Don't Fail Good Boards

Ted T. Turner
Hewlett-Packard, Loveland, CO

By understanding false failures in board testing and the tradeoffs involved in eliminating them, you can improve both the speed and quality of your testing.

When you test analog PCBs, good boards can fail a test. This syndrome has many names such as "no defect found" (NDF), "no trouble found" (NTF), "false reject," and my favorite, "false failure." You must learn how to detect false failures and how to set test limits to avoid them.

False failures can be caused by many different fixturing and manufacturing problems and can occur during any type of testing using any test system. It's possible, though, for you to learn to recognize false failures, determine why they occur, and minimize them. To illustrate the process, I will describe errors that an in-circuit test system can contribute to the measurement of a basic passive component.

I will use a 100-Ω resistor with ±10% tolerance as the subject of this analysis. By definition, this resistor is "good" if its actual value is anywhere between 90 Ω and 110 Ω—otherwise it's "bad." Given this "good" range, it might seem natural to set the test limits at 90 Ω and 110 Ω, but you can't do that because of *test uncertainty*. No tester is perfect, and all testers will pass some bad resistors and fail some good ones. Test uncertainty, therefore, requires you to think in terms of passes and failures rather than in terms of a good part or bad one.

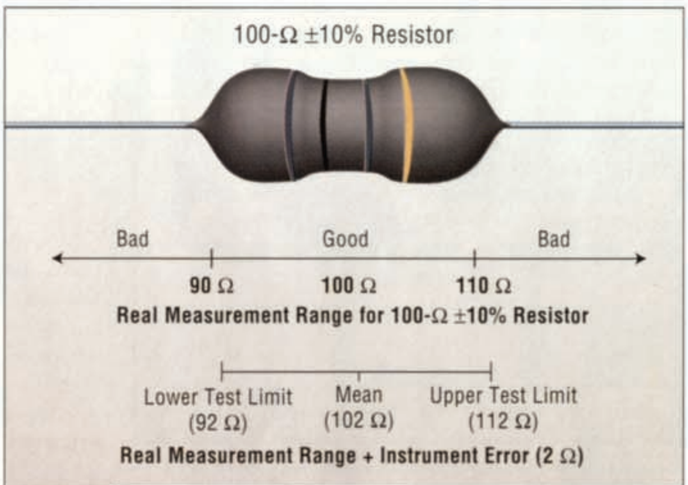

FIGURE 1. Test system inaccuracies such as average measurement error, in this case 2 Ω, will skew (or shift) the mean of all measured values of a component. Fortunately, you can correct this type of error by modifying the test program.

Several types of errors contribute to test uncertainty. If you test a perfect 100-Ω resistor and the tester indicates 102 Ω and then test a perfect 110-Ω resistor and the tester indicates 112 Ω, then your tester has an *average measurement error* of +2 Ω for this type of resistance measurement. **Figure 1** shows how the average measurement error skews the readings of a good resistor.

Whether the error is due to the design of the circuit board or to a test system error, the test programmer can correct by either inserting a –2-Ω correction factor for resistance measurements or increasing the test limits by 2 Ω. If the programmer doesn't correct this average measurement error, the 110-Ω resistor—which is the nominal 100-Ω resistor at its upper spec limit (110 Ω)—would fail because the tester would measure 112 Ω, and it would be a false failure.

You also need to consider another source of test uncertainty: the standard deviation of measurements that are inherent to test system hardware. You can't correct this type of error so you must learn how to live with it.

If you take many measurements of the same 110-Ω resistor and plot the results, you will get a curve like the one in **Figure 2**; this is known as the normal distribution curve. The averaged variance from the mean of the measurements is called standard deviation. To set test limits for the worst-possible measurement situation, you need to center the mean of the measurement error on the device tolerance limit—in this example, 110 Ω.

When centered on 110 Ω, the distribution curve shows that 50% of all measurements fall to the left of the 110-Ω mean; even though these readings are technically wrong, they still correctly indicate a good 100-Ω ±10% resistor, as long as the reading is 90 Ω or greater. But to the right of the mean, the same type of deviant values incorrectly indicate a bad resistor.

Statistics Can Help

You can correct this normal-distribution problem by setting proper test limits, and statistics can help. Statistics indicates that for a normal distribution of data, 34% of all measurements fall between the mean and +1 standard deviation, or σ. Adding the 50% of measurements that fall to the left of the mean, 84% of all measurements fall somewhere between 1 σ to the right of the mean and ∞ to the left of the mean. If you choose a 1-Ω σ for measurements of the 100-Ω resistor, you can then set the upper test limit to 111 Ω to partially account for the resistor's ±10% tolerance and the tester's uncertainty.

