Swept Receiver Dynamic Range Testing in the ARRL Laboratory

Testing DR at spot frequencies is useful, but a more complete picture comes from swept DR tests.

By Ed Hare, KA1CV, 225 Main Street, Newington, CT 06111, email: ehare@arrl.org

The ARRL Laboratory is sometimes envisioned as a dark, mysterious place where RF gurus cause technical conclusions and results to appear at will. In actuality, those who have visited the Lab know that it is modern, well-lit and staffed with people who obtain results the same way every one else does — ingenuity, hard work, a few false starts and no small amount of good luck. (Of course, anyone who has seen the collection of microwave circuits and projects on Zack Lau's bench knows that there is a bit of magic left in the hallowed halls of ARRL Headquarters, but *QEX* has done what it can to unveil the mysteries through the bimonthly *QEX* "RF" column so aptly written by our beloved Senior Laboratory Engineer.) Of late, the Lab staff has taken to writing *QEX* articles as a way of documenting some of the work we do and showing our fellow hams and technical types some of the steps we found necessary to get from a gleam in a Laboratory Engineer's eye to completed work.

Over the years, we have all marveled at the changes in technology that affect almost our entire lives. Most of these changes are obvious, such as the differences in bells and whistles between my old Heathkit DX-100 and today's modern transceiver marvels. Many of the changes are less obvious, like those in the internal design and components that make our modern receivers much more sensitive and effective in the presence of strong, off-channel signals than even the best receivers used by hams a few decades ago. In the mid 1960s, when I first learned the fun of ham radio, it was common for HF receivers to need a preamplifier on the 15- and 10-meter bands to improve sensitivity. In all the articles I so avidly devoured in those wild days of my youth, I do not recall any discussion of things like inter-modulation, blocking or dynamic range. (Perhaps there were a few that I don't remember because I didn't understand them well enough for them to make an impression!)

Fortunately, technology has improved, and today's hams are more likely to have at least heard the term "dynamic range." Although not all hams can fully explain it, most understand that dynamic range is a measure of receiver performance that indicates the ability of a receiver to function well in the presence of strong signals. Over the years, one of the "hidden" technology improvements has been a steady increase in the dynamic range of receivers.

Receiver Testing Basics

This article is about the dynamic range of receivers. In particular, it is a discussion of how and why a receiver's dynamic range varies with the frequency spacing between the desired signal a receiver is tuned to and the off-channel signal or signals used to determine dynamic range. Although all of the receiver engineers among our readers understand the concepts pretty well (or think they do, anyway, speaking entirely for myself), let me outline a few necessary receiver testing concepts and terms for those aspiring to be receiver engineers.

First, expressed in simple terms, the dynamic range of a receiver is the difference between the weakest on-channel signal a receiver can hear and the strongest off-channel signal a receiver can tolerate without degradation of the received signal. Normally, dynamic range is expressed as a ratio in decibels between the measured noise floor of a receiver and the level of off-channel signal that causes a specific amount of receiver degradation.

Noise Floor and Receiver Sensitivity

Noise floor is a measurement of how much internal noise a receiver creates. The internal noise of a receiver sets a limit on the weakest signal a receiver can hear. Noise floor is sometimes called "minimum discernible signal" (MDS), although this is somewhat of a misnomer. Some skilled operators can detect by ear signals that are as weak as 10 dB or so below the noise floor.

One might assume that a "perfect" receiver is noiseless, but, alas, that is not true. At any temperature above absolute zero, even a lowly resistor is a source of *thermal noise*. At room temperature, a 50-Ω resistor is a noise source of –174 dBm/Hz, meaning it generates a noise signal that is 174 dB below 1 mW in a 1-Hz bandwidth. Of course, useful receivers almost always have a bandwidth greater than 1 Hz. Not surprisingly, each hertz of bandwidth adds another –174 dBm of noise.

So, a hypothetical perfect receiver with a 2-Hz bandwidth would have twice as much noise power, or 3 dB greater than –174 dBm, for a noise floor of –171 dBm; a 10 Hz bandwidth would have 10 times as much noise power, or –164 dBm, etc. A more typical

500-Hz bandwidth in our perfect receiver would produce a noise floor 27 dB greater than the hypothetical 1-Hz receiver. Its theoretical sensitivity is, therefore, –147 dBm.

Receiver sensitivity can be expressed in a number of ways. One can express, as we have done above, the input noise floor power of a receiver. We have expressed it in dBm, but it could be expressed in watts or volts and still be valid. It may be less useful, however, at least in the context of using the input sensitivity to calculate dynamic range.

All of the numbers I have used above are for our perfect receiver. In reality, all receivers add additional noise over and above the thermal noise. This is due to a number of effects within the components used to build the receiver. Many modern receivers approach within a few dB of theoretical; some are within a few tenths of a dB.

"ARRL Standard" Noise-Floor Test Conditions

In the ARRL Lab, most of our receiver-sensitivity measurements are made of the receiver's noise floor in the CW mode. To help facilitate the comparison of one receiver-sensitivity measurement to another, we perform most of our receiver tests using "ARRL standard" test conditions. In the case of CW sensitivity, we select the available receiver bandwidth closest to 500 Hz, and to facilitate calculation of dynamic range, we express the sensitivity in dBm.

A block diagram of the ARRL Lab's receiver noise-floor test setup is shown in **Fig 1**. The generator currently in use is a Marconi model 2041. It is capable of generating signals from 10 kHz to 2.7 GHz at levels of –144 dBm to +13 dBm. The step attenuator is nominally set to 10 dB to ensure that the generator always sees a reasonable 50-Ω load no matter what the actual input impedance of the receiver under test might be. The audio voltmeter is normally set to take a relative reading of the output of the receiver. Over the years, we have used several different meters to do this test, ranging from an HP-339A audio test set to a GPIB-controlled HP-3478A multimeter to our newest acquisition, a National Instruments AT-2150A16-bit digital acquisition (DAQ) card.

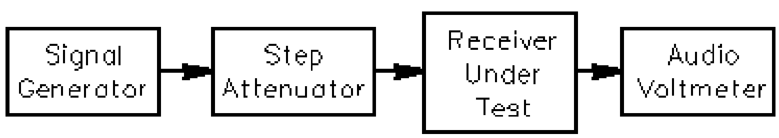

Fig 1—Test set-up for receiver noise-floor testing.

The noise-floor test is performed by setting the receiver to nominal operating conditions and a 500-Hz bandwidth. The generator is turned off and the receiver volume is set to a comfortable listening level. The audio meter is then used to measure the output level of the receiver's noise. The generator is then turned on, tuned to the same frequency as the receiver, and usually set to a level about 10 dB greater than the expected sensitivity of the receiver. The receiver tuning is adjusted for maximum receiver output, after which the generator amplitude is adjusted until the output meter level drops to a level that is 3 dB greater than the previously measured receiver noise. At this point the total output power of the receiver has doubled, indicating that the generator power and the noise power are equal. (Note that it is not necessary to measure the *absolute* level of the receiver output; we are looking for a *relative* receiver output level change.) The generator power can then be read from the generator, subtracting the 10 dB of the step attenuator, and logged in dBm as the receiver's noise-floor sensitivity.

Dynamic Range Measurement

I started this discussion about dynamic range by talking about receiver sensitivity because the dynamic range is based on a comparison of measured receiver sensitivity with measured receiver degradation in the presence of strong signals. There are actually two types of dynamic range, blocking dynamic range (BDR) and two-tone, third-order intermodulation-distortion dynamic range (TTTODR). Each type of dynamic range measures a specific type of receiver degradation. **Fig 2** shows the basic dynamic range test setup used in the ARRL Lab.

