

Ray Stata of Analog Devices Speaks Out on What's Wrong with Op-Amp Specs

Few products have more specs per cubic inch than operational amplifiers. It's no wonder that this is an immensely fertile field for Murphy's Law to raise havoc. Considering the number of variables to be juggled, and the relative inexperience of many op-amp users, it's not surprising that the circuits don't always play the first time around.

When manufacturers add the smoke screen of foggy specs, either deliberately or inadvertently (and I guess we are all guilty of this at one time or another), the poor user is really in trouble.

There's a bewildering array of numbers and most are just approximations.

Op amps are darned complicated. There are really lots of specifications, and most are only approximations of what's happening in the black box. Most parameters like voltage drift, current drift and open-loop gain are nonlinear functions of temperature; others, like common-mode rejection and common-mode impedance, are nonlinear functions of input voltage. And almost all parameters depend on supply voltages. Therefore, when you use a single number to specify a parameter, you must qualify the conditions of the measurements. Comprehensive graphs would be necessary to define performance completely.

Confronted with 50 more-or-less mysterious numbers, most engineers tend to select an op amp in terms of familiar values and to forget about the rest. An engineer who wants to replace a sensitive relay with a low-cost amplifier might simply concern himself with the output-current rating and neglect such factors as drift or gain.

And there are no standards.


Then there's the lack of standards for op-amp specs. Though various efforts have been made to unify terms throughout industry, this has not yet been done, so manufacturers have loopholes for specsmanship.

Along with real, honest-to-goodness specsmanship, we find the inadvertent errors of omission. The holes in the spec, or usually, the lack of additional information, is not revealed until the engineer has assembled umpteen amplifiers into his product and the whole batch is waiting to be shipped. Very sad.

We were attending our first IEEE Show when a group of indignant engineers advanced on us and used us as whipping boys for the whole op-amp industry. They accused us of being con men. This we take in our stride when it's leveled

at the competition, but these guys included us in their blast.

They had purchased a diff amp with a $5 \mu\text{V}/^\circ\text{C}$ drift spec. The data sheet said that this number meant maximum voltage drift. What the sheet neglected to say was that the specs held only for steady-state temperature conditions. What happened during thermal transients was another kettle of fish that no numbers covered.



Though we were able to show that this point had been covered in an early application note, we certainly hadn't referred to it in spec sheets for any amplifiers. Actually we didn't know how to specify this mode of operation quantitatively. And some competitors had omitted the point.

Now, let's take a closer look at some of the major specifications and see how these problems come up.

There's more to voltage drift than meets the eye.

It can help to go back to first principles. We all know that a transistor's base-emitter voltage varies at roughly 2400 microvolts for every degree C change. This variation develops an output, or offset, exactly as if a true input of equal magnitude were driving the transistor. The use of differential pairs enables the net voltage offset to be drastically reduced because both transistors

can be matched so their offsets track within a few microvolts over the temperature range. This is how the amplifier comes to have a drift spec of $5 \mu\text{V}/^\circ\text{C}$ instead of $2400 \mu\text{V}/^\circ\text{C}$; the transistors track within $5 \mu\text{V}$ for every $2400 \mu\text{V}$ of base-emitter drift, that is, for every 1°C rise.

But there's a flaw here. What happens if the base-emitter junctions are not at the same temperature? The spec states a figure for maximum drift, but this is really a tracking spec for base-emitter junctions at the same temperature.

Obviously, thermal gradients caused by adjacent heat-dissipating components can make nonsense of such specifications by making one junction hotter than the other. It only takes 0.1°C differential between the two junctions to develop $2400/10 = 240$ microvolts offset. To say the least, this is a substantial offset for an op amp with a $5 \mu\text{V}/^\circ\text{C}$ max drift spec.

In fact, offsets are caused not only by such obvious temperature-gradient sources as adjacent heat-dissipating elements and room-air drafts, but also by changes in the amplifier's own load current. Altering the output current from 2 mA to 20 mA produces an internal temperature transient that develops an offset due to unequal heating of the input transistor pair until temperature equilibrium is re-established. And when the amplifier has settled to its new output level, the input signal will invariably set the load current back to 2 mA, starting the problem all over again.

Offsets due to warmup and changes in operating conditions can be particularly annoying when an instrument or system is being adjusted. It takes only finger heat on one side of a differential amplifier to produce a distinctly-measurable offset error. How does one put performance specifications on such nebulous factors?


