BOB PEASE

What's All This Logarithmic Stuff, Anyhow?

Everybody knows that the V_{BE} of a transistor is a sort of logarithmic function. After you get the transistor up to the right threshold voltage—about 0.6 V for silicon transistors at room temperature—the transistor's collector current rises exponentially whenever the voltage increases slightly. In other words, you might say that the base-emitter voltage rises *logarithmically* as a function of the current.

Great! Now, we can cover a wide range of signals by looking at a finite range of voltages. The familiar slope of 60-mV output per decade of input seems nice and handy. You might just put your current into a diode or transistor. Or, you might put the transistor in the feedback loop of an operational amplifier *(Fig. 1a)*.

BOB PEASE
OBTAINED A BSEE FROM MIT IN 1961 AND IS STAFF SCIENTIST AT NATIONAL SEMICONDUCTOR CORP., SANTA CLARA, CALIF.

But, this "exponential" characteristic is really not a realistic, *usable* curve *(Fig. 1b)*. Why not? Well, the familiar "diode equation" is not usually properly written in terms of temperature. I_S is NOT temperature invariant. It varies wildly. If the temperature of a diode or transistor changes, you see the old familiar –2 mV/°C sensitivity. If the junction current is very small, that tempco may be as big as 3 mV/°C. So, if the V_{BE} changes by 18 mV, was this caused because the sun went behind a cloud and the circuit cooled off 9°? Or, was it because there was a 2:1 change of current? Kind of hard to tell. So, the raw V_{BE} of a transistor is NOT a good, useful logarithmic detector.

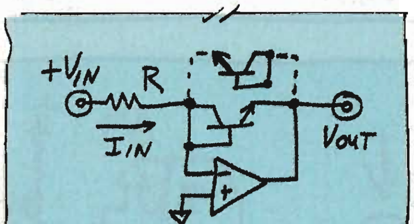

Fig. 1a

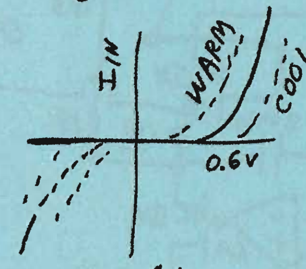

FIG. 1b

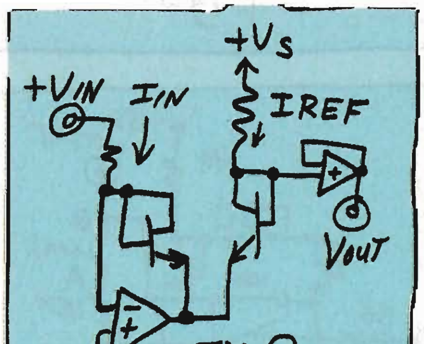

Fig. 2a

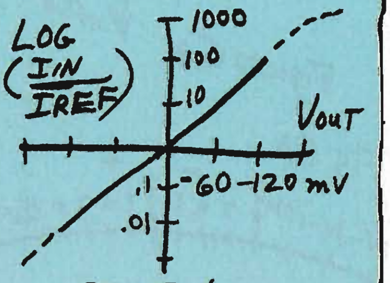

Fig 2b.

What, then, IS usable? There have been good log detectors around for well over 30 years. You don't just use one V_{BE}; you use a *difference* between two of them. This is much more predictable, and its temperature sensitivity isn't as bad *(Fig. 2a)*. When the currents through the two transistors are *equal*, the output is zero. That is stable with temperature. The output has a slope of 60 mV/decade at room temperature *(Fig. 2b)*. This circuit can easily cover a range of ±2 decades. But, this logger still needs some more improvements.

The seminal papers written by Horn and Gibbons were done in the late 1960s, showing that the logarithmic fidelity of the *collector* current of the transistor is more accurate than the *emitter* current. That's much better than most *diodes*. In Figure 2a, all of the input current flows through the npn's emitter. In Figure 3, all of the input current flows through the *collector*, which makes for a much more accurate log function—especially when the transistor's beta may fall off at low currents.

