

An Optimized QRP Transceiver

A rig doesn't need to be complex to work well. This 40-meter cw transceiver, designed for performance, ease of operation and low power consumption, is a case in point.


By Roy W. Lewallen,* W7EL

Many QRP rigs have been described in Amateur Radio publications over the years. The distinguishing characteristic of *this* transceiver is that it was designed and then optimized for high performance. It's relatively easy to build something that works, but it usually requires a great deal more effort to produce something that works really well. That effort has gone into this rig, and the result is a 40-meter cw transceiver with the following characteristics: full electronic break-in; clean keying and smooth, quiet transceive operation; *stable* VFO coverage from 7.0 to 7.15 MHz; receiver incremental tuning (RIT); single 12-volt supply operation; two-watt power input, 1.5-watt output into a 50-ohm load; receive current drain less than 20 mA; reasonable transmitter efficiency; high-performance direct-conversion receiver; and small size (1-1/2 × 2-1/2 × 3-1/2 in. [40 × 70 × 90 mm]).

This is *not* a step-by-step construction article. Rather, the purpose of this article is to share some of the many things I learned from designing, building and perfecting the transceiver. Very little of the article is devoted to mechanical packaging and, since there are no printed circuit boards in my rig, none are available from the author. I hope that this article will help potential designers of such gear to avoid some of the pitfalls I've encountered, in addition to provoking thoughts about how to make *good-quality*, simple rigs.

Some Underlying Philosophy

"High-performance direct-conversion receiver" may seem to be self-contradictory. After all, direct-conversion (DC) receivers are so simple they can't possibly compete with a good superhet, right? Wrong! DC receivers have only *one* significant disadvantage when compared to superhets: the presence of an audio im-



This diminutive QRP transceiver is a joy to operate. It features a high-dynamic-range receiver, smooth break-in operation, RIT and a host of other high-performance features.

age which doubles the amount of noise and interference heard.

The only other inherent disadvantage is the inability to generate other than audio-derived agc. The same careful attention to detail and potential problems is required in designing the DC receiver as is required for a top-quality superhet, if comparable performance is to be realized. This last point is frequently overlooked, and that may be one reason why the DC receiver is often looked upon as a mediocre performer.

All other problems can be overcome with careful design, and even the two inherent disadvantages can be overcome to some extent. On cw, narrow af filters may be used, reducing the image bandwidth along with the desired signal bandwidth. RIT helps also: When an image signal produces the same beat note as the desired signal, adjusting the RIT will move one up in pitch and the other down, thus separating them.

As for agc, this rig does without, and I've hardly missed it. This receiver is on a par with all but the best superhets for any type of operation, except perhaps during contests in conjunction with a high-power transmitter, but at a fraction of the com-

plexity. Note also that to use a superhet in a transceiver, an additional oscillator and mixer must be added to the *transmitter* to convert the VFO to the transmit frequency. In a transceiver using a DC receiver, the required shift is only a few hundred hertz, and can easily be accomplished by pulling the VFO.

Since the rig was designed for portable use, current drain was a major consideration. My experience indicates that many solid contacts may be had using simple antennas and operating during the night with 2 watts on 40 meters. This power level is also more than adequate for short-range daytime operation. Power drain is low enough that the rig will run for about a week of evening operation from one charge on ten NiCad "A" cells (660 mA-h).

The small size precludes wide-range antenna impedance matching — a necessity for field use — so a Transmatch was built in a separate box. The circuit for the Transmatch was taken directly from the reference (page 167).¹

While crystal oscillators have

¹Hayward and DeMaw, *Solid State Design for the Radio Amateur*, ARRL, 1977.

advantages for certain types of operation, a VFO is preferred in a rig which is intended primarily for ragchewing (and, I can't resist, a miniscule amount of DX-ing) and Field Day. Full electronic break-in was taken on as a challenge, and the convenience it offers is well worth the effort. RIT was originally left out of the design for the sake of simplicity. I later decided that RIT is a necessity in a transceiver, no matter how simple it may be, so an RIT circuit was added.

