

CRYSTAL SETS TO SIDEBAND

©2002 by Frank W. Harris

Chapter 11

Building a VFO for the higher bands (PMOs)



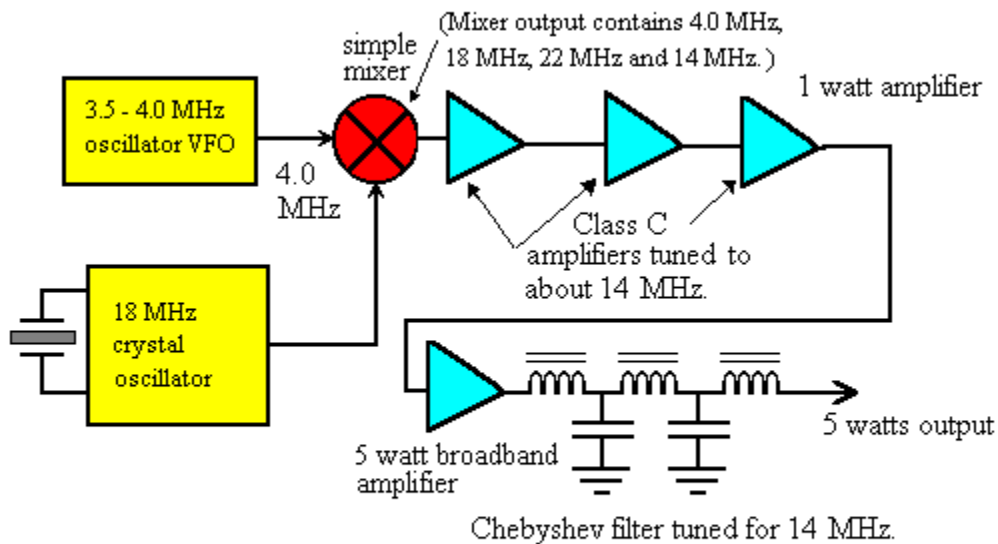
A QRP module for 30 meters is shown above. In many ways this module resembles the crystal controlled QRP module described in chapter 6. However, it uses a crystal-controlled premix oscillator (PMO) to convert the low frequency VFO sinewave up to the desired ham band. This particular module receives an 80 meter VFO signal and converts it to cover the 30 meter ham band, 10.100 to 10.150 MHz. The VFO signal and DC power come in at the back. The 5 watts of RF output exits from the phono jack at the front on the heat sink. The telegraph key port is at the right. The Chebyshev output filters are visible in the foreground. The filter amplifiers are at the rear. It would be more professional if it were enclosed in a metal shield, but I like to see all the components! With a CW QRP, you can get away with not using shielding. Later you will find that shielding is essential for SSB.

You can't multiply frequency anymore

In the old days, it was customary to build a VFO for 1.8 to 2.0 MHz or 3.5 to 4.0 MHz. Then for higher frequencies, we ran the signal through successive frequency multiplier amplifiers to get 7, 14, 21 and 28 MHz. A frequency multiplier was simply an amplifier tuned to the second or third harmonic of the input frequency. By using an amplifier tuned to multiples of the base frequency, the desired harmonic could be selected. For example, the tapped coil tuned amplifiers described in chapter 6 work well for this purpose.

If your VFO oscillator is crystal-controlled, then frequency multiplication is still practical. However, if your VFO drifts more than about 2 Hz, you may get complaints on the upper bands. For example, if you have an 80 meter VFO you will have to multiply the frequency eight times to raise the frequency up to 28 MHz. But, if your VFO drifts 5 Hz, then the multiplied signal will drift 40 Hz at 28 MHz.

Fortunately, carefully built high frequency crystal oscillators can be quite stable even up to 30 MHz. The solution to the drift problem is to “add” a low frequency VFO to a stable high frequency crystal oscillator. These crystal oscillators are called *Pre-Mix Oscillators* or *PMOs*. A “mixer” performs the frequency addition by literally combining two sinewave signals. The composite signal contains, not only the original signals, but also the sum and difference between the two signal frequencies. Filters follow the mixer to extract and amplify the desired frequency component. The process is illustrated by the block diagram of a 20 meter QRP transmitter shown below:



20 Meter 5 Watt QRP Transmitter Block Diagram

PreMix Oscillator method of frequency translation

In the diagram above, an 80 meter VFO is “converted” to 20 meters. The 80 meter sinewave is mixed with the output from an 18.000 MHz crystal oscillator. When the VFO is set to 4.0 MHz, the output from the mixer is a messy looking waveform that contains several frequencies, namely - 4.0 MHz, 18 MHz, 22 MHz and 14 MHz. By tuning the next three amplifier stages to 14 MHz, the “contamination” goes away and we get a pure 14.0 MHz sinewave tunable up to 14.5 MHz. The crystal oscillator may contribute a Hz or two of drift, but basically, the drift on 20 Meters is the same as it was on 80 Meters. The mixer is comparable in function to those used in superhetrodyne receivers, but PMO mixers are much less critical. Low noise and extreme image canceling aren’t necessary because both input signals can be as large as you like.