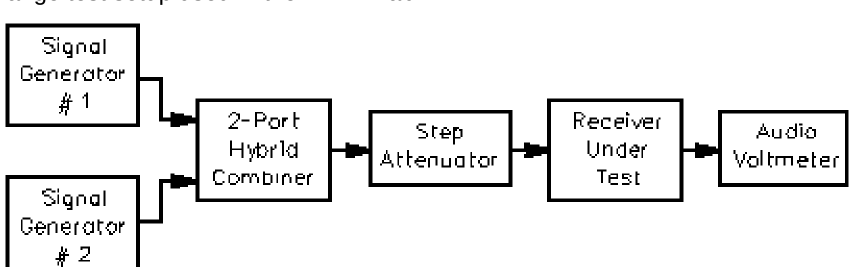

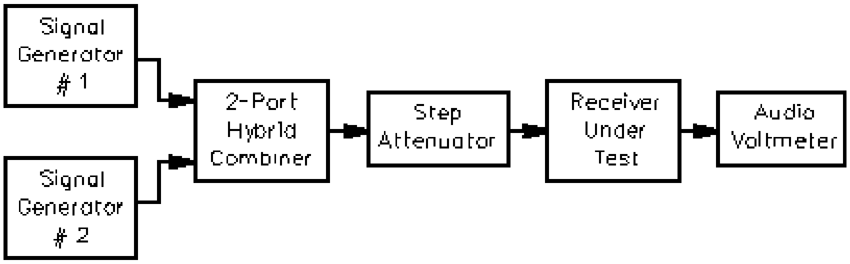

Fig 2—Test set-up for receiver dynamic-range testing.

Blocking Dynamic Range

Blocking dynamic range is a comparison between the noise floor of a receiver and the level of an off-channel signal that causes a specific measured reduction in the on-channel sensitivity. (This phenomenon is often termed "desense," especially in the VHF FM world.) The ARRL Lab uses the setup of **Fig 2** to manually test BDR. We use our Marconi generator for generator #1; for generator #2 we are using an older HP-8640B (if anyone wants to donate another Marconi 2041 generator to the ARRL Lab, call me here at HQ—collect!)

In this test, generator #1 is set to the desired frequency—the frequency the receiver is set to receive—usually standardized at 20 kHz above the lower edge of the amateur band being tested. The receiver under test is properly tuned to this test frequency and the receiver audio output is set to a comfortable level with the second generator turned off. The relative voltmeter is adjusted for a reading that is approximately midscale. Generator #2, which generates the *blocking* signal, is nominally set to a frequency that is 20 kHz higher in frequency than the test frequency, at a level of less than –100 dBm. (If a receiver were to block with an off-channel input signal of –100 dBm, we would initially set the generator to a lower level. If this occurred at 20-kHz spacing, this would be a particularly *bad* receiver.) The amplitude of generator #2 is then increased in 1 dB steps. When the output level of the receiver drops by 1 dB, the strong off-frequency signal is blocking the receiver. The difference between the noise-floor level and the level of signal that causes blocking is the blocking dynamic range (BDR). BDR is usually expressed in dB.

Signal Levels

The desired on-channel signal level must be within the linear range of the receiver. On the low end, it certainly must be higher than the noise floor. In fact, it should be *much* higher, to minimize the effects of receiver noise on the measurement. In practice, the desired-signal level should be as great as possible to minimize the effect of receiver local-oscillator noise; it takes more local-oscillator noise to mask a stronger signal. (More about noise later.)

There is an upper limit, however: at some point, the receiver becomes nonlinear. At saturation, one or more internal amplifier stages in the receiver is incapable of producing an increase in output corresponding to an increase in receiver input. The receiver is said to be in "gain compression." If the desired, on-channel signal level were at or above the gain-compression point, we would not be able to make a meaningful measurement. So, we have to keep the desired-signal level within a range limited on the low end by the needed signal-to-noise ratio and on the high end by gain compression.

"ARRL Standard" BDR Test Conditions

To help *QST* readers compare the relative performance of different receivers featured in "Product Review," the ARRL Laboratory uses standard test conditions for all units tested. In the BDR test, the 1-dB compression point of the receiver under test is measured with the receiver AGC turned off, if possible. The desired-signal input level is set to a point 10 dB below the 1-dB compression point. This ensures that the receiver is being operated in its linear region, using as strong a signal as can reasonably be accommodated by the receiver, yet not operating the receiver right at the limit.

In some receivers, the AGC cannot be disabled. In this case, we use a signal that is 20 dB greater than the receiver noise floor. This causes a few problems, however. First, the level of this desired test signal is usually not as strong as a signal that is 10 dB below the 1-dB compression point used to test receivers whose AGC can be disabled. This makes it more likely that we will get a noise-limited reading. In addition, we have found that changing the signal levels usually causes a receiver AGC transient response that temporarily desenses the receiver. It is then necessary to wait up to 5 seconds or so before taking a reading of the desired signal.

This isn't too much of a problem for manual testing, but when we do automated swept testing on receivers whose AGC cannot be turned off, we have to wait for the AGC to recover. This sometimes requires a few seconds between readings. If we are stepping a pair of generators in 1-dB amplitude increments over a 200-kHz range in 1-kHz frequency increments, a delay of 5 seconds per reading requires that the receiver spend about 10 hours on the test bench!

As I will show, the actual dynamic range of a receiver often varies with the frequency spacing between the desired and undesired signal(s). To help people readily compare receivers, an ARRL "standard" signal spacing of 20 kHz has been used for many years. In BDR testing, we use unwanted blocking signals that are 20-kHz above and below the desired signal, and report the worse result of the two cases in *QST*.

BDR—Best of the Best

No discussion of test results would be complete without answering the inevitable question: "What is the best radio?" To date, the best blocking dynamic range measured in the ARRL Lab has been that of the Yaesu FT-1000D (reviewed in March 1991 *QST*). The FT-1000D BDR was listed in the column as being ">154 dB." We listed it that way because we were putting about +20 dBm (0.1 W) into the front end of the receiver and it still was not blocking. I was the test engineer at the time, and I decided to stop the test because I was afraid to put more RF power into the receiver for fear that I could damage the receiver. Most excellent receivers have a BDR that is greater than about 130 dB or so.

Two-Tone, Third-Order Intermodulation Distortion Dynamic Range Measurement

When multiple signals are present at the input of a receiver (a normal situation), they can mix together and create additional signals that were not originally present. Some of these signals can cause interference to desired reception. All amplifiers and mixers are nonlinear to a degree. When two or more signals are present in the amplifier at the same time, additional signals will be created as the input signals mix together. These can range from simple harmonic generation or sum and difference products to more complex mixing involving the fundamental input signals and their harmonics.

Two-tone, third-order intermodulation distortion dynamic range is a real tongue twister. Unfortunately, we have had to learn to say it three times quickly here in the Lab, at least when we are discussing "Product Review" test results.

TTTODR is an expression of the difference between the noise floor of a receiver and the level of two off-channel input signals, spaced from the desired signal and each other in a third-order frequency relationship, that mix together in the receiver and create a third-order mixing product equal to the noise floor of the receiver. Like blocking dynamic range, TTTODR is usually expressed in dB.

Third-Order Frequency Relationships

The term "order" is often used in conjunction with intermodulation. "Order" is the number of terms in the equation used to characterize the frequencies involved. For example, mixing might generate signals at frequencies of $f_1 + f_2$ and $f_1 – f_2$. There are two factors in these expressions, so this is second- order intermodulation. (Second-order intermodulation includes cases such as $f_1 + f_1$, the mixing of a signal with itself to create the harmonic $2f_1$.)

Intermodulation can be more complex than these simple sum and difference second-order relationships. In a nonlinear circuit, harmonics of all the signals present are created, and those harmonics mix with all of the original signals plus those created by the nonlinearity. The equations used to describe these more complex mixing processes have more terms, so they are higher order.

In a receiver, third-order IMD is usually the most prevalent form of interference caused by intermodulation. This is because in-band signals can mix together to form third-order products that are also in-band. Two strong signals a few kHz apart up or down the band can create a third signal somewhere in the band. If the strong signals are at just the right (or wrong) frequencies, interference can be the result. Frequencies that result from third-order intermodulation are at $f_1 + f_1 – f_2$ and $f_2 + f_2 – f_1$.