Many of the circuits and concepts used here were taken directly, or with some modification, from the reference. The following discussion concentrates on the unique features of the circuits used, rather than on basic principles or those well covered in the reference.


The VFO

The oscillator, Fig. 1, is a good example of the principle that a circuit doesn't necessarily need to be complex to work well, if properly designed. This simple Hartley circuit exhibits less than 200-Hz warm-up drift, with about half of that occurring within one minute after turn-on. This drift performance is completely repeatable, as the building of many such oscillators has shown. The circuit used here is the result of a considerable amount of experimentation directed toward identifying the sources of drift in such VFOs. Without giving the details of the experiments, I'll summarize the results.

1) No part of the VFO circuit except the FET drain should be connected to a pc-board pad that is over a ground plane or near other pads. It's best to avoid pc construction (including ARRL "universal breadboard") of the VFO altogether, because the capacitances formed with the board material as a dielectric have extremely poor temperature and humidity characteristics. I prefer building VFOs using point-to-point wiring on standoffs above a ground plane of copper-clad board (copper side up).

2) Use NP0 ceramic capacitors. Commonly used polystyrene units are predictable, but have too strong a temperature coefficient to compensate a decent inductor. If a rather poor inductor is built, one might get lucky and have the considerable drifts cancel, as they are in opposite directions. But I don't consider that to be a good approach. I prefer to first reduce the temperature dependence as much as possible, then if necessary, compensate what's left. The NP0 ceramic capacitors have a much lower temperature coefficient than polystyrene or silver-mica types. "Dogbone" NP0 units have a black-painted end, and some NP0 disc capacitors are marked "NP0."

3) The gate diode is essential to minimize drift, for reasons put forth in the reference. A 1-megohm gate resistor provides better drift performance than the sometimes recommended 100-kilohm



Here's the complete W7EL QRP station, ready to pack for portable use. At the left, the box on top contains 10 "A" NiCad cells. Below it is a Transmatch. Below the transceiver itself in the center is a keyer and, at the right, for when the going gets rough and batteries are plentiful, is a 10-watt amplifier "brick."

value, possibly because of reduced tank loading.

4) The temperature of the FET itself has a negligible effect on this circuit. Therefore, circuits which more loosely couple the active device don't have any significant advantage over this one.

5) After the above recommendations have been followed, the only remaining significant source of drift is the inductor. Of the inductors I've tried, the best are those wound tightly on type-6 powdered-iron toroidal cores, with core size being relatively unimportant. A technique suggested by W7ZOI is to anneal the coil after winding, which I do by boiling it in water a short while, then letting it cool in air. This noticeably reduces drift, and this method was used to obtain the quoted drift.

If extreme environments with rapid temperature changes are to be encountered, you may want to compensate the VFO. This can be done by replacing part of the fixed capacitance with negative temperature coefficient (TC) capacitors, such as polystyrene or negative TC ceramic units.

It should be possible to make other oscillator types perform as well as, or better than, this one, as long as the above guidelines are followed. The secret, however, lies in the choice, rather than the number, of parts.

The rig had been used for a year without voltage regulation for the oscillator, and with no difficulty with chirp or hum. Supplies used have been a NiCad battery, an ac supply using a 3-terminal regulator, and fresh lantern batteries. This was possible because the sensitivity of the unregulated oscillator is

only 50 to 80 Hz/volt from 9 to 15 volts. A regulator was added when experiments showed noticeable modulation of received audio (and, presumably, transmitted rf) when a small amount of ac was purposefully introduced to the supply. It can now be used with poorer ac supplies or an automobile power system with the engine running.

The buffer, although designed for low current drain, is the major power consumer in the receiver, requiring 10 mA. The key to efficiency in this sort of buffer is to choose the transformer turns ratio to sustain as large a voltage swing at the output stage collector (or drain) as possible. Another potentially efficient approach is to use a complementary-symmetry stage. One was used for some time, but its temperature-stable, low-distortion design consumed as much power as the present buffer, and was more complex. Buffer voltage gain is approximately one half, providing about 2.5 volts pk-pk output.

