8 Hints for Making Better RF Counter Measurements

900,000,050 Hz

4,500,925,000 Hz

65,733,860

MHz 215,73

GHz

Agilent Technologies

The not-so-simple world of RF counters

At first glance, an electronic counter seems like a fairly simple instrument. You connect a signal to the input, and a digital readout tells you the frequency or some other parameter.

While we'd like to say that counters are a plug-in-and-go sort of instrument, the technical realities of quartz-crystal time-bases demand some care and attention on your part. The counter, its timebase options, your signal, and the way you set up the measurement all affect the quality of the results. In fact, so many factors play into the equation that the accuracy specification for a good counter is literally that—an equation.

The good news is that a little attention goes a long way toward better measurements, and these eight hints are a great place to start.

Selecting the right counter for the job

Selecting a counter is the first bit of confusion to clear up, since an array of related products perform a variety of counting jobs at various frequencies:

- **Universal counters.** Both frequency and time interval measurements, as well as a number of related parameters.

- **RF frequency counters.** Precise frequency measurements, up to 3 GHz and beyond.

- **Microwave frequency counters.** Precise frequency measurements, up to 40 GHz and beyond.

- **Time interval analyzers.** Optimized for precision time interval measurements.

- **Modulation domain analyzers.** Designed to show modulation quantities, such as frequency versus time, phase versus time, and time interval versus time.

This brochure focuses on RF frequency counters. For information on other counters and signal analyzers, please contact Agilent, your local Agilent Technologies sales office, or the Agilent website at **www.agilent.com/find/gp**

Contents

Hint 1: Understand the effects of counter architecture

Hint 2: Recognize the difference between resolution and accuracy

Hint 3: Schedule calibration to match performance demands

Hint 4: Choose the most appropriate timebase

Hint 5: Adjust sensitivity to avoid noise triggering

Hint 6: Configure your counter for low- frequency measurements

Hint 7: Smooth out jumpy displays

Hint 8: Use statistics to characterize signals

A hint about reading these hints

We suggest you read these hints in order from Hint 1 to Hint 8 because they start with the most general concepts and proceed to increasingly specific techniques.

Understand the effects of counter architecture

RF counters fall into two basic categories: direct counting and reciprocal counting. Understanding the effects of the two difference approaches will help you choose the correct counter and use it correctly.

Direct counters simply count the number of times the input signal crosses zero during a specific gate time. The resulting count is sent directly to the counter's readout for display.

This method is simple and inexpensive, but it means that the direct counter's resolution is fixed in Hertz. For example, with a 1 second gate time, the lowest frequency the counter can detect is 1 Hz (since 1 zero crossing in 1 second is 1 Hz, by definition). Thus, if you are measuring a 10 Hz signal, the best resolution you can expect for a 1 second gate time is 1 Hz, or 2 digits in the display. For a 1 kHz signal and a 1 second gate, you get 4 digits. For a 100 kHz signal, 6 digits, and so on. Figure 1 illustrates this relationship. It is also interesting to note that a direct counter's gate times are selectable only as multiples and sub-multiples of 1 second, which could limit your measurement flexibility.

Reciprocal counters, in contrast, measure the input signal's period, then reciprocate it to get frequency. Thanks to the measurement architecture involved, the resulting resolution is fixed in the number of digits displayed (not Hertz) for a given gate time. In other words, a reciprocal counter will always display the same number of digits of resolution regardless of the input frequency (Figure 2). Note that you'll see the resolution of a reciprocal counter specified in terms of the number of digits for a particular gate time, such as "10 digits per second."

Figure 1. The number of digits displayed by a direct counter versus frequency (for a 1 second gate time).

Figure 2. For the same gate time, a reciprocal counter's resolution is constant for all frequencies (the counter shown has a resolution spec of 10 digits/second).

Using the counter industry's benchmark of a 1 second gate time, Figure 3 compares the resolution of direct and reciprocal counters. You can see that the reciprocal counter has a considerable advantage over the direct counter in the lower frequencies. As an example, at 1 kHz, a direct counter gives a resolution of 1 Hz (4 digits). A 10 digit/second reciprocal counter gives a resolution of 1 µHz (10 digits).