To give a specific example for the 20-meter band, if f_1 is 14.04 MHz and f_2 is 14.06 MHz, a nonlinear system will create third-order products at 14.02 and 14.08 MHz.

Third-Order Amplitude Relationships

In addition to having a third-order frequency relationship, TTTODR has a third-order amplitude relationship. **Fig 3** shows two curves. The upper curve is the linear, on-channel amplitude response of the desired signal. Until compression occurs, the system is linear—a 1-dB increase in input signal results in a 1-dB increase in output signal.

The lower curve shows the third-order IMD amplitude response. In a classic third-order response, a 1-dB increase in signal levels will result in a 3-dB increase in output level. In addition to the on-channel and third-order amplitude curves, **Fig 3** shows the receiver noise floor. As can be inferred from the graph, the third-order response of the receiver is inaudible until the input signals are substantially higher than the noise floor. Once the third-order response becomes audible, it becomes loud very quickly as the amplitude of the two input signals increases more.

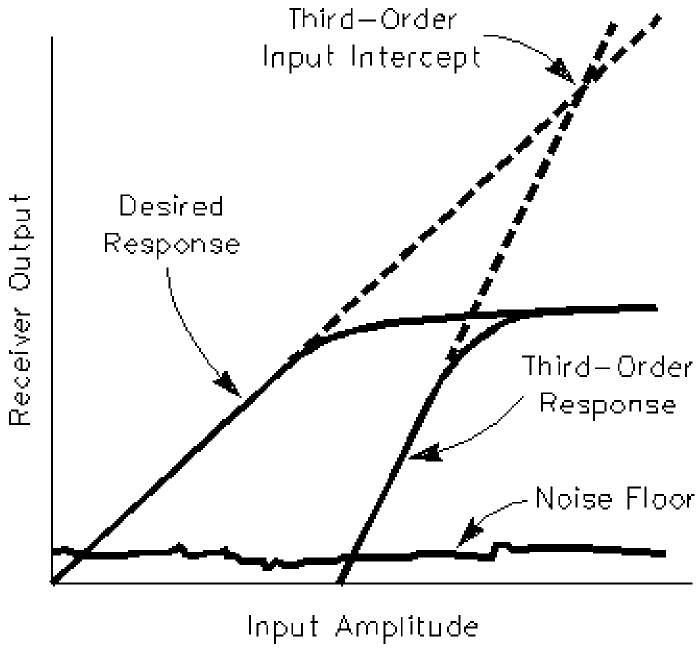

Fig 3—A plot of a receiver's on-channel and third-order IMD responses.

TTTODR Testing in the ARRL Lab

The test set-up of **Fig 2** is used to test TTTODR in the ARRL Lab. The receiver is tuned to a desired frequency, usually 20 kHz above the lower band edge. The two generators are each set to a level of –20 dBm, at frequencies 20 kHz apart. The generator frequencies are selected to produce a third-order response at the receiver frequency. The step attenuator is initially set to a high value, then decreased in 1-dB steps until a third-order IMD product appears in the receiver output. When this product is at the receiver's noise floor, the level of one of the input signals to the receiver is calculated from the generator and attenuator settings. The difference between the noise floor and the level that results in intermodulation distortion is the TTTODR, usually expressed in dB.

TTTODR also varies with the signal spacing of the test tones, as I'll show. To facilitate comparison of receivers, the ARRL Laboratory has standardized on a 20-kHz spacing between the test tones.

Test-Fixture IMD

So far, we have been discussing IMD created by the nonlinearities in the amplifiers and mixers in a receiver. IMD can also be created in the test fixture being used to test a receiver, corrupting the test results. If the signal from one of the generators enters the other generator's output, it can mix in the output amplifier with the generated signal, causing formation of intermodulation products. This happens if the two signal generators are not sufficiently isolated from each other. Like most signal generators, the Hewlett Packard and Marconi generators shown in **Fig 2** have internal amplifiers and a calibrated internal attenuator. When the generators are set to a low amplitude level, the internal attenuators provide quite a bit of attenuation between the internal amplifier and the external equipment that amplifier is connected to. Since the internal attenuator of each generator attenuates the incoming signal from the other generator, not much mixing can occur. And what little intermodulation signal may be generated will be further attenuated on its way back through the attenuator to the output, minimizing this "reverse" IMD.

In the test set-up of **Fig 2**, the two-port coupler provides a fair amount of port-to-port isolation between the two generators. At 14 MHz, if the coupler is properly terminated in a 50-Ω load, the MiniCircuits coupler we use offers a measured isolation of >35 dB. In addition, per our standard test procedure, the generators are set for a –20 dBm signal level, which places 20 dB of attenuation internally in the generator between its output amplifier and its output jack. This results in a total isolation of >75 dB between the two generator output amplifiers. This significantly attenuates generator signals capable of generating IMD in the test fixture.

The amplifiers used inside the signal generators are highly linear. Even if a reverse signal is present on another frequency, it is not likely to create measurable IMD products. One easy test to determine if there is IMD occurring in the test fixture is to change where attenuation is applied to the signals and note the effect. For example, consider a receiver with a noise floor of –140 dBm. If it generates IMD when the signal generators shown in **Fig 2** are set to –20 dBm with the step attenuator is set to 17 dB, the off-channel signals are –40 dBm at the receiver input (–20 dBm, minus the 17 dB of the step attenuator minus the 3-dB loss of the hybrid coupler), so the TTTODR is 100 dB . If we change the generators' output levels to –30 dBm, and change the step attenuator to 7 dB, the off-channel signal levels are still –40 dBm, but there is now an additional 20 dB of attenuation between the two generators. If the IMD being measured was created in the signal generators, it would then be significantly reduced. Such a change in the measurement indicates that there is a problem with test-fixture IMD, but if the measured IMD remains the same we can be confident that we were actually measuring the IMD of the receiver under test.

The ARRL Lab staff has determined that our test fixture is capable of measuring TTTODR of approximately 108 dB or so. To eliminate any questions about test-fixture IMD, we always perform the above procedure of moving attenuation around to verify the measurement accuracy for receivers that measure better than 100 dB of TTTODR.

TTTODR—The Best of the Best

As of press time, the best TTTODR measured in the ARRL Lab has been that of the IC-775. It delivered an impressive TTTODR of 106 dB! In practice, excellent receivers have a TTTODR of better than 100 dB or so.

Oscillator Noise

Oscillator noise can interfere with any receiver measurement, whether the noise is generated by the test- signal generator(s) or by the local oscillator of the receiver under test. The Marconi and HP generators we use are much less noisy than the local oscillators found in commercial amateur radios, so the limiting factor is usually the receiver under test.

All oscillators have some noise sidebands present at frequencies around the oscillator frequency. These noise sidebands are produced by a composite of amplitude and frequency variations. These off-frequency components of the receiver's local-oscillator noise can mix with the off-channel signal being used to test blocking, creating additional noise in the output of the receiver. In severe cases, this additional noise power causes an increase in the output level before the receiver starts to block.

We watch for this effect during BDR testing. If we see a 1-dB increase in the measured output level of the receiver, we deem the test to be *noise limited*. In recent *QST* product reviews, we have reported the level at which the test was noise limited, expressed in dB over the measured noise floor.

Noise can also affect TTTODR tests. This is not as common as it is for BDR tests because the TTTODR of a receiver is usually about 20 to 40 dB poorer than its BDR. That means it's likely that an IMD product will be audible before noise affects the measurement. This is somewhat offset by the fact that TTTODR tests are made at the noise floor, but in most cases the TTTODR wins the noise battle and valid results can be obtained. Receivers that have particularly poor noise characteristics *can* be noise limited on TTTODR measurements, however. In the case of TTTODR, we assume that if the receiver output (noise plus any IMD signal present) rises by 3 dB, we are noise limited at that level.