For the given 100 Ω ±10% resistor, the tester will return a value between 0 and 111 Ω 84% of the time; this means 16 out of 100 measurements will be above 111 Ω and will falsely fail the good resistor. If you move the upper test limit to 2 σ (112 Ω), 97.5% of all measurements will fall within the test limit; now 25 out of 1000 measurements fall above the test limit and falsely fail the good resistor. **Table 1** shows the false failure rates for different σ values. You can see the same information graphically in Figure 2.

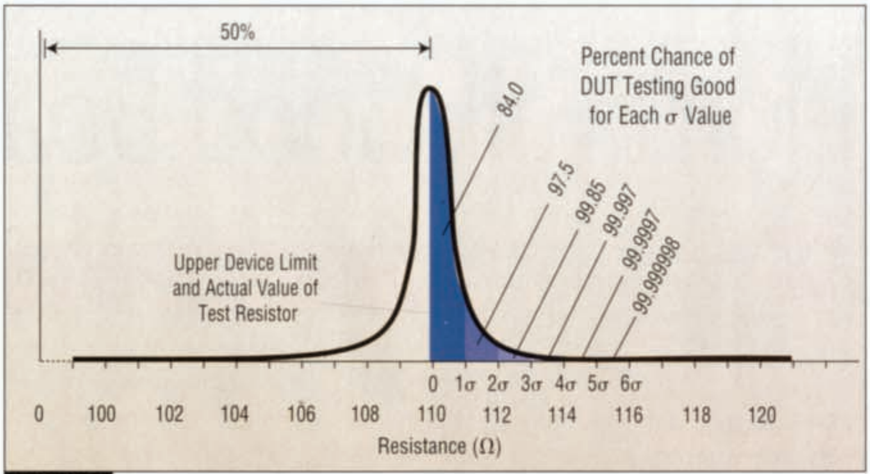

FIGURE 2. If you apply a tester's 1-Ω σ (standard deviation) for the subject resistor to the normal distribution curve, you can see how tester uncertainty affects false failures and how application of different σ multiples helps to eliminate those failures.

The tradeoff of false failures (failing a good component) against fault coverage (passing a bad component) can be a delicate balancing act. The standard method for eliminating false failures is to widen test limits. But if you do that, you reduce fault coverage with no immediately obvious effect, since bad components that pass a test are not visible until subsequent stages of test or until the product returns for warranty repair. Whether you set test limits statistically during test development to *avoid* false failures, or empirically over the course of production testing to *remove* them, you reduce fault coverage to some degree.

Many manufacturers have adopted 6 σ as their quality standard because it seems to be a good compromise between adequate fault coverage and excessive false failures. Applying the 6-σ standard to this example (the 100-Ω resistor and the 1-Ω σ for the tester), the closest you can set your test limits without false failures is ±6 Ω beyond the device tolerance limits. That translates to a ±16% tolerance for the 100-Ω ±10% resistor, giving an upper limit of 116 Ω.

If another tester had a σ of only 0.1 Ω, the 6-σ upper test limit would be 110.6 Ω. This gives the same false-failure rate (20 errors in 1 billion measurements) as the previous tester with a 1-Ω σ. But the fault coverage for the 0.1-Ω-σ tester is much better because you can set the test limits much closer to the device limits. **Figure 3** shows how much your tester's σ can affect test limits.

Table 1. Test results for different values of σ.

σ	False Failures/# of Measurements	% of Devices Testing Good
1	16/100	84.0
2	25/1000	97.5
3	15/10,000	99.85
4	3/100,000	99.997
5	3/1,000,000	99.9997
6	20/1,000,000,000	99.999998

Unfortunately, setting 6-σ test limits is like trying to hit a moving target, because the *random* variables are always varying. These random variables include:

- variations caused by the test system hardware design;
- test system aging;
- temperature and humidity (if your tester is in an uncontrolled environment);
- variation in the values of the components on the board;
- effects of no-clean manufacturing processes;
- inconsistent probing of test pads;
- oxidation of test fixture probes; and
- variations caused by test fixture maintenance.

You can model all of these variations mathematically, but your results will be inaccurate because you can't predict the variations over time. For example, when you set test limits to 6 σ based on samples taken over a very short time (such as during test development), those test limits may be only 1% of the value of limits computed with samples taken over longer periods of time.

This means the test limits you create during test development must be set *much higher* than 6 σ to achieve what are effectively 6-σ results on the production floor—it's the only practical method you can use to account for random variables over the long term. In fact, an automatic test program generator (ATPG) may set test limits of 100 σ to 200 σ, because it tries to eliminate false failures at all costs. This causes test engineers extra work if they must "pull in" test limits when debugging the test program.

