CIRCUIT CELLAR INK®
THE COMPUTER APPLICATIONS JOURNAL

Practical Analog Design

FEATURE ARTICLE

Bob Perrin

Time for a refresher on op-amps? If so, here's the article for you. Bob covers op-amps from the basics of common op-amp configurations and modeling to the more advanced aspects of instrumentation amplifiers and complex circuit designs.

Analog design is fraught with snares and pitfalls. In all the years I've been designing mixed-signal instrumentation, I've made my share of mistakes, but I've learned a bit, too. I want to share some practical circuit designs to make your experiences easier.

As you know, many texts cover op-amp modeling in detail. Why? Because op-amps aren't simple devices.

A typical SPICE model for a transistor has about 30 parameters to describe the transistor's behavior. An op-amp has 30 or more transistors. Imagine the interaction of all of the various transistor parameters.

OP-AMP BASICS

While real-life op-amps are complex assemblies, we can still get a lot of mileage out of simple models. I'll give you a refresher course on common methods for predicting the behavior of op-amp circuits.

To begin, an op-amp is a high-gain differential amplifier. The voltage difference between the noninverting (+) and inverting (–) inputs is amplified by the open-loop gain, A_o.

The model shown in Figure 1 predicts op-amp behavior reasonably well for most circuits. There are of course many second-order effects this model doesn't taken into account.

A_o is the single most important parameter influencing op-amp behavior. The open-loop gain is generally 10^6 or 10^7 V/V at DC. In the most common op-amps (i.e., those that are internally compensated), A_o typically begins a 20-dB-per-decade rolloff starting at a few hertz. Figure 2 shows A_o as a function of frequency.

Internally compensated op-amps have a dominant pole at f_o and trade open-loop gain for stability. These are the easiest op-amps to design with because they're stable over the widest range of

closed-loop gains and load impedances.

Noncompensated devices lack a dominant pole at f_o, enabling the amplifier to perform well at higher frequencies. However, the designer must ensure the system remains stable. Datasheets for noncompensated devices indicate the range of stable closed-loop gains.

Both compensated and noncompensated op-amps oscillate at some point when driving some types of reactive loads. A slightly capacitive load is usually the cause.

Noncompensated devices are most sensitive to load impedance. Unless you need good high-frequency characteristics, using compensated devices makes life much simpler.

When you're selecting an op-amp, especially for a high-speed application, it's imperative to know the behavior of A_o. Manufacturers talk about the gain bandwidth product (GBP), which is the product of A_o at unity gain (i.e., 0 dB or 1 V/V) and frequency f_t. The unit for the GBP is Hz × V/V or hertz (usually in the megahertz or gigahertz range).

In the good old days before marketers muddied the waters, you'd hear of GBP and think of the behavior shown in Figure 2. By dividing GBP by the intended closed-loop gain, you got a good idea of the bandwidth available for the application. But, not anymore.

The GBP in Figure 2 is constant anywhere along the slope of A_o. This feature is what makes GBP a useful parameter for figuring out the usable bandwidth. Not all devices have a constant GBP.

Current-feedback amplifiers (CFAs) have open-loop gain characteristics that depend on the closed-loop gain. Instrumentation amplifiers (IAs) have GBP characteristics that depend on their internal topology and the circuit gain.

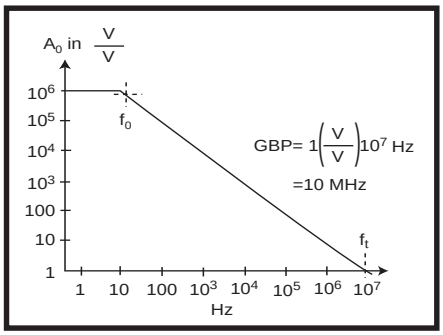

Figure 2—*In internally compensated op-amps, the open-loop gain has a dominant pole at f_o.*

GBPs for CFAs and IAs are horses of entirely different colors. When you select an op-amp for a high-frequency application, look at all device parameters, not just GBP.

For example, consider Burr-Brown's OPA643 op-amp. The datasheet for this useful, economical wide-bandwidth device boasts a GBP of 1.5 GHz, which means, qualitatively, it should be useful at relatively high frequencies.

The device is a voltage-feedback amplifier (versus a CFA). However, when you look at a graph of A_o as a function of frequency, the open-loop gain isn't as well-behaved as that in Figure 2.

The datasheet also indicates the OPA643 is a noncompensated device and only stable for closed-loop gains greater than 5 V/V. It can serve as a 10-dB video amp but is unsuitable as a unity-gain follower in a high-frequency measurement system.

In addition to the GBP, you need to consider the slew rate in high-frequency performance, which measures how fast an op-amp can swing its output voltage. The op-amp must be able to swing the output voltage as fast as the maximum derivative of the highest frequency component of the wave form. Slew rate is measured in V/µs.

Consider an amplifier designed to provide 10 V/V gain for input signals up to 1 Vp-p and 1 kHz. Maximum slew rate would occur on an output waveform of (1 Vp-p × 10 V/V) 10 Vp-p at 1 kHz, right at the zero crossing. So,

$$\left(\frac{d}{dt}\right)10\sin\left(2\pi 1000t\right) = 10 \times 2\pi 1000\cos\left(2\pi 1000t\right)$$

evaluated at t = 0, π = 3, results in 60,000 V/s or 0.06 V/µs. So, your chosen op-amp for this example circuit must have a slew rate of at least 0.06 V/µs.

Many common op-amps (e.g., LM741, LM324, OP497) have slew rates around 0.5 V/µs. Any of these would suffice here. By contrast, the OPA643 has a respectable slew rate of 1000 V/µs.

The output resistance of the op-amp, R_o, is the final parameter I want to discuss with respect to the model shown in Figure 1. R_o is often mistaken as the output resistance of the overall amplifier circuit, R_{out}.

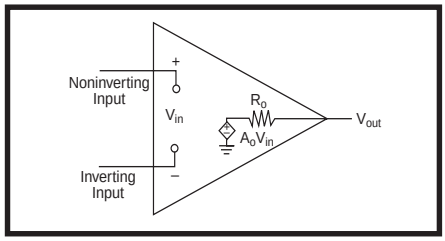

Figure 1—*The first-order op-amp model is simple yet remarkably useful.*

In typical amplifier configurations, the negative feedback desensitizes the circuit to R_o. Therefore, the output resistance R_{out} is not the same as the op-amp's R_o.

