electronic design

What's All This Long-Term Stability Stuff, Anyhow?

Electronic Design

Bob Pease
Bob Pease

Tue, 2010-07-20 (All day)

1 of Enlarge image

A few years ago, a guy called up and asked me about a new NSC amplifier: "What data do you have on this op amp for long-term stability?" Without even checking, I told him that we surely did not have any such data. This was just a fairly new version of a garden-variety op amp, and it was not trying to be a technical leader in low offset or drift. So we were fairly safe in not bragging. But the customer was not satisfied.

"Why don't you have full information on this?" I tried to explain that we can't afford to gather this data on every little product. We can't afford to take the data and analyze it. We also can't afford the delays in our time-to-market. Even if we released it soon and updated the datasheet later, people would look at the first datasheet and ask the same question.

The amount of manpower, engineering power, and technician power that's needed to do a proper study would be huge. It's rare to find a customer who has enough serious need for that info to make it worthwhile. Usually, if we have obtained good data on previous circuits using this process, and the changes are minor, such as a minor change of layout or amount of current, or a change of output stage, we assume that the new part will look very much like previous ones.

Now, we do perform a little drift testing just to make sure nothing stupid is happening. For example, we'll load up three boards, each with 30 parts from three different runs, to see if they look as good as expected. We'll compare the data from before 1000 hours of high-temp operation versus after. We may throw one of those boards back in the ovens to get data at 2000 hours. But usually this gets boring, real fast. And when we are all done, there's nothing to brag about. Nothing to write a glowing report about. Just "Passed."

The customer on the phone was still grouchy. "Okay, I'll just go over and ask some of your competitors what their long-term drifts are," he said. So I told him, "Be my guest," and I started to tell him the phone numbers of some of our competitors. I keep them memorized for just such occasions. But then I paused: "But then they are going to tell you the same story. They can't afford to do precision life tests on every good circuit they bring out."

WHEN DO WE MAKE A BIG DEAL?
When do we do a lot of testing and data-logging? When we're using a new process or a new circuit that's expected to provide superior performance. A new low-drift op amp? For sure. When National's new chopper-stabilized amplifiers came out a few years ago, we ran all sorts of life tests to make sure there weren't going to be any bad apples. And when there weren't any bad ones, how good would the good ones be? It took a lot of nitpicking to analyze enough data to state the typical "0.006 µV per month" such as on the LMP2012.

Oh, yeah? What do you expect to see for a year or two? Our standard rule is "For a time N × longer than 1000 hours, we expect the drift to be $n = \sqrt{N}$ × the 1000-hour figure." This usually makes the customer happy and/or shuts him up because (a) it is usually true, or close to true; and (b) if he wanted to get data, it's a lot of work for him! And he doesn't usually come back quickly. If at all. For example, for two years or 16,000 hours, $\sqrt{16} = 4$.

The LM199AH super-reference
When the new LM199AH came out about 38 years ago, it was designed as a new circuit using a new process to cancel out all probable causes for long-term drift. Its output tolerance was ±3%, but the long-term stability per 1000 hours was 0.0020% typical—and that's 20 ppm.

We put in lots of preliminary tests to screen out bad ones and then put in comparison circuits so we could use an excellent six-digit digital voltmeter (DVM) to compare several reference sources, such as ovenized standard cells, an ovenized band-gap reference, and several other fairly good zeners.

By using multiple references, we could avoid problems in case all of the devices under test (DUTs) seemed to drift at the same time. Was that caused by all the DUTs drifting? No, because the other references showed the same dip at the same time, meaning that the DVM's reference was to blame. And that effect could be "deducted," or at least ignored.

One day I got mad and grabbed a big double handful of these LM199AHs, soldered in a group of four, and averaged their outputs with small resistors (499 Ω?). This output seemed quieter and less drifty. Well, let's do it again. Soon I had four groups of four.

I compared the averaged output from eight LM199s to the other set of eight, and that was really good! Some tests showed less than 2 µV p-p for a limited bandwidth (4 Hz?). If I had averaged all 16, the output noise would have been even smaller! Most people don't need to make such low noise as that, but by averaging several circuits, you can get a square-root advantage. Until you run out of steam, space, and power.

Comments invited! czar44@me.com —or:
R.A. Pease, 682 Miramar Avenue
San Francisco, CA 94112-1232

Bob Pease obtained a BSEE from MIT in 1961 and was a Staff Scientist at National Semiconductor Corp., Santa Clara, Calif.