Being noise limited does not necessarily imply that a particular receiver is poor. For example, if a particular receiver actually blocks at 90 dB above the noise floor, it is not as good as a receiver that is noise limited at 150 dB! In the former case, blocking occurs before the receiver's local-oscillator noise could affect the test. In the second case, the blocking performance is so superior that even a small amount of receiver noise can affect the test. In essence, "noise limited" simply states that noise (from whatever source) won the match between noise and blocking or intermodulation distortion.

Dynamic Range versus Frequency Spacing

I have been alluding to the fact that in real receivers, dynamic range varies with frequency spacing. This could be somewhat misleading. In fact, dynamic range varies with the amplitude of the signals doing the mixing or blocking. But in real receivers, filters that are part of the receiver design serve to attenuate off-channel signals to one degree or another, lowering the amplitude of the unwanted signals within the receiver, and thus affecting the dynamic range as signals are varied in frequency in and out of the filters' passbands.

Fig 4 shows a receiver with no filtering. Actually, it is a direct-conversion receiver consisting of a preamplifier, a mixer and a post-mixer amplifier that passes both audio and RF (up to a point, anyway). Because there is no filtering, all signals present at the receiver input get amplified equally by the RF amp and presented to the mixer, causing sum and difference frequencies (among others) to appear in the mixer output. In a BDR test of this receiver, one of the applied test signals is set to the desired frequency, which for this receiver will be within a few hundred Hz of the local-oscillator frequency, and a second (blocking) signal is set to another frequency. Because there is no filtering in this receiver, any blocking-signal frequency will do as long as it is different than the frequency of the desired signal; the level at which blocking will occur is not dependent on the spacing between the desired and blocking signals.

Treating this receiver as a "black box," we really don't know where the blocking occurs, but for our purposes we also really don't care. All we need to know is that it occurs somewhere in the system. (In fact, it is likely that the blocking is occurring somewhere late in the chain, after the signals have been amplified by the RF and post-mixer amplifiers.)

A similar situation arises if we perform a TTTODR test of this receiver: the frequency spacing of the two off-channel tones will not affect the TTTODR of the receiver, as long as the signals are at the correct frequencies to generate a third-order product signal at the desired (received) frequency.

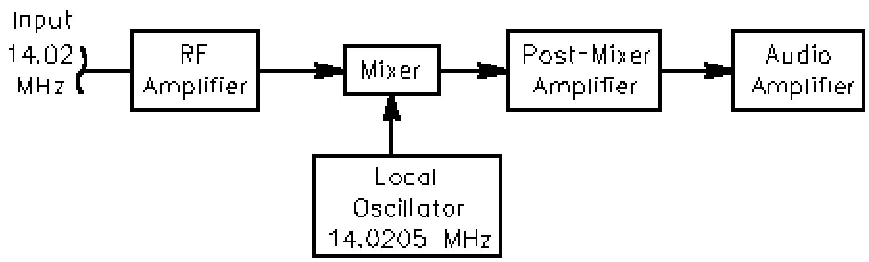

Fig 4—This direct-conversion receiver has no filtering. Its dynamic range is independent of test-frequency spacing.

Receivers with Filters

In the several receiver examples that follow, all filters are assumed to be lossless. **Fig 5** shows a receiver similar to that of **Fig 4**, except that we have added a front-end bandpass filter. This filter has the frequency response shown in **Fig 6**, from which it is apparent that the filter does not pass all frequencies equally to the rest of the receiver. The filter attenuates signal frequencies by differing amounts, depending on the amount by which the signal frequency differs from the filter center frequency. Correspondingly, the amount of off-channel frequency delivered to the receiver depends on the exact off-channel frequency and the characteristics of the filter.

In a BDR test of this receiver, the on-channel signal at 14.02 MHz will be passed through the filter with no attenuation, but a blocking signal at 14.06 MHz will be attenuated by 10 dB. This means that the level of the applied blocking signal will have to be 10 dB greater that it was in the receiver of **Fig 4** to achieve the same blocking effect; the BDR of this receiver, at 40-kHz spacing, is thus 10 dB better than that of the receiver of **Fig 4**. If the blocking signal is moved to 14.1 MHz, it will be attenuated by 30 dB. The BDR would be 30 dB better than the receiver of **Fig 4**.

The TTTODR of a receiver also varies with signal spacing, although the relationship between the filter characteristics and intermodulation performance is not quite as straightforward as it is for blocking. This is due to the third-order amplitude relationship between the input signal and the resultant products and the fact that in TTTODR tests, there are two input signals—at different frequency offsets from the frequency to which the receiver is tuned. The input and IF filters may attenuate each of these input signals by a different amount, making the two tones no longer equal in amplitude at the point at which IMD occurs.

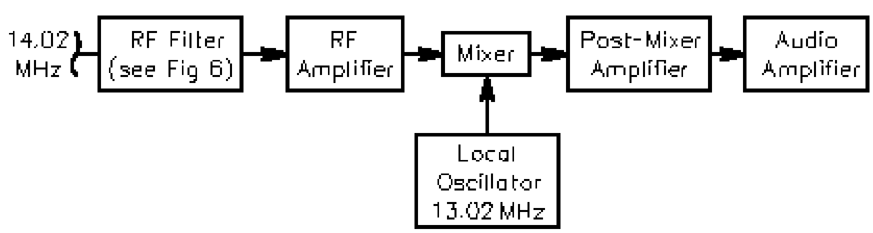

Fig 5—An RF bandpass filter has been added to the receiver shown in Fig 4. The characteristics of the filter are shown in Fig 6.

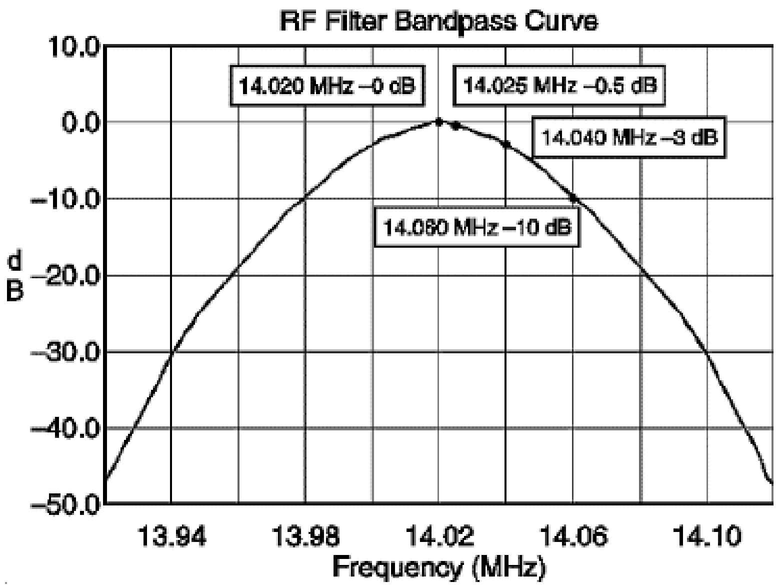

Fig 6—The characteristics of the RF bandpass filter used in several of the receiver examples.

A More Complex Case

In most real receivers, there are multiple stages and filters. Each of these can have a specific effect on how dynamic range varies with frequency. **Fig 7** shows a receiver with RF and IF amplifiers, a front-end filter, a relatively broad post-mixer IF filter (sometimes called a *roofing* filter) and a sharp IF filter near the end of the IF amplifier chain. The bandpass characteristics of the first IF filter are shown in **Fig 8**. The bandpass characteristics of the 500-Hz IF filter are shown in **Fig 9**. **Fig 7** could represent a real receiver, although most modern amateur receivers use multiple conversions and an even more complex filter arrangement.

For the sake of discussion, let’s assume that every filter is lossless and that each RF amplifier, IF amplifier or mixer stage will block or generate IMD if the input signal levels are at 0 dBm PEP or greater. In practice, the actual level at which distortion occurs will vary quite a bit, depending on the specific design of the stage. But this simplification allows us to easily calculate filter attenuation versus frequency and amplifier gain or loss to determine where in this hypothetical receiver blocking and IMD will occur for given input-signal levels and frequencies.