The RIT circuit uses a Zener diode as a voltage-variable capacitor. While Zener diodes are inexpensive and readily available, their nominal capacitances may vary a great deal with different manufacturers. An empirical procedure to adapt the circuit to an individual diode is to select a series capacitor (here 15 pF) to obtain a tuning range of about 1300 Hz with a diode reverse bias variation of about 9 to 4 volts. When the control is adjusted to the center of its range, the frequency shift should equal the center frequency of the receiver-audio filter (about 650 Hz). During transmit, or when the ZERO button is depressed, the shift is removed, causing the transmit frequency to be the same as that of a received signal peaked at the

audio-filter center and tuned to the correct side of zero beat.

The Transmitter

The transmitter is a fairly efficient (75%) Class C design. The Zener diode was added after twice blowing the output transistor by inadvertently transmitting with the antenna disconnected. The diode protects the output transistor from this hazard. Some caution is necessary when using a Zener diode at the output-stage collector, as many Zener diodes have a large amount of shunt capacitance. When adding the diode, the collector capacitance must be reduced by an amount approximately equal to the capacitance of the diode when it is reverse biased by the collector supply voltage. In this transmitter, the total capacitance at the collector should equal approximately 450 pF, including the fixed capacitor, the 51-pF receiver-pickoff capacitor, the Zener diode and the transistor (about 10 pF for this type). If the capacitance of the diode can't be measured, the 385-pF fixed capacitor should be made variable and adjusted for best transmitter efficiency.

The value of L2 is not critical, as long as it's not much smaller than the 10 μ H shown. Conventional solenoidal rf chokes will work fine also, but toroids are required in a tightly packed rig such as mine to keep mutual coupling acceptably low.


The Receiver

More time has been spent improving the receiver than any other part of the rig. The end result is no more complex than the first version, but the improvement has been great, again showing that complexity and performance don't equate. In the following discussion I'll relate why some types of circuits were chosen and others rejected.

Receiver signal pickoff is through the 51-pF capacitor from the transmitter output filter. When transmitting, the diodes protect the receiver and cause the 51-pF capacitor to become part of the transmitter output network. When receiving, the capacitor and L5 make up a fairly low-Q series resonant network to reduce signal attenuation by maintaining an approximately 50-ohm source impedance to the mixer. The additional filtering it provides is helpful also.

The mixer is a conventional doubly balanced type. Unfortunately, I didn't choose this by accident — it was selected after a good deal of frustration trying to use other kinds!

I'll digress here a moment to explain about a-m demodulation, a problem which is common in "simple" direct-conversion receivers (but not because they're simple!). Direct-conversion receivers have most or all of their gain at audio frequencies. Thus, if any device near the receiver input is nonlinear — such as forward- or reverse-biased diode



High component density is necessary to allow the author to squeeze all the circuitry of the transceiver into such a small package. Point-to-point wiring is also used. The transmitter circuitry is at the upper left in the photo above. The output transistor, which is bolted to the case, is hidden below the top layer of components.

junctions — audio from strong shortwave broadcast and a-m broadcast stations or ssb stations is detected. If passed through the mixer, this audio is amplified and appears as an annoying "din" in the background — or foreground, if severe enough! Leakage of the local-oscillator signal into the circuitry preceding the mixer definitely aggravates the problem, but I haven't attempted to isolate the (apparently) several phenomena involved. An often-overlooked point is that the audio amplifier itself will usually happily rectify any rf which reaches its input, amplifying the resulting detected audio.

A common solution is to use very selective tuned circuits at the receiver input, a solution not practical in this case. Another is simply to avoid using nonlinear elements ahead of the mixer — and following the mixer, if rf can get through it. This is difficult when using electronic T-R switching — this rig has the T-R diodes, Zener diode and PA transistor as potential culprits. Yet another solution is to use a balanced mixer which will, in theory, prevent detected audio from getting through the mixer. I took this last approach, first trying an MC1496 IC mixer. Try as I did, I was never able to obtain good rejection of a-m signals originating over a wide frequency band. The balance seemed to depend on the source impedance which, of course, changes with frequency when an antenna is the source — to say nothing of the transmitter output network and series-resonant network in the path. My attempts included different biasing and signal levels, and driving the inputs through baluns.