Even if you don't need micro-Hertz resolution, the reciprocal counter still offers a significant speed advantage: the reciprocal counter will give 1 mHz resolution in 1 ms, while a direct counter needs a full second to give you just 1 Hz resolution (Figure 4). Reciprocal counters also offer continuously adjustable gate times (not just decade steps), so you can get the resolution you need in the minimum amount of time.

The choice comes down to cost versus performance. If your resolution requirements are flexible and you aren't too concerned with speed, a direct counter can be an economical choice. For fast, high-resolution measurements, though, a reciprocal counter is the way to go.

You can figure out whether a counter is direct or reciprocal simply by looking at the frequency resolution specification. If it specifies resolution in Hertz, it's a direct counter. If it specifies resolution in digits-per-second, it's a reciprocal counter.

Figure 3. Comparing resolution for direct and reciprocal counters (for a 1 second gate time).

Figure 4. Here are the gate times needed to yield various resolutions with a 10 digits/second reciprocal counter.

Recognize the difference between resolution and accuracy

Wow, look at all those digits. Must be an incredibly accurate measurement.

Well, not necessarily. Equating resolution with accuracy is a common mistake They are related, but they are distinctly different concepts.

Resolution can be defined as the counter's ability to distinguish closely spaced frequencies (Figure 5). All other things being equal (such as measurement time and product cost), more digits are better—but the digits you see on the display need to be supported by accuracy. Digits can be deceptive when other errors push the counter's resolving ability away from the actual frequency. In other words, it's possible for a counter to give you a very precise reading of an incorrect frequency.

True measurement accuracy is a function of both random and systematic errors. Random errors, which are the source of these resolution uncertainties, include quantization error (the uncertainty surrounding the final count in the gate time window), trigger error (such as triggering on noise spikes), and short-term instabilities in the timebase. Systematic errors are biases in the measurement system that push its readings away from the actual frequency of the signal.

This group includes effects on the timebase crystal, such as aging, temperature and line voltage variations.

Compare the two counters in Figure 6. Counter A has good resolution but a serious bias error, so its displayed result in most cases will be less accurate than those of Counter B, which has poorer resolution but a smaller systematic bias error.

To paraphrase the mathematician John Tukey, it's far better to have an approximate answer to the right question than a precise answer to the wrong question. As you'll see in Hint #3, this is why calibrating your counter to match your accuracy needs is so important.

Figure 5. Resolution is the ability to distinguish closely spaced frequencies. However, resolution by itself is no guarantee of accuracy.

Figure 6. Simplified view of resolution vs. accuracy. Systematic errors related to the timebase "push" the displayed frequency away from the actual frequency. The random errors create a range of frequencies inside of which the counter can't distinguish different signals.

Schedule calibration to match performance demands

How often do you need to calibrate your counter?

The most precise, accurate answer is "it depends." It depends on what kind of timebase you have in your counter, the conditions the counter will be subjected to during the measurement, and, most importantly, how much accuracy you need from the measurement.

To understand why calibration is not a simple issue, you have to step back and consider the nature of counter accuracy in general. The quality of the answer you see on the display depends on four groups of factors:

1. Time-invariant performance factors related to the counter itself, such as the temperature stability of the counter's timebase (more on this in Hint 4)

2. Time-variant performance factors related to the counter itself, such as the aging rate of the crystal timebase

3. The nature of your signal, such as the presence of noise

4. Setup choices you make, such as gate time

This list explains why, when you try to look up the accuracy specification in the counter's data sheet, you see a rather complicated formula instead of a single, simple number. Let's focus on item #2, which is where calibration comes into play.

In spite of the fact that counters are electronic instruments measuring electrical signals, the quartz crystal that is the heart of every counter's timebase is essentially a mechanical device. As a mechanical device, the crystal is vulnerable to physical disturbances that can change the frequency at which it vibrates, which in turn affects the counter's accuracy. The cumulative effect of these various disturbances is known as *crystal aging*, and it is this aging that you are compensating for when you calibrate the counter (Figure 7).

