

CRYSTAL SETS TO SIDEBAND
© Frank W. Harris 2006, REV 10

Chapter 13

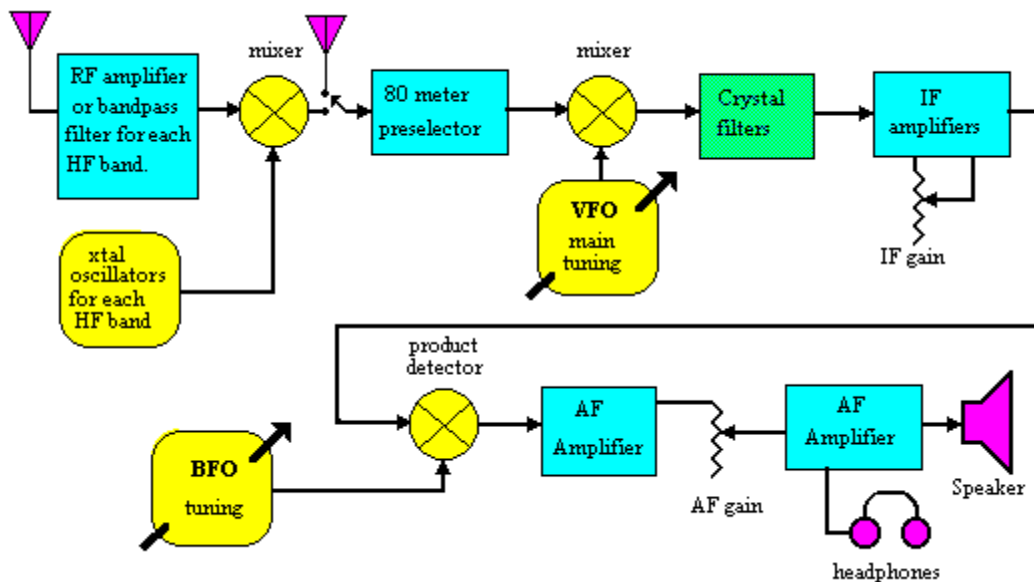
BUILDING A HOMEBREW HF RECEIVER

The Vanishing Art

The 1986 ARRL Amateur Radio Handbook reported that hardly anyone was building homebrew ham receivers. Since then, handbooks have had fewer and fewer practical details on building radios. Out of hundreds of contacts, so far I've only worked three guys, George, K7DU, Mike, NØMF, and Biz, WDØHCO, who were using homebrew receivers for the QSO. Two of these receivers were made from vacuum tubes. The third was an old-fashioned super-regenerative design. However, the super-regenerative was implemented with modern FET transistors and integrated circuits. All three of these receivers must have worked well since because they had no trouble hearing me on 40 meter CW. I talked to one other fellow, Gil, N1FED who told me he had just finished a vacuum tube receiver. Unfortunately, it was performing so poorly he was still using his modern transceiver on the air. Gil told me he didn't like transistors. I guess he found printed circuit boards and those pesky oscillations too much trouble.

In spite of this pessimism, *you CAN build transistorized receivers that work reasonably well.* I built mine because I was intrigued by mysterious circuits like “balanced mixers,” “product detectors,” “cascode amplifiers,” and “crystal ladder filters.” Before this project, I could recite the purposes of these circuits, but I had no “feel” for how they worked and why receivers are designed the way they are. What better way to learn than to build one?

What's a reasonable goal?



An “adequate performance” communications receiver

2. Chapter 13, Harris

My receiver is based on the “High Performance Communications Receiver” designed by W7ZOI and K5IRK described in most of the annual ARRL Handbooks in the 1980s. In my opinion “High Performance” is optimistic, but certainly “adequate performance” is realistic. I define *adequate sensitivity and noise figure* to mean that I can hear the DX and QRP stations that other stations are working. Before I built the receiver described here, I often had the impression I was hearing only the loudest signals. For me, *adequate selectivity* means that it's good enough for CW QSOs in the evening on 20 and 40 Meters. On these bands there are often dozens of narrow CW signals operating within a few hundred Hertz of each other. With a 10 KHz bandpass you may hear many stations simultaneously and not be able to copy any of them.

Adequate sensitivity will allow you to hear most QRP signals. I believe that 45 years ago hardly anyone had receivers that were “adequate” for QRP contacts. When I was a novice, my first transmitter was a 7-watt homebrew for 40 and 80 meters. It was a close copy of a design in the 1957 Handbook. I know it worked OK, because I talked to my novice buddies around town. Unfortunately I hardly worked anyone outside of town. It wasn't until I bought a 50-watt commercial kit, just like all the other novices, that I was able to talk to all the same stations my buddies were working. I was still using the same dipole antenna, so I can only assume the improvement was the extra power.


The sensitivity of the receiver described in this chapter is well under 0.5 microvolt on 80 meters and lower bands that have no RF preamplifiers. On the upper bands where the receiver has preamplifiers I could hear a calibrated signal source at 0.02 microvolts. Wow! No wonder I can hear those QRPs. In the old days sensitivity less than 1 microvolt was considered hot stuff.

Another issue is *adequate stability*. When your receiver is equipped with sharp crystal filters, it is vital that the VFO and crystal oscillators are stable enough that the signals you're listening to do not drift in and out of your passband. If you build a VFO like the one described in chapter 10, you will have no drift troubles.

Does it have to be so complicated?

Looking at the block diagram above, each one of those blocks represents one to three transistor stages. The front-end converter has three transistor stages for each separate HF band. That means you need to build about 20 transistor amplifier or oscillator stages for the converters to cover all the bands. You're probably wondering if there isn't some simple receiver you can build that will get you on the air rapidly. The best news is that you can build the above receiver in stages. The core of the all-band HF receiver is a quality 80 meter receiver. You can build that first, then at least you'll be in business on 80 meters. In the beginning you can also do without the loudspeaker and multiple crystal filters. After you have a functioning receiver, you can add features and the converters to hear the other ham bands.

Yes, you can build a less complicated receiver, but I doubt it will “adequate.” My direct conversion receiver in chapter 7 worked well, but was not selective enough. There are modern super-regenerative receiver designs with two or three integrated circuits that *MAY* be pretty good, but I haven't built one, so I'm not convinced. Chapter 14 describes a vacuum tube regenerative receiver that was great fun to build and quite good for listening to foreign short wave broadcast stations. Unfortunately, it was *NOT* selective and sensitive enough for ham communications. In summary, ***YES, a decent ham receiver does have to be complicated.***



A homemade ham receiver built in 1967

A homebuilt ham receiver from 40 years ago is shown above. It has 11 tubes, a simple crystal filter, and covers 80 through 10 meters. It doesn't cover the WARC bands. Yes, it works OK. But compared to the all-transistor receiver described in this chapter, it is insensitive, noisy, and has poor selectivity.

Being realistic, any receiver you build is unlikely to match the performance of high-end commercial rigs. But every time your receiver brings in DX on a new band or whenever you conquer one of the dozens of glitches you will encounter, you'll have a thrill and pride you'll never get from a commercial rig. If you decide to build your own version of the W7ZOI / K5IRK receiver, I recommend you find a copy of an old ARRL Handbook from the 1980s and Xerox their original descriptions. You'll find they built most circuit blocks differently than I did. Going back to the original description may give you some useful ideas. Perhaps their version will work out better for you.

PLANNING YOUR RECEIVER

Superhetrodynes offer crystal filters for CW

A superhetrodyne uses a mixer to produce a constant intermediate radio frequency (IF). This intermediate frequency signal is always the same so it can be filtered with fixed crystals or mechanical filters to establish bandpass widths for CW and upper and lower SSB. Before you commit to any design, make sure you can buy the critical parts you need, especially the crystal or mechanical filters for your IF. For example, many receiver designs use a 455 KHz IF. Unfortunately, I have yet to find an easy source for 455 KHz crystals for building filters and

4. Chapter 13, Harris

BFOs. Consequently, I have avoided this frequency. Among homebuilt ham receivers the most common IF frequency seems to be 9 MHz.

Why not single conversion?

I had always wondered why homebuilt all-band HF receivers are almost always dual conversion. It turns out that the fundamental challenge of homebuilt receivers and transmitters is building a stable VFO. Yes, you can build a reasonably stable VFO, but homebrew VFOs usually don't have much tuning range. 0.5 MHz is typical. And, in order to drift as few Hertz as possible, the VFO needs to be relatively low frequency. Homebrew VFOs are usually in the range of 2 to 7 MHz. The disadvantage of a low frequency VFO is that its harmonics will appear as one or two loud whistles on some upper HF ham bands.

Compared to the practical VFO range of 2 to 7 MHz, the HF spectrum is huge, 1.8 to 30 MHz. Right away one can see that a homebrew direct conversion 10 meter receiver is difficult because it needs a stable VFO that will tune 28 to 29.7 MHz. This problem can be solved by "converting" the VFO oscillator up to 28 MHz using a crystal controlled oscillator and a mixer plus 28 MHz filter/amplifiers. This complexity ruins the simplicity of direct conversion. If you're going to operate above 40 meters, you may as well build dual conversion like the rest of us.

How do modern receivers do it?

Modern receivers use integrated circuit frequency synthesizers to generate stable VFO signals anywhere they like. Sometimes modern HF receivers escape from artifact images and harmonics by using an IF frequency way up in the VHF range. In addition, after the initial mixer stage, some commercial receivers use multiple conversions to get the signal back down to an audio output. At each conversion stage, different kinds of filtering are applied. For example, the Yaesu FT1000MP has four down-conversions from an 89 MHz IF. (!) This includes the digital signal processor with its 32 KHz input.

In a superhetrodyne the VFO interacts with the incoming RF signals to produce an intermediate (IF) frequency. A 5 MHz VFO implies that the IF is going to be within 5 MHz of the band or bands it covers. Such a receiver might cover 28 MHz, but that would imply an IF of 23 MHz or possibly 33 MHz. The lower bands would be out of range unless the VFO could be tuned over many MHz. Consequently, a single conversion homebrew superhetrodyne can only cover one band well and can't possibly cover the whole spectrum.

In some old ham designs the VFO tuned 5.2 to 5.7 MHz. They used a 1.7 MHz IF and either subtracted or added the IF to the VFO frequency to cover either 80 or 40 Meters. Specifically, 5.7 MHz *minus* 1.7 MHz = 4.0 MHz and 5.3 MHz *plus* 1.7 MHz = 7.00 MHz.)



Homebuilt all-band dual conversion HF ham receiver.

Start with a single band, single conversion superhetrodyne

My "adequate receiver" is shown above. At the upper right is the preselector tuning for peaking or decreasing the sensitivity of the 80 meter IF. This control is useful as an attenuator when tuning in SSB stations. Above the frequency dial is a bargraph S-meter. The S meter is useful for zeroing in my transmitter VFO onto a station I wish to call. Unfortunately, the bargraph produces digital switching noises that often interfere with weak signals. When I substitute the analog meter on the left, the noise vanishes. Difficulties like this are typical of the real world of building your own equipment. Try to learn from these challenges and find pleasure in meeting them. Building perfect equipment is not a practical goal for us homebrewers.

Begin by building the best single-band receiver you can. The core of my receiver tunes 80 Meters. To cover the other eight HF bands, separate crystal controlled converters convert the signals down (or up) to 80 meters. This makes the receiver "dual conversion" on every band except 80 Meters. In the old vacuum tube days each homebuilt amplifier stage was designed to work on several bands. This meant that the front panel of a 1970 homebrew receiver usually had a bunch of tuning knobs to tweak each stage. Today, a transistor amplifier or mixer stage is physically tiny, so it's practical to build a separate converter and amplifier pre-tuned for each band. My experience was that building the converters to translate the bands to 80 meters was relatively easy. Unfortunately, building a decent 80 meter core receiver was much harder. Maybe with the help of this chapter, you'll find the core receiver easier to build than I did.

