ARRL Periodicals Archive – Search Results
A membership benefit of ARRL and the ARRL Technical Information Service

ARRL Members: You may print a copy for personal use. Any other use of the information requires permission (see Copyright/Reprint Notice below).

Need a higher quality reprint or scan? Some of the scans contained within the periodical archive were produced with older imaging technology. If you require a higher quality reprint or scan, please contact the ARRL Technical Information Service for assistance. Photocopies are $3 for ARRL members, $5 for nonmembers. For members, TIS can send the photocopies immediately and include an invoice. Nonmembers must prepay. Details are available at www.arrl.org/tis or email photocopy@arrl.org.

QST on CD-ROM: Annual CD-ROMs are available for recent publication years. For details and ordering information, visit www.arrl.org/qst.

Non-Members: Get access to the ARRL Periodicals Archive when you join ARRL today at www.arrl.org/join. For a complete list of membership benefits, visit www.arrl.org/benefits.

QST Issue: Aug 2004
Title: QST Product Reviews--In Depth, In English
Author: Michael Tracy, KC1SX

Click Here to Report a Problem with this File

By Michael Tracy, KC1SX

QST Product Reviews— In Depth, In English

Our lead test engineer describes product review testing to help readers make the best use of this popular column.

So why do reviews? Some of the reasons are obvious, some not so obvious. Most folks want to get the best possible equipment for their hard-earned cash. Others might be satisfied if they find something that fits their needs as long as it doesn't have any serious problems.

In one form or another *QST* has been "reviewing" Amateur Radio equipment since the early 1930s. The first investigations were pretty basic, giving block diagrams and circuit descriptions. In 1975, Recent Equipment saw the addition of test results from the ARRL Lab, and the column was subsequently renamed Product Review.

In the pages that follow, the process of Product Review will be described, and the Lab test data will be explained in a manner that will aid in providing a better understanding of these large collections of numbers.

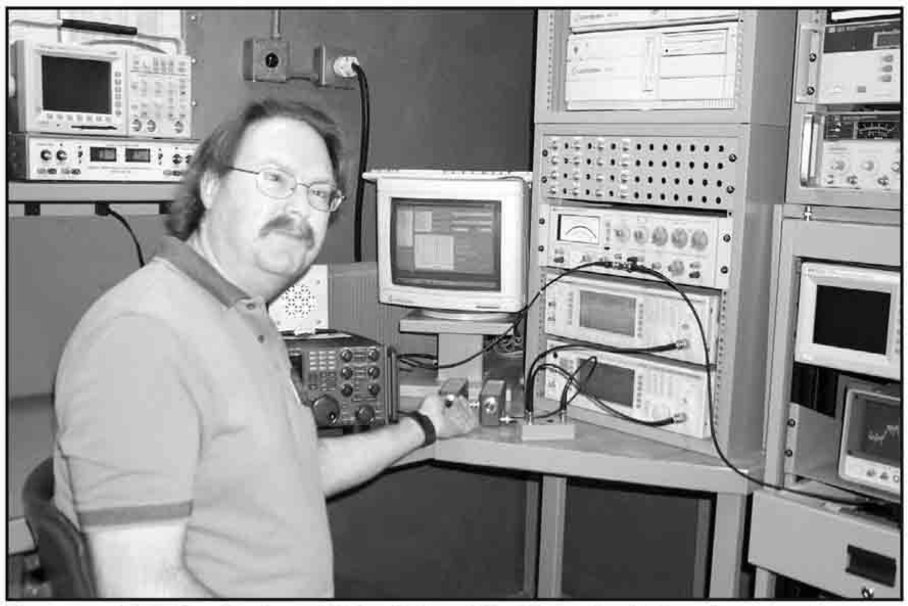

Figure 1—ARRL Test Engineer Michael Tracy, KC1SX, testing IMD performance in ARRL Lab screen room.

The Process

Selection

The Product Review Editor selects equipment for review, and selects an appropriate person to perform the review. The editor may choose to do the review himself or may select a writer knowledgeable in the field from among the licensed members of the Headquarters staff (with the exception of the Advertising Department), Technical Advisors and Contributing Editors.