As seen in Figure 3, the feedback network attaches to the output pin of the op-amp after R_o. Negative feedback drives the output pin to the required level. The circuit's overall output resistance is:

$$R_{out} = \frac{R_o}{1 + A_o b} \qquad [1]$$

The feedback factor, *b*, reflects the attenuation of the output voltage before it is fed back into the inverting input. For the inverting and noninverting configurations of Figure 3:

$$b = \frac{R_1}{R_1 + R_2}$$

Inverting and noninverting amplifiers have an extremely low output resistance. However, R_{out} increases as the frequency increases and A_o rolls off, which is a result of A_o in the denominator of equation 1.

R_o limits the maximum current the op-amp can source or sink. Datasheets most often give a maximum value for short-circuit current. This parameter can be used to back calculate R_o for modeling. Simply take the maximum output voltage and divide it by the maximum source current to obtain a value for R_o.

CIRCUIT-LEVEL ANALYSIS

Figure 3 shows schematics and equations for the two most common op-amp configurations—the inverting and noninverting amplifiers. To analyze these, sum the currents at the inverting node.

To simplify the analysis, assume R_o is zero. You greatly simplify the

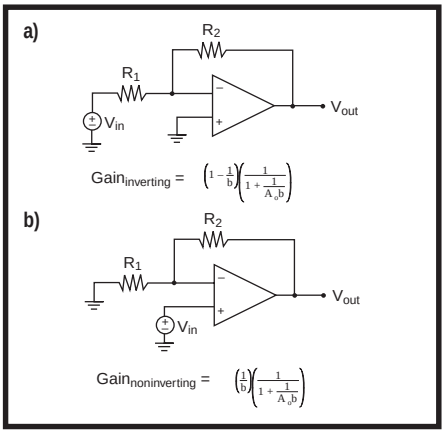

Figure 3—*Op-amps can be configured as inverting (**a**) or noninverting (**b**) amplifiers.*

analysis, yet still obtain useful results. Other nonideal characteristics can be analyzed separately and, through the use of superposition, can be integrated into an overall circuit model.

Figure 3 shows the equations deduced using the above technique. To further simplify, let A_o approach infinity to obtain:

$$\text{Gain}_{\text{inverting}} = \frac{-R_2}{R_1}$$

$$\text{Gain}_{\text{noninverting}} = 1 + \frac{R_2}{R_1}$$

Keep in mind that these two gain equations were derived with much simplification. They don't reflect the effects of the open-loop gain rolloff. However, these are the most commonly used equations for describing the behavior of the circuits in Figure 3.

When the frequency response of A_o is relevant (e.g., when high closed-loop gains are needed), use the equations in Figure 3. I recommend simulation tools like SPICE or MicroCap when you require highly accurate predictions.

SPICE models take into account second-order effects like bias currents, offset currents, and offset voltage. You can find detailed op-amp models for virtually every commercially available op-amp. Most manufacturers supply the models free of charge on the Internet or on CD.

A common concept used to analyze op-amp circuits is the idea of a "virtual ground" or "virtual short," which mean the same thing. The voltage on the op-amp's inverting node is forced by the negative feedback to be at the same potential as the noninverting node.

This concept assumes that A_o approaches infinity, therefore requiring the difference between the inverting and noninverting inputs to approach zero for V_{out} to be finite.

As with all approximations, there are times when this one is useful and times when it's not. Figure 2 shows A_o is generally pretty large at low frequencies. So, the virtual-short idea is most useful at low frequencies.

In the case of the inverting amplifier, the term "virtual ground" is applied to the inverting input of the op-amp. The noninverting node is held at ground. Negative feedback forces the inverting node to be at the same potential as the noninverting node (i.e., ground). And, the inverting input of the amplifier is said to be a virtual ground.

The phrase "virtual short" is used for the noninverting amplifier. The input signal is applied at the noninverting node, and the inverting input is forced to this same potential. Since the two nodes are held at the same potential by the negative feedback, and that potential is not necessarily zero, the inputs of the op-amps are said to have a virtual short across them.

SECOND-ORDER EFFECTS

Now that we have examined op-amps at the circuit level, let's go on to modeling and predicting behavior due to second-order effects. Developing an intuitive feel for these effects enables you to select devices best suited to your application.

Figure 4a shows how to model op-amp input bias and offset currents. I_{bias} is the tiny base current (or gate leakage) the internal transistors need for operation. This value is orders of magnitude smaller for FET-based input stages than for BJT input stages [1].

I_{offset} results from the mismatch between the input transistors. I_{offset} is typically one order of magnitude smaller than I_{bias} [1].

The polarity of I_{bias} is constant and predictable if the topology of the input stage is known. However, I_{offset} may be of either polarity.

To examine the effect of these currents, remove the stimulus from the circuit and look at the V_o generated by the offset currents. With the inputs grounded, the inverting and noninverting configurations become the same circuit. Figure 4b gives you a glimpse of the circuit, assuming that $R_{compensation} = 0$.

This configuration offers no hope for nulling offsets introduced by bias currents, which is why $R_{compensation}$ has been added. Analyzing the circuit with $R_{compensation} \neq 0$ yields:

$$V_o = \left(1 + \frac{R_2}{R_1}\right) \times \left[\left(R_1 \,||\, R_2\right) I_- - R_{compensation} \times I_+\right]$$

where $I_+ = I_{bias} + I_{offset}$ and $I_- = I_{bias} - I_{offset}$.

To compensate for the bias currents, simply set $R_{compensation}$ equal to the equivalent resistance of R_1 in parallel with R_2. However, this doesn't compensate for I_{offset}. The overall error reduces to:

$$V_o = \left(1 + \frac{R_2}{R_1}\right)\left[-2\left(R_1 \,||\, R_2\right) I_{offset}\right]$$

Since I_{offset} is typically an order of magnitude less than I_{bias}, adding $R_{compensation}$ (equal to R1||R2) is usually sufficient.