The voltage at the base of Q2 will move at 60 mV/decade. But, that is linearly proportional to Kelvin temperature. Not so good. So, how do we compensate for that? Take a look at the two resistors around A2. The input resistor is a special wirewound resistance specified at +3400 ppm/°C. There are several suppliers of resistors at this tempco. This compensates for the "gain" of the logger, giving pretty good overall gain.

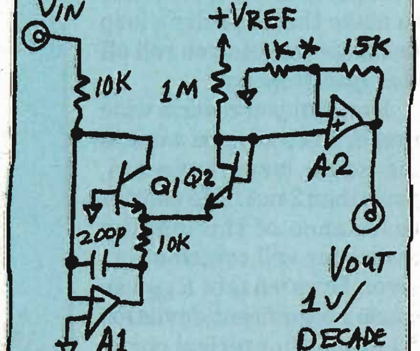

Fig. 3

Of course, for best results, you can't expect perfect accuracy if you get far away from room temperature or are taking extremely wide current ranges. The example in Figure 3 covers the range from 10 µA up a couple of decades (to 1 mA), and down to 0.1 µA. Of course, to maintain full accuracy, you would want an op amp that starts out with less than 1 mV of offset. That way, when you trim it down to 10 µV, it will stay there. Cheap grades of OP-07 are often suitable. BiFETs and CMOS amplifiers are sometimes pretty stable.

You should note also that the dc output error at the emitter of Q1 is not directly affected by the offset voltage of A1. If you have a 1.0-V signal input, the error caused by 1 mV of offset from A1 will only create a 26-µV error at Q1's emitter. That's because the emitter voltage is established by the fact that the base is tied to ground—NOT to the summing point. If you were using the circuit in Figure 2a, it would cause 1.026 mV of error. So, the Figure 3 circuit will cause less errors in some cases.

Note the feedback capacitor around the first amplifier. The gain of Q1 makes it necessary to add the resistor and capacitor at its output, making the loop stable. The reason is that the transistor adds so much gain to the loop. When you add extra gain to your feedback loop, it tends to make the amplifier's loop unstable—unless you roll off that gain properly.

How can you cover a wide range? Let's say you want to log some large currents, more than 2 mA. The emitter resistance of the logging transistor will contribute to error. Even 0.5 Ω of $R_{EE'}$ can cause a significant deviation from the theoretical curve (18 mV/octave). The improved circuit in Figure 4a can make good compensation. For example, 10 mA × 0.5 Ω = 5 mV. Here, the signal input is permitted to rise, permitting the transistor's base to be pulled up to compensate for the I × R drop in the emitter. The accuracy, improved up to 20 mA, is shown in Figure 4b.

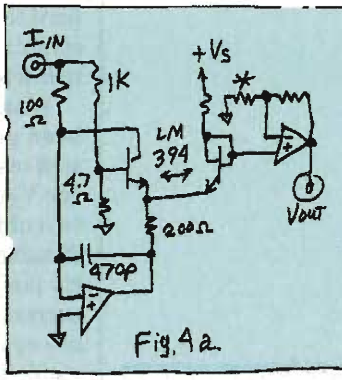

Fig. 4a.

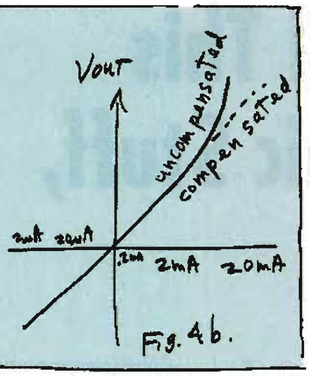

Fig. 4b.

