One Measurement

Problem: build a good resistive termination.

Dr. Howard Johnson

signal consulting inc

This short luncheon talk was delivered to a general engineering audience at Lockheed in 2013.

Measurements define the body of knowledge we call Signal Integrity. Master the technique of making proper measurements and you will become a guru of the art. Work on one type of measurement every quarter, or every year, until you fully grasp the relationship between circuit theory, simulation, and measurement.

Every careful measurement harbors subtle difficulties. All reveal sparkling gems of insight.

As an example, I'll relate to you one of my first experiences with a high-speed measurement: making a good signal termination.

Selecting Good Values

2

signal consulting inc

In 1983, my mentor, Martin Graham, PhD, had me build a Wheatstone bridge for measuring the common-mode impedance of certain twisted-pair cables, now known generically as Category 3 UTP (unshielded twisted pair). My setup required some closely matched carbon-composition resistors with values accurate to within ±½%.

The measurements would be taken at fairly high frequencies covering the 10- to 100-MHz range. At such frequencies, even the parasitic series inductances and shunt capacitances of the resistors must match. That requirement rules out the use of trimming potentiometers to meet the stringent accuracy requirement.

I needed a few bull's eye, hit-the-spot, on-the-money, perfect resistors. Looking at the lab stock available on that day, I found no high-precision carbon-composition resistors. There were some 2% metal-film resistors in stock, but I knew that the manufacturers of those parts sometimes etched serpentine patterns in the metal film to elongate the resistor, thus increasing the parasitic series inductance to levels unacceptable for my application. Carbon-composition resistors are made in a simple cylindrical shape that is ideal for high-frequency use.

The only carbon-composition resistors I found had a 10% tolerance. I decided that they might work if I hand-selected values good enough for my purpose. I reasoned that out of 100 parts rated at 10% tolerance, about 10 should fall within 1% of the advertised value and that even more would do so if the distribution were centrally clumped. From those 10 parts, I hoped to select a couple of pairs suitable for my setup.

Imagine my surprise when, after an hour of labor and after checking 300 resistors, I found that none-absolutely none-fell within my 1% initial selection window.

Distribution of 10% Resistors

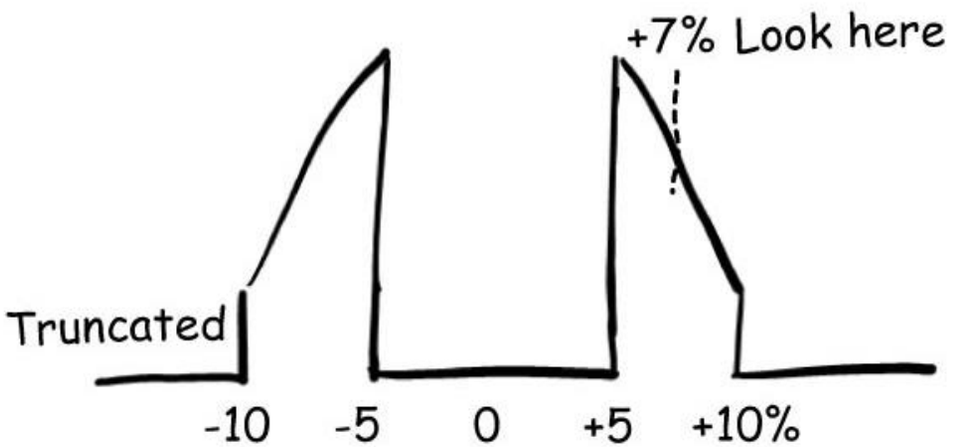

Use values that lie 7% away from the nominal standard values, either higher or lower, and you'll find plenty of those in the bins when you do your hand selection.

"Seven-Percent Solution," EDN, June 10, 2010

Mathematically, if the component values were distributed evenly across the whole tolerance band of ±10%, the probability of any one resistor's falling within a ±1% selection window should be one out of 10, or 0.1. The probability of any one resistor's falling outside the selection window then equals the complement of that value, 0.9. If you repeat the experiment 300 times, the probability of all the resistors' falling outside the selection window equals $(0.9)^{300}=1.8\times10^{-14}$.

It seemed to me inconceivable that such a low-probability event could ever actually occur in my sample of 300 parts. Monkeys striking random keys on a typewriter could compose a sonnet in the mean time between experimental failures of that magnitude.

Perplexed, I sought guidance from Martin. I found him in the company cafeteria enjoying a meatloaf sandwich. On a ketchup-stained napkin, he patiently drew this odd-looking curve.