Purchase

After an item is chosen, procurement must be made. To ensure that the equipment is as close to "typical" as possible, purchases are made from Amateur Radio dealers or indirectly by third parties. In effect, we purchase equipment the same way our members do. Indeed manufacturers are often not aware that a review is in process until the "wrap-up" (described later).

Laboratory Testing

When new equipment arrives, its first stop is in the ARRL Lab. The equipment is inspected for any possible shipping damage, inventoried for completeness, and then run through a series of performance tests (in most cases—some review items do not require any bench testing).

Hands-on Testing, Writing and Editing

After the Lab testing, the item is handed over (or shipped) to the designated *reviewer*, who is then responsible for putting it through its paces in real world situations. The reviewer is responsible for writing the actual review text. The Product Review Editor then edits the completed text in order to make it fit the available space and comply with *QST* style.

Afterward

Equipment that has been reviewed is generally auctioned off to members via the **prauctions** page on the ARRL Web site.

Lab Testing—Overview

Types of Testing

Although the types of equipment that generate the most interest are transceivers, receivers and amplifiers, the ARRL Lab often tests station accessories as well. These include items such as transverters, power supplies, SWR meters and just about anything else that can be tested on a bench. Antennas are not subject to Lab tests because the ARRL does not have the calibrated test range required to obtain proper gain and pattern figures.

Accessories and Specialty Items

The testing that is performed on accessories depends on the type of equipment. For example, SWR meters are checked for power accuracy and SWR, plus insertion loss, return loss (measures SWR of the meter input when the output is a proper load) and frequency range. These are all important for optimum operation.

Radio Equipment

Judging by inquiries to the ARRL Technical Information Service, transceivers create an enormous amount of interest. Of course, they are usually the first thing that comes to mind as soon as "the ticket" is on its way from the FCC. Some folks hold onto their first transceiver for decades; others trade rigs every few months. Most of us fall somewhere between these two extremes, but all of us seem to be interested in the new ones. Receivers also tend to create a lot of interest.

Lab Testing Up Close (and What the Numbers Really Mean)[1]

RECEIVER TESTING

Sensitivity

Sensitivity is a measure of a receiver's ability to make use of weak signals. One common measurement standard is called minimum discernible signal (MDS), although this is more aptly known as the receiver's *noise floor* because the human ear can often discern signals that are weaker.

The noise floor is the amount of power present in the receiver's internal noise, determined by measuring a signal level equal in power to that noise. The output of the receiver consists of the receiver's internal noise, plus the constant unmodulated tone from the signal generator. This is the "stick" by which all radios are measured for *QST*'s Product Review data tables. Typical noise floor figures for modern transceivers can be anywhere from –120 to –140 dBm. The term dBm refers to decibels relative to a milliwatt. If you think these are very small signals indeed, you are correct! Receiving a level of –120 dBm is somewhat akin to trying to view a 4 W night light at a distance of several miles without the aid of a telescope.

Unfortunately, it is impossible to duplicate "real world" conditions on the test bench—an approximation is all that can be achieved. To further complicate matters, real world conditions are different for everyone, so any attempt to duplicate a given set of conditions would only be useful to a fraction of hams. The best that can be hoped for is a consistent "yardstick" for making measurements that allow meaningful comparisons.

In the shack, the radio is connected to an antenna, and what you hear from the speaker is a combination of receiver noise and local noise (atmospheric and man-made). In some circumstances in HF and most in VHF, the receiver noise might dominate, but in the vast majority of cases, on HF the local noise predominates (even on a "quiet" band in winter). While atmospheric noise is very random (similar to the receiver internal noise), man-made noise sources are often pulse-type or otherwise very periodic in their characteristics.

The noise from the receiver that is heard by the ear is proportional to the bandwidth. Reducing the (effective or "noise") bandwidth from 500 Hz to 50 Hz will result in a 10 dB or ten times reduction in noise power. Narrow bandwidths can make a dramatic difference in how weak a signal can be copied by the ear. Most rigs do not have the capability to get that narrow, but for many to shift from a 2 kHz to a 500 Hz filter will make a noticeable improvement in the received signal to noise (S/N) for a CW signal. We use a 500 Hz filter for such measurements, when available.