Other methods to reduce error due to I_{offset} are to:

- keep R1||R2 small

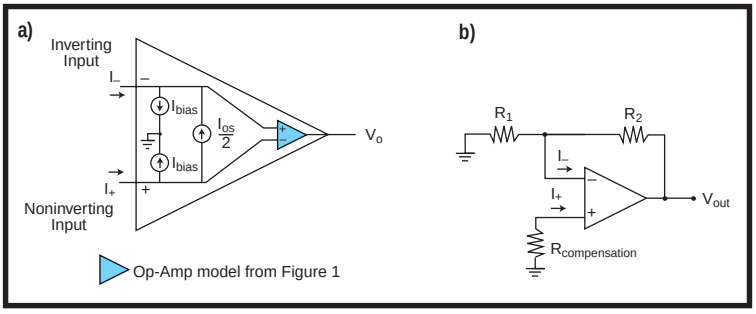

Figure 4a—*Input bias and input offset currents are modeled as separate independent sources.* ***b***—*$R_{compensation}$ provides a means of nulling the offset in V_{out} caused by I_{bias}.*

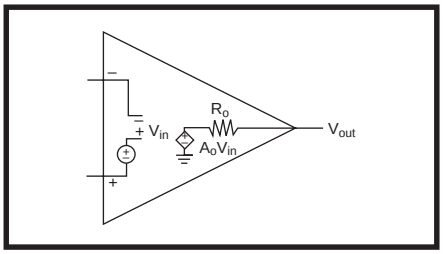

Figure 5—*Input voltage offset is modeled by the inclusion of an independent voltage source in the simple model shown in Figure 1.*

• keep R2 small

Both increase power dissipation. The tradeoff may seem simple, but a quick look at a graph of I_{bias} versus dice temperature reveals the insidious truth. Many op-amps, especially those with FET input stages, increase I_{bias} exponentially as a function of temperature.

Devices with superbeta input stages, like the OP297 and OP497 from Analog Devices, offer a good compromise between I_{bias} and the temperature coefficient. Superbeta input stages have moderately low initial I_{bias} currents and do not suffer from an exponential temperature dependence.

Increased power means increased dice temperature. Evaluate the tradeoffs carefully.

I_{offset} is related to input current noise. Intrinsic noise on the dice creates a small amount of fluctuation in I_{offset}. This current noise will develop into a voltage across R1||R2 and be multiplied by:

$$A_{noise} = \left(1 + \frac{R_2}{R_1}\right)$$

The noise gain (A_{noise}) is the same for both inverting and noninverting configurations.

The op-amp input current noise is typically very tiny—subpicoamps. However, the voltage noise at the output can become significant as current noise is developed into a voltage across 2× R1||R2 and multiplied by A_{noise}.

Op-amp noise is frequency dependent. Datasheets always list tabulated data, and better datasheets have graphs.

Another source of error in op-amps is the offset voltage (V_{os}). The cause of V_{os} is the unavoidable mismatch of devices and operating points on the dice. Figure 5 shows the model for offset voltage.

V_{os} is highly temperature dependent. The polarity is unpredictable and may reverse over temperature.

Using the model in Figure 5 with the circuit topology from Figure 4b (with $R_{compensation} = 0$), I derive:

$$V_o = \left(1 + \frac{R_2}{R_1}\right)V_{os} = A_{noise} \times V_{os} \quad [2]$$

Compensating for V_{os} can be accomplished via external nulling (i.e., introducing an external voltage at one of the op-amp inputs) or internal nulling (i.e., unbalancing one of the internal differential pairs).

External nulling can be accomplished by using a trimmer, digitally controlled pot, or DAC to dial in the required compensation voltage. Internal nulling is accomplished by connecting a trimmer to the offset pin or pins on the op-amp. The device datasheet has trimmer-value recommendations and suggested configurations.

A common problem occurs when multiple op-amps are chained in a circuit. Each op-amp contributes an offset to the system. The most common question asked is, "Do I null each op-amp's V_{os} or null the overall system at a single point?"

I lean toward nulling each op-amp using internal nulling. I definitely shy away from nulling the entire system offset using the internal nulling on a single stage. The latter is much simpler but unreliable over temperature.

If you use internal nulling to introduce a huge imbalance in a single stage to compensate for the total system offset, you are assured of temperature-tracking problems. The grossly unbalanced op-amp can't track the overall system offset.

If you null each stage internally, the system offset remains as stable as practically possible over temperature.

With the DC offset nulled, let's turn our attention to noise on V_{os}. All V_{os} noise, regardless of origin, shows up in the amplifier output multiplied by A_{noise} (see equation 2).

Intrinsic noise on the dice causes perturbations in V_{os}, which is referred to as input voltage noise (E_{os}). Datasheets give tabulated data for E_{os}.

The power supply induces another V_{os} noise. As the power-supply voltage fluctuates, the internal bias points of the op-amp shift. Because V_{os} is the result of internal bias-point mismatches, power-supply–induced noise is modeled as noise on V_{os}.

The power-supply rejection ratio (PSRR) is:

$$PSRR = \left(\frac{\Delta V_{os}}{\Delta V_{supply}}\right)$$

Another culprit introducing offset error is the common-mode voltage on the op-amp's inputs. As the voltage on the inputs is raised or lowered, the op-amp's internal bias points shift and additional V_{os} develops. Thus, common-mode noise can be modeled as noise on V_{os}.