One way to sidestep the temperature-gradient error is to avoid the differential amplifier. A chopper-stabilized amplifier corrects automatically for transient offsets of this kind. Its error-sensing circuit is independent of mismatches between the transistors. But sometimes the chopper unit won't win. Apart from being twice as big and costly as a differential amplifier, the chopper-stabilized unit has only a single input terminal. (The other terminal is "used up" in the stabilizing circuit.) A single-ended amplifier can't easily be used in noninverting or differential circuits. So what you gain on the roundabouts breaks the rope on the swings.

Fortunately, some diff amps are inherently less susceptible to temperature gradients. Single-

chip, dual-transistor front-end circuits cut down thermal inertia and reduce physical spacing between base-emitter junctions of each differential pair. Monolithic IC op amps are generally good on this score because of the small spacing between junctions.

Consequently, temperature transients never pull the two junctions more than a few hundredths of a degree apart. Such amplifiers offer a meaningful steady-state maximum voltage drift of about $1 \mu\text{V}/^\circ\text{C}$, with a short-lived and worst-case offset of about 75 microvolts for 40°C thermal shock.

While an amplifier with a dual input transistor exhibits a considerable improvement in offset stability, there is no way to discern this fact in comparing the usual published offset-drift specifications. A user can get some idea of comparative thermal-gradient performance by making thermal shock tests, like dropping the amplifier in an oil bath 40°C above ambient. He'll get a response like this.



Okay, so we've raised one straw man and then beaten it down. But aren't there other problems that only specsmanship has solved thus far?

Zeroing can hide some second-order pitfalls.

What happens when you adjust an amplifier's offset potentiometer to zero its output at the selected working temperature? Not surprisingly, there's more to the process than meets the eye. And it's not always easy to tie up the sources of error in neat, crisp numbers.

The offset pot is frequently a variable resistor in series with a collector load resistor. The amplifier's output is zeroed by altering collector



current to change voltage balance between the two front-end transistors.

So What? Well, it turns out that there's a second-order effect that causes interaction between the actual value of collector current and the rate at which the transistor's base-emitter voltage drifts with temperature. Adjust collector current in one transistor of a matched differential pair and it no longer tracks the other as temperature varies. True, the adjustment modifies the base-emitter drift by only $0.7 \mu\text{V}/^\circ\text{C}$ for each $250 \mu\text{V}$ change in emitter-base voltage or initial offset voltage.


But such differences can swamp the carefully-designed tracking specs of today's state-of-the-art differential amplifiers. For example, some chopperless differential amplifiers have better than $1 \mu\text{V}/^\circ\text{C}$ maximum voltage drift. To zero an initial offset voltage of 1 mV with the internal balance pot in one of these amplifiers would introduce a change in temperature drift of $2.8 \mu\text{V}/^\circ\text{C}$. And that really screws up the works.

Some manufacturers neglect to point out the second-order effects caused by trimmer adjustments, but as op amp specs improve these effects can no longer be ignored. One must really know the condition of the balance resistor to specify voltage drift uniquely. We get around this by eliminating provisions for an internal offset trim on low-drift amplifiers and by recommending external offset biasing circuits.


Average drift is a trap. Remember the man who drowned while wading through a stream with an average depth of only four feet?

Neither voltage drift nor current drift are linear functions of ambient temperature. Accordingly, the numbers published for voltage drift and bias-current drift can refer only to average values, or to values at specific operating conditions. For example, an amplifier's total change in voltage offset for a temperature excursion from -25 to $+85^\circ\text{C}$ might be 2200 microvolts (referred to the input). The average drift rate over this interval works out to $2200/110^\circ$ or $20 \mu\text{V}/^\circ\text{C}$. But when you look at the actual drift


curve you see that the drift rate at the extremes of temperature can exceed the specified average drift rate by a substantial amount.



You have a special problem when the drift curve changes slope—a real live possibility. Here the average slope calculated by subtracting end points is entirely meaningless. Nonetheless, some manufacturers have taken advantage of this golden opportunity for specsmanship.



Perhaps trickier to interpret than the voltage-drift figures are the nonlinear errors caused by exponential variation of bias currents in FET and varactor-bridge amplifiers. Usually, the bias current at a given temperature is specified, and the spec reader is reminded that the bias current doubles for approximately 10°C temperature rise as it does here.



If a value of bias-current drift is given, it is usually quoted at room temperature, where the slope of the bias current versus temperature curve (i.e., drift) is shallowest. But at 85°C, the bias current and drift slope will be 64 times worse than at room temperature, making the FET amplifier a worse choice for high-temperature applications than some bipolar transistor types. This is not really a question of specsmanship, but it certainly requires that the user know his way around a spec sheet.

How much output do you really get?