My 80 meter receiver has a 9 MHz IF. 9.00 MHz crystals are available for less than \$1

from Digi-Key and Mouser. The low price is important because, depending on your filter plans, you may need as many as 11 or more 9 MHz crystals. The downside of 9.00 MHz was that I had to wait months for the crystals to arrive.

An 8 MHz IF wasn't such a good idea

At first I used the more common 8.0 MHz crystals. Unfortunately, to receive 4.0 MHz, the VFO had to tune 4.0 MHz. I expected the 4.0 MHz VFO signal would be “a little birdie” that would mark the high end of the band. I thought this “edge-of-band marker” would be kind of convenient. Instead, the “birdie” was more like a screaming siren that overwhelmed the IF and made the upper end of the 80 meter ham band unusable. So when the 9.0 MHz crystals eventually arrived, I rebuilt everything for 9 MHz. Now the VFO (the big tuning knob on the above receiver) tunes 5.0 to 5.5 MHz to cover 4.0 to 3.5 MHz. That is, $5.0 \text{ MHz} + 4.0 \text{ MHz} = 9.0 \text{ MHz}$. Of course the BFO frequency also had to change from 8 MHz to 9 MHz.

An unusual adventure

Once your receiver begins to work, you'll have interesting glitches. Until I got my 80-meter preselector filter working, I usually heard rap music from my local 1190 AM radio. Also, the 31 meter shortwave band is just above the 9 MHz IF. Before it was aligned, I was hearing sermons from HCJB in Quito, Ecuador. Later, my 20 meter converter was overwhelmed by Dr. Scott, a Los Angeles evangelical minister, who preaches on 13.8 MHz. Once I had my modules tuned and sealed, Dr. Scott and his friends were silenced. Actually, I got a kick out of these problems.

Building a receiver revived my interest in shortwave listening. I've had shortwave radios since I was a kid. Some of them, like my Army surplus Collins R-388/URR, were excellent. In spite of this, I rarely listened when I wasn't actively hamming. But once my homebrew receiver(s) began to work, I found myself exploring the bands as never before. For instance, on 80 meters I was amazed to hear hams from all over the continent. I had heard about guys who work DX and earn WAS certificates (Worked All States) on 80 meters, but I never really believed it. I have even worked QRP stations on 80 meters. 80 meters is usually so noisy, I didn't know that was possible. Until I built this receiver, I had never heard “spy code stations” before. Some of them have an announcer reading what sound like random letter groups, but most of them are CW “spy stations” sending Morse code 5 letter groups, just like the WWII Enigma signals. I often hear them on 10 and 30 meters. As I built converters for each of the HF bands, it was like hearing those bands for the first time. The “WARC bands” were interesting because they were near to shortwave broadcast bands I hadn't listened to in years.

Building with modules

Aside from the need to shield circuit blocks from one another, a homebrew receiver with a single big board full of discrete components has another problem. If you build the whole thing at once without buying a kit and pre-cut board, I guarantee it won't work. ***To make homebrew stuff that works, you have to develop your own technology based on parts can get and circuits you understand.*** Thinking this way has always been difficult for me. Rather than “building a receiver,” I had to lower my sights and build one circuit at a time, e.g., “an oscillator,” “a mixer,” “an audio amplifier,” etc. Then I put the blocks together to complete my project. Some of these circuit blocks didn't work the first time so I had to build a new block. There were various

7. Chapter 13, Harris

reasons the modules didn't work. Usually, I wasn't able to buy the exact parts used in the circuits I was copying. Or my craftsmanship or shielding wasn't adequate. Sometimes I never did figure out why one version of a circuit block was superior to another.

On rare occasions my circuits didn't work because there were errors in circuit diagrams in QST magazine or in the handbooks. I found some serious errors in my 1979 Handbook and a minor one in my 1998 edition. Perfect editing is not possible, so we shouldn't expect it.

By building my receiver using separate little shielded modules for each circuit block, I could replace a circuit block whenever I managed to build an improved version. Otherwise, I would have ruined the whole big board. If painstaking R & D is new for you, prepare for a long battle. On the other hand, you'll learn a lot and victory will be especially sweet.

Mechanical construction

Homebrew receivers should be built in big cabinets. That way there's lots of room to add features and swap modules. The table in my shack is small, so I shoe-horned my receiver into a fairly compact package. In general, roomy boxes with lots of panel space are a better idea. For example, when my bargraph wasn't practical, there was no room on the panel for an analog meter. My long-range dream is to recycle my receiver modules into a larger cabinet. Eventually I want to add DSP, frequency counters, phase lock loops, and other goodies that intrigue me. With a big enough chassis, a receiver can grow and improve continually.



The receiver modules are interconnected with right-angle phono plugs and skinny coax.


Metal boxes shield each circuit block

A basic reason for building a receiver in metal-shielded modules is that capacitive coupling from one circuit block to the next can degrade performance. For example, my first 80-meter receiver module was built on one board. A crystal filter determined the selectivity. I made two "plug-in" crystal filters so I could use different filters for CW and for SSB. Using my 20

meter converter, I was tuning around 20 meters one day. I could hear lots of hams, but I was bothered by poor sensitivity and poor selectivity. I thought, "What's wrong with this receiver today?" I soon discovered there was no filter at all plugged into the 80 meter receiver board. What I was hearing was just stray coupling between the mixer and the IF amplifier. Amazing! So if you want band-pass filters with 50 dB skirt attenuation, you're going to need at least that much isolation between the stages. That means you need metal shields between all stages, coax interconnects and lots of bypass capacitors.

The metal-shielded modules could be small circuit boards mounted in commercial boxes. What I usually do is make shallow rectangular boxes out of pieces of two-sided circuit board material soldered together. The circuit is then carved into the floor of the box using small wood-carving chisels. The press-fit lid of the box is made from thin sheet aluminum folded over the corners of the PC board box.

Egg carton construction of the "mother board" of an 80 meter receiver



The compartment on the left is the product detector and audio amplifier. The large area in the center is the IF amplifier and AGC. The two modules on the right are the mixer and an optional RF pre-amplifier for the crystals.

If you plan on having more than one circuit block on the same board, you can isolate circuit blocks from each other using circuit board barriers soldered in place. The result is "egg-carton construction." Power can be routed between compartments using feed-through capacitors. If you suspect that a circuit block might not work, then wire your circuit on a separate square of PC board and then mount it on the floor of the desired compartment. All of these techniques are illustrated above. The IF amplifier in the center compartment is on a separate PC board.

Shielded modules and shielded cable interconnects

For connections between stages I use thin **RG-174 coax and phono plugs**. Right angle phono plugs are not designed for RF and they have too much capacitance. However, they're cheap, available, easy to wire, and don't take up too much room. I don't pretend that phono plugs are OK for 6 meters and VHF. I got some feel for their frequency limit when I discovered that switching from a phono plug to a UHF PL-259 on the input to my 50 watt transmitter amplifier


9. Chapter 13, Harris

vastly improved its performance on 10 meters. So far, I haven't seen any problems when working with lower frequencies and far lower power levels. Fortunately, receivers have tiny currents and low voltages. Most phono plugs have plastic bodies. That means that about 3/8 inch of the center conductor is not shielded. For my crystal filter module I used metal-bodied phono plugs that are at least a slight improvement over the plastic.

TV cable connectors are electrically superior to phono plugs, but they are dreadfully intermittent. Personally, I've found them unusable and I wish the TV industry would junk them. Proper RF connectors like BNC, SMA, or TMA cost \$2 - \$6 each. Some of them are hard to assemble and your receiver could easily contain \$300 worth of "proper" connectors. Also, most of these connectors are too long to fit gracefully in a small receiver. Right-angle phono plugs are quite short and quite cheap. Hurrah for phono-plugs!

Use plastic knobs

One odd little problem I encountered was that touching the metal control knobs or the front panel sometimes causes scratchy noise in the headphones when I listened on the higher bands. Yes, the metal panel was grounded and the chassis was wired to the station ground. The station ground is a heavy 12 gauge wire that grounds all the various metal boxes to a copper water pipe next to the station. I don't really have an explanation for the noises, but I switched from metal to plastic knobs and the annoying scratchy noises greatly improved. That was a shame - the metal knobs looked much more sophisticated.



Bottom view of the receiver


Band switching and power supplies

The precision power supply for the VFO is at the top right. The low drop-out regulated supply for the rest of the receiver is at the bottom right. These are the same circuits used earlier with the transmitter VFO and QRP transmitter modules. The band-switch is the multi-wafer ceramic switch on the left. The black wires on the left are skinny coax that interconnect the inputs and outputs of the converters for every ham band except 80 meters. It is desirable to cover the bottom with a metal plate to help keep stray signals out of the power leads.

80 Meter input preselector


The “front end” of the 80 meter receiver is a mixer. No RF pre-amplifier is needed on 80 meters because, if the receiver works well, then the atmospheric noise coming in from the antenna will be louder than the receiver internal noise. In this situation an RF amplifier won’t help. However, the mixer does need a sharp bandpass “preselector” filter to keep out the low frequency AM radio and limit the input signals to 3.5 to 4.0 MHz. It’s especially important to filter out AM broadcast stations. These stations put big voltages on the antenna and tend to overwhelm the mixer unless they are well attenuated.

The mixer subtracts the VFO frequency, (5.0 to 5.5 MHz) from the IF frequency, (9 MHz), to tune 80 meters, (3.5 to 4.0 MHz). The two inputs to the mixer are the VFO signal and the antenna signal. When I first examined the 1986 ARRL design, I was disappointed to see that the pre-selector had a primitive variable capacitor that the operator was supposed to tune for maximum gain for a particular part of the band. After all, the bandpass filters for the other HF bands were fixed and not accessible from the front panel. I attempted to build my own fixed bandpass filter, but my filters had too much attenuation (poor sensitivity) and sometimes let in AM broadcast stations. That is, it was like listening to a crystal set.



The recommended 80 Meter preselector filter for the core receiver mixer input.

So I returned to the ARRL design with the 365 pF variable capacitor. It had so much attenuation on 80 meters, I couldn’t hear a thing. I ran the ARRL circuit on a Spice program, and, according to Spice, it should have worked well. However, mine didn’t. I must have had some bad part somewhere, but I couldn’t find the problem. Using trial and error, I removed some parts and ended up with the circuit shown below. My filter works pretty well, although, according to my Spice program, it shouldn’t have. Sigh ...




My version of the 80 Meter preselector filter for the mixer input

I mounted the preselector in a shielded box up on the front panel where the peaking capacitor was accessible. Someday, when I figure out how to build a better preselector, I shall

replace the whole module.

The 365 pF variable capacitor acts like an attenuator

On the positive side I found that the variable peaking capacitor is quite useful as an attenuator for receiving strong single sideband signals. In other words, strong SSB phone signals are often much more intelligible when the preselector is mistuned and signal strength is decreased. If I didn't have this capacitor as an attenuator, I would have to build an attenuator.



The preselector is built in a little box up front behind the front panel.