Other types of sensitivity measurements include signal to noise (S/N), signal plus noise to noise (S+N)/N and signal plus noise and distortion to noise and distortion (SINAD). These are all ratios so they are measured in dB. For AM, sensitivity is often specified by manufacturers as 10 dB S/N, a level where the signal is 10 dB greater than the noise. In the Lab, we measure AM as (S+N)/N at a level of 10 dB so that all receivers can be readily compared.

On FM, the measurement standard is 12 dB SINAD. Although 12 dB might sound like a fairly high signal level, FM signals are difficult to discern with noise levels higher than this, so it is the level of minimum practical signal strength. This can be measured on a special instrument known as a SINADDER, but it can also be measured by looking at distortion on a sine-wave modulated signal because the noise is also distortion (relative to a constant-amplitude single-tone waveform).

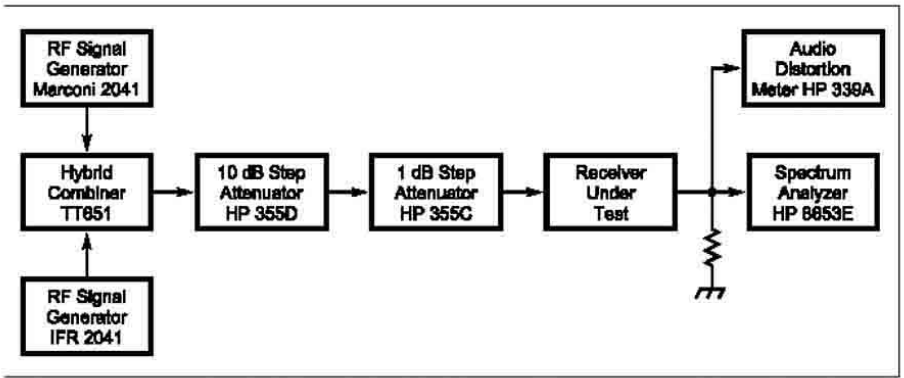

Figure 2—The test setup for measuring receiver dynamic range.

Dynamic Range

Dynamic range is generally the difference between the weakest signal that can be perceived and the strongest signal that can be present without adversely affecting that weakest signal.

Specific to receivers and transceivers, dynamic range is the difference between the receiver's noise floor and the level of strong signals that are close in frequency yet outside the receiver's passband (therefore assumed to be undesired). While receive dynamic range is a critical issue to contest operating, it can be important even to casual weak-signal DXers if they have to share a crowded band with strong local stations. Stations with high gain antenna systems are also prone to dynamic range issues.

When problems do occur, a rig's attenuator can be of help. If a particular receiver has a noise floor of –140 dBm and the local noise level is –130 dBm, adding 10 dB of attenuation will not make any difference in the weakest signals that can be perceived, yet it will reduce problems from the interfering strong signals.

Actually, it should be noted that –140 is typical only of a rig with a preamp on. It is preferable to turn the preamp off prior to adding attenuation because the preamp adds some noise of its own as well as generating undesired products. Under circumstances when the band

[1]Notes appear on page 36.

abounds with moderate signals (assuming they are ones you want to work), you can even increase the attenuation even more. A good example is operating the lower bands during the early portions of Field Day. Of course, it is usually better to have a rig with too much gain and capable of reduction rather than having not enough gain in the first place.

Blocking Dynamic Range

Blocking dynamic range (BDR) refers to a condition in which the weak signal is "blocked" or suppressed. You'll often hear this described as *desense* because the strong signal reduces the effective sensitivity of the receiver.

BDR as a lab measurement normally refers to the point at which the weak (presumed desired) signal is reduced by 1.0 dB ("blocked") by the presence of a strong (presumed undesired) signal at a frequency above or below the desired signal. The frequency difference between the two is the *spacing*. Thus, blocking dynamic range is a measure of the difference between the receiver's noise floor and the level of the signal that caused the blocking condition.