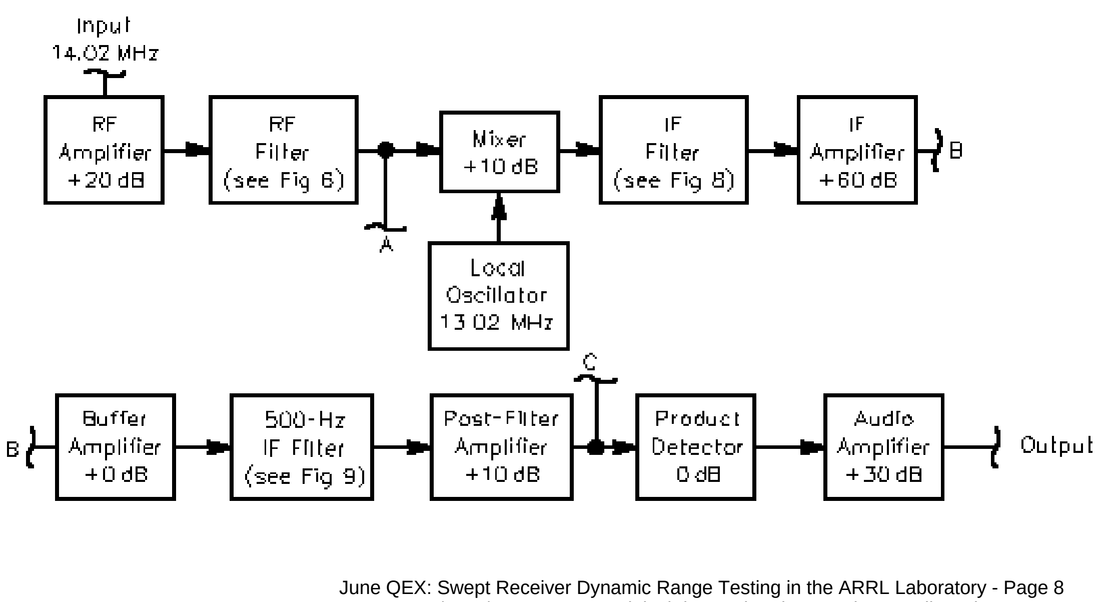

Fig 7—A more complex receiver using the RF filter shown in Fig 6, the IF filter shown in Fig 8 and the 500-Hz narrow filter shown in Fig 9.

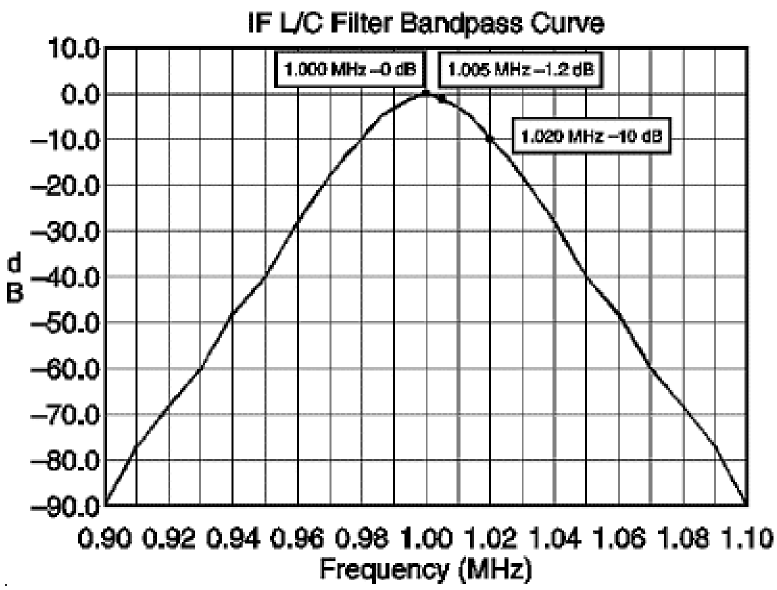

Fig 8—The characteristics of the IF roofing filter used in the receivers of Figs 5 and 7.

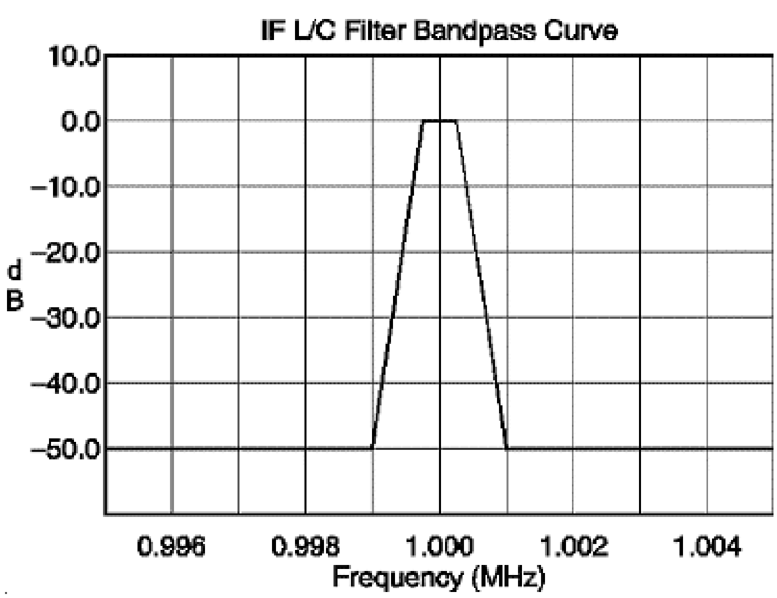

Fig 9—The characteristics of the 500-Hz narrow IF filter used in the receivers of Figs 7 and 11.

Let's look at BDR first. Assume that the receiver of **Fig 7** is properly tuned to an input signal of –120 dBm at 14.020 MHz. At Point A, this signal will be amplified to –100 dBm. By the time it gets to Point C, it will be at –30 dBm (P_1) and will have been mixed down to 1.000 MHz. –30 dBm is well below the defined distortion point of 0 dBm, so the receiver will operate normally.

Now let's add a second (blocking) signal of –120 dBm at 14.025 MHz. This is 5 kHz away from the signal to which the receiver is tuned, so it will not be heard in the output. The input filter will attenuate this signal by 0.5 dB and the 1-MHz IF roofing filter will attenuate it by another 1.2 dB, for a total of 1.7 dB of attenuation. At Point B, the level of the 14.025-MHz signal, now mixed down to 1.005 MHz, will be –31.7 dBm (P_2). Point B is in front of the sharp IF filter, so both the 1.000 and 1.005-MHz IF signals will be present at this point simultaneously. The peak RMS power of these two signals (P_1 and P_2) can be found as:

$$\mathrm{PEP} = 20\log\left(10^{P_1/20} + 10^{P_2/20}\right) \qquad \text{Eq 1}$$

where P_1, P_2 and the resulting PEP are all expressed in dBm.

Solving Eq 1, the PEP is –24.78 dBm, still less than 0 dBm, so the radio will continue to function without degradation of the desired signal; the –120 dBm signal on 14.025 MHz will not interfere with reception of the desired signal on 14.020 MHz.

But at –120 dBm, the undesired signal is a pretty weak signal. In a crowded band, it is likely that strong signals on other frequencies will be much stronger than –120 dBm. If the signal on 14.025 MHz were increased to –88.58 dBm, the total PEP of the two signals present simultaneously at Point B would be 0 dBm (adding the total of 90 dB of gain to each signal, subtracting the 1.7 dB of loss through the filters and solving Eq 1). At this level, by our previous definition of receiver performance, the desired signal would just start to block and intermodulation products would appear at other frequencies as the two products mixed together in the nonlinear stage. In this case, the receiver is not tuned to any of the intermodulation products, however, so the main concern is blocking. If we assume the noise floor of the receiver of **Fig 7** to be –140 dBm, the BDR *at this frequency spacing* is 51.42 dB, the difference between the noise floor and the –88.58 dBm of interfering signal that causes blocking.