CRYSTAL OSCILLATORS ARE STABLE, AREN’T THEY?

A few years ago I thought I had the VFO problem conquered. I had just enjoyed a nine

3. Chapter 11, Harris

month “VFO vacation.” During this time my signal was so stable, that no ever commented on it. I was quite proud of myself. Then I built QRP modules to get on 17 and 30 meters. Suddenly the complaints started up again and I was mystified. After all, I was using the same VFO. What had changed?

I checked out my VFO. I discovered that, when it was cold, it drifted downward 20 Hz the first minute. Then, after a few more minutes, it stabilized and the drift was plus or minus a 2 or 3 Hertz. Of course, by definition, whenever I start sending, the VFO is cold. Therefore, unless I send for minutes on end, it must always be drifting. But even so, that didn't explain the 100 Hertz complaints.

Surely, it couldn't be the frequency converter crystal oscillators, could it? Crystals drifting!!?? Blasphemy!! I checked out my 17 and 30 meter crystal oscillators. The 30 meter oscillator drifted downward 50 Hz the first minute, 25 Hz the 2nd minute and eventually stabilized 150 Hz below the starting frequency.

Use HC-49 or larger crystals

The problem with my 30 meter converter turned out to be the crystal. The crystal was a tiny weeny can, about 1/4” square and a 1/32” thick. It had come from my junk box and I don't know what size number it was. However, I've since concluded that all the tiny ones in my collection aren't as stable as the HC-49 or larger crystals. Little bitty overtone crystals are particularly bad. Yes, eventually they settle down and become reasonably stable. But by then you have switched the QSO over to the other guy. Now your oscillator is cooling off again so it will be ready to drift during your next transmission.

Not All tiny crystals are bad. I have some little half-size 9.00 MHz HC-49 crystals that work extremely well in my receiver BFO and IF filters. I guess the lesson is that you need to check out your oscillator stability during that first critical minute. The drift after 5 minutes is interesting, but not very important for a ham transmitter.

TTL oscillator blocks

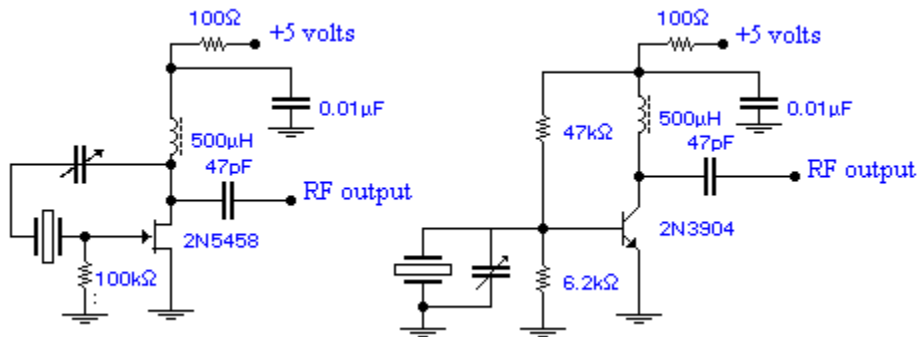
The 17 meter board had one of those sealed-in-a-can, TTL oscillator blocks. They're like an integrated circuit with the crystal and oscillator sealed in the same package. I used one because it was the right frequency and I happened to have it in my famous junk collection. My oscillator started out at the right frequency, but it ran surprisingly hot. Then it dropped at 25 Hz per minute. Although the drift slowed, the frequency never stopped sinking. I happen to have a bag of various frequency oscillator blocks and they all did that. Every one I tried was terrible! All except the really high frequency ones, like 50 or 100 MHz..... Those were really terrible. Some moved as much as 500 Hertz a minute. The only good news is that they're consistent. They all drift downward.

Solutions to crystal drift

I could leave the crystal oscillators running continually. That might work with the ones that stabilize, but then I might be forced to listen to a harmonic of the oscillator signal in the receiver. No, thanks! I already have a couple whistling artifacts in my receiver. Notice that vacuum tube oscillators might actually have an advantage in the warm-up problem. Since the tube filaments run continuously, a vacuum tube oscillator is always hot and a few milliamperes

of plate current aren't going to change its temperature much. In the old days, temperature controlled ovens were used to hold crystal oscillators at constant temperature. I don't know about you, but that's too drastic for me!

All oscillator circuits are not equal



Two common crystal oscillator circuits


The drawing above shows two common crystal oscillator circuits that I used in some of my first QRP PMO converters. The variable capacitors are used to trim the frequencies to the exact Hertz. **Both oscillators have the crystal connected to the base or gate.** (Remember that feature and you'll know which oscillators to be wary of.) In my experience, these base-connected crystal oscillator circuits drift downward when first turned on. Finally, after a minute or two, they stabilize.

Plan your frequency converters so the drifts cancel.