From a user's perspective, the details aren't terribly important, but it is definitely important to know that crystal aging occurs. The good news is that aging is fairly easy to predict and even easier to compensate for through calibration.

Now that the role of crystal aging is clear, you can find out if you need to calibrate by looking at the aging rate specification in your counter's data sheet. Let's say the rate is 4×10^{-8} per day. If it's been 300 days since calibration, aging will add a timebase error of 1.2×10^{-5} into the overall accuracy calculation. If this uncertainty ($\pm$ 12 Hz on a 1 MHz signal), plus the other uncertainties inherent in counter measurements, is acceptable for your measurements, you don't need to calibrate. Otherwise, you need to calibrate. There, it wasn't that complicated after all (but keep in mind that aging is only one of several factors that affect overall accuracy).

Figure 7. Calibration compensates for the effects of crystal aging, which affects the crystal's long-term stability. In this example, the aging rate is 4×10^{-8} per day. (The short-term deviations are the result of phase noise and other factors.)

Choose the most appropriate timebase

Measurement accuracy in frequency counters begins with the timebase because it establishes the reference against which your input signal is measured. The better the timebase, the better your measurements can be. (Notice the "can be" part here; you still have to calibrate and take care of your counter to maximize performance.)

The frequency at which quartz crystals vibrate is heavily influenced by ambient temperature, and time-base technologies fall into three categories based on the way they address this thermal behavior:

• **Standard.** A standard or "room temperature" timebase doesn't employ any kind of temperature compensation or control. While this has the advantage of being inexpensive, it also allows the largest frequency errors. The curve in Figure 8 shows the thermal behavior of a typical crystal. As the ambient temperature varies, the frequency output can change by 5 parts per million (ppm) or more. This works out to ± 5 Hz on a 1 MHz signal, so it can be a significant factor in your measurements.

• **Temperature-compensated.** One way to deal with the crystal's thermal variation is to make sure that the other electronic components in the oscillator circuits have complementary thermal responses. This approach can stabilize the thermal behavior enough to reduce timebase errors to around 1 ppm (± 1 Hz on a 1 MHz signal).

• **Oven-controlled.** The most effective way to stabilize the oscillator output is to simply take the crystal off this thermal roller coaster. Counter designers do this by isolating the crystal in an oven that holds its temperature at a specific point in the thermal response curve (Figure 8). The result is much better timebase stability, with typical errors as small as 0.0025 ppm (± 0.0025 Hz on a 1 MHz signal).

There is more to this story than just temperature-related accuracy, however. Oven-controlled timebases also help with the effects of crystal aging, which means you don't have to take your counter out of service for calibration as often. For example, the per-month aging rate of a standard Agilent 53181A RF counter is < 0.3 ppm (± 0.3 Hz on a 1 MHz signal). The optional high-stability oven reduces this to < 0.015 ppm (± 0.015 Hz on a 1 MHz signal) per month. In other words, the standard timebase ages 20 times faster than the high-stability model, and will therefore require calibration more frequently to maintain your required measurement accuracy (see Hint 3).

Frequency deviation (ppm)

Figure 8. The frequency output of an unprotected crystal can vary widely in response to ambient temperature. Putting the crystal in a controlled thermal environment (an oven) helps maintain a stable output frequency.

Adjust sensitivity to avoid noise triggering

Figure 9. The two small peaks (spurious signals in this case) generate unwanted triggers at point 1 and point 3 because the trigger band is set too narrow.

Figure 10. Lowering the trigger sensitivity by expanding the trigger band produces the desired count.

The good news is that high-quality counters are broadband instruments with sensitive input circuits. That's also the bad news.

To a counter, all signals basically look the same. Sine waves, square waves, harmonics, random noise—they all just look like a series of zero crossings as far as the counter is concerned. A counter figures out the signal's frequency by triggering on these zero crossings to measure frequency. If your signal is clean and uncluttered, the process works quite well. Noisy signals, however, can trick the counter into triggering on spurious zero crossings. When this happens, the counter doesn't count what you think it's counting.