Let's analyze BDR at other frequency spacings using the same –120-dBm 14.02-MHz signal we used in the previous example. A –120-dBm signal at 14.04 MHz (20-kHz spacing) will be reduced 3 dB by the RF filter. It will then be mixed down to 1.02 MHz where it will be attenuated 10 dB by the IF roofing filter. By the time both signals get to Point B, the 14.02-MHz signal will be –30 dBm (at 1.0 MHz) and the 14.04-MHz signal will be –43 dBm (at 1.02 MHz). Solving Eq 1 we find the total power will be –28.25 dBm and the receiver does not block.

If we increase the input power at 14.04 MHz to –77.28 dBm, subtract the attenuation of the two filters, add the gain of the various receiver stages and solve Eq 1, we find that the power at Point B is 0 dBm and the receiver has just started to block. The BDR can then be calculated as 62.72 dB for 20-kHz spacing by subtracting the noise floor of the receiver from the level of input signal that causes the receiver to block. If we do the same thing for blocking signals at other frequencies, we'll find that the BDR varies with the frequency spacing between the desired and unwanted signals.

Things get worse when the unwanted signal gets inside the narrow 500-Hz filter. Assume that the receiver is properly tuned to a signal at –120 dBm. Let's now put a –100.92-dBm signal at 14.02025 MHz, with both the 14.02-MHz and 14.02025-MHz signals squarely within the narrow filter. This case will result in a total PEP of 0 dBm at point C, so that stage will block or generate IMD, with a resultant dynamic range of 39.08 dB.

Note that the RF filter is *after* the broadband RF amplifier. So by our previous definition that the receiver becomes nonlinear if the signal level reaches 0 dB at the input to any stage, if the input signals are at 0 dBm, the receiver will block or IMD will be created. So the dynamic range of this receiver is never better than 140 dB at *any* frequency spacing, regardless of the amount of attenuation that may be offered by the RF or IF filters.

A graph of the calculated BDR of our hypothetical receiver is shown in **Fig 10**.

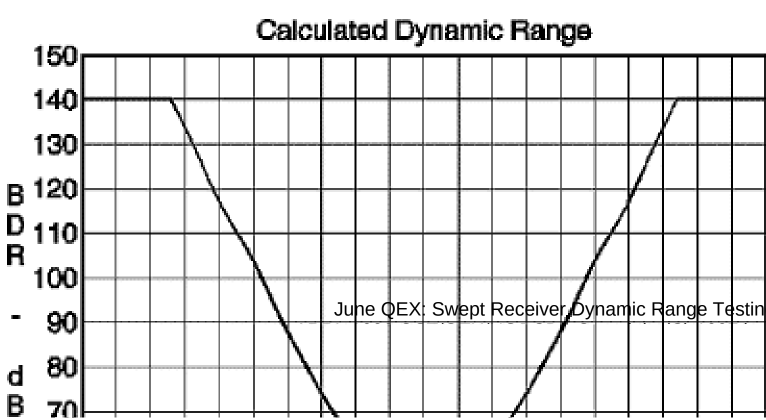

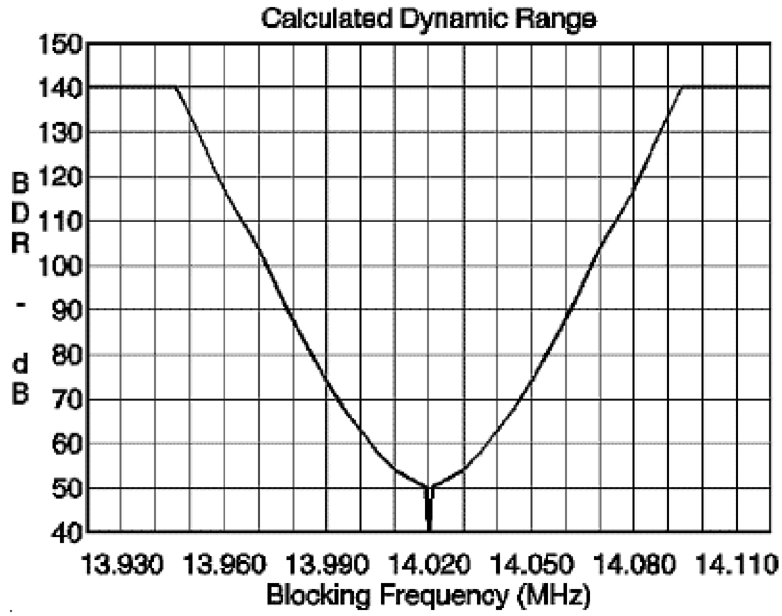

Fig 10—A graph of the calculated BDR of the hypothetical receiver shown in Fig 7.

One More Receiver

Let's analyze one more receiver—the one shown in **Fig 11**. In this receiver, the selective IF filter has been moved up much earlier in the IF amplifier chain. Just as in the previous examples, the input signals will be attenuated by the RF and IF filters. However, signals that are outside the passband of the 500-Hz IF filter will be attenuated significantly relatively early in the receiver stages. For most frequency spacings, it will take a correspondingly much greater off-channel signal to obtain 0 dBm anywhere in the amplifier chain than it did in the previous examples.

For the example of two signals inside the passband of the 500-Hz filter, the receiver's dynamic range will be the same as in the example for the receiver of **Fig 7**. But for the example of a blocking signal at 20-kHz spacing, the receiver of **Fig 11** will perform much differently. Assume that the receiver is properly tuned to a signal on 14.02 MHz. If we introduce a blocking signal at 14.04 MHz, it will be attenuated 3 dB by the RF filter, then another *50 dB* by the 500-Hz filter. This is 40 dB more attenuation than was obtained by the IF roofing filter in the receiver of **Fig 7**. The gain distribution of this receiver is a bit different, however, with an IF gain of 70 dB after the filter, as opposed to the receiver of **Fig 7** having a gain of 60 dB after the IF roofing filter and then another 10 dB of gain after the narrow 500-Hz filter. Working out the various gains and losses then solving Eq 1, we can calculate that it takes an input signal of –37.91 dBm to create a PEP of 0 dBm at Point C—a BDR of 102.09 dB if we assume the receiver noise floor to be –140 dBm. Clearly, receiver design has a strong effect on the blocking dynamic range of a receiver versus frequency spacing.

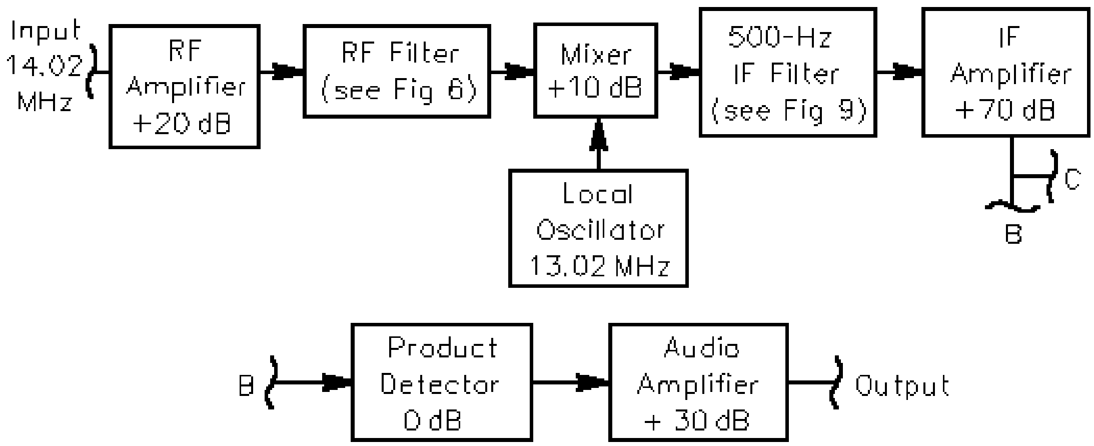

Fig 11—In this receiver, the 500-Hz narrow IF filter has been moved to a point early in the IF-amplifier chain.