The variable frequency oscillator (VFO)

The receiver VFO is the same design as the 5 MHz transmitter VFO discussed in chapter 10. The big tuning knob controls the VFO. Actually, in superheterodynes the VFO is usually called a *local oscillator* or *LO*. The range and stability of the VFO determine what VFO and IF frequencies are practical. Like a transmitter VFO, a receiver VFO should be stable to less than 5 Hz/ minute drift, although less than 20 Hz will probably be OK.

Unfortunately, if the VFO frequency is too low, it probably won't span enough Hz to cover the bands you're interested in. Notice that 10 meters is so huge, 1.7 MHz wide, that you may have to cover it with multiple converters. So far, my receiver just tunes the first 500 KHz of 10 meters, which includes all the CW activity. I guess I'll worry about the rest when I get on 10 meter SSB. And that will probably happen when the sunspots return. The VFO frequency and its harmonics should avoid the ham bands as much as possible. Frequencies with round numbers are best because they make calibration easy. Put all these factors together and you'll find your choices are fairly limited.

In fact, once you've built separate transmitters and receivers and used them on the air, the disadvantage of having to tune two VFOs makes it obvious why most modern rigs are transceivers. When planning a homebrew receiver and transmitter, you might consider using a common VFO module. A common VFO would be a great help on the air. But, if you use just one VFO, you will have to master the 500 to 800 Hz send/ receive frequency-offset problem. Also, an isolation amplifier is needed to keep the cable connecting one to the other from loading down the receiver VFO. And finally, when used on the upper bands, each crystal oscillator in the converter and each transmitter PMO must be on the corresponding frequencies in order for the receiver to listen on exactly the same frequency used by the transmitter. In other words the

converter oscillators need to be shared too. All this represents lots of work. It makes the most sense to build a transceiver from the beginning, starting with a gigantic mainframe cabinet.

How much VFO signal do you need?

As explained in chapter 10, the stability of a VFO is partly dependent on generating the minimum heat possible inside the VFO box. Therefore, unlike me, you should decide out how much VFO signal voltage you will need before you build the VFO. The required VFO voltage is the level needed to drive your mixer. MOSFET and JFET mixers only need 1 or 2 volts peak, so if you're going to use one of these, there is no good reason to generate 5 volts, then throw away most of it in a potentiometer at the input to the 80 meter mixer. Instead of running the VFO on 12 volts or 5 volts, you can reduce the supply voltage externally to 3 volts or whatever minimum voltage you actually need.



A varactor tuned 5 MHz VFO. Tuning is accomplished with the big round potentiometer.

Mixer magic

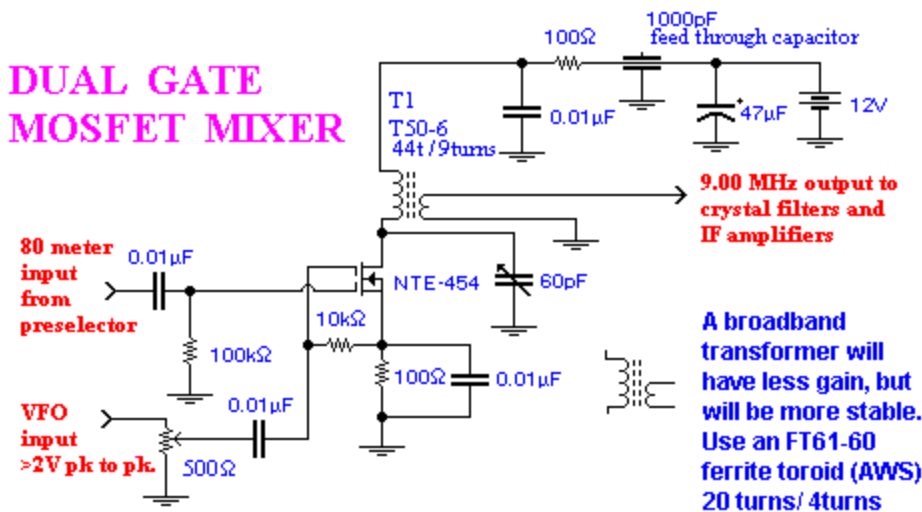
The purpose of a mixer is to translate the frequency of an incoming radio signal to a constant intermediate frequency (IF) that can be amplified and filtered more easily. Mixers combine a local oscillator sinewave with the incoming radio signal to make a composite signal. The new signal contains the original frequencies, plus the new sum and difference frequencies. Mixers intended for moving a VFO up to a high band were described in chapter 11. Mixers for this purpose can be quite crude and will work well. Unfortunately, receiver mixers are much more difficult because the signal input can be so tiny.

One way to look at mixers is that a big local oscillator sinewave keys the incoming RF signal on and off, cycle by cycle. The lesson is that the local oscillator (the VFO tuning knob) must be a big signal while the RF input signals may be arbitrarily small.

The ARRL Handbook presents six or eight different mixer designs made with discrete diodes, inductors, and transistors. However, in most ARRL receiver designs since the 1980s, the mixer is an integrated circuit or little canned assembly labeled “mixer.” I guess everyone else was having mixer trouble too, so they resorted to integrated circuits. Last year there was a receiver project in QST that used an IC that contained both the mixer and the VFO. I’ll bet these marvels work fine, but the contents of these ICs are a mystery. Use one if you want.

Mixers will give you lots of static... and squeals, howls, and squawks

So far I've built five different mixer designs from discrete parts. First I built a classic balanced mixer with ferrite cores and a hot carrier diode ring. When I turned it on, I heard loud, roaring static in the headphones. "Oh goodie!" I thought, "Listen to all that atmospheric static! It must be working!" I soon figured out that the static was coming from the mixer and the IF amplifiers and had nothing to do with the outside world. I had just learned a basic truth about mixers: Mixers aren't just prone to generate "a little background noise." They often produce gigantic Niagara Falls noise that obscures everything coming in the antenna. However, once I had proper mixer input levels and resonant circuits tuned up as best I could, the noise disappeared and I began to hear stations. Unfortunately, as I tuned across the band, there were loud whistles like marker beacons every few KHz. In between the whistles, I could sometimes barely hear strong stations. So much for diode ring mixers.



A practical mixer

My favorite mixer is shown above. Most of the others suffered from noise and "birdies" and usually insensitivity as well. Unlike diode mixers, the operation of the MOSFET mixer is obvious. It is essentially an ordinary tuned transistor RF amplifier. The radio signals come in on one control gate. This modulates the large current passing from drain to source of the transistor. The small voltage on the control gate controls the large drain current thereby amplifying the original signal. A second input gate amplifies the local oscillator signal. This means that 2 volts of VFO signal is plenty. The local oscillator signal is so strong it turns the drain to source current totally on and off, "chopping" the input RF signal into tiny segments. The big output current from the transistor becomes an amplified "mixture" of the two input signals.

I originally used a tuned version of this mixer in which the transformer primary was tuned with a capacitor to resonate at 9 MHz as shown above. That works OK, but is rather tweaky and subject to noise and birdies. I now prefer an untuned version (not shown above) which has slightly less gain, but isn't so easily misaligned. Instead of the powdered iron core and 60 pF trimmer, just use a CWS (Amidon) FT50-61 ferrite core with 20 turns primary and 4 turns secondary.

Because of the high gain of the MOSFET, the VFO sinewave signal can be low amplitude, typically 2 volts peak, and still fully chop the radio signals on and off. In contrast, a diode ring mixer needs a big local oscillator signal, 12 or more volts peak, to chop the signal. Other transistor mixer designs use junction FETs or bipolar transistors. These designs use the emitter resistor as the VFO input port. The emitter or source has no gain, so these designs also need big local oscillator signals.


A small local oscillator signal is especially helpful for building the converters you will need for each of the other nine HF bands. Each of these converters must work through its own mixer. If the dual gate mixer were not used in the converters as well, each of the eight local oscillators you will build would need separate amplifiers to get the local oscillator signal up to 12 volts peak. I used a common, broadband pre-amplifier for the crystal oscillator inputs to a common untuned mixer. That is, the mixer has a ferrite, broadband transformer as described above.

All dual gate MOSFETs are not equal


Alas, a dual gate MOSFET mixer isn't a guaranteed success either. When I first built a MOSFET mixer, I couldn't buy any of the transistors recommended in the handbook. I first tried a generic part, the NTE221 transistor. This produced the usual oscillations and insensitivity. I was getting discouraged, but I tried the similar NTE454 and *IT WORKED!* The only obvious difference in the specifications was that the gate shut-off voltage was smaller. In other words, the NTE454 was more sensitive. Since then, I've discovered the NTE222 seems to work as well as the NTE454. The NTE455 seems too sensitive. In my circuit it produced whistles, birdies, and noise. On the other hand, the NTE455 worked great as a product detector (Chapter 7).

MOSFET TRANSISTOR CASES

DUAL GATE MOSFET
TO -72 CASE
e.g., NTE454, NTE 221, etc.



MOSFET POWER SWITCH
TO -220 CASE
e.g., IRF9140, SMP20P10,
MTH25P05, etc.



I first tried to use two dual-gate MOSFETs in a "semi-balanced" configuration. The handbook said this is superior to the simple mixer shown above. Balanced designs are supposed to help cancel images, by that I mean, birdies. My semi-balanced mixer produced the usual oscillations, but at least the noise and sensitivity were adequate and I began to hear weak 80 meter stations with my receiver.

Reading up on mixers, I learned that mixers are only happy when they receive the exact input levels. That's why I put a pot on my VFO drive to inject the optimum level. As I turn up the VFO drive to the mixer the output signal strength rises abruptly then levels off. Higher levels

of VFO contribute only slightly more gain, but much more noise. I adjust the VFO input to where the gain first begins to level off. (Note: My VFOs are designed to work into a 500 ohm load, hence the 500 ohm pot.)

The sensitivity of mixers to ideal levels explains why most modern transceivers have input attenuators so that they can be adjusted to tolerate strong signals. I got a QSL card from a guy who wrote, "Sorry about the 529 signal report. After we signed off, I discovered I had the attenuator on." As mentioned earlier, the 80 meter preselector filter may be deliberately mistuned so that it acts as an attenuator to limit signal strength.

Note: Reception on 80 meters and 160 meters is best with a tuned transmatch

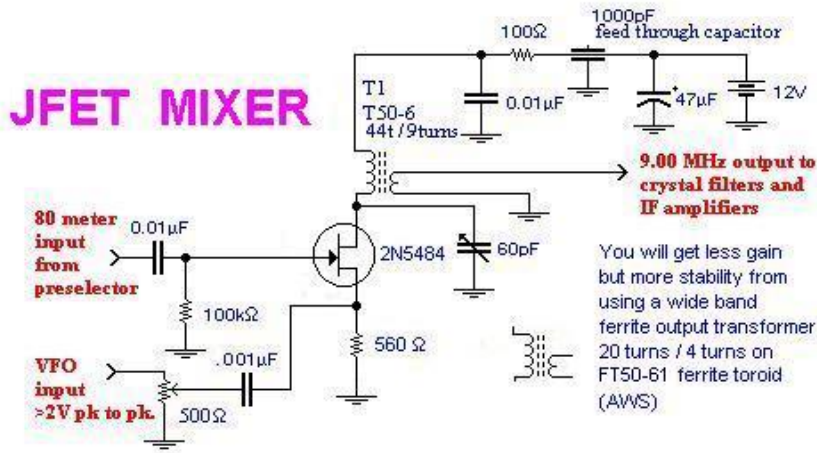
By accident I discovered that reception on the two lowest HF bands is much better when the receiver is sharing the antenna with the transmitter and the antenna is tuned with the "T" antenna coupler described in chapter 9. In my neighborhood at least, the signals from the local AM radio stations are so high that they tend to overwhelm the 80 meter mixer. This results in a lack of audible signals on 80 and 160. I didn't realize I had a problem because I wasn't hearing the AM stations in the headphones. However, when I tuned up the transmatch, suddenly numerous ham signals appeared. The obvious conclusion is that my receiver preselector filter is not selective enough. Even my old Collins receiver is greatly improved by a tuned antenna coupler.