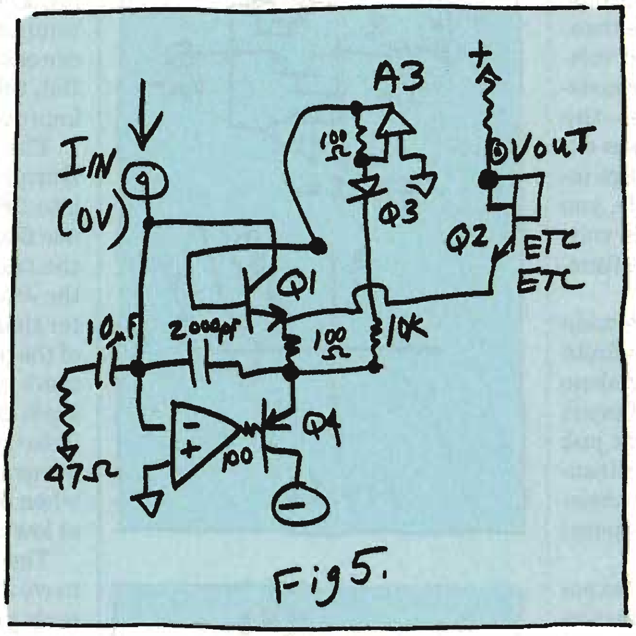

Fig 5.

Another version is shown in Figure 5, where the current input (a real summing point) is held strictly at 0 mV. A dummy current is generated through Q3 and this lets A3 pull up the base of Q1. I tried this one at 100 mA, using a big power transistor. It seemed to work OK, though the self-heating in the transistor was not negligible.

Now, we indicated way back at the beginning that somebody might like to do the logarithms of both positive and negative signals. But, just using two diodes in the feedback of an op amp—as shown by the dashed lines in Figure 1a—is a lousy way. Yet, a thoughtful application of the circuit of Figures 2a and 3a can lead to Figure 6a.

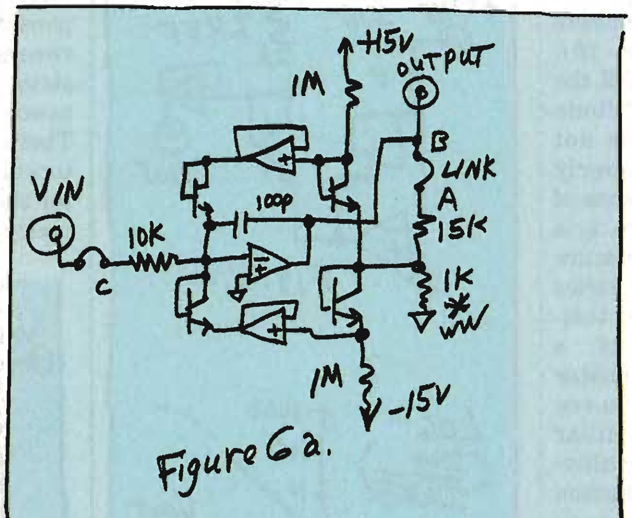

Figure 6a.

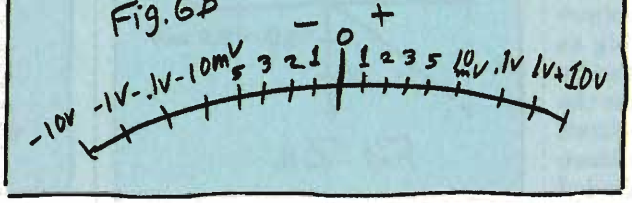

Fig. 6b

This circuit was originally designed as a log null indicator—a semilogarithmic null. It simplifies how you null in a circuit as you trim it. It will have very good resolution for small signals, but you can still see what is happening with small changes of a large null error. First designed as a small, round, epoxy-potted module, it was to be bolted onto the back of an analog meter and called a "meter-mate" by the Nexus guys back in '67. Cute idea.

Then, somebody discovered that this ac null function—$\sinh^{-1}$ function—can be used as a logarithmic *compressor* function for audio signals. Plus, when you take a second one and connect it as an *expander*, it makes a *compander* (that is, a compressor-expander) function. One of these days I