A measurement that is *noise-limited* is one in which the undesired blocking signal caused an increase in receiver noise output before the desense effect was observed. Usually, this is caused by interaction of the signal with the phase noise of the receiver's internal oscillators. It is often the case that a transceiver that has high transmit composite noise will be noise-limited on receive since the same oscillators are used for each. Some consider this to mean that a *real* BDR measurement cannot be made for that rig. The ARRL Lab considers that the effective blocking dynamic range on a noise-limited measurement is the point at which the noise increases by 1.0 dB. That point results in the same change in the signal to noise ratio as would occur had the desired signal decreased by 1.0 dB.

So, for a receiver where the noise floor is –140 dBm and the 20 kHz spacing blocking dynamic range is 125 dB, the level of signal that caused blocking effect would have been –15 dBm. To relate that to something that may be observed on a receiver, it is convenient to use S-units for discussion purposes. However, it should be noted that few transceivers follow the established S-meter standard. In that standard, S9 = –73 dBm (or 50 microvolts, for a 50 Ω system). Therefore, –15 dBm would be close to S9+60—quite a strong signal, but certainly a level that might be observed under the right conditions. Also, it should be mentioned that many radios do not have a blocking dynamic range that is that high.

Two-Tone Third-Order IMD Dynamic Range

Intermodulation describes the effect of two or more signals mixing (modulating) each other, if you will, thereby creating undesired signals on other frequencies. These signals are referred to as *intermodulation distortion* (IMD) products and they are most often created in the amplification or mixer stages of a receiver, although they can be generated in any non-linear element. In the ARRL Lab we simulate this with two carefully selected signals, as described below.

IMD dynamic range is the difference between the receiver's noise floor and the level of the unwanted signals that caused an undesired signal to appear right on the listening frequency. The process of mixing makes the largest of such signals appear at a frequency spacing equal to the difference of the two signals. For example, if the receiver is tuned to 14,020 kHz and there are strong signals at 14,040 and 14,060 kHz, a false signal may be heard because the second harmonic of 14,040 (28,080 kHz, generated in a nonlinear stage) beats with the 14,060 kHz signal to produce a difference signal at 14,020 kHz, right where we are trying to listen. Because the signal is the result of a product of a second order term and a first order term, it is referred to as a third order response. If the receiver were tuned to 14,080 kHz, it would also hear the other third order combination.

As with blocking dynamic range, IMD dynamic range can be noise-limited. In this case, the effect on the frequency that the receiver is tuned to is created entirely by the interaction of the nearest frequency strong signal and the receiver's phase noise. This results in a noise "signal" that is equal in strength to the receiver's noise floor. In this case, when the more distant signal is removed after the IMD noise is observed, then the noise would still be there. In the case of non-noise limited measurements, if you remove either signal the intermodulation ceases.

It is important to note that these lab measurements don't duplicate real-world conditions because unmodulated carriers are used for the measurements. On the air, there are usually many more than two undesired signals for the receiver to contend with. However, these tests provide an excellent means by which to compare different receivers.

Intercept Points (Third Order and Second Order)

Third-order intercept is related, as you might expect, to two-tone, third-order IMD. Now, if receivers behaved in an ideal fashion, the signals that you intend to listen to would produce a linear receiver response—that is, as the signal gets stronger, the output would get louder and as it gets weaker, the output would decrease, exactly in proportion. A 3 dB change (a doubling or halving) of the input signal power would produce the same 3 dB change in the output. Of course, real receivers don't behave quite this way. In fact, the whole purpose of automatic gain control (AGC) is to prevent changes in output with sudden input changes, helping to preserve the listener's hearing. Nonetheless, a significant portion of the receiver's response is indeed intended to be linear.

This applies only to desired signals within the passband of the receiver. Because IMD products are created by a non-linear mixing process, they change at a faster rate than the desired signal. As the undesired signals go up, the third-order distortion products also go up, but three times as fast. Likewise, when the undesired signals get weaker, the distortion products decrease three times as fast. Sharp readers will conclude that this response change is also linear—indeed that is so, but this line (if plotted) would have a slope three times the response plot of the desired signal. If these two responses were plotted on the same graph, the two lines would intersect at a point. This point is known as the third-order intercept. Actually, because the value stated is the input signal level, this is technically the third-order input intercept.

