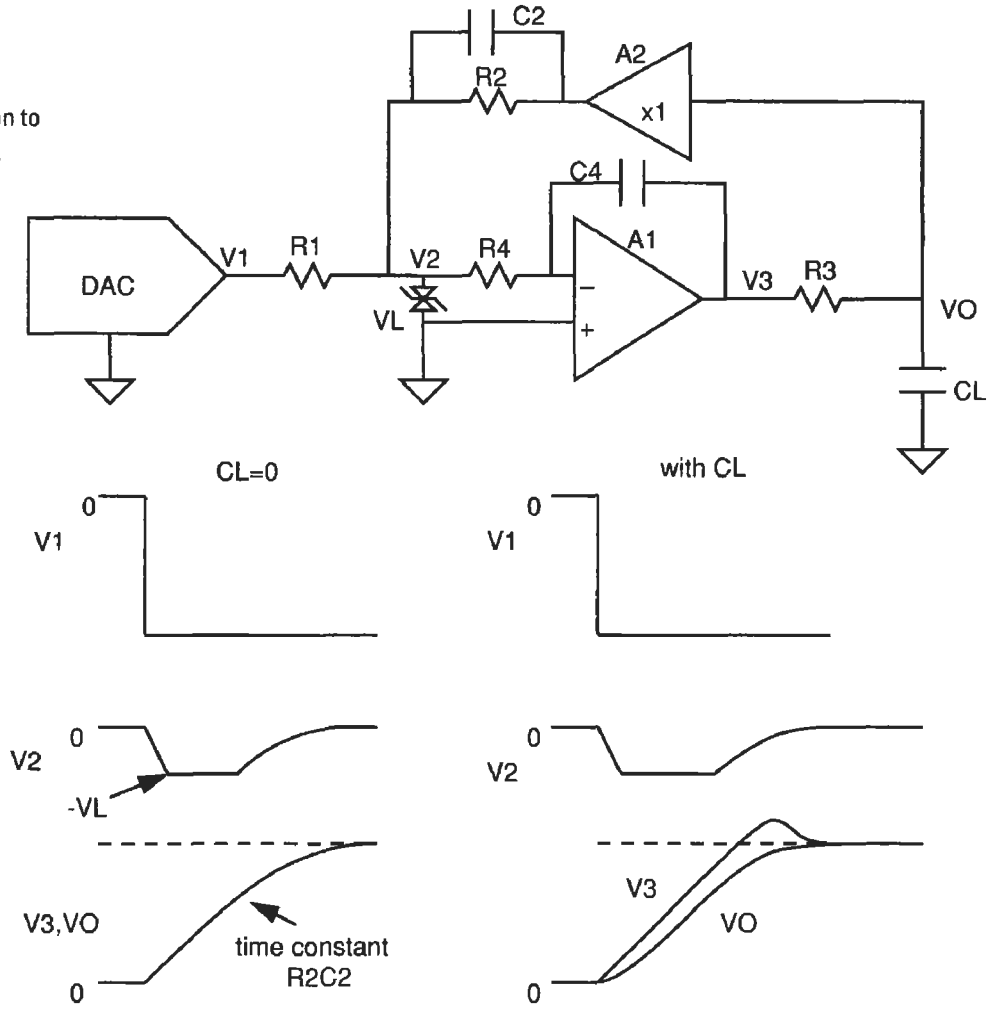

Figure 22-11.
Adding anticipation to reduce overshoot.

Looking ahead to the real DPS, the closed loop gain from DAC output to source output will be programmable. With other values for R1, C2 would increase loop gain by a different amount, even though the anticipation was the same. I want to make the frequency response independent of range, so I will put C2 in an inverting op-amp circuit. This way, the compensation will behave more nearly ideal. My goal is to have simple, independent circuit blocks that can be implemented in the final circuit (see Figure 22-12).

I have added R6 in series with C2 because the real A3 circuit would have a gain limit at high frequency. In the model, R6 allows us to use an ideal infinite gain op amp for A3 and still set a realistic maximum gain for this stage. For frequencies up to several times the loop UGF, I want the gain of each stage to be set by the passive components and not limited by amplifier gain. That way, the frequency response is sufficiently predictable and the amplifier could be changed to another type without serious effect.