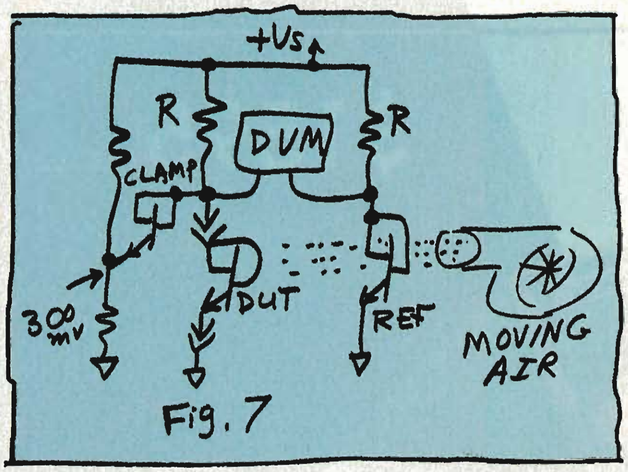

Fig. 7

may even connect one of these up in my phone system.

Note: the definition of sinh(x) is: $1/2(e^{x} - e^{-x})$. Thus, the current in this circuit *(Fig. 6a)* increases in magnitude *exponentially* as you get away from the null. Conversely, when you look at the inverse function, $\sinh^{-1}(y)$, it has a flat slope near null. But, as you get far away from null, it looks more and more like a log function, with a scale per Figure 6b.

Just what we wanted! That's what the compressor does—the circuit of Figure 6a. To use this circuit as an expander, apply the input signal to point A, and then connect the output of the op amp from B to C.

Back in the Vietnam era, some guys were proposing to insert this compressor function right after a soldier's microphone to cut down the effect of the sounds of gunfire. These guys were figuring they could bugger around with the transfer function and cut down on the noise of the shots. I don't know if they ever got that to work. But, it was a cute concept. Recently, a guy had a similar requirement: to cut down the impact noise of bowling pins falling!

Of course, for all these logging applications, transistors should be matched in pairs for V_{BE} at a nominal temperature. Specifically, pick up the transistors with tweezers. Insert them into a socket that's adjacent to a reference transistor running at the same current *(Fig. 7)*. Any deviation between the V_{BE} of the DUT and the reference transistor's V_{BE} is used to grade the transistors into bins such as 1 or 2 mV wide. Then, when you take the transistors out of that bin, they are well-matched. Oops! I almost forgot to say that you have to blow a big air blast—a steady flow of room-temperature air—over the DUTs. In fact, the DUTs waiting to go into the test should be kept in that same moving air.

After these transistors are matched, it's important to install them properly to keep them at the same temperature. Use some epoxy and some metal or junk as a thermal mass to keep the transistors at a fairly constant and equal temperature, along with the temperature-compensating resistor.

Another way to get matched transistors is to buy matched pairs, such as LM3046 or LM394. The '3046s are typically matched to within 1 mV, but the specs are about 3 mV max. You get a quad of transistors at this price. The LM394s are a bit more expensive for a dual of WELL-matched transistors. But, you get a spec of 300 µV max for the LM394CH.

So, as you can see, there are many games to play in the area of logarithmic and exponential functions. Stay tuned. I might share some more of them with you. Let me know if you'd like to see more on this subject. But this should give you enough to think about for now.

All for now. / Comments invited!
RAP / Robert A. Pease / Engineer
rap@galaxy.nsc.com—or:

Address:
Mail Stop D2597A
National Semiconductor
P.O. Box 58090
Santa Clara, CA 95052-8090

IDEAS WANTED

Be Creative. Send your Ideas for Design submissions to:

Ideas for Design Editor
Electronic Design
611 Route 46 West
Hasbrouck Heights, NJ 07604
Fax: (201) 393-6242

BOB PEASE, *Contributing Editor*
• • • **rap@galaxy.nsc.com**

WHAT'S ALL THIS LOGARITHMIC STUFF, ANYHOW? (PART 2)

The I_C versus V_{BE} of modern transistors has some excellent log characteristics, as we have discussed.* If you ground a transistor's base and compensate for its V_{BE} with a matching V_{BE}, you can do some good logging over a wide range, from 1 mA to 1 pA, and really quite accurate from 100 pA to 100 µA.