Some companies set test limits to be between 20 σ and 100 σ before

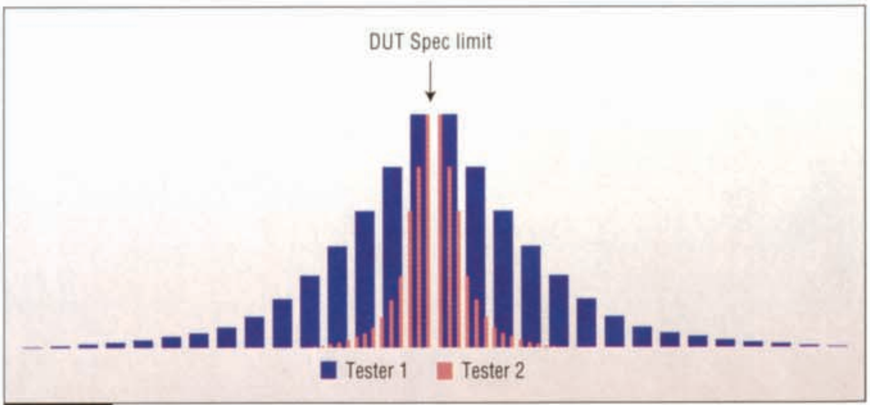

FIGURE 3. The better a tester's σ, the better its test coverage—its ability to find out-of-spec components as opposed to just gross problems such as wrong or missing components. If you compare the 6-σ test limits of Tester 1 and Tester 2, you can see that Tester 2 will produce more useful data than Tester 1 because its test limits are closer to the DUT's spec limits.

they release test programs to production. Although these limits seem very high, they appear to provide the same degree of stable testing in long-term production as 6-σ limits do in the short term.

If you applied this test coverage approach to the 100-Ω ±10% resistor example, making the standard deviation 1 Ω and the test limits 50 σ, you would test the resistor to a ±46% tolerance! About all you can tell from this gross a test is whether or not there is a resistor present, or perhaps if there is a short, an open, or an incorrect resistor at that location. You could probably speed up testing without reducing fault coverage by eliminating such a test altogether. If the tester's σ were 0.1 Ω, you would be testing to ±10.5%. This test *could* detect grossly out-of-tolerance resistors and would probably be acceptable for in-circuit test.

Find Your Tester's σ

Before you can start testing boards in production, though, you must determine your tester's σ for each component; then you'll know how accurate your board tests *really* are.

To determine your tester's σ, you need to perform a series of tests that are not influenced by any external variables. Test only one board to eliminate variations between UUTs. At the same time, make certain that you account for any variations in the tester itself. One way to do this is by cycling the test fixture between sets of measurements. Capturing every measurement you make is important; whether the devices pass or fail doesn't matter.

Make as many measurements as possible, but limit the measurements to a "small-enough" number of samples so the tester's repeated probings don't damage the board. Between 20 and 100 sets of measurements should be adequate to create a small-sample-size distribution.

Assign a test descriptor to each measurement and store all measurements in a datalogging file. Copy the logged measurements into a spreadsheet and calculate the standard deviation of each group of measurements (each device you've measured). Multiply each standard deviation by a constant of your choice, preferable between 20 and 100. You now have a measurement-variance-error value for this tester that you can add to the test limits for each device you've characterized to establish a production test that approximates 6 σ.

To compute the tester's sigma, subtract the device's tolerance limits from the measurement-variance limits and divide the result by 2.

Then divide the result by the standard deviation that you calculated for that device's data in the sampling experiment. The result is the σ of the tester for this device component. If, for example, we arbitrarily chose test limits for the 100-Ω resistor of 85.6 Ω and 114.2 Ω, while its spec limits are 90 Ω and 110 Ω, the measurement variance error is:

$$((114.2\ \Omega - 85.6\ \Omega) - (110\ \Omega - 90\ \Omega)) / 2 = (28.6\ \Omega - 20\ \Omega) / 2 = 4.3\ \Omega$$

If you calculated the standard deviation of the 100-Ω ±10% resistor used for your many measurements to be 0.56 Ω, the σ of your tester for the resistor is:

$$4.3\ \Omega / 0.56\ \Omega = 7.68\ \sigma$$

Now that you know how to trade off fault coverage against false failures, you need to minimize the standard deviation of your measurements if your test system's accuracy is unacceptable. Make sure your fixtures undergo the proper routine maintenance. Verify that your test setups are correct for each measurement. Sometimes you can reduce measurement error by changing how you make the measurement.

When you have eliminated all possible sources of error in your tester, you are left with the noise floor of your test hardware. If *that* is too high to give you the measurement error you need, the only solution is to buy quieter hardware. *T&MW*

FOR FURTHER READING

Scheiber, Stephen, *A Six-Step Economic-Justification Process for Tester Selection*, Quale Press, Haydenville, MA, 1997.

Ted T. Turner *is an engineer at Hewlett-Packard's Manufacturing Test Division. He has been working in the test industry for more than 15 years and has experience in manufacturing, applications development, R&D, and marketing.*

For quantity reprints of this article, contact Carol Bien at 847-390-2235.