As stated earlier, real receivers are not linear over their whole input range. As a result, the third-order intercept can never be reached (or measured) because the receiver always goes into gain compression or desense before that can happen. For that reason, the third-order intercept is, strictly, a theoretical point. While its usefulness may not be immediately obvious, this figure gives a good indication of a receiver's overall strong signal performance.

The second-order intercept is similar to the third-order intercept. Second-order IMD products are produced directly from the sum and difference of the undesired signals. So while the third-order products are produced by signals that are near the desired frequency, the second-order products are often quite distant in frequency. For example, if a receiver is tuned to 14,020 kHz, then the frequency of two signals (note that there are many more possibilities) that would cause a second-order response at 14,020 are 6020 and 8000 kHz. The rate of change in the second-order products is twice that of the desired signal. The second-order intercept is then the point at which the second-order response plot would intersect the desired signal response.

IF and Image Rejection

As if the effects of multiple undesired signals were not bad enough, receivers also can experience problems created by the influence of external signals over frequencies that are intentionally present within the receiver. One such internal frequency is the receiver's first intermediate frequency (IF). In general-coverage HF receivers, this is usually a frequency higher than 30 MHz, such as 45 MHz. Even with robust filtering before the first conversion stage, some energy from strong external signals that coincide with the receiver's first IF can still find its way into the first mixer. To measure IF rejection in the lab, a signal generator set to the receiver's IF is connected to the antenna jack, and the generator output is increased until a signal appears at the receiver output that is equal to the receiver's noise floor. The difference between the noise floor and the generator level is the amount of rejection.

In many receivers, good IF rejection can be provided by sharp filter skirts at the RF stages. In wideband VHF and UHF receivers, however, particularly handheld units, the IF is often within the receiver's normal operating range or very close to it. On bands that are close to the IF, the rejection is often poor because of the modest rejection provided by the filter skirts at close frequencies.

Another example of the influence of external signals on internal ones is *image rejection*. One of the characteristics of mixers is that they produce many different products in addition to the intended one. Filtering following the mixer is intended to attenuate these undesired products, leaving only the desired IF signal. Some RF frequencies can produce images in the first mixer such that the images coincide with the IF and are therefore not attenuated after the mixing process. These signals are measured in the test of image rejection.

Testing image rejection is much the same as testing IF rejection—a signal at the image frequency is dialed up on the signal generator, and the level is adjusted for a noise floor signal on the output of the receiver. Because images are usually far removed from the tuned frequency, image rejection is often excellent, perhaps 80 to 100 dB or more. On the higher UHF bands, however, the image rejection in a handheld wideband receiver may be poor because of the very broad front-end filtering often used at those frequencies.

Other Tests (Audio Output, IF/AF BW, etc)

The tests described so far cover the "meat" of a receiver's performance, and they are usually given the most weight when comparing different models. Of course, there are other receiver performance issues that interest different folks, and these are covered in the comprehensive set of tests performed for *QST*'s Product Reviews. The audio output test gives information about the transceiver's audio performance—useful to know if you plan on using the receiver in a noisy environment. The IF/AF bandwidth test gives the net bandwidth of the receiver's cascaded IF and AF stages using its nominal filter widths. The squelch sensitivity test tells the strength of a signal that will "break through" the squelch at its minimum setting (called the threshold). The S-meter test notes the strength of a signal that indicates S9 on the receiver's S-meter. This reveals how different S-meters can be on various receivers.

TRANSMITTER TESTING

Power Output

Power output, the most straightforward of transmitter tests, gives an easily understood result. The aim in this test is simply to determine the actual power output from a transmitter in watts. While most MF/HF transceivers are designed for a nominal output of 100 W, they will sometimes exceed this figure by a few watts, or in some cases, fall just shy of the mark. Maximum output often varies from band to band as well.

Those who like to dabble in low power (QRP) operating from time to time will also want to know a transmitter's minimum output power. The ARRL Awards program defines QRP as 5 W or less power output. Many transceivers can be "throttled back" to less, but some exceed that level and that is useful knowledge for this type of operating.