In chapter 7 I described a high pass filter for the direct conversion receiver to attenuate the AM broadcast signals. I have found that this little filter also improves the 80 meter receiver board performance. Because I had built it in a little module with phono plug connectors, I could try it out first without building anything new.

A JFET mixer

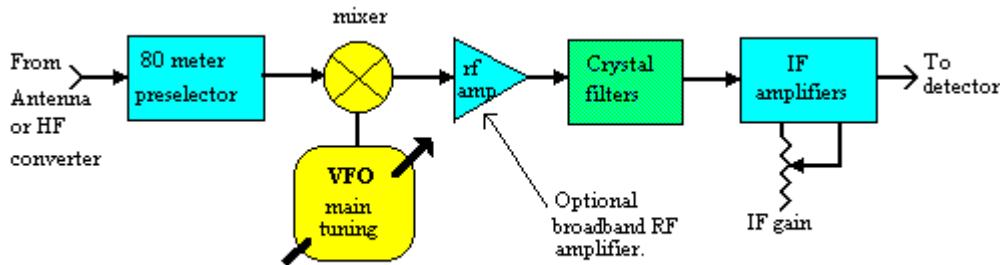
Since this book was first written, dual gate MOSFETs have become increasingly expensive and hard to find. One part you might try is the RCA 40673 available from Mouser for \$7. I haven't tried them, but my friend Jayram, JU2JN in India, has used them in receiver mixers with great success. Dual gate mixers are his favorite design too.

In the event that you need a substitute design, here is a JFET circuit that works but in my opinion, isn't as sensitive. It is practically the same circuit, but instead of introducing the local oscillator signal into a separate gate, it is introduced across the source resistor. Surprisingly the optimum local oscillator signal input level for this circuit is only about 1 volt peak-to-peak. For this reason the signal is first passed through a 500 pot. I would have expected that, for a circuit like this, the optimum peak voltage would have approached the supply voltage so that the transistor would be entirely turned on and off. I guess the answer is that this is a *depletion type* MOSFET that is already half turned off with no bias. Consequently, it doesn't take as much drive to turn it either full on or full off. I also attempted to use the dual JFET PMO mixer design described in chapter 11, but it worked poorly in my receiver.



Crystal filters and BFOs

Crystal IF filters give you the selectivity you will need for working CW. They eliminate interference from nearby stations and also eliminate a great deal of the atmospheric noise. The output from the superhetrodyne mixer is a weak, broadband, IF frequency signal that needs amplification and filtering before it is ready to be detected. The bandpass filtering is usually done right after the mixer.



Location of crystal filters in a superhetrodyne


The filter could be a "mechanical filter" if you are using a low frequency IF like 455 KHz. But if your IF is 9.0 MHz like mine, then you'll need one or more crystal filters. Before I describe building crystal filters, I'll discuss the Beat Frequency Oscillator (BFO). You will probably need the BFO as a tool to select the crystals for your filters.

The Beat Frequency Oscillator (BFO)

A beat frequency oscillator is an RF oscillator that operates on the intermediate frequency of a superhetrodyne. The BFO mixes with the IF signal to make CW and single sideband transmissions audible and/ or understandable. Without the BFO CW signals would be inaudible or just thumping noises at best. Single sideband phone would be unintelligible "Donald Duck" sounds. In single sideband the transmitter filters out the basic carrier frequency leaving just one of the modulation sidebands. The BFO serves to restore the carrier sinewave, in effect returning the sideband signal to its original amplitude modulation.


During detection, the audio signal passed on to the loudspeaker is the difference between

the IF frequency and the BFO frequency. For example, when listening to CW signals, the IF frequency might be 9.000,000 MHz. The BFO frequency might be 9.000,700 MHz. What you would hear in your headphones is a musical tone of the difference frequency, 700 Hz. If that pitch is too high for you, adjust the BFO frequency to say 9.000,500 MHz to produce a musical tone of 500 Hz. For the musical tone to stay constant, the BFO oscillator must be quite stable. Therefore, we use a crystal oscillator and pull the frequency up or down using a variable capacitor, just like we did in the crystal controlled QRP in chapter 6.



The Beat Frequency Oscillator

This BFO is directly from the W7ZOI and K5IRK receiver in the 1986 ARRL Handbook. Its unusual feature is that its DC power supply rides on the same line as the RF output. This makes it easy to install the BFO in a little metal box up on the front panel remote from the main receiver board. A variable capacitor on the front panel “pulls” the BFO crystal frequency above and below the nominal frequency. The BFO, together with the crystal filter, allow you select upper or lower sidebands. The BFO is connected to the main board with a length of thin coax jumper cable. If you wish to use this oscillator for matching 9 MHz crystals for filters, I suggest you install the BFO crystal in a small IC or transistor socket.




The Beat Frequency Oscillator (BFO) module

The BFO frequency tuning range should extend above and below the bandpass of your crystal filters. *When the BFO frequency is below the center of the filter bandpass, you are listening to the upper sideband. When the BFO is tuned above the filter bandpass, you're hearing the lower sideband.* Adjusted for the upper sideband, tuning down the CW band will cause the whistle pitch of a signal to start high, drop down to a low pitch, and then disappear. Adjusted for the lower sideband, tuning down the band will cause the pitch to start low then climb up high and disappear. When the BFO is tuned to the center of the filter bandpass, the tone will start at a medium pitch, drop to a low pitch, briefly become inaudible, then climb back up to medium pitch and disappear.

When tuned to sideband phone signals, the signal is harder to make understandable if you select the wrong sideband. Consequently, it will be important to calibrate the BFO tuning knob so that you will know when you are listening to upper or lower sideband.

Ladder filters

Building crystal filters was easy, once I figured out how. Most commercial transceivers use modular crystal filters that have specific bandwidths and are sealed in little cans, something like an integrated circuit. I made my filters from discrete crystals.



One, two, and three crystal “ladder filters”

A ladder filter is just two or more crystals in series with capacitors bypassed to ground at the nodes. The bandpass width is inversely proportional to the number of crystals and the

capacitance to ground. *In general, the lower the capacitance on the nodes, the wider the bandwidth.* Large capacitors, like 50 ohms reactance, will give narrow bandwidths and higher attenuation. Long “ladders” of course produce even narrower bandwidths and more attenuation. If all the crystals are all identical, then the “skirt attenuation” on either side of the peak passband becomes steeper and steeper as more crystals are added.

Crystal filters with two or three crystals are selective enough for “CW after dinner.” By that I mean you can work the busy CW bands early in the evening with adequate signal separation. A single crystal filter made from just one 9 MHz crystal is good for single sideband phone (SSB). With three or four crystals the bandwidth becomes so narrow that SSB is totally unintelligible. With strong CW signals and lots of QRM, triple or quadruple crystal ladders are extremely useful. If you tune down to the bottom of 20 meters in the evening, you’ll usually hear a cacophony of CW stations all trying to work the same 5 or 6 DX stations in exotic countries. With only one crystal switched in, you’ll hear practically everyone at once. Switch in two and three crystals and suddenly, you’re not only listening to one clear station, but most of the background noise has gone too.

How many crystals can you use in a ladder?

One limit on how many crystals you can use in series depends on the precision with which you match the crystals. I didn’t match my first set of filters and they produced more attenuation than filtering. I then proceeded to waste time getting more gain out of my IF amplifier. Then, after I achieved the gain, the selectivity wasn’t much better than single crystals. Finally, I tested my crystals one by one by putting them in the BFO and measuring the frequency with a frequency counter.


I didn’t expect that the frequency filtering characteristics would be exactly the same frequency as when the crystal is used in an oscillator. However, I figured I could at least select sets of crystals that were similar. When I ran them in the oscillator, I was shocked to discover that the crystals were as much as 2.5 KHz different! No wonder it worked so poorly. I had put 9.001 MHz crystals in series with 9.003 MHz crystals. I had built a “crystal barrier,” rather than a crystal filter.

Luckily I had bought twenty 9.000 MHz microprocessor crystals. That sounds extravagant, but they’re under \$1 each at Digi-Key or Mouser. Because I had a wide selection of 9.000 MHz crystals, I was able to match up two crystals that were within a few Hz of each other. And I was also able to match up a group of 3 crystals that were within 50 Hz. This time when I put matched crystals in my ladders, the improvement was dramatic. When I switched from a single crystal to a double crystal, the signal strength barely dropped. With three crystals, the signal strength just dropped slightly more.

In theory, you can make perfectly matched crystals by placing tiny trimmer capacitors in parallel with each crystal. Then you would trim the crystal and capacitor assemblies in the oscillator one at a time so that each crystal/capacitor assembly oscillates on precisely the same frequency.

Another limit on the number of crystals in a filter is the shielding and RF isolation between the mixer and the IF amplifier. If the shielding is poor your IF amplifier will “hear” the signal from the mixer without the signals ever passing through the crystal filters. In my receiver,

5- or 6-crystal ladders aren't worth building.



220 pF (3)

Four matched 9.000 MHz crystals

Switch in your filters with a rotary switch

I first built my filters as “plug-ins,” but I soon discovered it was too hard to change them in the middle of a QSO. Eventually I wired them to a rotary switch in a shielded box. As I became more experienced with using the filters, I began to use the triple filter more and more. Finally I built a quadruple filter and now I use it routinely. I find it works well with the S meter (strength meter) as a way to tune my transmitter to zero beat with another ham's signal. I just sweep the transmitter VFO across the band until the S meter jumps up to maximum. This occurs when the VFO frequency converted into the IF matches the frequency of the four matched crystals. The offset of the BFO and the tone of the Morse code signal take care of themselves. That is, if the fellow is on the upper sideband, then the S meter only responds to my VFO when I am on the same BFO offset that he is using. This happens because only one sideband is audible at a time with 3- or 4-crystal ladder filters.

By the way, one of my 20 crystals behaved erratically. While I watched, the frequency wandered about +/- 200 Hz. Consequently, I didn't use it. With my bare board oscillator, I found that *all crystals wandered about 2 to 5 Hz*. This implies that, ideally, all crystal oscillator circuits should be treated with the same respect needed for a stable VFO. In other words, crystal oscillators should be housed in metal boxes with regulated power supplies.