The common-mode rejection ratio (CMRR) is:

$$CMRR = \left(\frac{\Delta V_{os}}{\Delta V_{common\text{-}mode}}\right)$$

Both PSRR and CMRR generally roll off with frequency. Some devices—especially chopper-stabilized amps—have a notch in the rolloff. Always refer to

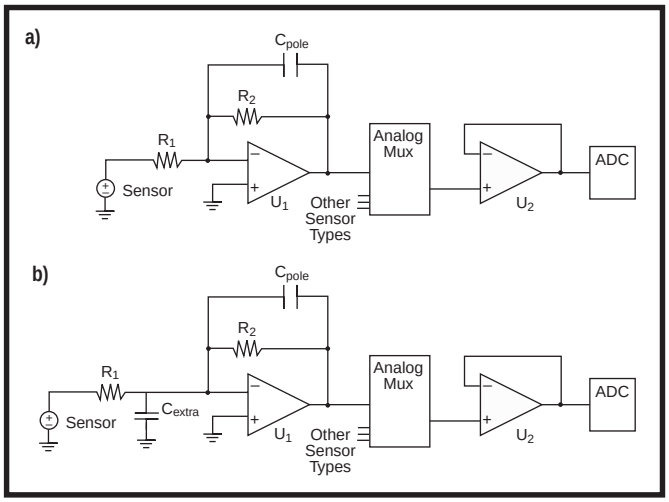

Figure 6a—*R_2 and C_{pole} form a single-pole filter.* ***b—****A virtual ground exists at the inverting input of U1, so C_{extra} has a virtual short across it.*

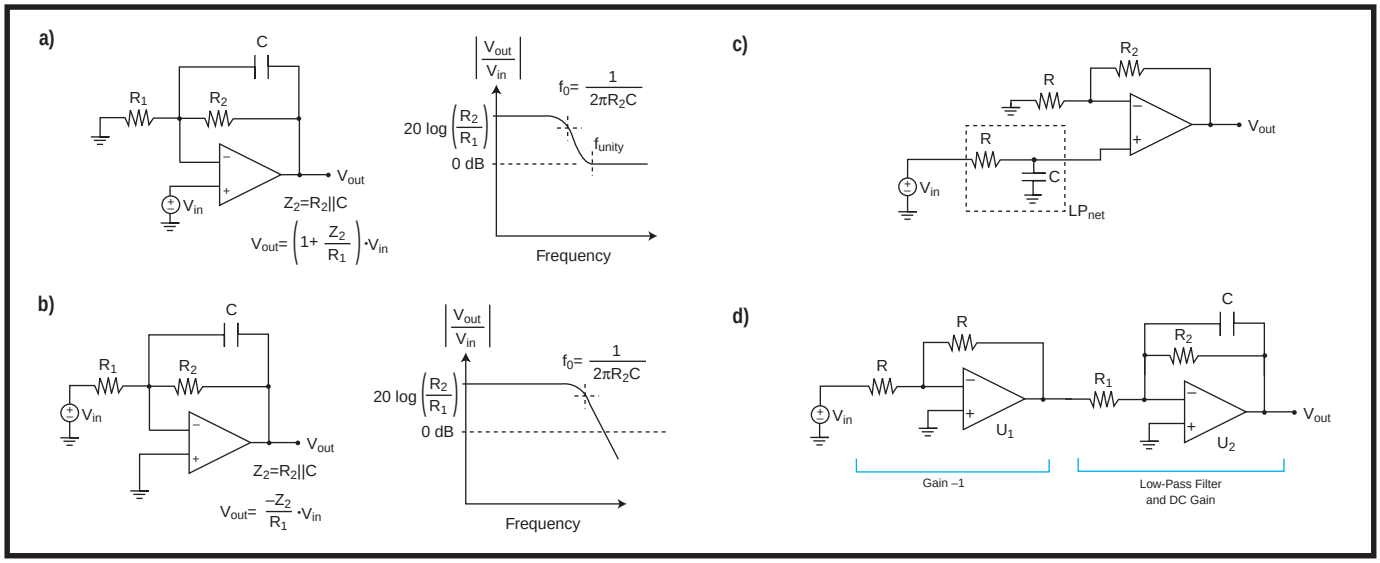

Figure 7a—*The noninverting configuration has an intrinsic zero and does not allow attenuation below 0 dB.* ***b***—*The inverting configuration permits attenuation below 0 dB.* ***c***—*LPnet continues to attenuate noise from V_{in} independently of the op-amp gain stage.* ***d***—*A two-stage antialias filter can provide high gain, low noise, low R_{out}, and excellent frequency rolloff.*

graphs of these parameters, and don't forget that CMRR and PSRR are referenced to V_{os}.

It's incorrect to reason that if you have *x* V_{RMS} on your rail and the PSRR is –60 dB, your output will only see *x*/1000 V of noise. Wrong!

Instead, you need to think, "If I have *x* V_{RMS} of noise on my rail, my noise gain is 40 dB, and my PSRR is –60 dB, then my output will see *x*/10 V of noise."

Well, the basics review is over. Now, let's look at some real-life examples.

WATCH OUT FOR VIRTUAL SHORTS

A number of years ago, I was developing a hand-held instrument with optical sensors on a meter-long probe. The signal paths to the sensor formed long loop antennae. The functional geometry for the instrument required this suboptimal electrical configuration.

In my signal path, I had the single-pole low-pass filter/inverting amplifier shown in Figure 6a. I needed an additional pole to provide sufficient band limiting of noise.

I added a capacitor as shown in Figure 6b. I expected R_2 and C_{pole} to form the first pole, and R_1 and C_{extra} to form the second pole. On testing the circuit, I found it behaved as only a single-pole (R_2 and C_{pole}) circuit.

After much wailing and gnashing of teeth (and a little algebra), I realized that C_{extra} is attached between a virtual ground and ground. C_{extra} had a virtual short across it and was therefore superfluous. I added my second pole with a passive RC on the output and obtained satisfactory results.

As it turns out, C_{extra} isn't entirely superfluous. The virtual-ground concept is predicated on the assumption that $A_o >> A_{closed\ loop}$. As A_o rolls off and the virtual ground degrades, C_{extra} enters into the equation. The best way to analyze these kinds of second-order effects is with SPICE.

ANTIALIASING FILTER TRADEOFFS

The noninverting amplifier has an interesting feature—an intrinsic zero. By virtue of the circuit topology, there is a zero that prohibits the closed-loop gain from dropping below 1 V/V (0 dB).

The easiest way to see this is to review at the transfer function for a noninverting amplifier:

$$\frac{V_{out}}{V_{in}} = 1 + \frac{Z_2}{Z_1}$$

Regardless of what Z_2 and Z_1 do (short of going negative), the gain is always larger than unity. Keeping this in mind, let's compare the inverting and noninverting topologies for use as antialiasing filters.

Figure 7 shows four circuits used as antialiasing filters. Figure 7a allows all high-frequency noise beyond f_{unity} to be aliased back into the pass band without attenuation. The circuit in Figure 7b has inverted gain but is vastly superior to the circuit in Figure 7a at attenuating high-frequency noise.