Spectral Purity

FCC rules have strict requirements for spectral purity on the HF and VHF bands—these are outlined in *The ARRL's FCC Rule Book*,[2] and they are also described in detail in *The ARRL RFI Book*.[3] In addition to rules compliance, it is useful to know the amount of a transmitter's harmonic and spurious output to prevent interference to other radio services and other amateur bands—a chief reason that our allocations are generally harmonically related.

Two-Tone IMD

Transmit two-tone intermodulation distortion, or two-tone IMD, is a measure of spurious output close to the desired audio of a transmitter being operated in SSB mode. This spurious output is often created in the audio stages of a transceiver, but any amplification stage can contribute.

If you've ever heard someone causing "splatter," the noisy audio that extends beyond a normal 3 kHz nominal SSB bandwidth, then you have heard the effects of transmit IMD. Frequencies close to the transmit signal are affected the most, but depending on the amount of IMD, large portions of the band can suffer from one poor transmitter.

Carrier and Unwanted Sideband Suppression

One of the main benefits of single sideband operation is that the required frequency spectrum is greatly reduced compared to AM. It allows stations to operate close together without interfering with each other. This assumes that the reduction in the carrier and opposite sideband is sufficient to prevent interference. Thus, it pays to know the amount of suppression instead of just taking it for granted. The level of suppression is measured relative to the desired sideband. In the ARRL Lab, this is done by feeding a sine wave at a known audio frequency into the microphone input, and adjusting the amplitude level until the transceiver is operating at its rated output. Although having more suppression is almost always better, 45-50 dB or so is generally adequate.

Keying Waveform

The CW keying waveform can tell quite a bit about the way your transmitter will sound in someone else's receiver. The ARRL Lab test for this is performed using a custom-built keying generator (basically a precision timing circuit with a switching transistor output). The generator is set up to send a string of dits at a rate of 60 WPM, and the output of the very first dit and second dit are captured on a storage oscilloscope. Subsequent dits are usually identical to the second dit. This test shows whether there is any dit shortening in break-in (QSK) operation (usually there is some), it also shows the waveform shape (which can indicate a tendency to produce key clicks) and indicates the keying delay—the time from when the key is depressed until RF actually starts to appear.

The top trace in these photos is the voltage on the transceiver's key jack, as determined by the transceiver itself (since the keying generator does not put out any voltage). When the transceiver is key down, the voltage on this line will be close to zero (as it would be if you were using a straight key). When the transceiver is key up the voltage goes up to whatever value the transmitter's circuit produces while in receive. Sometimes this key line voltage can be oddly shaped (such as having a curved rise time on key up), but this is not of any consequence.

Turnaround Time Tests

The turnaround time test measures the delay between receive and transmit, and the delay between transmit and receive. This test is performed in the SSB mode (important for folks who like to operate digital modes), and in the FM mode (important for packet operators and in some cases for FM repeater operation). A transmit-to-receive delay of 35 ms or less in SSB indicates that a rig is suitable for digital operation. In FM, the receive-transmit delay determines the appropriate TNC settings for packet. If the delay (either T-R or R-T) on FM is long enough (200 ms or more), it starts to become noticeable to folks operating on repeaters—the T-R delay can cause loss of the first syllable (or part of it) of some words. The R-T delay can cause the loss of some syllables unless you remember to add a short pause between PTT and start of speaking. Long R-T delays can also lead to "doubling" in group conversations if other listeners are not aware of it.

Composite Transmitted Noise

In some receiver dynamic range measurements, you'll see a "noise-limited" figure, as discussed above. Often this is the result of an internal oscillator (such as the primary VFO) that is "noisy." All oscillators have some minor variations in their output that can be in either amplitude or frequency or both. This variation, which results in noise appearing close to the oscillator's intended frequency, is referred to as *phase noise* because it is manifested as short-term changes in the phase of the oscillation frequency.

Measuring the transmit phase noise can be done at any frequency by comparing to the output of a low noise signal generator. *QST* Product Reviews include a performance figure of composite transmitter noise. The majority of this is usually also the receiver's phase noise, but since other noise sources can also contribute, the name is a little different. If the noise level of a transmitter is high enough, it can even show up in a receiver that is close in frequency; however, the receiver's dynamic range performance is often affected at lower signal levels.