Fortunately, all good counters offer a way around this problem. First, they require the signal to pass through both lower and upper hysteresis thresholds before they register a zero crossing. The gap between these two levels is called *trigger sensitivity*, the *hysteresis band*, the *trigger band* or something similar.

Second, good counters let you adjust this band to minimize unwanted triggering. Figure 9 shows a signal with some spurious components that are causing trouble with the count. The trigger band is fairly narrow, so both the unwanted noise (at points 1 and 3) and the real signal (at points 2 and 4) cause the counter to trigger. What is really just two cycles of the signal get counted as four.

By adjusting the trigger band to make the counter less sensitive, you can avoid these spurious triggers. In Figure 10, the trigger band is wide enough (which is to say the sensitivity is low enough) that the spurs don't get counted as zero crossings. The counter registers two valid zero crossings and goes on to compute the appropriate frequency.

If you think your signal might have some noise problems, try switching your counter into low-sensitivity mode. If the displayed frequency changes, chances are you were triggering on noise.

Configure your counter
for low-frequency measurements

Hint 5 discussed the problem of triggering on unwanted components in your signal. This problem can be even more acute with low-frequency signals (roughly 100 Hz and below), since the chance of spurious triggering on irrelevant high-frequency components is that much greater. In addition, the signal's slew rate affects trigger accuracy; the lower the slew rate, the more chance there is for error.

Here are three quick steps you can take to help improve the quality of counter measurements on low-frequency signals.

Invoke the counter's low-pass filter

Assuming your counter has one, use the low-pass filter. The Agilent 53181A RF counter, for instance, has a 100 kHz low-pass filter that you can switch into the signal path. This reduces the chance of triggering on harmonics and high-frequency noise.

Use manual triggering

When a counter is set to use auto triggering, it estimates the peak-to-peak level of the signal and computes the midpoint to establish a trigger level. This approach generally leads to good results but can cause trouble on low-frequency signals. The problem is that the auto trigger algorithm can take less time than the signal takes to transition between its minimum and maximum values. As a result, the auto trigger can wind up following the signal level up and down, rather than setting a single trigger level based on a consistent estimate of the minimum and maximum values. The solution is to turn off auto trigger and set the trigger level manually.

Use dc coupling

Many counters offer a choice between ac and dc coupling on their primary input channel. This setting works the same way on a counter as it does on an oscilloscope: ac coupling removes any dc offset from the signal, whereas dc coupling admits the entire signal, offset and all. The trouble with ac coupling is that it also attenuates lower frequencies. In fact, your counter's performance probably isn't even specified below a certain frequency when ac coupling is used. To ensure better results all the way down to fractions of a Hertz, use dc coupling instead.

HINT
7

Smooth out jumpy displays

A jumpy display, where the last several digits fluctuate rapidly, can be a challenge if you're trying to adjust a circuit in real time or perform some other task based on the counter's display. Depending on your counter's capabilities, you have several options for taming these unruly digits.

• **Reduce the number of displayed digits.** The first option is simply to reduce the number of digits displayed using the "Fewer Digits" button or whatever that function is called on your counter. While this can quiet the display, it might hide information you need to make decisions about circuit behavior. Note this is strictly a display function that doesn't have any effect on the actual measurement.

• **Use limit testing.** Another possibility is to use limit testing with a visual indicator, if you just need to know whether the signal is within a certain band of frequencies. The catch here, of course, is that your counter needs to have this feature before you can use it.

• **Use signal averaging.** Averaging (labeled "mean" on many counters) is a good option to consider any time your signal is jumping around. Unlike simply reducing the number of displayed digits, of course, averaging actually improves the quality of your measurements. By reducing the effects of random variations in the signal, it reduces the number of display changes.

HINT
8

Use statistics to characterize signals

A good counter can do more for you than simply count. Built-in statistical functions help characterize signals and identify trends, two tasks you'd have to perform on a PC otherwise. A basic set of statistical functions includes arithmetic mean, minimum, maximum, and standard deviation.

For example, a wandering carrier signal is a great application for counter statistics. Min, max, mean and standard deviation can give you a clear picture of the carrier's behavior, without resorting to a spectrum analyzer.

Moreover, the counter will give you better frequency resolution than you could get with a spectrum analyzer.