The circuit in Figure 7c uses the op-amp only as a gain block. A passive network (LPnet) provides the antialias function. This configuration doesn't attenuate the noise introduced into the system by the op-amp, but other noise is band limited by LPnet. The circuit in Figure 7b is again superior to Figure 7c because V_{os} noise is attenuated.

The weakness in the circuit in Figure 7c is the LPnet preceding the gain block. You might argue that LPnet could follow the gain block for improved noise attenuation.

But this newly proposed configuration doesn't present a low impedance to the ADC. The ADC may load LPnet, thus introducing more error.

For example, the ADC's sample and hold (S/H) may gulp current out of

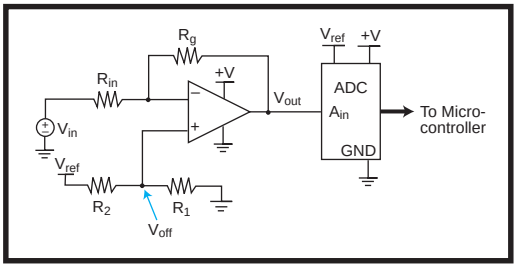

Figure 8—*This versatile interface maps an arbitrary input range into 0–V_{ref}.*

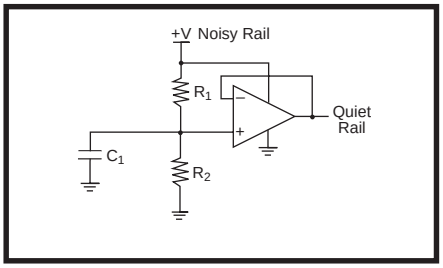

Figure 9—*A spare op-amp can be used as a low-cost voltage regulator.*

LPnet, and enter hold mode before LPnet can deliver sufficient charge to bring the voltage sampled up to the proper value. A good rule of thumb is to have an op-amp driving the input of the ADC's S/H.

If the V_{os} noise (intrinsic, power-supply induced, and common mode) in the circuit shown in Figure 7c multiplied by A_{noise} is suitably small for the application, then this circuit is an acceptable solution.

If signal inversion is important and high gain is required, you can precede the circuit of Figure 7b with a unity-gain inverting stage (see Figure 7d). The new stage introduces a small amount of noise. A_{noise} for a unity-gain inverting stage is only 2 V/V.

The filter portion (U2 in Figure 7d) band-limits the newly introduced noise. As well, another pole can be added to the filter by placing a capacitor in the feedback loop of the unity-gain inverting stage.

When you're designing filters around a noninverting amplifier, always keep in mind the zero that exists as an artifact of the topology.

SINGLE-ENDED ANALOG STAGE

Figure 8 illustrates a versatile single-ended analog front end. The op-amp and ADC can be powered by a single rail supply while accepting bipolar or unipolar inputs. Tiny signals can be expanded to fill 0–V_{ref}. Large input signals can be compressed. The input signal need not be center around zero.

For example a 4–20-mV range can be mapped into 0 – V_{ref}. If you select the resistor values and V_{ref} properly, this circuit serves almost any application.

Input resistance versus gain is the major tradeoff to consider with this circuit. R_{in} must be large enough so the circuit doesn't load the sensor, but small enough to allow the ratio R_g/R_{in} to supply sufficient gain.

Designing with this circuit is simple:

- determine the sensor's output voltage range
- determine the ADC's input range
- determine the required gain (or attenuation). Deduce ideal values for R_g and R_{in}.
- determine required V_{off}. Deduce ideal values for R_1 and R_2.
- select standard resistor values
- verify input range maps to output range with the selected resistor values
- repeat the last two steps until satisfied

Let's walk through the procedure. The governing equation is:

$$V_{out} = \left(V_{off} - V_{in}\right)\frac{R_g}{R_{in}} + V_{off} \qquad [3]$$

The first step is to figure out what the application gain needs to be:

$$|g| = \frac{V_{ref}}{V_{inmax} - V_{inmin}}$$

From equation 3, you see that if V_{off} is 0, the circuit gain is:

$$g = -\frac{R_g}{R_i} \qquad [4]$$

Now, select R_g and R_{in}. R_{in} should be as high as possible to avoid loading the sensor.

Next, determine V_{off}. Note that the circuit is an inverting amplifier. When V_{inmax} is applied, the desired V_{out} is zero. Thus from equation 3, you obtain:

$$0 = \left(V_{off} - V_{inmax}\right)\frac{R_g}{R_{in}} + V_{off} \qquad [5]$$

Using equation 4 and rearranging equation 5, you can determine V_{off}:

$$V_{off} = V_{inmax}\left(\frac{|g|}{1+|g|}\right)$$

Now, pick R_1 and R_2. R_1 and R_2 form a voltage divider that dials in V_{off}. You can select R1 and R2 by using:

$$V_{off} = \frac{R_1}{R_1 + R_2} \times V_{ref}$$

Finally, pick standard resistor values and use equation 3 to verify that the design maps the V_{sensor} range in to the V_{ADC} range.

If a pole is needed, you can add one by placing a capacitor in the feedback loop.

The circuit in Figure 8 is shown as an input interface. However, the same topology can be used to condition the output of a DAC. A bipolar output is possible if a negative rail is available.

You can use this circuit to map any arbitrary input range to any arbitrary output range. This is restricted only by practical limits of available components.

COST-EFFECTIVE LOW-NOISE RAIL

To supply clean power to low-noise analog circuits, you can add a linear regulator just for the analog section. Modern micropower amplifiers don't need huge amounts of current. Regulators are bulky, and dropout voltage is always a concern. In many cases, a dedicated regulator is overkill.