That's six decades where the limitation of R_{EE}' on the high end is the major limitation. Input leakage currents on the low end are not really a limitation, these days, as good CMOS amplifiers have I_{IN} smaller than 1/1000 of the 100-pA signal. Of course, this works best around room temperature. The compensating resistor R_x (1-kΩ wirewound resistor at +3500 ppm/°C) works fairly well for a moderate temperature range, and it works best around room temperature. Figure 1 shows a standard log ratio circuit, found in AN-29 and in many books. Its output is −1 V per decade for inputs larger than 0.2 µA.

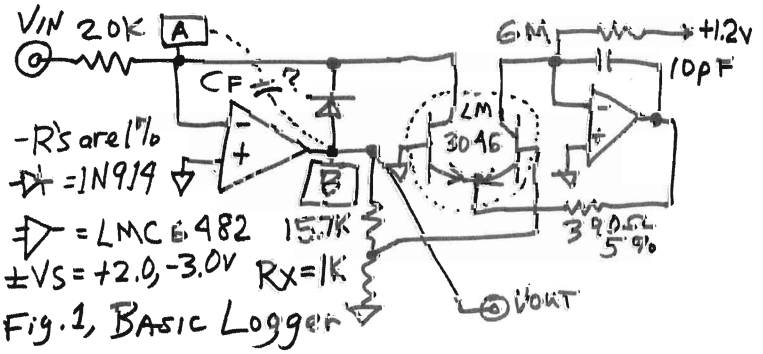

Fig. 1, Basic Logger

However, there are definite limitations in speed of response. A large input current requires a large feedback capacitor, such as 500 pF, because the transistor adds so much gain to the loop that it ruins the loop stability. When I_{IN} is as large as 1 mA, the transistor has a gain of 800—too much gain to add to an op amp. (Op amps are happy with attenuation in the feedback loop, but they don't like gain added.) So we have to add a large C_F, such as 500 pF, to make the loop stable at high frequencies.

Now when I_{IN} is decreased down to 1 µA, the loop is very stable and very slow. The 500 pF is much too big. The bandwidth falls below 1 kHz. A feedback capacitance of 2 pF would be plenty. Some of our customers needed a logger with good audio bandwidth over a wide range of input currents. What to do?

I remembered that some engineers back in the 1960s had this problem and used some diodes to solve it. But I never saw how they did it. So I went back to the scene of the crime. All I needed was a feedback capacitance that was big when the signal was big, and small when the signal was small.

I figured out that the circuit of Figure 2 might do it. All I had to do was build it up and try it out. Do a cut-and-fit on the capacitance sizes. Sure enough, it worked quite well. The effective feedback capacitance increases as the output goes more negative, and the diodes start to conduct. I used a fairly spacious layout, so the stray feedback capacitance (C_O) was only 1 or 2 pF, due mostly to the amplifier's socket and the 1N914. If you made a really good layout, you could get the strays even lower than that and get fast response below 1/10 µA.

I kept a bandwidth of 12 kHz, from 1/4 µA up to 0.4 mA. I didn't really try hard to optimize it further. I stopped, as it was working entirely well enough for me. It did overshoot a bit, but it did not ring at any signal level. You have my permission to optimize it as needed. You might use three or four diode-R-C networks to cover a wider range or to get cleaner response at all levels. So there's an old trick that hasn't been seen for about 40 years, but when we have to reconstruct it, it works pretty well. ●●

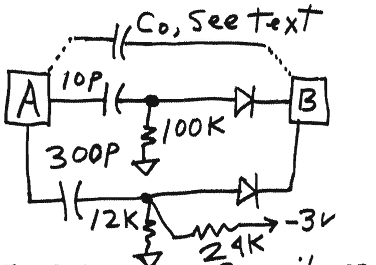

Fig. 2, Adaptive Capacitance

ED Online 11301

Comments invited! *rap@galaxy.nsc.com —or:*
Mail Stop D2597A, National Semiconductor
P.O. Box 58090, Santa Clara, CA 95052-8090

---

* "What's All This Logarithmic Stuff, Anyhow?" *Electronic Design*, June 14, 1999, p. 111; ED Online 6068 at *www.elecdesign.com*

BOB PEASE *obtained a BSEE from MIT in 1961 and is staff scientist at National Semiconductor Corp., Santa Clara, Calif.*