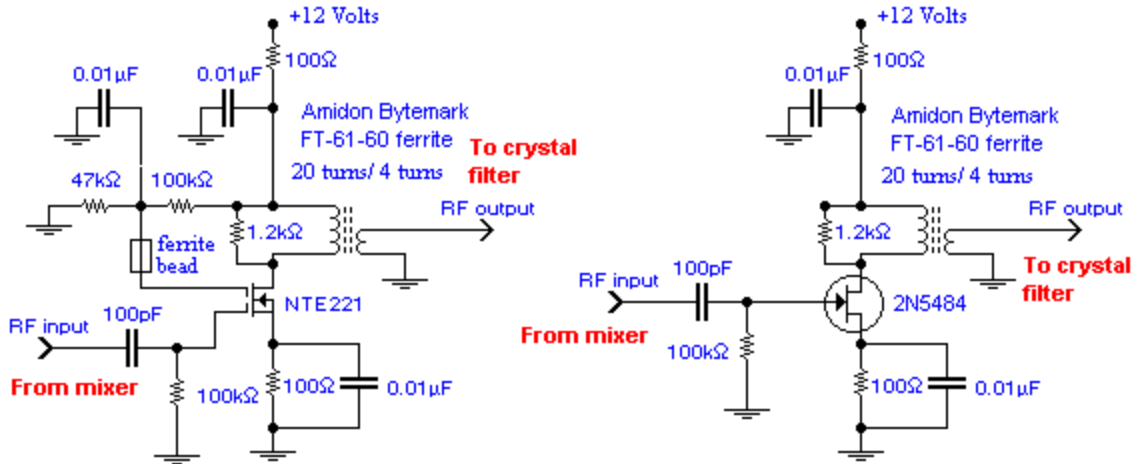
Series and parallel cut “XTALS”

There are two kinds of simple crystals, series and parallel cut. As I understand it, the difference is the oscillator circuit for which they are designed to be used. For example, a series-type crystal is intended to be used in series with a specific capacitance in an oscillator. When this exact capacitance is used, it will oscillate at the rated frequency, for example, 9.000 MHz. In contrast, if you use a parallel cut crystal in the same circuit, it might oscillate at 9.004 MHz. You may use either kind of crystal, but your filter frequency may not be exactly 9.000 MHz. If you like, you may place trimmer capacitors in series or parallel to get exactly 9.000 MHz.

All 9.000 MHz crystals aren't equal

In my experience larger crystals, size HC-49 or bigger, work well. The tiny surface mount or half-size HC-49 crystals need more signal strength from the mixer output to push the signals through the filter. The broadband “optional” amplifier shown below will give you extra gain if you need it. It's the same design that will be used later as a RF amplifier for the higher HF bands. If you are short of MOSFETs, a JFET alternate circuit is also shown below.

RF Amplifiers, MOSFET and JFET versions



Broadband RF amplifiers Place one between the RF mixer and the crystal filter.

I also observed a difference between different brands of crystals. My crystals from the ICM company were extremely similar to each other. For ladder filters, this would be the brand to buy. My crystals from the ECS company had much wider frequency variation from one crystal to the next. This isn't helpful for building ladder filters, but if you need crystals that are slightly different from each other to build wider pass-band filters, or perhaps SSB oscillators that must operate 2.5 KHz above and below the nominal frequency, then buy the ECS crystals. I found them both useful.

By trial and error I found that the untuned impedance step-down transformer shown in above pre-amp circuit performs better than leaving the transformer out and simply using an inductor on the drain then coupling to the crystals with a capacitor. In other words, this implies that the crystals look like roughly 50 ohms, but must be matched to the high impedance output of the MOSFET transistor. There's a design in my ARRL handbook that uses impedance step-up transformers to match the filter. I find it hard to believe that design is optimum unless the crystals they used behave differently than mine.

The second input gate in the above amplifier is used to set the DC bias and make the amplifier class A. A voltage divider delivers about 4 volts DC to the gate. The *ferrite bead* is a tiny inductor (RF choke) that helps insure that the MOSFET doesn't oscillate. The ferrite bead is literally a 1/8 inch cylinder with a tiny hole through the center. For example, you could use a CWS type (Amidon) FB43-101 bead. The type isn't critical. I used several different kinds of beads and have had no trouble with oscillation. If it does oscillate, remove the 0.01 μF bypass cap from the 100 ohm source resistor. The resulting negative feedback should kill the oscillation at the expense of a small amount of gain.

The IF amplifier

The IF amplifier is another tricky part of a superhetrodyne. It's a high-Q amplifier that must handle signals with a range of 100 dB or more without oscillation or noise. This is a huge

dynamic range. The gain on the IF amplifier stages should be adjustable using an IF gain control. Too much gain and you will have noise and squeals. Too little gain and you can't hear weak DX stations.

Moreover, if you used miniature half-size HC-49 crystals to build your bandpass filters, you will need even more gain to stuff signals through the significant attenuation of the filters. In the last section I described a simple RF amplifier that can be placed between the mixer and crystal filter to overcome this difficulty.

Oscillations in an IF amplifier come in several flavors. As you tune the LC circuit of an IF amplifier stage you will hear squeals, harsh roaring, silent dead spots, and gentle static. The setting that brings in the loudest signals is surprisingly noise free. The first time I turned on my receiver, I quickly learned that *most of the receiver noise is coming from the mixer and IF amplifiers*, not from the outside world. The noise comes from a maladjusted mixer or too much IF amplifier gain.

Although I was able to tune up my 80 meter receiver using a signal generator, real 80 meter ham signals worked the best for me. There's no point in simulations when you have the real thing. A problem with tuning up the IF on 80 meters is that 80 may be dead during the day. During the summer, 80 may not be so hot in the evening either. As a result, you might consider building a converter for 20 meters early in the project. 20 meters is usually full of signals anytime, day or night, all year long. Since your 80 meter receiver may not yet be working, you can adjust your converter by feeding the output into a commercial receiver tuned to 80 meters. Then, after you have the converter working, you'll be confident that there are lots of real signals for your 80 meter receiver to hear.

Impedance matching the crystal filter to the IF amplifier

Looking at examples on crystal filter circuits in handbooks from various years, I found circuits that seem to assume the filters are low, medium, or even high impedance. In most of my attempts, I seemed to get the best gain when my filters were assumed to be relatively low impedance, say 50 to 100 ohms. That's why the optional amplifier described above used a step down transformer output. I tried step-up, step-down and no transformer to feed the signal into the IF amplifier shown directly below. Step up worked best on the two versions as shown below.

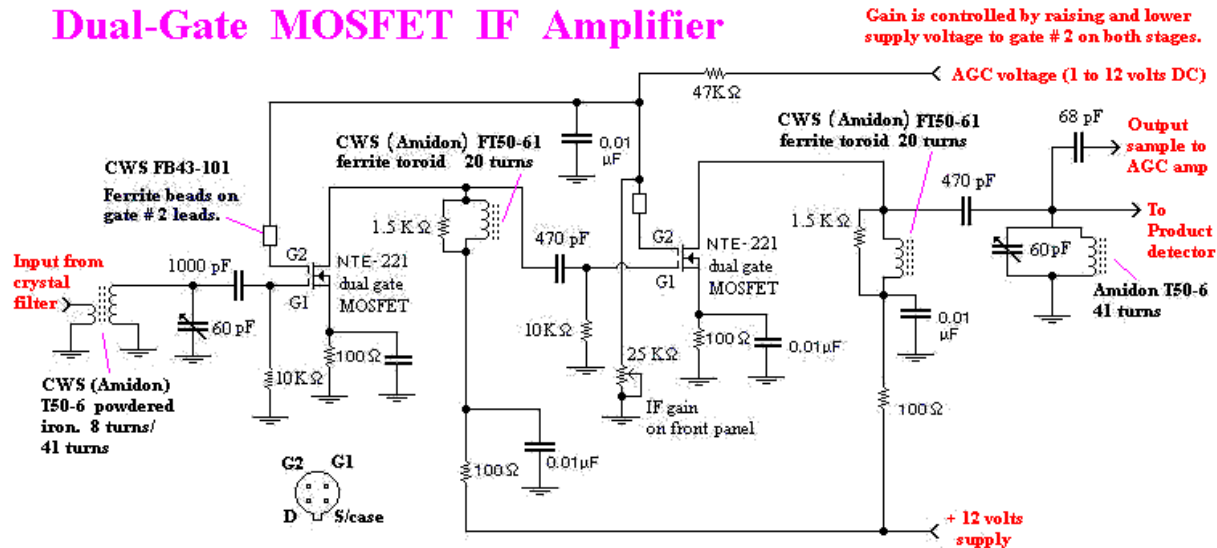
A dual-gate MOSFET IF amplifier

My first versions of the IF amplifier used dual gate MOSFET amplifiers, similar to the crystal filter preamplifier described earlier. Just what you need - another circuit made from 7 dollar transistors! The gain of each MOSFET transistor can be controlled by varying the DC bias on one of the two control gates. This control voltage can be generated by either the IF gain knob or by the automatic gain control circuit. In short, the dual gate MOSFET looks ideal for IF stages. Unfortunately, I had lots of trouble with squeals and noise.

With struggle, I eventually managed to get the dual-gate version working. Originally it consisted of two dual-gate tuned amplifiers in series. Each amp was similar to the dual gate amplifier on the previous page. Instead of the fixed 4 volt bias on the second gate, those gates were biased with the variable IF gain/AGC voltage. When I started, I used tuned LC circuits as the impedances between the + supply and the two MOSFET drains. This approach was quite unstable and prone to oscillations. I decreased the instability but retained enough of its

sensitivity by replacing the tuned drain LC circuits with ferrite cores - that is, RF chokes. I moved both tuned LC circuits to the input and output respectively. These LCs must be tuned carefully for the maximum signal strength without squeals. The circuit below is the result.

Dual-Gate MOSFET IF Amplifier



The main lesson I learned from this circuit was that high gain amplifiers shouldn't be tuned unless the input signal is continuous. Whenever there is little or no signal to amplify, a high gain tuned amplifier will amplify any noise that may be present and may begin to oscillate at the tuned frequency. To prevent this, you are better off doing your tuning function separate from the amplification. The two amplifiers above are wideband but relatively low gain.

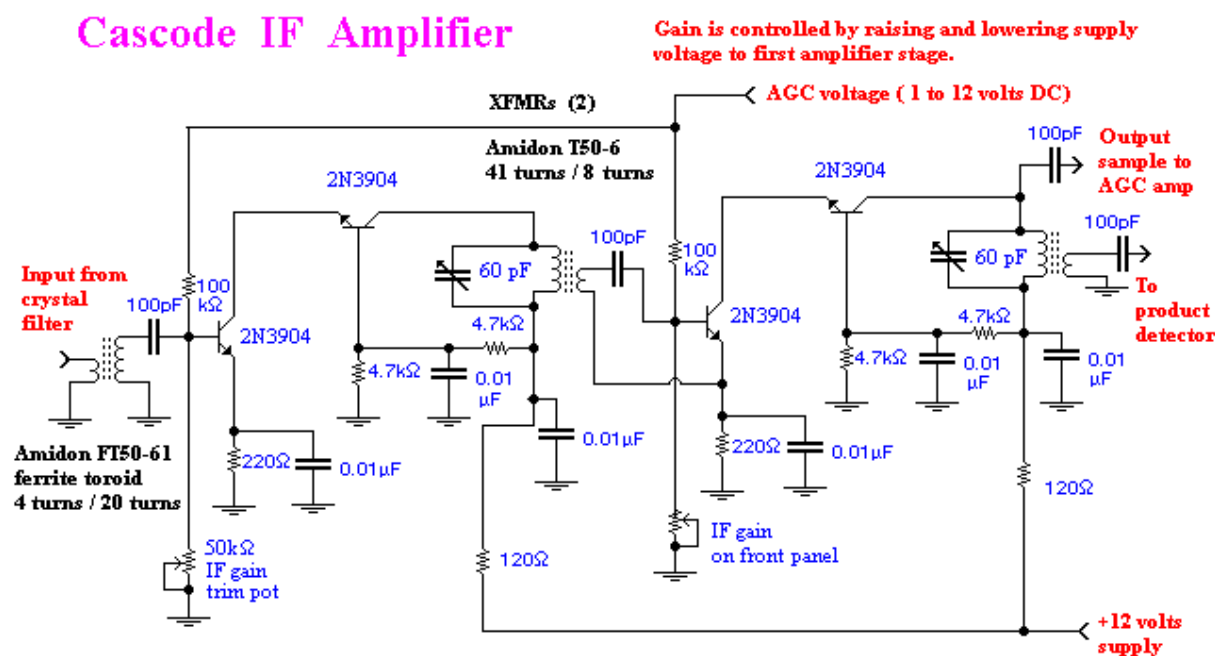
I suspect that a similar circuit would work OK with ordinary bipolar transistors, but I haven't tried it. The IF gain would control the forward bias level of the two transistors, that is, the bias that makes them class A. A gain/ bias system like this is used in the cascode amplifier shown below.