Back To The Frequency Domain

At this point, we have an idea that improves the time response of the circuit, but we have ignored loop stability. To get an idea of the small-signal stability, we need to

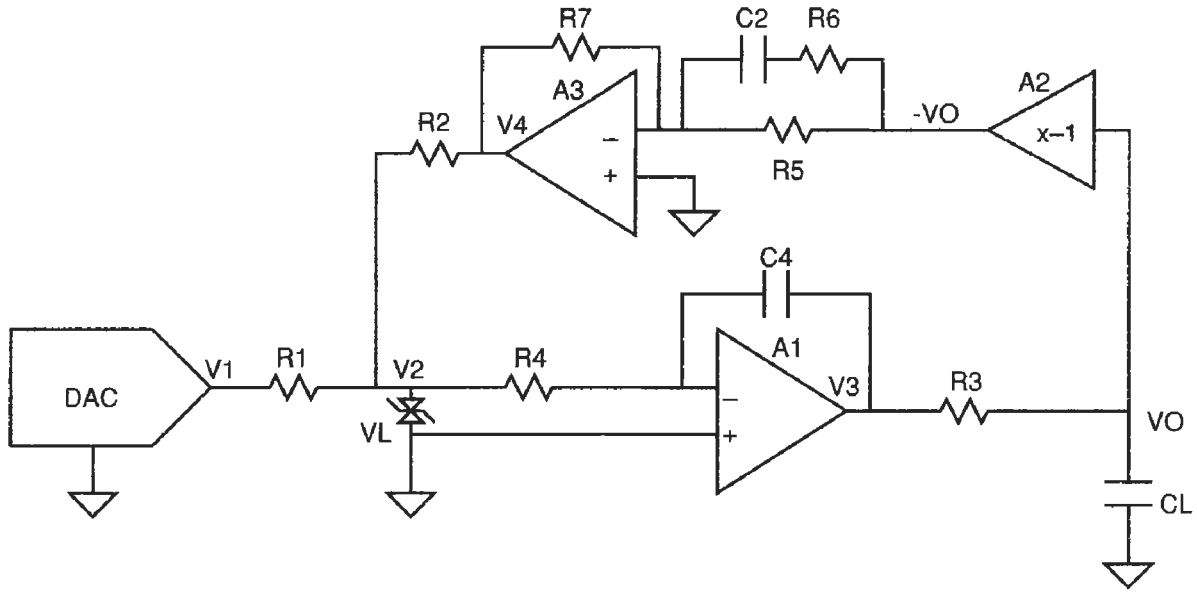

Figure 22-12. C2 moved to be in series with R6.

make Bode plots of the new model, as shown in Figure 22-13. From a Bode plot, I can estimate phase shifts according to the sketch in Figure 22-14.

My criteria for stability of this circuit are:

1. Loop UGF shall be 80 kHz or less.
   Although the Bode plots don't show it, the real circuit will have many poles above 100 kHz. If the loop UGF were higher than 80 kHz, there would be lots of phase shift added by these poles causing serious instability.
2. Phase shift at UGF is no more than 135° to 150°.
   At 135° the step response has a small amount of ringing. Much more phase shift would increase the ringing unacceptably.

Applying the stability criteria above to the drawing of compensated Bode plots, I conclude:

When CL is zero the UGF is too high.
When CL is small both the phase shift and UGF are too high.
When CL is just right, the loop will be stable.
When CL is too large the phase shift is too high (this case is not shown in drawing, but similar to uncompensated).

Is this a problem? I'm not surprised that making CL too large would cause instability. This compensation scheme has a limited range. The big problem is the excessive UGF for small CL. Overcompensation is expected to result in slower than optimum time response but should not cause oscillation!

To reduce the UGF when the loop is compensated, but CL is small or zero, we need to reduce the loop gain. This gain reduction has to be done in the forward part of the loop at A1. Changing the gain in the feedback patch would change the closed loop gain which sets the full-scale range. Increasing R4 is the best way to reduce the loop gain. Changing C4 would also change loop gain, but using solid state switches to switch capacitors is likely produce side effects. Increasing R4 to reduce the loop gain by a factor of 4 shifts the Bode plot down, decreasing the UGF to an acceptable 80 kHz.

By reducing the gain of the compensated loop we have achieved stability when CL = 0. However, the settling time is still seriously affected by the compensation