In some cases, the nifty circuit in Figure 9 may be an almost free solution. It uses the following concepts to provide a low-noise power rail [2]:

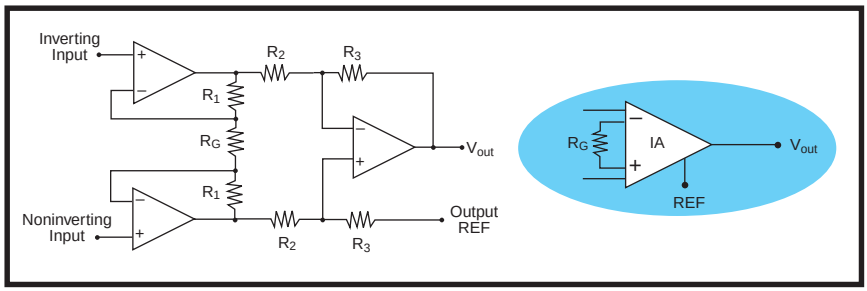

Figure 10—*Today, $3 can buy a single monolithic chip with a full-blown instrumentation amplifier.*

- PSRR characteristics of the op-amp
- R_o desensitivity
- the high input impedance of the noninverting node
- low A_{noise}

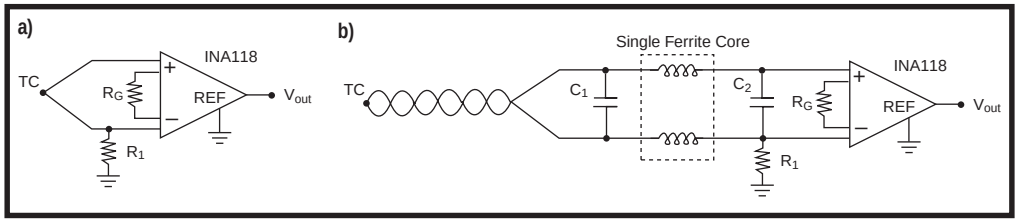

Figure 11a—*A low-cost instrumentation amplifier can provide a nearly ideal thermocouple interface.* ***b—****A common-mode filter and twisted-pair sensor cable clean up induced common-mode noise.*

The op-amps are capable of sourcing several milliamperes, which is often enough current to drive many micropower analog stages. An op-amp in a SOT-23/5 package and a couple 0402 passives may occupy less board space than a bulky linear regulator.

If the circuit already has an extra op-amp (e.g., in an existing dual or quad package) powered from a noisy rail, the two resistors and capacitor are all you need for a low-noise power rail for your quiet analog section.

INSTRUMENTATION AMPLIFIERS

Instrumentation amplifiers (IAs) are differential. They have a schematic symbol similar to op-amps, but are not op-amps. Figure 10 shows its symbol and a three op-amp implementation of it. IAs can be built from discrete components or purchased as single chips.

The voltage between the inverting and noninverting inputs is amplified by the IA's gain. This gain is usually set by a single resistor, R_g. Single-chip IAs are often capable of 1–1000 V/V gains.

The IA's two input op-amps are configured as noninverting amplifiers. They provide the IA's high-impedance input characteristics that it's known for. The output op-amp is configured as a differential amp. The overall IA gain is:

$$IA_{gain} = \left[1 + 2\left(\frac{R_1}{R_g}\right)\right]\left(\frac{R_3}{R_2}\right)$$

Although other implementations of IAs exist [1], the one in Figure 10 is the most widely known.

The advent of the single-chip monolithic IA is a boon to instrumentation engineers. In the not-so-good old days, a circuit designer had to build an instrumentation amp from discrete op-amps and resistors. It looked pretty good on paper, but in practice, matching the discrete resistors was difficult and performance suffered.

For an IA to achieve CMRRs on the order of 120 dB, the resistors must be precisely matched in value and temperature coefficient. The mismatches associated with 1% resistors in discrete designs were often unacceptable in high-precision applications. Also, careful consideration was needed to keep the resistor pairs isothermal.

The single-chip IAs offer superbly matched components. With all the parts on a single die, everything stays isothermal. Laser trimming of internal resistors guarantees maximum precision. Monolithic IAs save board space, reduce component count, save money, and maximize performance.

Some manufacturers have single-chip IAs with chopper-stabilized front ends. These devices are designed to minimize V_{os}. For example, the LTC1100 from Linear Technologies has an admirable V_{os} of 10 µV. This device has an internal oscillator and internal caps to support the chopper stabilization.

Figure 10 shows that both inputs of the IA have a high impedance. Notably, 10 GΩ is a fairly common equivalent input resistance for an IA. This, coupled with high CMRR and a wide range of gains, makes the IA an easy-to-use and versatile tool.

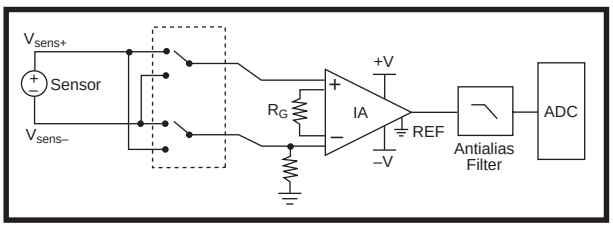

Figure 12—*Chopping the input signal enables the nulling of system offset in software.*

IA THERMOCOUPLE CONDITIONER

A thermocouple is a junction of dissimilar metals that produces a voltage proportional to temperature. Thermocouples typically have an absolute accuracy of only a degree or two Celsius. However, thermocouple measurements are extraordinarily stable and repeatable.

In a calibrated system, temperature Δs of 0.001°C can be reliably measured with thermocouples. Figure 11a shows a simple thermocouple (TC) interface. R1 prevents the high-impedance inputs from developing sufficient charge to drift outside of the supply rails.

Figure 11b offers a remedy for noise induced on the sensor leads. The ferrite bead and two capacitors form a common-mode filter. The twisted pair ensures noise is picked up equally on both conductors.

Briefly, let's explore how much loading error is introduced into the measurement by the IA and common-mode filter in Figure 11b. The leakage resistance of an X7R dielectric ceramic capacitor is about 10 GΩ. The addition of the two capacitors cuts the DC input resistance seen by the thermocouple by two-thirds, putting about a 3-GΩ load on the thermocouple.

Thermocouples are often modeled as a voltage source with a series source resistance, R_s. The size and type of junction determine the value of R_s. For my needs, 10 Ω is a reasonable value for R_s.

The load resistance (3 GΩ) and R_s (10 Ω) form a voltage divider, with the tap of the divider being the actual voltage observed. The loading error is only about three parts in a billion. This is a negligible effect for thermocouple applications.

Minimal loading of the TC, high CMRR, high PSRR, and differential front end are needed to measure temperature changes on the order of millidegrees Celsius. IAs provide a simple and economical solution for applications measuring small temperature changes.