In the circuit above, the input LC attempts to limit input noise to just 9 MHz noise. Similarly, the output noise is similarly clipped with the second LC. Chapter 15 discusses PMOs for single sideband generation, which have a similar problem. Every time there is a pause in the speech being transmitted, a tuned amplifier will usually oscillate. Consequently, all the frequency filtering is done in separate passive filter blocks, rather than being incorporated into the amplifiers.

In summary, the above amplifier stages are broadband and use ferrite transformers that pass on any radio frequency signal to the output with little or no filtering effect. On the other hand, ferrite transformers have lots of internal inductance and capacitance. They will resonate at lower frequencies, perhaps even audio frequencies. Therefore 1.5K Ω resistors were placed across the transformer primaries to lower the Q still more. If you are still plagued with squeals and motorboating oscillation, you might try removing the 0.01 μF emitter resistor bypass capacitors.

Cascode amplifiers - variable gain with constant Q

I had heard of *cascode amplifiers* but didn't have a clue why they were wonderful. I built two other IF strips before I settled on the circuit shown above. Reading in an old handbook, I spotted the IF amplifier shown above. The handbook said that simple transistor amplifiers were poor for IF amplifiers because, when you tried to change the gain of a single transistor, the Q of the output tank circuit changes and you get squeals and noise. "Yes!! Yes!!" I cheered. "That's my problem!" The next circuit uses two (cheap) bipolar transistors in each stage in a "cascode" configuration.



The input transistor is wired as an ordinary grounded emitter amplifier with its high input impedance. The clever part is that the second transistor is wired to the first in a grounded base configuration. This gives the amplifier a super-high output impedance which supposedly makes it immune to changing the DC bias on the first stage. Besides, the phrase "cascode amplifiers" sounded cool and I wanted to use some. This cascode amplifier worked well for me.

It's interesting to see what happens when one tunes up an IF amplifier with a scope probe on the IF amplifier output. As expected, the audio signals ride on the IF frequency signal, just like amplitude modulation should. When the amplifier is tuned for optimum signal reception, the scope shows that the amplifier is producing the most modulation on the IF signal. But when the output is tuned slightly differently to produce the largest 9 MHz signal, the reception is OK, but not the best. I had not realized that these two attributes aren't the same thing.

How to tune an IF amplifier

When you first turn on your 80 meter receiver module, there will be several modules with variable capacitors and pots that all need to be adjusted. Let's assume that you have checked out the VFO and it is providing enough voltage sinewave to drive your particular mixer. You should also have a look at the waveform of your BFO after it has arrived at the receiver module. It should be a smooth, clean sinewave. A clean sinewave is more important than a high amplitude.

You may want to readjust the trimmer capacitor on the BFO oscillator for the best waveform.

To align the IF you could just start twiddling pots and trimmers and hope you hit a favorable combination of settings. I've tried that - it doesn't work very well. First, clip an oscilloscope probe onto the output capacitor going from the IF amplifier to the product detector. At this point your BFO comes to your rescue. Clip a test lead onto the connector where the BFO plugs into the 80 meter receiver board. This signal line also has 12 volts riding on it, so connect other end of the test lead to a small capacitor, say 100 pF. This will filter out the DC. Use the other end of the capacitor is a signal probe to inject a 9 MHz test signal into your IF amplifiers. Since the oscilloscope probe is at the output of the IF, start by aligning the trimmer and input pot of the last amplifier.

Next move the BFO test signal down the to input of the first amplifier and adjust the trimmer capacitor and input trim pot. Peak the 9 MHz signal at the output as before. Now remove the BFO signal test probe but leave the oscilloscope probe in place.

At this point, if there are strong 80 meter signals at the input to the mixer, you should begin to be able to hear them. Now tune each stage carefully for maximum sound level. Notice that maximum sound level is not the same setting as maximum 9 MHz carrier in the output.

Automatic Gain Control (AGC) is not a luxury

The automatic gain control is a receiver feature that holds the signal level relatively constant while tuning in signals of varying strength. Before I built one, I thought an AGC was in the same category with digital readouts and beautiful cabinets. Why do I need one? Am I too lazy to turn the IF gain up and down? It turns out that an AGC has many advantages. The main one is that it helps you achieve the gigantic signal strength dynamic range (100 decibels) that you need in practical ham receiver. After I built an AGC, I realized it was it was also a big help in getting rid of the noise and oscillations.

Although I had been happy with the performance of my IF without an AGC, I could never get rid of the "noise zone" in my IF gain control. That is, I had to keep the IF gain below a certain level or it would produce a roar of receiver-generated static. Apparently, IF amplifier stages are only happy when they are processing signals of a limited range of amplitude. Noise and oscillations happen when the signals in the final IF amplifier are too large. With an automatic gain control, it was easier to tune the IF so that the IF gain control acts like a "volume control" without a noise zone.

The S meter and other uses for the AGC


A benefit of an AGC is that, when I switch in higher selectivity crystal filters, the AGC compensates for the filter attenuation to a large degree. Also, when you put a meter on the AGC signal level, you have made an *S meter* - in other words, a "strength meter." The S meter taught me that what you hear in the headphones doesn't always correlate with the signal strength in the IF strip. In other words, the S meter is reacting to big IF frequency signals, not the level of modulation on those carrier signals.

The best use for an S meter is to tune the transmitter VFO to match the receiver. In other words, if you are answering a CQ, you can tune your transmitter right in on top of the fellow you want to call. First, you need to switch in a 3- or 4-crystal filter. Then, as you slew your

transmitter VFO across the frequency, the S meter will soar when you are lined up right on top of him. Without using this technique, “zero beating” the VFO is time consuming. Modern transceivers don’t have this synchronization problem because the receiver and transmitter and using the same VFO.

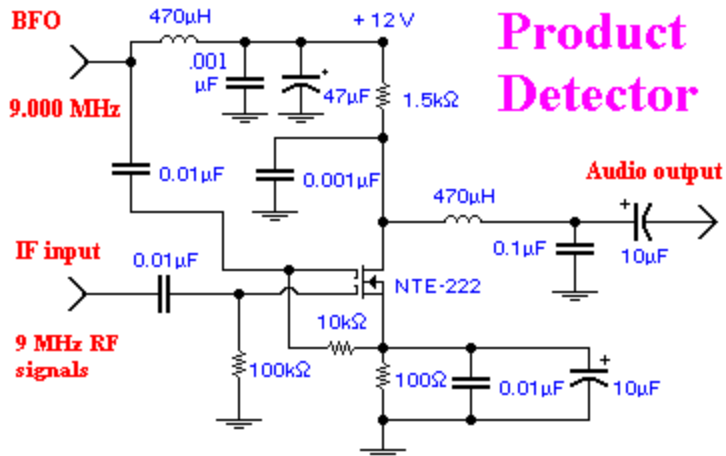
I originally used a digital bargraph S meter that looks very racey. Unfortunately, like almost everything digital, it makes a hissing radio noise when it changes level and I don’t need that. I tried hard to filter the bargraph circuit, but as usual, I couldn’t get rid of the hiss. I finally replaced it with an old-fashioned analog meter and the noise vanished.

Automatic Gain Control (AGC)



An AGC works by sampling the output level of the last IF amplifier stage. Signals are detected like a crystal set using a diode and averaged with a capacitor to produce a DC level proportional to the IF signal strength. This DC level is then amplified and used to bias the IF amplifiers. For example, the above circuit can deliver the positive voltage bias on IF amplifier stages made from dual gate MOSFETs. Or if the IF amplifiers are made from bipolar transistors, the same circuit can put Class A bias current into the bases of the transistors. For big signals, the AGC automatically turns off the bias and runs the transistors “Class C.” When signals become weak, the bases are biased “ON” so that the signals don’t have to exceed the 0.6 volt input barrier.


The product detector



My product detector is basically the same circuit I used as my mixer. Product detectors are “direct conversion mixers” that mix an RF “beat frequency” (BFO) signal with the IF frequency to produce a difference frequency which is the audio signal. A 470 microhenry RF choke keeps the RF out of the audio output. To say it another way, the choke keeps the .1 μF cap from shorting out the RF while letting audio frequencies pass on to the AF amplifier.

Notice that the 12 volt DC power supply for the BFO oscillator passes through another choke and goes out to the BFO oscillator box on the front panel. That is, the DC power input for the BFO and 9 MHz RF output from the BFO share the same wire. The 470 microhenry choke prevents the 9 MHz signal from shorting to the power supply line.

Product detectors are exactly what're needed for CW or SSB. However, when you tune in an AM broadcast station, it will have a whistling overtone on it until you tweak the BFO perfectly to get rid of the whistle. If you plan to listen routinely to short wave AM broadcast stations, you'll probably want to replace the IF crystal filter with a short circuit. I have a setting on my crystal filter rotary switch that is a simple short. Otherwise, the 3 KHz width of a single crystal will be too narrow and the sound will be “low-fidelity.” Another change you might consider is to put in a switch to bypass the product detector and use an ordinary diode detector for AM signals. Any of the four dual gate MOSFET transistor types mentioned earlier will work well for a product detector. Here is a JFET version that works OK.




With product detectors, anything works at least a little

In my experience, receiver RF mixers that produce a quality IF output are extremely finicky and are often plagued with low sensitivity and oscillations up and down the band. In contrast, a product detector is amazingly uncritical. I haven't tried to make one out of wood shavings or pebbles from the driveway, but I wouldn't be surprised if I could still hear signals.

For example, I built a new IF strip and product detector in hopes of improving the noise problem. It worked, but I was slightly disappointed with the sensitivity. I was inspecting my dual gate MOSFET product detector when I realized that I had soldered the MOSFET in 90° out of alignment. In other words, the drain was connected to the RF input gate, the source was connected to the drain circuit, and the BFO input was connected to the source. Delighted that I had discovered my problem, I correctly soldered in a new transistor. When built properly, it worked better - but not dramatically better.

In another experiment I disconnected the RF input so that the input to the product detector was just stray coupling from the IF strip. Signals were weak, but it still worked amazingly well! Finally, I disconnected the BFO input. I was relieved to confirm that it no longer tuned and received ham-band signals. Instead it worked like a crystal set and received the loudest signals on or near the ham band input. For example, on 17 meters, it brought in the Deutsche Welle (Radio Germany) loud and clear.

The AF amplifier



The output from the product detector is an audio signal that needs to be amplified before it goes to the earphones or speaker. Most ARRL designs use integrated circuits marked “audio amplifier.” The LM386 is a typical one-chip audio amplifier. I’ve used these and they usually work great. But of course I didn’t learn anything from the experience. So this time I built my audio amplifier out of discrete parts from an example in my 1986 handbook. It looked like two straight-forward “R-C coupled amplifiers” in series. But the design had extra filter components I didn’t understand. Every part that I didn’t understand, I left out. That was my education. The audio amplifier was dead as a doornail when I first turned it on.