Automated Testing

As I have shown, the dynamic range of a receiver varies with the spacing between the desired-signal frequency and the off-channel signals used to perform the test. The hypothetical graphs shown in various figures are similar to real receiver performance. Several years ago, back when I was the Lab test engineer, I spent a few hours in the ARRL shielded screen room manually testing a few receivers at a number of test-frequency spacings. Unfortunately, we had to decide that taking this much data was just too time consuming to do on a regular basis. But we did continue to investigate ways to automatically test receiver dynamic range at multiple spacings. To aid in this process, we acquired our Marconi GPIB-controlled signal generator, GPIB-programmable step attenuators and an HP-3478A GPIB-controlled RMS voltmeter. This equipment was first controlled with some *QuickBasic* software written by Jon Bloom, KE3Z. In essence, the software stepped the Marconi generator over a frequency range, changed the programmable attenuators in 1-dB steps and measured the total output power of the receiver under test. When the receiver output had either decreased by 1 dB, indicating that the receiver was blocking, or had increased by 1 dB, indicating that the test was noise limited, the BDR data was logged and later plotted.

In late 1995, we purchased a copy of National Instruments *Lab Windows CVI*, a well-known Windows-based instrumentation package. Using this, I wrote a C-language program that performed and graphed the same measurements as the *QuickBasic* software. To use this software, we manually set our (nonprogrammable) HP-8640B signal generator to the desired frequency (14.02 MHz in most of the examples discussed in this article) and allowed the software to step the Marconi 2041 in small increments (usually 1 dB) across a swept frequency range, simultaneously using the HP-3478B to measure the receiver's output signal voltage. It sure was slow—about 1 second per reading best case—but it could be set up and run without a test engineer present. (Okay, I must admit that Jon and I watched the first tests being run for hours on end).

In early 1996, we received our National Instruments AT-2150A 16-bit digital data-acquisition (DAQ) card. We now use this to measure the output level of the receiver under test. The DAQ card has a distinct advantage over the HP-3478A: it takes readings a *lot* faster. It can be programmed to sample at rates ranging from 4 kHz to 51.4 kHz. I wrote a short program that takes 256 samples at a sampling rate of 48 kHz and graphs the data in the time domain. The program also uses some of the National Instruments Advanced Analysis Library DSP functions to calculate the frequency domain representation of the sampled signal, which is also plotted. Even with all these calculations, the software is able to take and graph as many as 10 measurements per second. At first blush, it appeared that we could do swept dynamic-range testing in the blink of an eye. Unfortunately, we found that receiver AGC responses and noise sometimes prevent taking really fast swept dynamic-range data.

Noise and Averaging

Noise appears in many receiver tests. To measure a receiver's noise floor you have to (surprise!) measure noise. The desired signal is at the same power level as the noise—a signal-to-noise ratio of 0 dB. The random nature of noise makes it likely that a short-term measurement will be inaccurate, so to obtain a valid measurement in the presence of noise it's necessary to perform some sort of averaging to smooth the result.

In manual measurements, the averaging is more or less automatic. The best way to do this is to use an analog audio voltmeter to measure the receiver output level. If the meter indicator is steady, the measurement is not being affected by noise. You can then proceed to the next step in the test (usually to increase the level of one or more of the signal generators) and again observe the receiver output level. If the analog meter level is unstable, the measurement is being affected by noise. In that case, a skilled test engineer is usually able to mentally average the results by visually noting the level around which the signal is varying.

The computer can do this too, more or less. I say more or less because the human interpretation quickly determines if a reading is noisy and doesn't try to average it unless it is required—the computer isn't that smart. Still, if the software is written to average a number of readings, the effects of noise can be significantly reduced. With simple software, however, there is a price to pay: you may need to average 10 or more readings to get a reasonable measurement of a really noisy signal. Our initial software used the HP voltmeter to make the measurements. This allowed about one reading per second. To obtain a smooth graph of dynamic range over 200 kHz, the program must vary the signal amplitude in 1-dB steps over a 50-dB range and the frequency in 2-kHz steps over a 200-kHz range. At one reading per second, this would result in a total test time of almost 14 hours.

Averaging can work. Refer to **Fig 12**. This figure shows two BDR tests that were performed on an Icom IC-765. One of the tests is quite noisy, the result of performing the test with no averaging of results. The other curve shows the result of averaging out noise over 10 readings for each measurement taken. (These data were manually verified at several points across the swept range, showing excellent correlation with the test methods that have been proven over the years in the ARRL Laboratory.) But the averaged measurement was slow. (In this graph, the lower line has been reduced in level by 10 dB so facilitate easy comparison of the shapes of the two curves.)

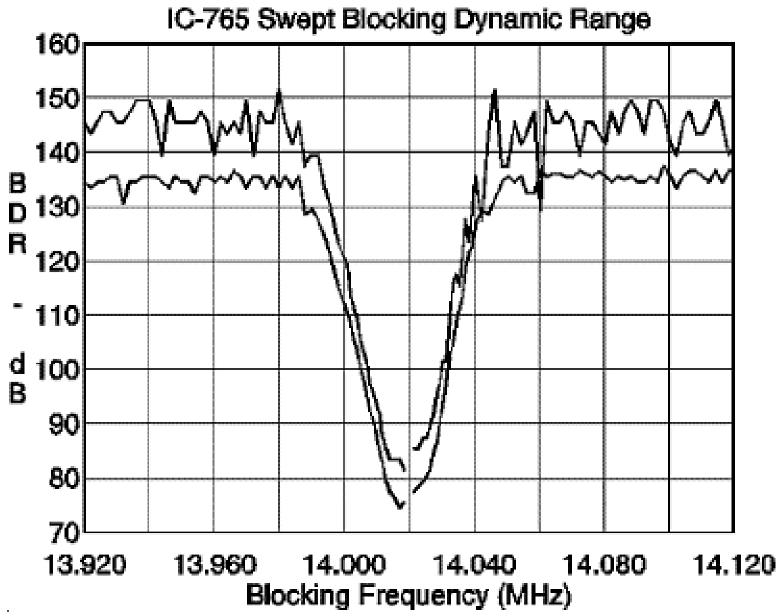

Fig 12—The upper curve shows the noisy test results that result when averaging is not used. The lower curve shows the same receiver with averaging used as the measurements are taken, to eliminate the effects of noise. The lower graph has been offset by 10 dB.

Enter the DAQ Card

The major limiting factor in the above scenario is the measurement time of the hardware; the HP-3478A is capable of about 1 ac voltage reading per second. The National Instrument DAQ card is capable of taking readings quite a bit faster. As soon as the Lab received the new DAQ card, I rewrote the software to use it. The first radio I tested was a small, homebrew QRP transceiver. This is a simple radio with no AGC. It blocked before it was noise limited, so noise was not really a factor. When I wrote my software to take about 4 readings a second, stepping in 2-kHz steps and in 1-dB increments, it ripped through the test in about a half hour.

Enter the Real World

The next radio I decided to test was my old Ten Tec Omni D. This radio uses a permeability-tuned VFO instead of the phase-locked loop used in most modern rigs. The PTO is very low-noise, so the Omni D was not noise limited during the testing. However, its AGC cannot be disabled. As the Marconi generator’s internal attenuators were switching, they were generating loud transients, resulting in strong transient AGC responses in the receiver. When this happened, the radio would go momentarily deaf, giving much more than the 1-dB drop in output power the software was looking for. The first time through, the software told me that the radio had a dynamic range of about 40 dB, jumping all over the place on the graph. Of course, I knew this to be wrong; the Omni D is much better than that! This is where I learned the easy C function called *Delay*().