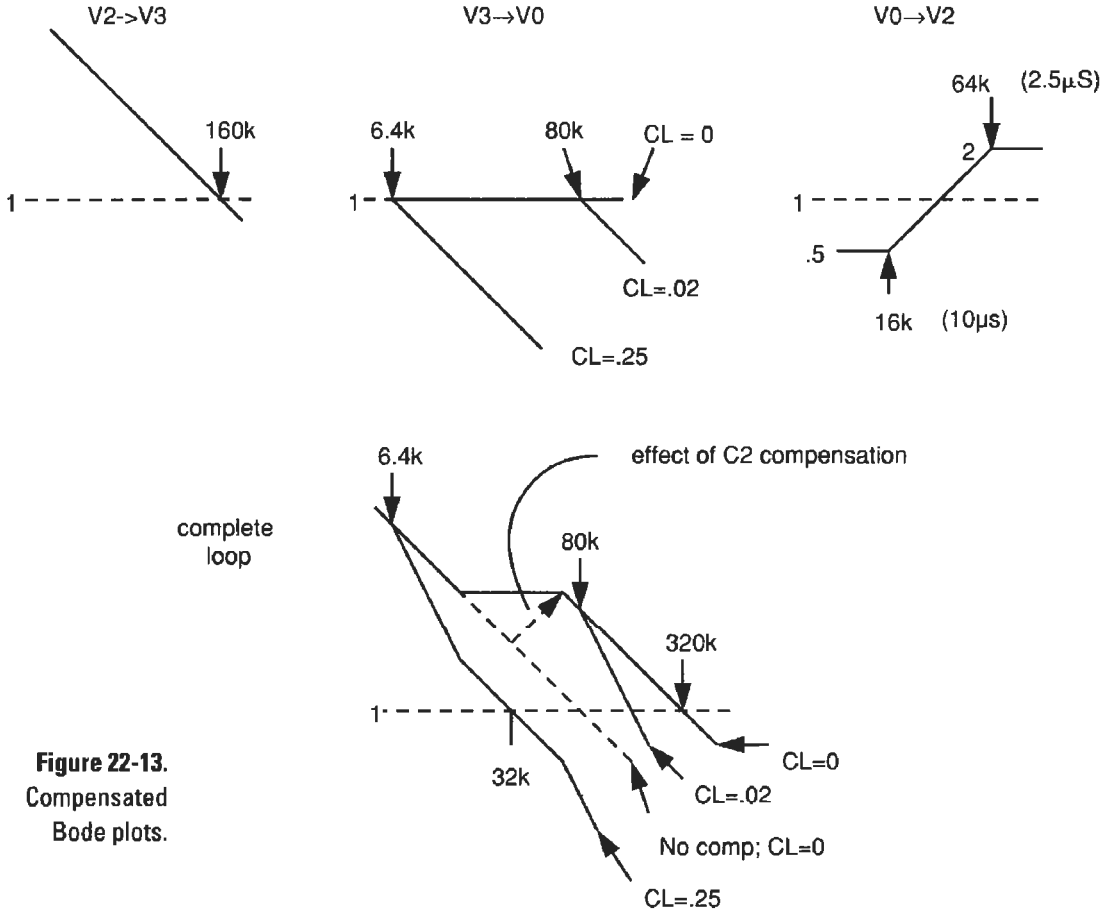

Figure 22-13. Compensated Bode plots.

because there is now a pole well below the loop UGF. The compensation should be programmed or switched to an appropriate value according to the load capacitance. In Figure 22-15 I have drawn a circuit to show switching one value of compensation.

S1 adds the anticipation or frequency compensation and S2 reduces the loop gain. To achieve both stability and the elimination of overshoot, these functions must be located in the different parts of the loop as shown.

Range of Compensation Required

The basic stability problem of capacitive loading is caused by the time constant of R3 and CL. R3 is expected to range from 1 Ω to 1 kΩ. CL could reasonably range

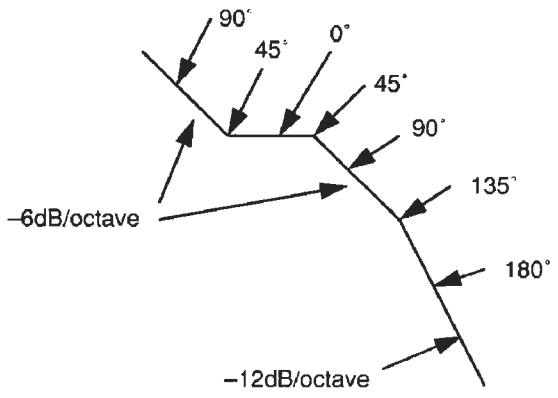

Figure 22-14. Bode plot phase shift estimates.