An audio Automatic Gain Control (AGC)

I was particularly puzzled by the low frequency feedback link, R1, R2, and C1. I couldn’t understand what sort of “low frequency filtering” the designer was trying to accomplish. But, when the amplifier seemed completely dead, I put these mysterious components back in the circuit. Voila! The earphones came to life. It turns out that this loop biases the amplifier “on” for weak signals and biases it “off” for loud signals. It’s a sort of audio AGC circuit.

Remember that for a bipolar transistor to turn on, the input signal must be greater than 0.6 volts or no current will flow into the base. In a “Class A” amplifier a DC signal is added to the base. This increases the base voltage above 0.6 volts so that it’s always turned on. In this way a class A amplifier can amplify signals much smaller than 0.6 volts. The low frequency feedback adjusts the bias for weak and strong signals. When the signals are weak, the second transistor is turned off, so it’s collector voltage is high and unchanging. This big collector voltage is leaked into C1 to provide a forward bias for its own base, biasing it on and raising the sensitivity. Conversely, when the signals are strong, the collector has a big current flowing but a low average DC voltage from the collector to ground. This lower voltage biases the transistor more “off.”

Protecting your ears from strong signals

The audio amplifier is able to blow your ears off when you encounter a strong signal.

Therefore it's essential to add a clamp circuit to limit the voltage to the headphones to less than about a volt. I first did this with back-to-back 5 volt Zener diodes across the headphone jack. In practice, with sensitive modern 8 ohm headphones, I found that less than one volt peak is plenty of volume for me. Eventually I put in two ordinary silicon 1N914 diodes "shorted" in opposite directions across the headphones. This limits the positive and negative sound peaks to just +/- 0.6 volts and my ears haven't been blasted since.

How Hi-Fi should it be?

The original circuit was also sprinkled with 0.1 microfarad bypass capacitors as if the designer were trying to kill all higher frequency sounds and shunt most of the audio to ground. Since I was trying to get more gain wherever I could, I left out the bypasses. The amplifier worked well without them, but the sound of the static had an obnoxious, piercing high pitch that irritated my ears. I put the bypasses in and, as I expected, the audio sounded more "base" and became somewhat weaker. However, getting rid of that piercing, hissing static was well worth the loss of gain. Experiment!

The original design also had no emitter bypass capacitor, the 10 microfarad capacitor across the 220 ohm resistor. Not having this bypass capacitor reduces the gain because some of the audio voltage signal is wasted across the 220 ohm emitter resistor. Since I needed gain, I put in the capacitor and my gain jumped up noticeably. This bypass has no disadvantage that I could detect.


Sharp audio filters

Many receivers have audio filters that limit the audio frequency of signals passed onto the headphones. This can be useful for separating CW signals that are nearly on the same frequency. If I didn't have my multiple crystal filter selections, I would definitely want audio filters. But in practice, when QRM (interference) happens to me, the guy that is interfering usually has the same audio tone pitch as the fellow I'm trying to listen to. Obviously in this case an audio filter would not help. But if you want to add one later, it is never too late. Unlike IF crystal filters, audio filters can be added external to the receiver. There is an example of a 700 Hz audio filter in chapter 7.

Driving a speaker

If you don't need a speaker, you don't need a third amplifier stage. By the same token, an 8 ohm speaker plugged into the 8 ohm headphone output is much too faint. Also, 0.6 volts peak is not nearly enough to drive a loudspeaker.


In the original handbook AF amplifier design, the third stage was an emitter-follower for driving either a speaker or low impedance headphones. The advantage of this design was that the emitter-follower drove the speaker directly and there was no need for a high-to-low impedance transformer. The speaker was located between the emitter and ground while the collector was connected to the positive supply. It looked straightforward to me, so I tried it. Unfortunately the emitter-follower distorted the sound and "motorboated" on strong signals. That is, the sound ran in "putt-putt" bursts. I tried several modifications to solve these problems but was never able to fix it. I gave up on the emitter-follower and used another impedance step-down audio transformer to drive the low impedance speaker. I happened to have a handful of tiny speaker transformers in my transformer junk box, so for me this was an easy solution.



An optional extra amplifier for driving an external speaker

You'll find that a big speaker sounds much better than a little one. A speaker small enough to fit in the receiver itself will sound "tinny." I eventually wired my speaker output to a remote, 12 inch wide speaker.

Converters for the other HF bands



Block diagram of a converter for an HF ham band

I used the RF amplifiers and crystal oscillators out of the W7ZOI and K5IRK receiver. I built these modules close to what was described in the handbook and they worked right away. For my mixer module I used the same dual gate MOSFET circuit I developed for the 80 meter receiver. I had some difficulty with the low frequency preselector filters, so I used other designs as will be described.

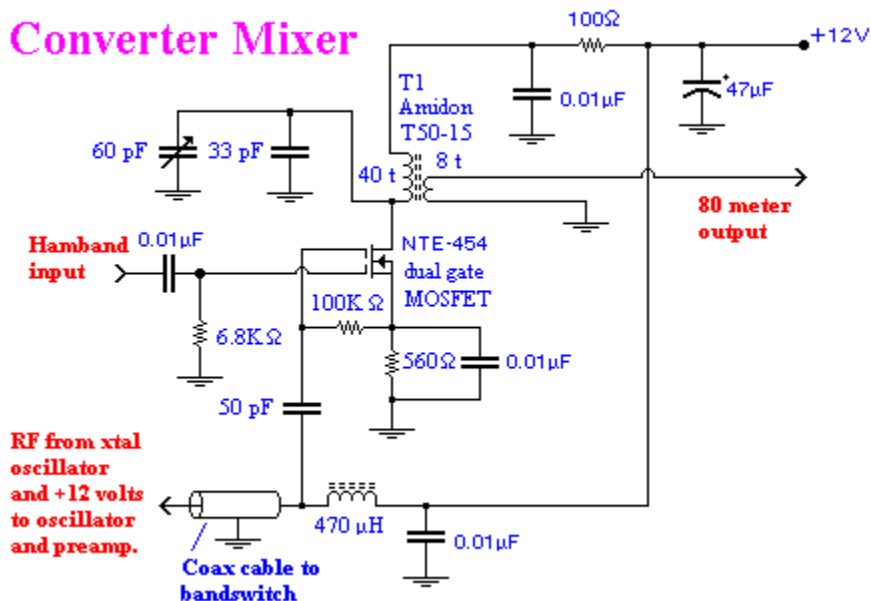
In my receiver all converters for bands other than 80 meters share the same dual gate MOSFET mixer. Band switching would be easier if each converter had its own mixer. On the other hand, those dual-gate MOSFETs are pricey transistors, so do what you like. Each band needs its own crystal-controlled oscillator and a pre-tuned bandpass filter or "preselector" to limit the input to the desired band. Bands above 30 or 40 meters need an RF amplifier. Below 20 or

30 meters, the signals and noise are stronger and an RF amplifier on the antenna input shouldn't be needed. I have had some trouble with weak signals on 30 meters, so maybe I should have added an RF amplifier on 30 meters after all. On the other hand, maybe the signals are just plain weak.

Each ham band needs its own HF converter

In the old vacuum tube days, tubes were physically large and it would have been extravagant for a homebuilder to use a separate converter for each band. My old homebrew vacuum tube receiver had a single multi-band converter that had to be tuned manually for each upper band. Every time I tuned in a weak station, I tweaked three variable capacitors and three separate gain controls.

Transistors and powdered iron cores are tiny, so today we can easily house a complete HF converter into a few cubic inches. Moreover, each converter runs on just one low DC voltage. This simplicity means that the power can be routed to the converter using the same coaxial line that receives the output from the converter. Since each converter is optimized for just one band, it can be tuned up once and forgotten.



All my converters share the same converter mixer shown above. A rotary bandswitch brings in the filtered and amplified hamband RF signal on the upper input. The lower input takes in the RF from the local crystal oscillator for each band. In addition, the lower "input" is also an output that delivers 12 volt power to the oscillator and the preamplifier for that band.

Hamband Converter Preamplifier

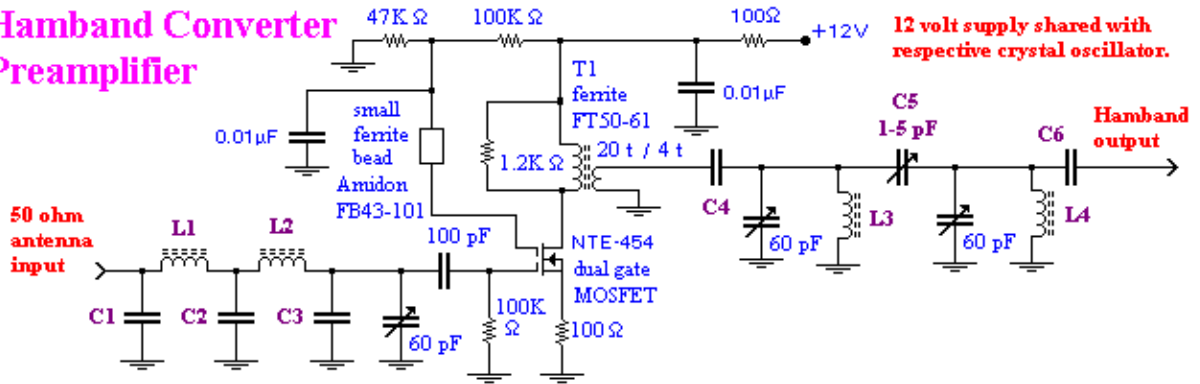


Table of values for converter preamplifier and filters

Hamband	C1	C2	C3	C4	C5	L1	L2	L3	L4	
	(pF)	(pF)	(pF)	(pF)	(pF)	(μH)	(μH)	(μH)		
30 meters	300	680	33	33	4.1	0.68	3.36	1.16		
							T50-6 CWS (Amidon) toroids	(13t)	(29t)	(17t)
20 meters	220	500	22	27	4	0.58	2.50	1.16		
							T50-6	(12t)	(25t)	(17t)
17 meters	180	390	none	22	3.9	0.40	1.94	1.44		
							T50-6	(10t)	(22t)	(19t)
15 meters	150	330	none	20	3	0.40	1.60	1.44		
							T50-6	(10t)	(20t)	(19t)
12 meters	120	200	none	12	2.1	0.26	1.44	0.58		
							T50-6	(8t)	(19t)	(12t)
10 meters	110	250	none	12	1.6	0.26	1.16	0.40		
							T50-6	(8t)	(17t)	(10t)

Preselectors for the lower hambands

As explained earlier, hambands lower than about 30 meters shouldn't require an RF preamplifier. The only purpose of the preselector is to limit the input signals to the particular hamband. This design shown below is the one recommended by W7ZOI and K5IRK .

Hamband Preselector Filter





Table of values for preselector filter (without amplifier)

Hamband	C1 C3 (pF)	C2 (pF)	C4 C10 (pF)	C5 C8 (pF)	C6 C9 (pF)	C7 (pF)	L1 L2 (μ H)	L3 L4 (μ H)
30 meters	300	600	32	180	50	4.1	0.68 (13 turns)	1.16 (17 turns)
		CWS (Amidon) T50-6						
40 meters	430	860	42	180	50	4.6	1.16 (17 turns)	2.50 (25 turns)
		CWS (Amidon) T50-6						
160 meters	1720	3440	none	250	120	20	4.64 (29 turns)	30 (79 turns)
		CWS (Amidon) T68-2						

My 160 meter filter is similar to the original handbook design and was extrapolated from their 40 meter values. For 30 and 40 meters I used the following filter design that I adapted from an article in QEX:

Preselector Filter for 30 and 40 Meters



- * All variable caps are 60 pF trimmers
- * C7 and C8 are .001 μ F (1000pF) feedthrough caps augmented with 0.001 μ F disk ceramic caps.