The drawing complete, Martin said, "A 10% carbon-composition resistor is made in a somewhat slipshod manner. The manufacturer tries to get it right, but some of the variables are just too difficult to control. They make up a batch, test them all, and then throw away the bad ones. What's left is a distribution of values truncated on either side at the ±10% limits. The other main feature of the distribution is the big gap-toothed section in the middle. That's where they pulled out all the good parts and sold them at a higher price with a ±5% tolerance. How else do you think they make 5% resistors?“

My jaw hit the floor when I grasped how perfectly his explanation matched my results. He paused and then passed along another point of wisdom: "Design your circuit to use values that lie 7% away from the nominal standard values, either higher or lower, and you'll find plenty of those in the bins when you do your hand selection."

Try Measuring Nothing

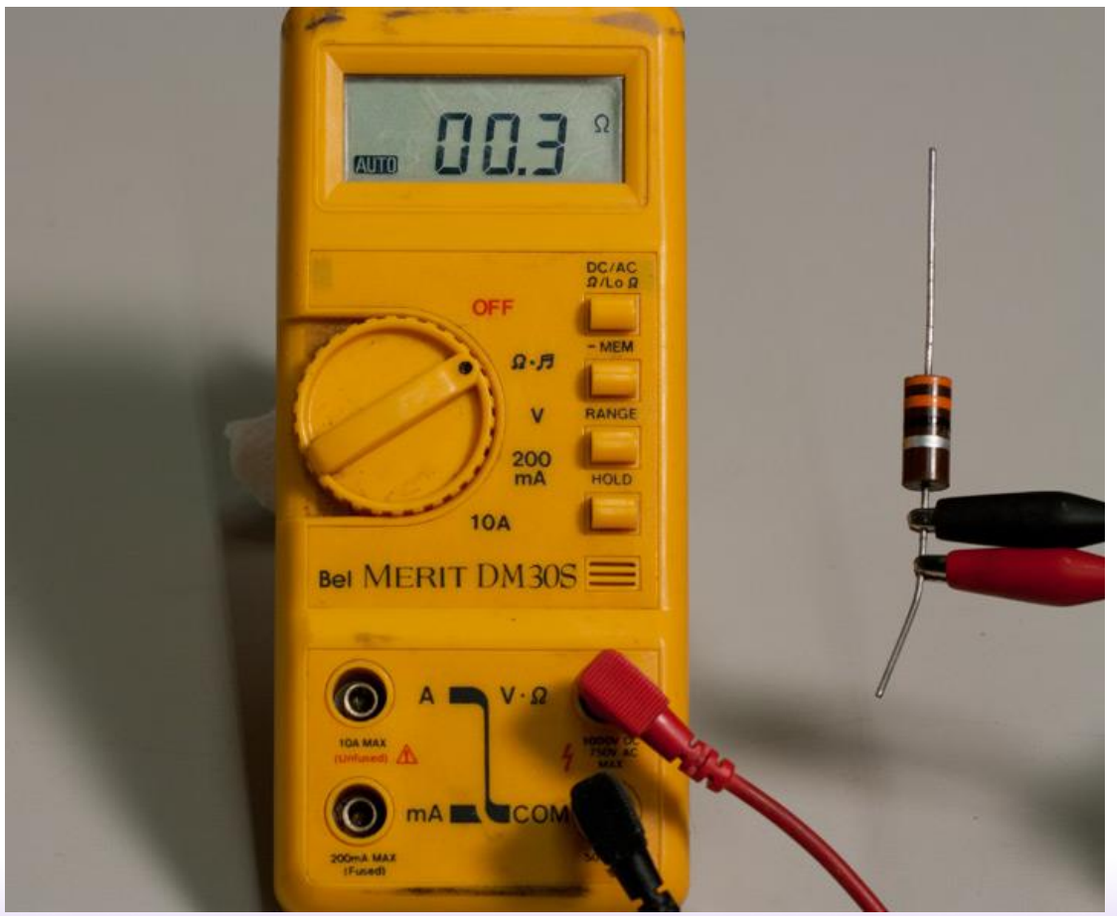

4

OK, so I re-designed my experiment so that it required accurate *35-ohm* resistors and started looking for a set of such resistors.

Martin noticed my work, and came over to the bench. He placed the measurement probes *both on the same side* of the resistor. The result should, naturally, show ZERO, but it didn't. It never does.

Every probe picks up extraneous noise and bias. Some of that noise is self-generated, and some may be generated by the system under test. In my case, the test leads, the banana jacks they plus into, and the clip-on connections to the resistor leads all had a certain amount of intrinsic series resistance. Even if the meter were perfect, it would have to report the sum of those little values, and you cannot tell from a single measurement how much is the "real" measurement and how much is the intrinsic bias.