I next tried a singly balanced diode mixer, with and without two extra diodes for improved balance. I used this for quite a while, and it was quite satisfactory after I replaced the T-R diodes with an MOS transistor switch. The noise figure was

marginal, however, and there was still some background a-m interference when propagation was good. While trying to improve that situation, I performed quite a few experiments using the mixer with and without the extra diodes, with the input ports exchanged, with Schottky and

Fig. 1 — The circuit of the W7EL 40-meter QRP transceiver. Resistors are 1/4 or 1/8 watt, 5%. All ferrite cores are available from Amidon Associates. When winding the inductors that use BLN-43-2402 cores, the wire should be passed once through both holes of the core for each "turn" specified. See the illustration at far right.

- C1 — 1- μ F, 3-V non polarized ceramic.
- D2, D3, D5-D13, incl. — Silicon general-purpose/switching diode; 1N914, 1N4152 or equiv.
- D4 — Zener, 33-V, 400-mW; 1N973 or equiv.
- D14 — Zener, 10-V, 400-mW; 1N961 or equiv.
- L1 — Approx. 3 μ H; 26 turns on a T-44-6 core. Tap at seven turns from ground end.
- L2 — Approx. 10 μ H. 43 turns on a T-50-2 core.
- L3, L4 — 1 μ H; 19 turns on a T-37-6 core.
- L5 — 9.4 μ H; 58 turns on a T-37-6 core.
- Q1, Q11 — Silicon n-channel JFET, 300 mW, 2N4416.
- Q2, Q3, Q10 — General purpose, silicon npn, 310 mW, 2N3904.
- Q4 — General purpose, silicon npn, 1.8 W, 2N2222.
- Q5 — Rf power, silicon npn, 7 W, 2N3553 or 2N5859.
- Q6, Q7 — General purpose, silicon pnp, 310 mW, 2N3906.
- Q8 — General purpose, silicon npn, 310 mW, 2N4124 or 2N3565.
- Q9 — General purpose, silicon npn, 310 mW, 2N3565.
- RFC1 — 100- μ H subminiature choke, wound on a 1/4-watt-resistor-sized ferrite form. Dc resistance is approx. 8 Ω .
- T1 — Primary 15 turns, secondary 3 turns. Wound on a BLN-43-2402 core.
- T2 — Primary 39 turns (approx. 6.7 μ H), secondary 5 turns. Wound on a T-44-6 core.
- T3, T4 — Five trifilar turns on a BLN-43-2402 core.
- U1 — Op amp, LM301.
- U2, U3 — Dual op amp, LM358N (one section of U3 unused).

conventional diodes, and with various VFO source impedances. None were satisfactory with respect to both a-m demodulation and noise figure.

These problems virtually disappeared when I replaced the mixer with a doubly balanced type. Additional improvements were significant when I designed the present preamplifier and input diplexer. Now, a huge a-m signal is required to cause interference, and none has been heard since it was implemented. The noise figure is very good, with a minimum-discernable signal level of less than 0.1 μ V. Other balanced mixers, such as CA3028 IC or discrete balanced JFET mixer, might match the a-m demodulation characteristics of this receiver, but they probably won't match its signal-handling capability. A signal 50 kHz away must be greater than 100 mV in amplitude (120 dB above the minimum discernable signal) to have any noticeable effect on a medium-amplitude (30- μ V) received signal. Try that test with *your* station receiver!


Following the mixer is a diplexer. Its purpose is to provide a wide-band 50-ohm termination for both rf and af, while preventing rf energy from getting into the af amplifier and preventing af energy from being wasted in the rf termination. The rf termination consists of the 0.1- μ F capacitor and 51-ohm resistor; RFC1 and the 0.47- μ F capacitor form a low-pass filter which prevents any residual rf from reaching Q10, thus greatly enhancing immunity to a-m. Q10 presents an input impedance of approximately 50 ohms for maximum power transfer.