An amplifier output rating of ± 10 volts, ± 20 milliamps implies that the user can drive a 500-ohm load at the full signal swing of 10 volts, 20 milliamps. For discrete-component op amps this is usually what such specs mean. But, integrated circuit manufacturers seem to have different ideas. "Sure you can get 10 volts output; certainly it will develop 20 milliamps," they say, but they often forget to add, "so long as you don't ask for them simultaneously." In fact, it is not unusual for an IC amplifier labeled as having a ± 10 V, ± 5 mA output to have a maximum power rating of only a fraction of the product of the VI figures given.

If an engineer needs an op amp with a relatively high output, he generally wants to know what gain he's getting at that current level. If 20 mA is the amplifier's full-load rating, it would be nice and simple if the manufacturer stated the amplifier's open-loop dc gain at this full load. Not all manufacturers do.

The op-amp user should know his amplifier's roll-off curve in order to build a circuit with adequate gain stability over the working frequency range. The manufacturer may be perfectly justified in departing from the conventional 6 dB/octave frequency compensation to achieve desirable features like fast settling time, high slew rate, fast overload recovery, or increased gain stability over a wide range of frequencies.

But to obtain these improved features generally requires fast roll-off characteristics and therefore a propensity toward oscillation. Key to preventing instability, of course, is knowing that you have this kind of amplifier. You can then use one or more well known circuit techniques to tame the oscillations.

The great crime occurs when manufacturers use fast roll-off compensation to obtain improved published specs without giving an open-loop response curve or some other indication of what's going on.



Let's look at common mode, the specs we'd all like to forget.


To many users the common-mode rejection ratio of an amplifier is a rather mysterious number that they'd like to forget. Many op-amp manufacturers feel the same way. In fact, if a particular amplifier has a poor common-mode spec, some manufacturers thoughtfully omit it from the data sheet.

Low-cost FET amplifiers are the worst culprits, with typical common-mode rejection ratios around 1000. This is exactly the kind of number that manufacturers would like to lose. They often do. For a noninverting amplifier circuit, the common-mode error is $1/\text{CMRR}$, which works out to 0.1% for an amplifier with CMRR of 1000, or 60 dB. Recently, new circuit tricks have enabled manufacturers to overcome this fundamental limitation of low CMRR, but a FET amplifier with 100,000 CMRR tends to cost more than \$100.

Picking a typical CMRR spec from measured op-amp data is great sport. The numbers vary over an enormous range — from 500:1 to as high as 100,000 for FETs — all seemingly at random. A reasonable way for a manufacturer to select a typical common-mode figure for his data sheet is to pick a value that is met by 70 or 80% of all units of that particular type. Some manufacturers, with considerable ingenuity, average all the test numbers to find a "typical" value. Tweaking up a few samples to get 100,000 CMRR can do great things for averages of this kind and can, of course, lead to very respectable-looking common-mode rejection figures.

As with drift, an amplifier's common-mode-rejection performance can vary with operating conditions, notably with the value of common-mode input voltage as we can see in the sketch. This leaves room for some really fancy footwork. For example, some well known FET types boast a common-mode voltage range of ± 10 volts. But the common-mode rejection figures are specified for a ± 5 -V common-mode

range. It's possible for the CMRR to degrade by as much as a factor of ten when the applied common-mode voltage is raised from 5 to 10 volts.



Another difficulty with CMRR is that, since it's a nonlinear function of input common-mode voltage, a single spec number can at best give an average value over the test-voltage range. For small input-signal variation about some large common-mode voltage, the specified "average" CMRR gives little indication of the actual errors you can expect due to the steeper error slope at high voltages.

How full is the full-power response?

We all know that an amplifier rated at, say, 10-MHz unity-gain bandwidth, doesn't give full output-voltage swing at this frequency. Invariably, distortion caused by internal slew-rate limiting induces the manufacturer to specify the maximum full power frequency several decades lower.

Generally, an amplifier with 10-MHz unity-gain bandwidth would have a full-power response of about 1 MHz or maybe as low as 100 kHz. At the "small signal" unity-gain bandwidth, the achievable output-voltage swing is usually related to the swing at the full-power bandwidth by the ratio of the full-power response, f_p , to the unity-gain bandwidth, f_t . Thus, for example, you get only a 100-mV swing from a 10-V amplifier at unity-gain bandwidth for an amplifier with an f_p of 100 kHz and an f_t of 10 MHz.

One large problem with full-power-response specs is that no one really says what he means by the published minimum number. For one thing full-power response has nothing to do with amplitude vs frequency as the term "response" normally implies. Instead it is a measure of output distortion caused by slew-rate limiting.