For example, consider a microcalorimeter. The maximum expected Δ is

10°C. The typical output of a type-E thermocouple is 60 µV/°C. If the system has a 2.5-V full-scale ADC, then the gain on the IA must be:

$$\frac{2.5}{10 \times 60 \times 10^{-6}} = 4166 \text{ V/V}$$

This is achievable with a single $3 eight-pin IA from Burr Brown—the INA118.

CHOP TO NULL SYSTEM OFFSET

Offset voltage is still a caveat. IAs have an offset voltage much like op-amps. For many IAs, this offset is around 100 µV. In high-gain systems like a microcalorimeter, V_{os} can be a major problem (4166×100 µV = 0.42 V of offset at the output).

Chopping the input is a technique for reducing offset voltage. Two measurements are made of different quantities. Since the offset voltage shows up in both measurements, a subtraction drops out the offset voltage.

Many chopper-stabilized ICs make a measurement of the input voltage, then internally measure a ground potential. The two measurements are subtracted and presented to the output. All of this goes on internally and is transparent to the external circuit.

For high-gain applications, selecting an IA with low V_{os} is a good start. Two low-offset nonchopper-stabilized IAs are Analog Devices' AD620B and Burr Brown's INA118 with a V_{os} of 50 µV. The INA118's V_{os} depends on the gain setting and can be as much as 300 µV at unity gain. The Linear Technology LTC1100 is an internally chopper-stabilized IA with a V_{os} of only 10 µV.

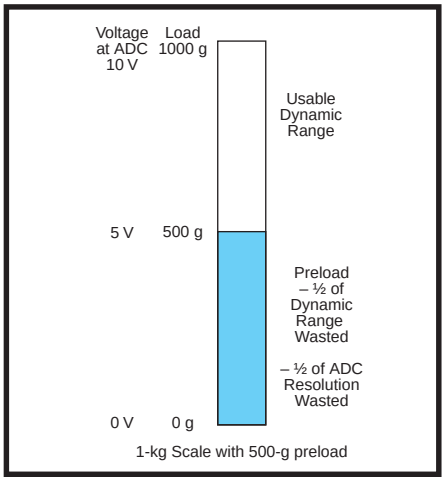

Figure 14—*ADC resolution and dynamic range are expensive. Don't trade them away in a simplistic nulling scheme.*

Chopping externally to the IA is an outstanding and time-proven technique for reducing channel offset. Figure 12 shows a signal-conditioning circuit that uses chopping to null the system offset. This technique nulls the offset associated with the entire measurement channel, not just the IA V_{os}.

The sequence of events in a measurement cycle is:

- the input mux is set such that V_{sens+} routes to the noninverting input and V_{sens-} routes to the inverting input
- a measurement is taken with the ADC: M1 = (V_{sensor} + V_{os}) Gain
- the input mux is reversed such that V_{sens+} routes to the inverting input and V_{sens-} routes to the noninverting input.
- a measurement is taken with the ADC: M2 = (–V_{sensor} + V_{os}) Gain
- compute:

$$V_{sensor} = \frac{M1 - M2}{2 \times \text{Gain}}$$

When $M_1 - M_2$ is computed, the V_{os} introduced into the measurement drops out.

For chopping to work, measurements M_1 and M_2 should be taken as closely together in time as possible. Normally, the limiting factor is the settling time of the system's antialiasing filter. Burr-Brown has an excellent app note on fast-settling low-pass filters [3].

A major advantage of the chopped input technique is the ability to null the system offset with each measurement. In real-world systems, system offset voltages drift with time and temperature. This technique enables the system to compensate for the unpredictable drift of offsets.

IA RESISTIVE BRIDGE SENSOR INTERFACE

High input impedance makes the IA ideal for measuring resistive bridge sensors. Strain gages, pressure transducers, RTDs, and load cells are common examples of resistive bridge sensors.

Bridge transducers are balanced or unbalanced. Balanced bridges use feedback to force the voltage on the differential nodes to be the same by tweaking one or more resistive elements in the bridge. Unbalanced bridge systems simply measure the voltage on the differential nodes.

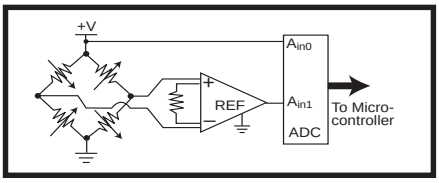

Figure 13—*Instrumentation amplifiers can measure unbalanced bridge transducers without significant loading.*

Both systems rely on precise measurement of the voltage across the differential nodes. This work demands a high-impedance differential amplifier—an IA.

Figure 13 shows a measurement system using an IA as the bridge interface. The IA amplifies the voltage across the differential nodes of the bridge. The ADC measures both the output of the IA and the excitation voltage on the top of the bridge.

From these two quantities, the controller determines the magnitude of the physical stimulus on the bridge. The exact determination of the value depends on the bridge characteristics.

Typical strain gages have bridge elements between 120 and 350 Ω [4]. The IA has an input impedance a hundred million times greater, and the loading effect is negligible.

In a strain-gage or RTD application, the full-scale differential output voltage may only be 10 mV. IA gains of 60 dB (1000 V/V) are therefore required to bring the signal up to typical working levels, but that's not a problem for IAs like the AD620.

As with high-gain thermocouple applications, V_{os} must be considered in the system design. For example, the 50-µV V_{os} of the AD620 times 60 dB yields an output-voltage offset of 50 mV.

If the full-scale output voltage is 10 V and measured with a 12-bit ADC, the offset introduced by the IA is 20.4 codes out of 4096. This error is manageable and can be removed in software.

Bridge transducers are often located on long cables, which tend to pick up common-mode noise. Multiples of 50 and 60 Hz are the most common.