Receiver Audio

Because the receiver audio gain exceeds 100 dB, great care must be taken to prevent feedback or amplification of power supply hum. This receiver uses an active decoupling circuit consisting of Q9 and associated parts to avoid these problems in the input stage, where the sensitivity is greatest.

Following the preamp is the active audio filter. This one is a peaked low-pass type with a Q of five — low enough to keep ringing unnoticeable. It is simple, noncritical and adequate for general operating. The peak frequency is about 650 Hz, which corresponds to the transmit-receive frequency difference with the RIT control centered. An LM301 is used because of its low noise and relatively low current drain. A TL071 or TL072 should give comparable performance, and one section of an LM358 may be used with a 2-dB increase in noise figure, an amount I feel is quite acceptable.

The last two stages are conventional amplifier stages, with frequency response rolled off outside the range of about 150 to 1500 Hz. The gain distribution (31 dB in the first stage, 52 dB in the second) is unusual, but not for any special reason — it just evolved that way.



The output spectrum of the Optimized QRP Transceiver. Vertical divisions are each 10 dB; horizontal divisions are each 5 MHz. For a transmitter of this power level, FCC rules require that spurious outputs be suppressed at least 30 dB. The 14-MHz second harmonic can be seen at about -37 dB. The pip at the extreme left is the zero-frequency reference, generated within the spectrum analyzer.

U3 is a sidetone oscillator. The reference suggests keying just the bias resistor, but this doesn't work with the LM358, as the negative supply voltage (ground) is an acceptable input. It will still oscillate (at a very low frequency) with both 10-k Ω resistors grounded! Therefore the IC supply line is keyed also. Sidetone injection level is set by the 100-k Ω resistor at pin 6 of U2B; this may be varied to suit individual taste.


Keying and T-R

Three events must occur when this type of transceiver is keyed, and they must occur in the proper sequence if operation is to be clean. They are, in order: (a) receiver mutes, (b) VFO shifts frequency, and (c) transmitter keys. These events must occur in reverse sequence when switching back from transmit to receive. The sidetone oscillator must also be keyed, but its timing isn't as critical.

Attention to this sequence and proper transmitter waveform shaping makes the difference between a poor-sounding rig and a really clean one. Many people have been surprised to learn my power input — because "it doesn't *sound* like a QRP rig." Only a few parts are required to accomplish this. In addition, it's easier to copy a clean weak signal than a poor one, so good keying and freedom from chirp, clicks and roughness are particularly important for QRP transmitters.

The receiver is muted by Q11 which acts as a series gate. Q10, U1 and U2A are all driven to saturation for a while when the transmitter is keyed, and again when the key is released. This is caused by the relatively large rf signal appearing across the T-R diodes, as well as stray rf pickup by the receiver. Such a disturbance is impractical to eliminate, as it would require a T-R switch with very high attenuation, and extensive shielding. I solved the problem by other means: Q11 is turned off immediately when the transmitter is keyed, then turned back on after the disturbance

is over, about 60 ms after the key is released. The diodes around U2A prevent the output of U2A from swinging to ground during the disturbance, a condition which turns Q11 on when it should be off. I find the 60-ms delay to be ideal, as it removes distractions between dits and dahs at medium speeds while being short enough to provide essentially instantaneous break-in. The disturbance (hence, required delay) could possibly be reduced further by limiting the swing of either or both Q10 and U1, or biasing U1 and U2A outputs closer to the positive supply voltage.



The keyed CW waveform of the Optimized QRP Transceiver. The horizontal divisions are each 5 ms. The upper waveform indicates the actual key-down time. The rise and fall times of approximately 2 ms result in a crisp sounding but clickless signal.