There is only one way to resolve this problem, and that way, if you embrace it, leads to remarkable insights about noise, grounding, and the nature of digital systems.

The only way to directly observe the noise and bias in your measurement system is to attempt to measure nothing. That's what Martin did. His configuration is called a null experiment. Ideally, you should see zero, zip, nada, or, as the English call it, "naught."

What you actually observe is your own noise floor, a plethora of noise sources, a whole ecosystem of interferences all superimposed. In an RF measurement, creative use of your trigger circuits combined with vertical averaging can often pull apart these tiny effects, deeply buried in a sea of foam, for close inspection. You can learn a great deal measuring nothing.

In my case, the offset from bias factors amounted to about 0.3 ohms, roughly 1% of the value I was trying to measure. Worse yet, wiggle any of the connections and the result changes.

In order to characterize parts down to ½% and below, I needed a much better test setup.

Use Appropriate Equipment

Resolves to 0.1 nH. Operates at 1 MHz.

Enter the HP 4271B Digital LCR Meter.

This instrument uses a 4-terminal measuring process, which involves injecting a known current through a device under test using two leads and then, separately, with two other leads, measuring the voltage across the device under test.

We don't have time to go into all the details, but the bottom line is this: it can resolve resistances down to 0.001 ohm, *including* the inaccuracies contributed by the probes.

This is an adequate instrument for doing my part selection.

Notice something else about the picture. The instrument reports a value of parasitic series INDUCTANCE at the same time it reports RESISTANCE. The measurement is being made at a frequency of 1 MHz.

Consider Parasitic Effects

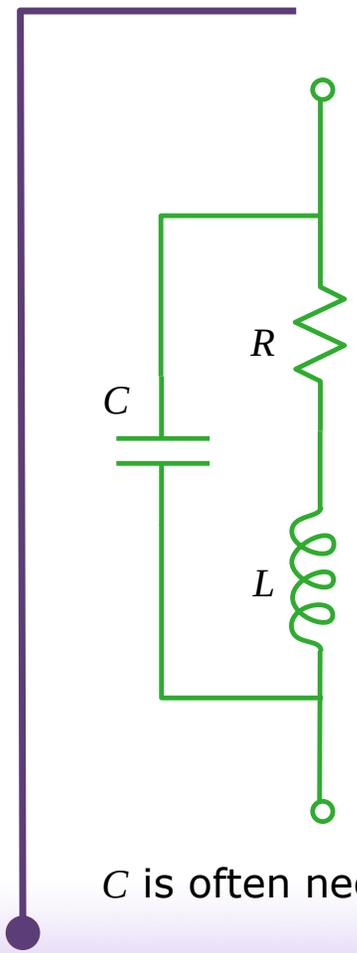

Exaggerate the effects you wish to measure.

$$Z(\omega) = R + j\omega L$$

Measure at HF to see this effect

Use a tiny value of R so it won't affect your result

C is often negligible in low-impedance applications.

All component suffer from series inductance. In the case of a terminating component, this inductance can affect its performance.

I was planning to use my resistive components at frequencies up to 100 MHz, but my test instrument only works at 1 MHz.

Once you think about that, you may realize that, because inductive effects scale in proportion to frequency, the net effect of inductance in my actual circuit (as 100MHz) would be 100x greater than the effect of inductance at 1 MHz.

Things that could cause difficulty at high speeds may be invisible at lower speeds. If you can't make your measurement at the intended speed of operation, then you're going to have to do something to **exaggerate the effect you want to see** at the frequency where you want to measure it.

So, here's what I did (another of Marty's suggestions). I measured the inductance of a 1-ohm resistor.

In this new measurement, the (jwL) term now looms much larger than the resistance (R=1 ohm), making the inductive effect easy to see, and easy to accurately measure.

Question: is the inductance of a 1-ohm resistor (which I can see and test at 1 MHz) the *SAME* as the inductance of a 33-ohm resistor?