The high CMRR of IAs effectively

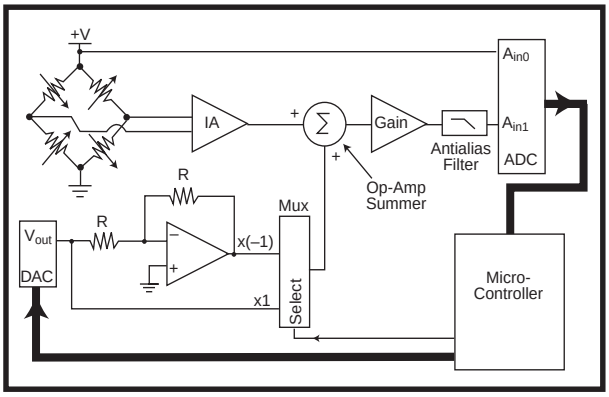

Figure 15—*Preload compensation can be accomplished without sacrificing dynamic range.*

reduces this unwanted noise. Common-mode RF can be significantly attenuated by ferrite beads (see Figure 11b).

The single-chip IA greatly simplifies the job of the analog instrumentation engineer. Measuring bridge circuits is still considered somewhat of an art. For most of us, precanned IAs reduce the problem to the paint-by-numbers level.

PRELOAD COMPENSATION

A sensor preload is a load that is undesired but necessary for the measurement. For example, a load cell in a digital scale may use a bowl to hold the product being measured. The bowl is a preload.

One way to null a preload is to sample the sensor with the preload and subtract the preload from subsequent measurements in software. At first glance, this software nulling technique seems like a good solution. But, the tradeoff for simplicity is dynamic range.

For example, consider a system with a 0–10-V ADC input corresponding to a 0–1000-g stimulus on the load cell. If a 500-g bowl is placed on the load cell, the ADC sees a 5-V preload. The remaining 5–10-V range is still available for measurement, but half the system's dynamic range is lost to the preload (see Figure 14).

If a 1000-g limit is determined by the system gain settings and not the load cell's capacity, you can restore the dynamic range of the system via a preload-compensation circuit. For Figure 15's circuit to work, the load cell must measure the maximum expected preload plus the maximum load due to product.

For the sake of discussion, let's say the system needs to measure 0–1000 g of product with up to a 1000-g preload. So, the load cell must be capable of measuring 2000 g.

The microcontroller sets the DAC and mux to generate a compensating voltage equal to the preload voltage on the sensor (but of opposite polarity). The summer adds the sensor output and compensating voltage. Once the DAC section is set up, the output of the summer is zero when the preload is present.

The ADC now sees 0 V when the bowl is preloading the load cell. As product is placed in the bowl, the load on the sensor increases and the summer's output increases. The ADC sees a voltage corresponding only to the product weight, and the full 0–10-V range is available for measuring product. The system's usable dynamic range is restored.

This technique differs from gain-switching schemes because the product placed on the sensor is always measured in a zero to full-scale range for the ADC. So, regardless of the preload, the scale always has the same resolution. Using the full range of the ADC for product measurement ensures maximum resolution.

There are limitations to this technique based on sensor linearity and DAC versus ADC resolution. Although I've simplified the technique, you can get a good start with the topology shown in Figure 15.

Gain-switching schemes can be combined with this nulling technique to provide an overall system capable of incredibly fine measurement resolution while accommodating a huge preload.

If the Gain block in Figure 15 is programmable, we could have a high-gain mode and a low-gain mode. For example, low-gain mode, may map 0–1000 g of sensor stimulus to 0 to full-scale on the ADC. High-gain mode may map 0 to 1g of sensor stimulus to 0 to full-scale on the ADC.

Consider again the case of the 500-g bowl on the sensor. To compensate for the preload:

- the microcalorimeter selects low-gain mode
- the 500-g bowl is measured, and the DAC section is set to null the preload
- the microcalorimeter selects high-gain mode
- the microcalorimeter tweaks with the DAC until any remaining preload disappears (this step assumes the DAC is very high resolution, which is a simplification)

Next the user places product—e.g., 25 mg of salt—into the bowl. In high gain mode, the system gives the user maximum resolution in the 0–1-g range.

If the user pours another 750 g of salt into the bowl, the system switches to low-gain mode and gives a resolution corresponding to the 0–1000-g range.

The technique in Figure 15 is a versatile and powerful method of preload compensation. For applications where 100% of the ADC's dynamic range must be used, Figure 15 offers a good starting point. For less-demanding applications, a combination of gain switching and pure software nulling suffices.

INFO YOU CAN USE

I've reviewed op-amps, as well as examining some interesting mistakes and some excellent interface circuits. I didn't discuss PCB-layout considerations, but this topic is well-covered elsewhere [5].

As you've seen, analog design is full of tradeoffs. Models help, and SPICE is a great tool. But, nothing beats gaining insight and intuition by examining other people's successes and failures.

Bob Perrin currently works for Z-World. He has designed mixed signal instrumentation for agronomy research, food science, harsh-environment industrial computing, and embedded control. You may reach Bob at bperrin@zworld.com.

REFERENCES

[1] S. Franco, "Practical Op-Amp Limitations," *Design with Operational Amplifiers and Analog Integrated Circuits*, McGraw-Hill, New York, NY, 188–241, 1988.

[2] C.D. Motchenbacher and J.A. Connelly, "Capacity Multiplier Filter," *Low Noise Electronic System Design*, John Wiley & Sons, New York, NY, 309–310, 1993.

[3] Burr-Brown, *Fast Settling Low Pass Filter*, App Note AB-022, 1994.

[4] K. Hoffmann, "Selection Criteria For Strain Gages," *An Introduction to Measurements Using Strain Gages*, Hottinger Baldwin Messtechnik GmbH, Darmstadt, Germany, p. 58, 1989.

[5] B. Baker and J. Graeme, "Systematic Approach Makes Op-Amp Circuits Resist Radiated Noise," *EDN*, **40:15**, p. 93, 1995.

SOURCES

OP297 and OP497 op-amps, AD620B instrumentation amplifier
Analog Devices
One Technology Way
Norwood, MA 02062-9106
(781) 329-4700
Fax: (781) 461-4261
www.analog.com

OPA643 op-amp, INA118 instrumentation amplifier
Burr-Brown Corp.
P.O. Box 11400
Tucson, AZ 85734-1400
(520) 746-1111
Fax: (520) 741-3895
www.burr-brown.com

LTC1100 instrumentation amplifier
Linear Technology Corp.
1630 McCarthy Blvd.,
Milpitas, CA 95035-7487
(408) 432-1900
Fax: (408) 434-0507
www.linear-tech.com