Table of component values for alternate preselector filter

Hamband	C1, C2 (pF)	C3, C4, C5, & C6 (pF)	C7, C8 (pF)	L1, 2, 3 (μ H)	L4 (μ H)
30 meters	710	33	2000	3.5	4.6

			CWS (Amidon) T50-6	(30 turns)	(34 turns)
40 meters	1000	47	3000	5.0	6.6
			CWS (Amidon) T50-6	(35 turns)	(41 turns)

The above design is a series of 4 series LC circuits, each in its own egg carton compartment soldered onto the PC board. The signals are fed from one little chamber to the next using .001 μ F feedthrough capacitors. Extra capacitors were soldered on to bring the values up to the values listed above. These filters tuned up well and I have no complaints about the receiver performance on 40 meters and above. 80 and 160 meters are more difficult, as discussed earlier and below.

160 meters is hard

160 meters presents two problems. First, standard broadcast signals are huge and just below 1.800 MHz. If you don't attenuate the broadcast band signals, you may be plagued with the loudest local AM stations appearing in your IF. You may find the AM broadcast filter described in chapter 7 useful. It worked best when placed in between the 160 meter preselector and the input to the 80 meter receiver board. It was not so helpful in series with the antenna jack outside the receiver. My 160 meter preselector was primarily a high pass filter.


The second problem with 160 meters is its proximity to 80 meters. When I first got my 160 meter converter working, I immediately heard hams and thought I was successful. It was a several evenings before I discovered that some of those hams were actually on 80 meters. Oops! When I tuned in a ham, I then switched to 80 meters to see if he was still there. If he disappeared, then I knew he was on 160 meters. Not really up to snuff. Since my core receiver is designed for 80 meters, the bandpass filtering on the 160 meter converter must be quite selective to keep out both the broadcast signals and the 80 meter signals. I'm still not satisfied with my present filter. As discussed elsewhere, 80 and 160 meters work much better when you listen to them using your transmitter "T-match" antenna tuner.

It can be hard to find an appropriate, cheap crystal for 160 meters. 5.5 MHz is available as a standard frequency and looked ideal. Unfortunately, this produces a large whistling artifact on 2.00 MHz. 5.6 MHz worked OK because the artifact is now on 2.1 MHz, completely out of the 1.8 to 2.0 MHz ham band.

Still another difficulty with 160 meters is that, if you build a multistage filter like those shown for 40 and 30 meters, the variable capacitors will be physically quite large. I didn't have room for such a large preselector, but if you started out with a big enough chassis that shouldn't be a problem. To scale the 40 meter filter to 160 meters, multiply the value of each capacitor and inductor times four. Good luck!

Crystal oscillators for the converters

These oscillators are nearly identical to the BFO oscillator described earlier. The 12 volt power for each oscillator is delivered via the cable from the bandswitch.



Converter Crystal Oscillator(s)

Data table for crystal oscillators

Hamband	Crystal freq. (MHz)	T1 primary total turns	Primary tap turns	Secondary turns
160 meters	5.600	39	8	7
40 meters	11.000	30	7	6
30 meters	13.800	26	6	5
20 meters	18.000	28	5	4
17 meters	21.900	20	4	4
15 meters	25.000	20	4	4
12 meters	24.800	17	4	4
10 meters	32.000	14	3	3

All toroids are CWS (Amidon) T50-6 powdered iron

Since I am converting to 80 meters, my receiver uses the same local oscillator frequencies I used in my CW transmitter, which has an 80 meter VFO. For most bands I use microprocessor crystals cut for frequencies 4.0 MHz above the bottom of the desired ham band. For example, I used 11 MHz for 40 meters, 18 MHz for 20 meters, 25 MHz for 15 meters, and 32 MHz for 10 meters.

You could also use local oscillator frequencies below the desired band. For example, to cover 15 meters, you might use a 17.5 MHz crystal oscillator. This will convert the 15 meter band to range from 3.5 MHz to 3.95 MHz. Notice that the direction of VFO tuning is reversed from using local oscillator frequencies above the desired band. Whatever you do, it's good to be consistent so that direction of tuning and the calibration can be consistent. Also, if the local oscillator frequencies are round numbers, like 11, 18, 25, and 32 MHz, then decimal calibrations on the VFO will be identical. Unfortunately, to get quality crystals for the WARC bands that didn't drift, I eventually had to have custom crystals ground for me. Maybe cheap isn't so

important after all.

Mechanical construction of converters

I built my converters on two-sided PC board material. Strips of PC board were soldered onto the main board to provide the walls and partitions in the egg carton construction. A folded aluminum, press-on cover keeps out most stray RF. The assembly shown below contains four converters for 20, 40, 30, and 12 meters. It looks daunting, but I built them one band at a time over several months. I only built another converter when I was ready to listen to a new band.




Converter modules for 20 meters, 30 meters, 40 meters, and 12 meters



This module contains converters for 160 meters and 17 meters.

Band switching



The very first module in my receiver that worked properly was the converter for 15 meters. I checked it out by using it with my old receiver tuned to 80 meters. I immediately noticed that it worked better than the old receiver tuned to 15 meters. Later, when the new 80-meter receiver began to work, I was able to listen to 15 meters by plugging the 80 meter receiver to the 15 meter converter directly with phono plug cables. Next I built converters for 10 meters and 40 meters. Whenever I switched bands, I moved my cables around like plugs on an old-fashioned telephone switchboard.

Eventually, as I built more and more converters, this process became increasingly cumbersome. Finally I gritted my teeth and spent a Saturday afternoon wiring up the bandswitch shown above. To switch from one band to another, a rotary switch selects the desired input filter/ RF amplifier and local oscillator for each band. Like the BFO design described earlier, the 12-volt power for each local oscillator and RF preamplifier arrives riding on the local oscillator RF signal to the mixer. In other words, a single switching circuit carries both the local oscillator signal and the 12 volt power. Even with this reduction in complexity, the 9 band rotary switch becomes a confusing pile of RG-174 coax spaghetti. Be sure to label each plug and socket. I used a fold of white plastic tape on each cable and labeled it with a fine-point permanent marker pen.

Receiver muting

While transmitting it is extremely distracting to hear your own signal blasting your ears in the receiver. Even with the antenna switched over to the transmitter, the receiver will still be overwhelmed by your own loud signal. A CW signal will usually sound distorted and will interfere with your sending. I have found it best to turn the receiver off altogether while transmitting. That is why I built a Morse code sounder into my electronic bug so I would have a clear tone to listen to while I am sending.

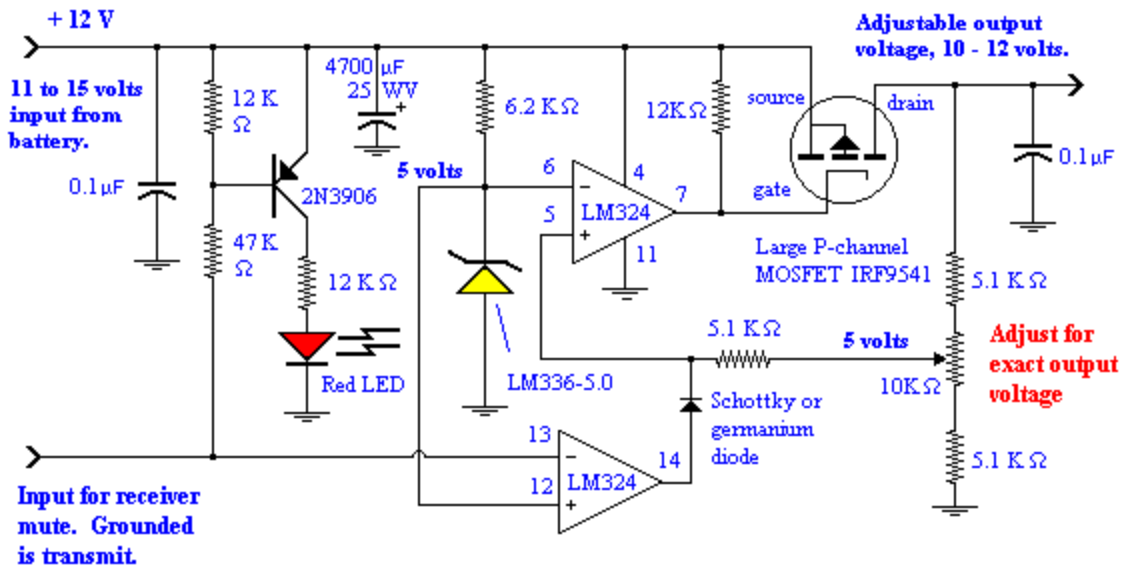
Turning the receiver off by hand is too slow, so you will need to build a muting circuit. I trigger mine with the same transmit/ receive signal line used to turn on the transmitter. I have

found that the simplest solution is turning off some or all of the power supply lines. I first tried just turning down the audio and IF gain pots by pulling the pot wipers down to ground using the collectors of transistors. This was inadequate, so I have since gone to keying the receiver power supply on and off, just as I did in the transmitter. I leave the receiver VFO supply on constantly to prevent any drifting due to having the VFO circuitry cool off and warm up repeatedly.

Receiver power supplies

The VFO power supply used in the receiver is the same one described in chapter 10. You could also use the line powered regulated supply described in chapter 8. I run my station on a 12 volt battery, so I just needed the low drop-out regulator shown below. The power supply requirements for the receiver are similar to the low power stages of the transmitter. I used the same supply that I developed for use with my QRP modules in chapter 8. This supply is over-designed for a receiver and can supply at least a few amperes.

The mute capability was implemented using an unused op-amp as a comparator. The mute input lead comes from the transmitter. Or in my case, it comes from the transmit/receive switch on my homebuilt bug. During transmit the mute line is pulled low to ground. During receive, the 12 K and 47 K resistors pull the mute line up to 12 volts.



Low Dropout 12 Volt Power Supply with Mute Capability

Referring to the lower op-amp, when the positive input wired to the 5 volt reference is higher than the mute line, the op-amp output, pin 14, rises up to 12 volts. This causes current to flow down through the diode to the positive input of the upper op-amp, pin 5. When the positive input (pin 5) is higher than the 5 volt reference on pin 6, the upper op-amp output (pin 7) rises to 12 volts, thereby turning the power MOSFET off and shutting off the current supply to the receiver.

Red mute indicator light

While the receiver is being muted, the base of the 2N3906 transistor is pulled down

thereby turning on the transistor and lighting the red LED. My LED was an extra bright type that only needed 1 milliampere to be adequately visible. You may need a resistor smaller than 12 K for good brightness.

Use linear regulators, not switching regulators

I tried to use switching supplies for this receiver, but reception was always impaired by switching supply noise. Switchers cause buzzing and hissing on every ham band. RF filtering was a great improvement, but it was never good enough. Building a low-noise receiver is hard enough without fighting a built-in noise generator.

In summary,

If you've already built homebrew transmitters, you'll find that building a receiver is easier than building your multiband, VFO-controlled CW transmitter. The hardest parts are the 80-meter pre-selector, the 80 meter mixer, and the IF amplifier. After that, it's pretty straightforward. Receivers don't need RF power gain stages, so the currents are low and the circuit board layout is forgiving.

In closing, remember that when you build a transmitter, you will hear no end of complaints about your signal until you get *ALL* the bugs out of it. But when you're on the air with your homebrew receiver, your contacts will never complain about your receiver. Or at least you won't ever hear them complain.