Understand Your Components

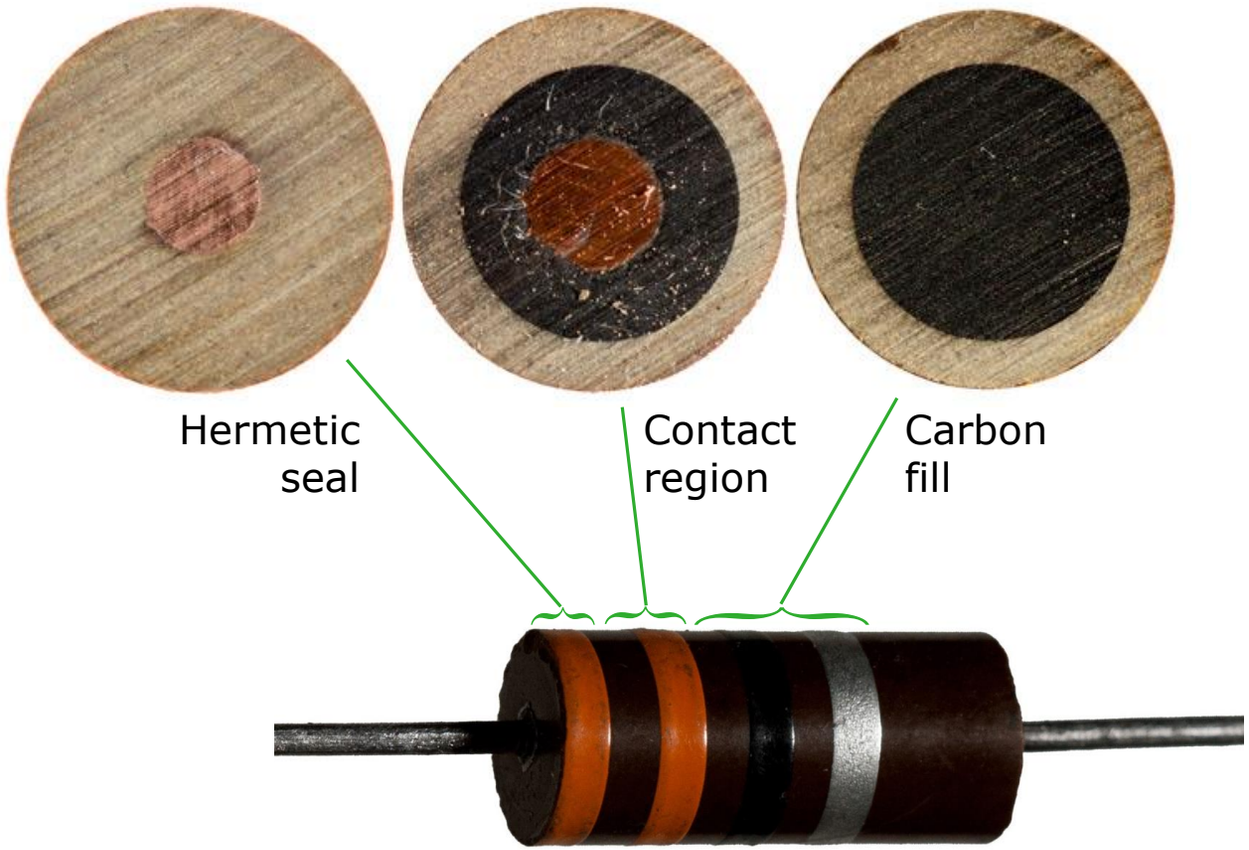

7

signal consulting inc

Inductance is a property of where the current flows within the body of a component.

In my case I was using carbon-composition resistors which have a particularly simple path: current spreads out into the carbon plug that fills the center of the component and flows straight down to the end.

Cutting open a few components verified that the construction of the 33-ohm resistor and the 1-ohm resistor were the same. They just used carbon materials with different resistivities.

I could therefore assume that the inductance measured on the 1-ohm sample would reasonably represent the inductance of my 33-ohm (hopefully 35-ohm) samples.

Understanding your components matters.

Had I used metal-film (MF) resistors, or wire-wound resistors, the constructions might have been different, because those types of resistors incorporate different serpentine patterns, in the case of MF resistors, or different number of turns of wire, in the case of wire-wounds, so the inductance measured at 1-ohm would **NOT** have represented the inductance of my 33-ohm samples.

I was beginning to understand why Marty had stipulated carbon-composition resistors in the first place.

An Circuit Model Is Not Your Goal

What you want, is a good terminator.

Finally, remember that a good, non-inductive 35-ohm resistor is note really our goal.

What we really want, is a good terminator.

Sometimes a part meets our design specification but still doesn't work in the actual circuit because of some flaw in our reasoning.

We have to consider ourselves as part of the measurement problem, and design ways to make tests that overcome our own intrinsic biases.

The best way to do that, is the make tests in an environment that, as closely as possible, resembles the actual working environment of your system.