Debouncing

At first, I put delay into the program, waiting a specific amount of time each time I changed generator levels, just in case the change generated a generator transient. This brought me right back to square one, however, with about 1 second per step. After a few false starts, I settled on using the C function shown in **Listing 1** to “debounce” the reading by waiting until the receiver’s output level is stable enough to give two readings that are of the same approximate level. (I must confess, I am an old Applesoft BASIC programmer, so if there is some resemblance between Applesoft and the following C code, don’t blame me: old habits are hard to break).

Listing 1 is a bit simplified from my actual code. Some of the variables that get passed into this function, and then on to the *measure_2150* function, have been left out or hard-coded with #define statements to unclutter the presentation of the concept. In essence, this code represents the fastest approach I could take to measuring the receiver output, by taking readings every 100 ms and waiting until two are the same (within the tolerance called out for in *ACCURACY*). If the receiver is recovering from an AGC transient, the receiver output is not stable. The function will keep trying until the receiver recovers and its output becomes stable and steady, then return a measurement. If the receiver is stable, only two readings, spaced by a small delay, will be necessary to obtain a valid reading, as opposed to waiting for a fixed delay, whether it is required or not.

Measurements with Noise

This function works great for noiseless signals that are changing in response to a receiver transient. But in most cases, especially with better receivers, receiver noise is increasing as the level of undesired, blocking signals is increasing. As a result, the receiver output signal contains a fair amount of noise in addition to the desired signal. The function in **Listing 1** can't distinguish between two readings that are different because the receiver is recovering from a transient and those that are different because of the presence of noise. So with noisy signals, the function may have to take a number of readings before the laws of probability give two noisy readings that are within *ACCURACY* of each other. This is not necessarily a bad thing, but it does add quite a bit of time to the test; I have seen up to 23 readings on moderately noisy signals.

Digital Signal Processing

Digital signal processing (DSP) can rescue us here. *Lab Windows CVI* has an extensive signal-analysis library. Among the routines are one—*AutoPowerSpectrum*()—that sorts the acquired data into frequency spectrum *bins* and one—*PowerFrequencyEst-imate*()—that scans those bins to determine both the frequency peak and the total power contained in the frequency peak. If the noisy signal has a significant component at one frequency, as would the receiver output when it is tuned to a signal, these functions can be applied to the acquired data to discard most of the noise (except the small amount actually contained around the frequency peak). Thus, the actual signal voltage can be obtained for even a noisy signal. If this noisy signal is still recovering from a receiver transient, the software knows and keeps waiting. If the actual signal is stable but also contains broadband noise, the software knows that and does not continue to take readings.

Averaging

As handy as *Lab Windows CVI* DSP functions are, they have their limits. As signals get noisy, it is still difficult to accurately measure the desired signal without averaging. Fortunately, though, DSP gives us a way of determining if we *need* to do averaging on a particular reading. If we compare the results of the ac-voltage estimation function with the measurement of power at the frequency peak, we can calculate the signal-to-noise ratio. We then can only use averaging if the signal contains noise. This whole process allows us to take two readings (to debounce) and simultaneously determine if the signal is noisy. If it is noisy, we can then take a number of readings, average them and return the averaged value, essentially without noise. If the averaged value of these frequency- selective readings has dropped by 1 dB, the receiver is blocking, even if it is noisy enough for the total power to be essentially unchanged because the noise is rising just as fast as the desired signal is blocking.

In the upper curve of **Fig 12**, that is exactly what was happening to make it so noisy at the upper end of the BDR scale. As the level of the off-channel signal was being increased, the receiver was starting to block, but the noise was increasing at almost exactly the same rate. The *total* power that was being measured as an ac voltage stayed approximately the same over a 10-dB input-signal range. However, noise being noise, if the results were not averaged, on some readings the software would determine that that the signal had just increased by 1 dB; on the next reading, however, the software would not make that determination until 10 dB more input signal had been supplied. So the BDR appears to vary over a 10 dB range, primarily due to the unpredictable nature of noise itself. The lower curve (offset by 10 dB) shows the way the proper use of frequency-selective measurement and averaging gives us a reasonably smooth curve in spite of some serious testing difficulties.

Comparison with Manual Testing

While we have always been able to average out noise, until now we did not have an easy way of making frequency selective measurements on the desired signal. We could manually use our spectrum analyzer to look at the receiver's output, but we quickly decided that this wasn't practical for "production" testing. While I can't predict how this new ability will affect out Product Review testing in the future, I will point out that the smoothed graph of actual BDR shown in **Fig 12** could not have easily been done using manual test methods. We used earlier versions of the swept-BDR program to create the graphs featured in the expanded test-result reports that are available for the IC-775DSP, TS-870S, IC-706 and FT-1000MP HF transceivers. As we incorporate the software changes discussed in this article, especially those that relate to our ability to accurately measure noisy signals, we will end up with dynamic-range results that are a bit more exact, more accurately showing the level at which the desired signal blocks by 1 dB or the level at which the total power rose by 1 dB, indicating that the test was noise limited.

An Actual Receiver

The test results from our ARRL Lab IC-765 are shown in **Fig 13**. This graph is a classic, showing BDR decreasing with decreased frequency spacing as the test tone gets inside various internal filters in the receiver. While this receiver uses a multiple-conversion design, most of the ultimate selectivity is late in the IF amplifier chain. The result of this design is the fairly wide area of reduced BDR evident on the graph.

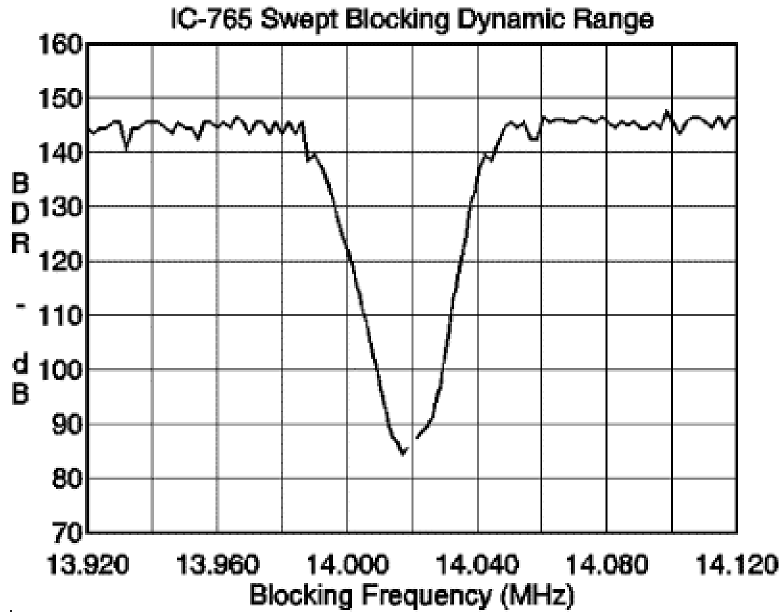

Fig 13—The results of a swept BDR test taken on the ARRL Lab's Icom IC-765.

Listing 1

```
double debounce_2150 (int number_of_samples, double delay)

/* This function gets the ac voltage from successive DAQ measurements and compares them. When it
has determined that one of them is within ACCURACY of the previous reading, it assumes that the
receiver has recovered its composure after an AGC transient and returns a double-precision
floating point value of the ac voltage. */
  {
   int i = 0;
   double attempt1;
   double attempt2;
   #define MAXIMUM_TRIES 256
   #define ACCURACY 0.025
   attempt1 = measure_2150 (number_of_samples, VAL_AC);
   /* This is my function that acquires the data and returns its ac voltage */
   while (i < MAXIMUM_TRIES)
  {
     /* prevents an endless loop if two readings are close to being the same */
     Delay(delay); /* delay is usually 0.1 seconds or so */
     attempt2 = measure_2150 (number_of_samples, VAL_AC);
     if (attempt1 < attempt2 * (1 + ACCURACY) && attempt1 > attempt2 *
(1 - ACCURACY))
       {
```

```
      i = MAXIMUM_TRIES;
      attempt1 = attempt2;
     }
     i++;
  }
  return ((attempt1 + attempt2) /2);
}
```