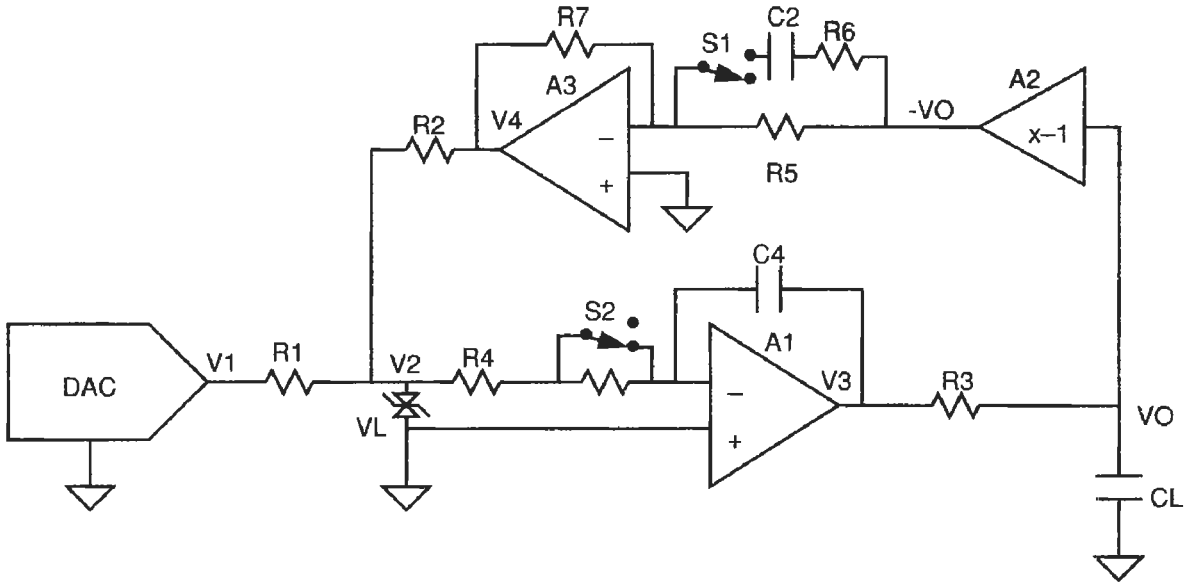

Figure 22-15. Switchable compensation.

from 1 nF to 1000 µF. That is a time constant range of 1 nsec to 1 sec. Time constants up to about 1 µsec will have little effect since the loop time constant is 2 µsec. Time constants above 100 msec are probably too long to be useful in an ATE system. Therefore, we have a time constant range of 100,000:1 to compensate for. The circuit above has only a narrow range of good compensation. There would have to be too many programmable choices of compensation networks to cover this wide range. Wouldn't it be neat if there were a way to cover a wider range with one compensation network?

Phase Margin Approach to Loop Compensation

We have been looking at the frequency domain. Now let's consider the capacitive load problem from a phase margin point of view. If the loop frequency response is a single pole, then the loop phase shift is 90° at UGF. A capacitor load adds, depending on frequency, from 0° to 90° additional phase shift to the loop. Phase shift approaching 180° is a problem. What if the no-load loop frequency response were –3 dB/octave instead of –6 dB/octave? That would give the loop a 45° phase shift at no load. Adding a capacitive load would increase the phase shift from 45° to a maximum of 135°. Anywhere in that phase range, stability would be acceptable. This

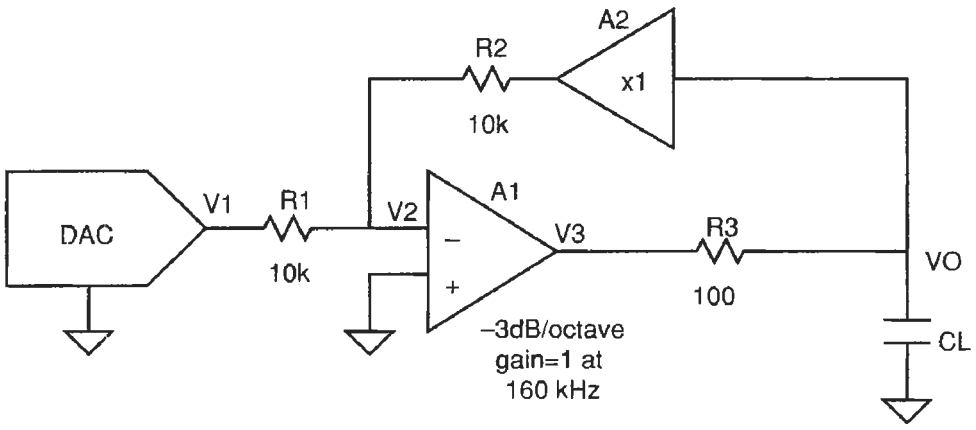

Figure 22-16. –3 dB/octave loop gain.