Expanded Testing (and reporting)

In the latter half of 1995, the ARRL Lab staff considered a number of ideas on how to give ARRL members more value without changing the way *QST* Product Reviews were presented. The result of this was the introduction of an expanded set of Lab tests, with the results to be included in special Expanded Test Result Reports. These are available on the ARRL Member's Web Pages (or by mail for those without Web access). The expanded reports include data on all the bands for which it is taken (*QST* reviews only report worst case figures), and includes some new tests that were not previously performed. In addition, some background on the test methods are given. More information on these expanded reports can be found in the April 1996 *QST* article, "Under the Microscope—The ARRL Laboratory's Expanded Test Result Report," by Dean Straw, N6BV. A copy of this article appears on the Product Reviews section of the Member's Pages. Because of their time-intensive nature, only some of the products that go through Product Review are selected for the expanded testing.

Hands-On Testing and Writing

After the ARRL Lab has put a piece of equipment through its paces on the bench, it goes to the reviewer. The reviewer must become familiar with the equipment by checking out all the features and functions in order to assess its ease of operation. Next, the reviewer will "put it through its paces" in real world situations, usually at a home station, to see how the equipment behaves in a practical sense. A reviewer must be thorough, and use the equipment in as wide a range of operating conditions as possible. Although the idea is to attempt to replicate the same situations that most readers will encounter there is, of course, a limit to the degree that this is possible. It is one of the reasons some Product Review items are evaluated by multiple reviewers.

Once a reviewer is finished with the new equipment, he or she must actually write about it. The reviewer must be thorough here, too, touching on all aspects of the equipment and documentation. Reviewers must be as objective as possible, avoiding the bias of personal preferences or opinion. At the same time, the reviewer may add some creativity and style (and anecdotal experiences) to make the review more readable than dry technical text.

Editing and Wrap-Up

The editing and wrap-up phases take place after the equipment has been completely evaluated. However well written, every Product Review must still undergo some amount of editing. Aside from typographical or grammatical correction, the text must also be double-checked for completeness and technical consistency.

Before the finished review is approved for publication, a copy is provided to the equipment manufacturer. This is done so that any technical errors or omissions that may have been missed in the earlier review stages can be corrected or issues resolved. Manufacturers do not have a free hand, however—only objective comments are considered for inclusion.

The last step before publication is the final editing—this is where the graphics, figures, tables and text all come together to form a "final" version of the review. This is the job of ARRL's Graphics, Production and Editorial staff (of course, they prepare *all* articles for *QST* and other publications).

What Happens Afterward

Equipment used for reviews is retained in the ARRL Lab for at least 30 days after publication of the review in *QST* to allow a retest if needed. It may then be retained by the ARRL, but most often the items are sold on the basis of competitive bids. A minimum bid is established, below current market price, and invitations to bid are published at **www.arrl.org/members-only/prodrev/prauctions**.

Where Do We Go From Here?

Over the past two decades, the Product Review column has continued to expand and improve, and testing in the ARRL Lab has followed suit. This process continues even today. As the ARRL Lab receives input on test methods from professionals in the field, we try to incorporate the latest measurement techniques while maintaining as consistent a process as possible to allow meaningful comparisons.

Notes

[1]For details of the ARRL Lab test procedures, see the ARRL Lab Test Procedures Manual at **www.arrl.org/members-only/prodrev/testproc.pdf**.

[2]*The ARRL FCC Rule Book*, Thirteenth Edition Chapter 4, pp 44-48. Available from ARRL dealers or the ARRL Bookstore for $12.95 plus shipping. Order number 9000. See **www.arrl.org/shop/** or call toll-free in the US 888-277-5289, or 860-594-0303.

[3]*The ARRL RFI Book*, Chapter 17, pp 9-11. Available from ARRL dealers or the ARRL Bookstore for $24.95 plus shipping. Order number 6834. See **www.arrl.org/shop/** or call toll-free in the US 888-277-5289, or 860-594-0303.

Michael Tracy, KC1SX, is an ARRL Test Engineer. He can be reached at **mtracy@arrl.org**.