Measure In-Situ

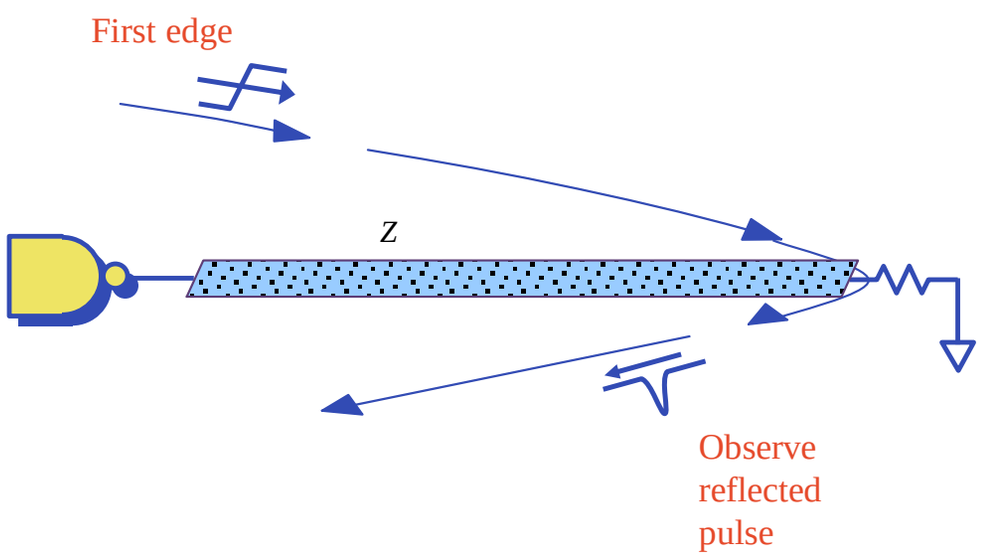

The more closely your test setup resembles your actual application, the more meaningful the test results.

In the case of a terminating component, that may mean building a transmission line, applying the component to the end of the line, stimulating the front end, and measuring the reflection that returns.

This measurement takes into account not only parasitic effects within the terminating resistor itself, but also the layout of the pads and vias used to connect that resistor to the transmission line.

Never Make Just One Measurement

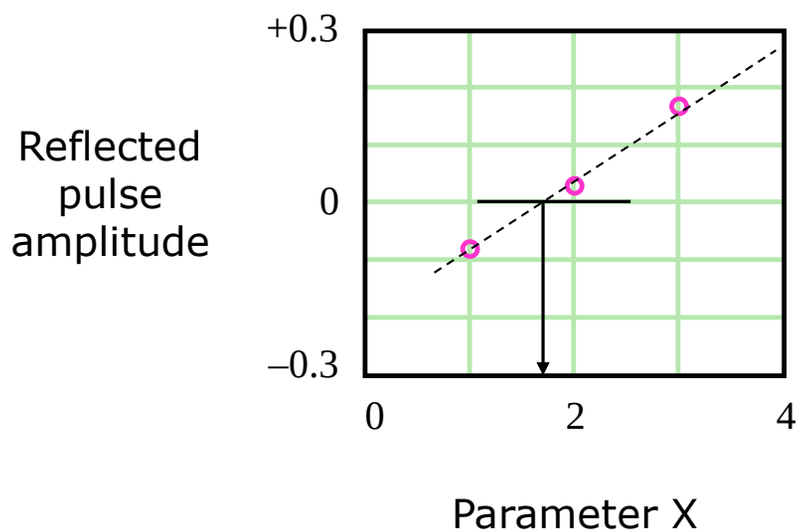

Measure what you need to know.

Ignore other factors.

If you are going to bother to make a test setup, I can suggest that you include in the setup some parameter that you can adjust that you know will affect the outcome.

For example, if I am concerned about the inductance contributed by a via used to ground the terminating resistor to the underlying reference plane, I might lay out three test circuits: one with a normal via, one with a skinny via, and one with a larger, fatter via.

In testing I can compare the results of the three tests and should get an outcome something like this chart.

If the vertical axis represents the size of the inductive pulse returned in response to a step edge, ...

and the X-axis represents a numerical parameter (via hole diameters, in this case) ,...

and you wanted to minimize that amplitude of that spike,...

you now have the test data necessary to make an informed judgement about the via hole diameter required to do the job (1.7 units).

The excess inductance is coming from at least three place: the layout where the transmission line touches the resistor, the resistor itself, and the way the resistor is grounded. This measurement in no way separates those three effects – they are all jammed together – but what happens is, you learn what you have to do to overcome the effect.

That's all you need to do effective engineering.

If you are just starting out in an engineering field, I hope you find a mentor like Marty who can show you the ropes.

Dr. Howard Johnson

Dr. Johnson no longer teaches seminars but he hopes these brief comments are useful and wishes you the best of success in your next project.

www.signalintegrity.com

signal consulting inc