To move beyond these basic statistics, consider one of the connectivity software packages such as Agilent IntuiLink, VEE, BenchLink Meter or National Instruments LabView® designed to post-process measurement data. Depending on the software you choose, you can easily add documentation and graphics, archiving, powerful analytical functions and test automation.

LabVIEW is a registered trademark of National Instruments Incorporated.

Agilent 53181A RF Counter

- Resolution of 10 digits per second; frequency ranges up to 12.4 GHz

- Built-in limit testing and statistics simplify your measurements

- IntuiLink software transfers data and images into your standard PC applications with little or no programming

- Choice of economical and high-precision timebases

The 53181A RF counter offers 10 digits/second resolution at up to 225 MHz on one channel, with an optional 1.5, 3, 5, or 12.4 GHz second channel.

The counter's reciprocal architecture delivers both high resolution regardless of input frequency and the flexibility to choose any gate time you want.

Built-in limit testing with a unique in-limit indicator simplifies real-time circuit adjustments.

Real-time digital signal processing boosts measurement throughput by analyzing data while simultaneously taking new readings. While other counters are stuck in processing "dead time," the 53181A has already moved on to the next measurement. With continuous GPIB data transfer rates of up to 200 measurements per second, and the standard IntuiLink connectivity software, you'll get the job done in a hurry.

Automated limit tests, instant recall of test setups and one-button access to the features you need most simplify your work. Plus, you can simultaneously measure and track average, min/max and standard deviation.

Agilent offers a complete line of counters, including the 53131A and 53132A universal counters and a wide selection of microwave frequency counters, time interval analyzers and modulation domain analyzers. The engineers at Agilent or your local sales office can provide more information. Or visit our website at **www.agilent.com/find/gp**

Agilent 53181A 225-MHz RF Counter

Measurements	Frequency, frequency ratio (with optional CH 2), period, peak voltage
Analysis	Automatic limit testing, math (scale and offset), statistics (min, max, mean, standard deviation)

Measurement characteristics

Frequency range (std.)	CH 1: dc –225 MHz
Frequency resolution	10 digits/s
Measurement speed	Up to 200 meas/s

Input conditioning

Impedance, coupling	1 MΩ or 50 Ω, ac or dc
Low pass filter	100 kHz, switchable
Attenuation	X1 or X10

External timebase reference input

	1, 5, 10 MHz

Trigger	Trigger on rising/falling edge; set level by percent of signal level or absolute voltage; set sensitivity to LOW, MED, or HIGH
Gating and arming	Auto, manual (set gate time or number of digits of resolution); external; delay
Interfaces	GPIB (IEEE 488.1 and 488.2) with SCPI; talk only RS-232
Power	100 to 120 Vac ± 10% 50, 60 or 400 Hz ± 10% 220 to 240 Vac ± 10% 50 or 60 Hz ± 10%
Net weight	3 kg (6.5 lbs)
Size (H x W x D)	103.6 x 254.4 x 374 mm (4.1 x 10.0 x 14.8 in)
Warranty	3 years

Ordering information

53181A	10-digits-per-second 225-MHz RF counter
34812A	BenchLink Meter software for Windows®
34131A	carrying case
34161A	accessory pouch
Opt. W50	additional 2-year warranty
Opt. 1CM	rack mount kit
Opt. 001	medium-stability timebase
Opt. 010	high-stability timebase
Opt. 015	1.5-GHz channel 2 with BNC connector
Opt. 030	3-GHz channel 2 with BNC connector
Opt. 050	5-GHz channel 2 with type-N connector
Opt. 060	rear terminals
Opt. 124	12.4-GHz channel 2 with type-N connector

Windows® is a U.S. registered trademark of Microsoft Corporation.

Agilent Technologies' Test and Measurement Support, Services, and Assistance
Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Support is available for at least five years beyond the production life of the product. Two concepts underlay Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise
Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you use Agilent equipment, we can verify that it works properly, help with product operation, and provide basic measurement assistance for the use of specified capabilities, at no extra cost upon request. Many self-help tools are available.

Your Advantage
Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.

Agilent Technologies