Oscillator frequency shift is obtained by changing the bias on the Zener diode (used as a voltage-variable capacitor) in the VFO circuit when going from transmit to receive or vice versa. The timing is provided by Q8, which comes on fast when the key is closed, but goes off some five milliseconds after the transmitter output drops to zero following the release of the key. Shaped transmitter keying is provided by Q7 and associated components. I found that simultaneous keying of the base and collector circuits of the driver stage was required to give the desired rise and fall times of a few ms at the transmitter output. The sidetone oscillator is keyed from the same line.

Construction

I discourage others from attempting to duplicate the construction of my unit. To do so requires access to subminiature parts and several no-longer-available items, a good understanding of potential crosstalk, shielding and ground problems, and a large amount of patience. There are, however, a few points which may be of interest to those wishing to build similar gear.

A great deal of information is available regarding bypassing, decoupling and layout techniques. If the potential builder isn't familiar with these basics, construction of a similar unit may cause a great deal of frustration indeed. I would suggest, as a minimum, that construction be over a ground plane, as shown in many

QST articles and in publications such as the reference, or *The Radio Amateur's Handbook*.

I am certainly no expert on miniaturization, and this rig doesn't by any means approach the ultimate in that regard. My only general advice is to begin with the box and build the rig into it, rather than the other way around, and get a good idea of the placement of controls, connectors and large components before you begin. Since it's difficult to troubleshoot or modify such a rig once built, ideally a larger breadboard version should first be constructed, perfected and operated. When I got really pressed for space, I found that building the circuitry on small pieces of perfboard and mounting the boards vertically allowed very dense packing. It helps a great deal to mount components on both sides of the board, and to ignore the usual conventions of placing parts in neat rows. The use of 1/8-watt resistors saves a surprising amount of space compared to 1/4-watt units. Another great space-saver is the use of tantalum, rather than aluminum, electrolytic capacitors. Small parts are nearly always more expensive and less available than their larger counterparts, so each builder must decide if the trade-off is a good one.

Adjustment and Operation

The only adjustments required are the

VFO trimmer, used to set the VFO frequency at the lower band edge; the drive level pot, used to set power input at 2 watts (although no major problem will arise if driven at higher or lower levels, efficiency may drop slightly); and the transmitter rf-amplifier tank circuit, which is peaked at the center of the frequency range. None should require re-adjustment once set.

Operation is, by design, simple. The only point worth noting is that, as with any direct-conversion receiver, signals must be tuned on the correct side of zero beat so that the transmitter will be on the same frequency as the received signals. Guest operators have picked this up in a few minutes, so the SPOT button is seldom used. When the rig is new, however, it's nice to have the assurance of knowing just where the transmitter will be when the key is pressed.

I do want to emphasize that this isn't just a "paper design," but a rig which has undergone a good deal of operation at W7EL and, on many occasions, from portable locations, including Field Day operation. The first version was built about two years ago. Nearly all states, as well as a few DX stations, have been contacted using simple antennas. I enjoy ragchewing, and countless enjoyable QSOs have been had with this rig. It's a pleasant experience anytime to operate a stable, clean, full-QSK, essentially crush-

proof rig. And, to my taste, to do this from a backpacking tent or cabin at the beach enhances the pleasure even more. The very best part, however, was best stated in the closing paragraph of the reference: *That* is "where it's at."

Closing Remarks

I hope that this article has illustrated a few important points: that simplicity and performance aren't mutually exclusive, that a well-designed direct-conversion receiver is a good receiver indeed, and that really good designs don't generally just happen. I also hope that some readers are moved to question the statements I've made, and those which have been made elsewhere, so that more of the subtleties of simple solid-state gear can be widely understood. Most of all, I hope that this will be of help to people who were puzzled, as I have been, by some of these phenomena.

To the extent that time permits, I will be glad to answer questions. An s.a.s.e. would be appreciated for inquiries. No circuit boards, board layouts, parts kits or parts availability information are available from the author. I'm too busy working on my *next* rig! I wish to thank Wes Hayward, W7ZOI, for his comments, criticism and encouragement during the design and testing of this rig, and the writing of this article. 