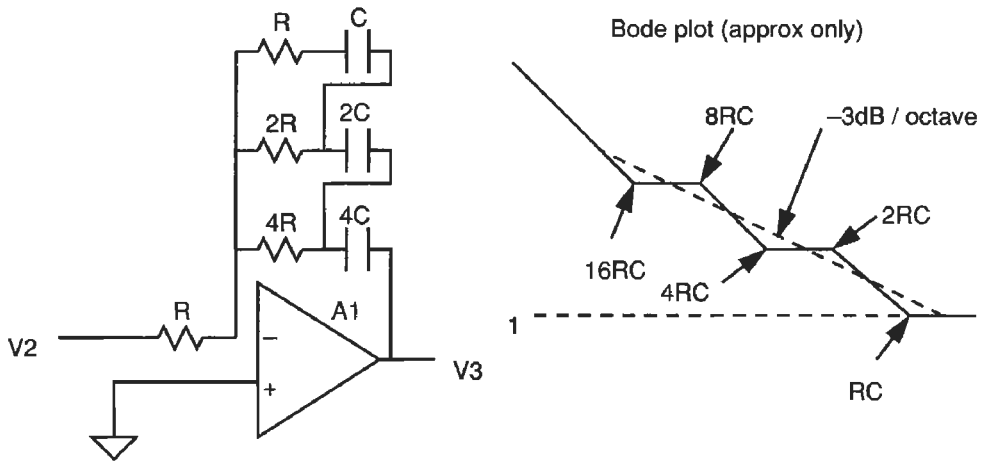

Figure 22-17. –3 dB/octave amplifier.

idea sounds like it would work for any value of capacitive load without any switching.

How would we get a loop gain of –3 dB/octave? One way would be to give the A1 stage –3 dB/octave response and the feedback path a flat response, as shown in Figure 22-16.

Replacing C in the inverting op-amp model with a series-parallel RC network comes close to –3 dB/octave and 45° phase shift over a range of frequencies, as shown in Figure 22-17.

This approach to a –3 dB/octave loop does not fix the overshoot problem since it does not allow for the anticipation network we need in the feedback path. A better approach to –3 dB/octave loop response is to leave the forward gain (A1) at –6 dB/octave and make the anticipator circuit +3 dB/octave. Adding more series RC's to R6 C2 will give the anticipator the desired +3 dB/octave over a range of frequencies. The more RC's, the wider the range. However, each RC adds a pole below the loop UGF, and these poles increase settling time. The compensation network does cover a wide range, but it still needs to be switchable to minimize settling time.

At this point, analysis by inspection became less reliable for me, so I used SPICE to simulate the multiple RC compensation network. I ran a SPICE analysis of a triple RC network. Phase analysis is easy with simulation but hard to do by inspection or on the bench. I found it tough to get close to a 45° phase lead out of a small number of components. I decided to shoot for 30° phase lead. SPICE showed me I could get 30° from a double RC network over a 50:1 frequency range. That covers a sufficiently wide range of capacitive load without making settling time too bad.

LTX Device Power Source (DPS) Performance

The LTX DPS turned out pretty good. The LTX Cadence programming language includes a DPS statement to specify the expected load capacitance and a choice of modes: minimum risetime, minimum settling time, or minimum overshoot. The operating system takes into consideration the voltage range and the current measure range, then selects one of four compensation networks (or no compensation) and the corresponding loop gain settings. The three modes are just different degrees of compensation. To minimize overshoot takes the greatest compensation and results

in the longest settling time. The compensation works well for all practical values of load capacitance. Overcompensation (specifying a larger capacitance than the actual) makes the settling time longer but causes no stability problem.

The DPS contains other functions, more stages, and many more details than I have outlined here. Each of the analog functions I designed using the same method. All the stages can be grouped together into blocks which match very closely the simplified blocks I based the design upon. The DPS behaves very much like the block diagram and simplified models would predict.

Summary of My Method

By way of an example, I have shown the method I like to use for analog design, especially for feedback loops. Here's an outline:

(Simplify!)

1. Draw a "front panel" of the instrument to be designed. "Try out" its functions.
2. Make a simple circuit model for one function or aspect of the instrument. The model should emphasize that one aspect and deemphasize other aspects.
3. Make simplifying assumptions and analyze the circuit by inspection where possible. Go back and forth between time domain and frequency domain analysis. Check your assumptions.
4. Change the model and analyze again until the results are acceptable.
5. Repeat steps 1–3 for other aspects of the instrument.
6. Design the full circuit with circuit blocks that behave like the ideal blocks in the models.
7. Test a prototype of the instrument to see if it behaves like the models.

Simple, isn't it?