But there can be monstrous difference depending on whether you set 1% or 10% as the acceptable distortion level. We have evaluated some amplifiers where the output looks a triangle at the specified full-power response. Where do you draw the line on the acceptable distortion level?

Distortion is only one consideration in setting the criteria for the full-power-response spec. A subtle, but very often more important side effect, is dc offset error due to rectification. Feedback signals developed by an unsymmetrical, distorted output, will not counterbalance the input signal at the amplifier's summing junction. So, the error signal is rectified by the amplifier's input stage, generating an undesirable offset voltage.

But how can such application factors be reasonably prevented? Should each data sheet be turned into a thesis, or may the manufacturer assume that his customers are already aware of the difficulties awaiting them? There is no easy answer for either the manufacturer or the customer.

Slewing rate for the most part is just another way of looking at rate limiting of the amplifier's circuitry. The slewing-rate specification applies to transient response while full-power response applies to steady state or continuous response. For a step-function input, slewing rate tells how fast the output voltage can swing from one voltage level to another. Fast amplifiers will slew at up to 300 V/ μ s, while amplifiers designed primarily for dc applications often slew at 0.1 V/ μ s or less.

Maximum slew rate, S , is related to full-power response f_p by $S = 2\pi f_p E_o$. As the voltage swing is reduced below the peak output, E_o , the operating frequency can be proportionately increased beyond f_p without exceeding the slewing rate. For many amplifiers the slewing rate may differ in the inverting and noninverting configurations — a fact that is not always apparent from published specs. Moreover, there is almost always a difference in slewing rate between fall time and rise time, and between positive and negative output signals. Opportunities for specs-

Who is Ray Stata?

A soft-spoken man with a wry sense of humor, Ray Stata doesn't publicly expose the fact that he's a dynamo. He has packed more experience in his 33 years than many men twice his age. He earned his BS and MS in EE at MIT, then worked at MIT's Instrumentation Lab before he went for some industrial experience at HP.

In 1961, just three years after he obtained his MS, he and Matt Lorber started Solid State Instruments on a shoestring. They sold that company and used the proceeds to start Analog Devices in January 1965. In just three


years, they brought the sales level to a \$5-million annual rate.

Ray enjoys writing on op-amp theory and applications as much as he enjoys building companies. A Boston Symphony enthusiast, he retains an active interest in music and in his record collection.

His greatest joy however (and the joy of his wife Maria) comes from his four-month-old son, Raymond Paul, who has not yet developed an interest in operational amplifiers, symphonic music or building companies.

manship arise here since many possible slopes can be measured.

A recurrent problem in our applications department is the irate telephone call from a customer claiming that our amplifier does not meet its slewing-rate spec. Closer inquiry typically shows that the customer is trying to obtain $100 \text{ V}/\mu\text{s}$ slew rate with a circuit like this one.



The problem here is that the inevitable 1 pF of stray capacitance must be charged by the $10\text{-}\mu\text{A}$ signal current. This limits the output slew rate to $10 \text{ V}/\mu\text{s}$ regardless of how fast the op amp may be. With an input impedance of $10 \text{ k}\Omega$ or less, the customer would obtain the desired $100 \text{ V}/\mu\text{s}$ slewing rate. In other words, you cannot get fast response from an inverting amplifier when using high input impedance owing to the slugging effect of stray capacitance.

Settling time is a parameter of increasing interest. This spec defines the time required

for the output to settle within a given percentage of final value in response to a step-function input. Common accuracies of interest are settling time to 0.1% and 0.01%. Heretofore, engineers have been forced to use slewing rate and unity-gain bandwidth as rough indicators of relative settling-time performance when comparing or choosing amplifiers, since no other data are given. As it turns out, these two specs have little bearing on settling time, particularly to 0.01%.

Manufacturers now realize this and are beginning to publish settling-time figures. The hooker here is that settling time is really a closed-loop parameter (while all other op-amp specs are open-loop parameters), and therefore depends on the closed-loop configuration and gain. Fortunately, even when published for only one gain, usually unity gain, it serves as a realistic yardstick for comparing amplifier performance.

The few examples I've shown illustrate the tremendous difficulties in specifying op-amp performance and the many pitfalls confronting the user. In the long run the user and manufacturer are on the same side in seeking unambiguous communications. One needed step is for an industry group to establish standards on definitions and test circuits. The use of op amps is growing very rapidly and this move is long overdue.

There's one further point that can be most important of all. An engineer should breadboard his critical circuits instead of relying totally on paper designs. Better yet, he should contact the applications department of an op-amp manufacturer to review his problems. Few engineers take advantage of this free service or give proper weight to this important aspect of picking a vendor.

EEE