Suddenly I understood why I received no complaints while using my 40, 20 and 15 meter crystal-controlled frequency converters. They used oscillators with the crystals connected to the bases like those above. However, the crystal frequencies were 4 MHz **above** the target band. As the oscillators drifted downward for the first minute, typically at 20 Hertz per minute, my 4 MHz VFO was also drifting downward at the same rate. For example, (25 MHz - 20 Hz crystal drift) minus (4 MHz - 20 Hz VFO drift) = 21.000,000 MHz. The result was a relatively constant frequency and no complaints. After a few minutes the drifting stopped and the crystal oscillators were slightly more stable than the VFO.

Notice that if these converter crystal oscillators had been **below** the target hamband, then the drifts would have been added instead of subtracted. It turns out that my receiver was designed this way. OOPS! Oh well, no ham has ever complained about my receiver during a QSO. And of course receiver oscillators can run non-stop, so initial drift isn't a big deal.

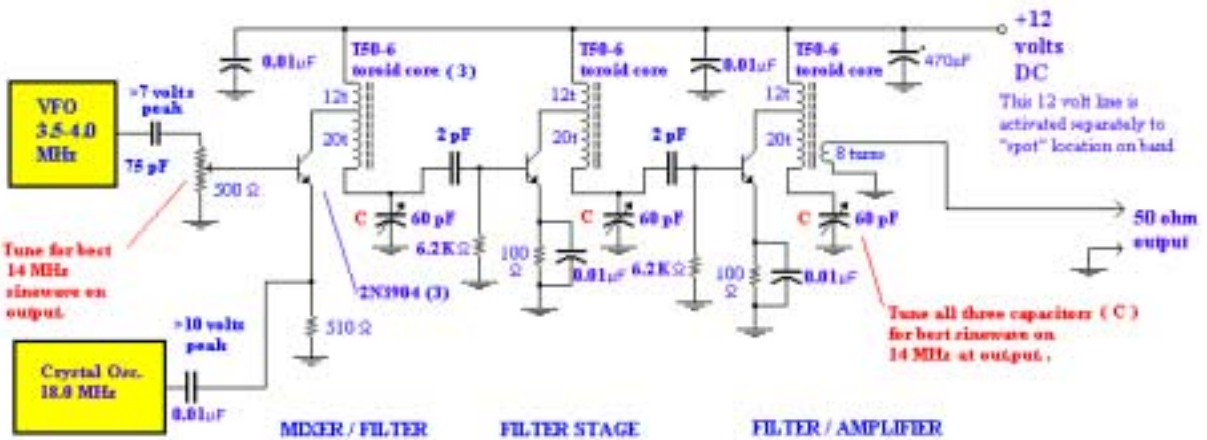
Butler is better



A Butler crystal oscillator

I checked out every oscillator in my rig and discovered that some of them didn't drift when turned on. The stable ones used the Butler crystal oscillator above. Notice that *the crystal and its capacitor are in parallel with the emitter resistor*. The equivalent FET oscillator circuit works as well and maybe better. I don't really know why, but this circuit is stable the moment you turn it on. Maybe it's because the crystal isn't connected to the P/N base junction while it warms up. Anyway, Butlers typically drift no more than a hertz or two per minute. Two of my oscillators showed zero Hertz drift during the first minute. This is the same oscillator I recommended in chapter 6. Depending on your application, Butler also has the advantage that the series capacitor can pull the frequency lower than the base-connected oscillators above.

A VFO controlled QRP



QRP DRIVER FOR 20 METER CW (14 MHz)

My "standard design" for a QRP driver is shown above. Unfortunately each QRP driver covers just one band. However, once it's tuned and working, it covers the entire band without further tuning or fussing. I should mention that the same filter train can be designed so that it

could be tuned up for several different bands, for example 20 meters through 10 meters. But of course, changing bands would mean changing the crystal oscillator and retuning the entire chain for the new band – not exactly convenient band switching.

So far, I've built eight versions of this design covering 80 through 10 meters. Taken as a whole, the entire circuit shown above can be considered as a 14 MHz VFO. In other words, all those parts accomplish is to generate a stable sinewave on 14 MHz. A simple crystal oscillator using 14 MHz crystals generates the same result, but of course it will only tune a few KHz at best. Life is hard for us homebuilders in the 21st century.

So why didn't I build one driver that works on every band? Back in the vacuum tube days, that was easy to do. However, the spectral purity and stability of our signals back then was awful. Also tubes are much easier to work with. If you're a lightly-equipped tinkerer like me, you'll find that getting even one band working to modern standards is much harder than it looks. I suspect that's why hardly anyone does this kind of scratch-built homebrewing anymore. I strongly recommend you *start simple*.

Changing the direction of tuning

In the 20 meter QRP driver above, the 80 meter VFO is mixed with an 18 MHz local oscillator. Note that the oscillator could also run at 10.5 MHz. As an experiment I ran my 20 meter QRP with both 18 MHz or 10.5 MHz crystals. All I had to do was change the crystal and it worked fine. The filter remained tuned up without any adjustment. The operational difference is that *the direction of the VFO tuning reverses*. As explained in the last chapter, this can be useful if you are using a varactor VFO and need to have the high frequency end of the VFO range tuning the lower end of a ham band.

The mixer needs a big local oscillator drive signal

The mixer stage in the above converter is just an RF amplifier made out of a bipolar 2N3904 transistor, much like the amplifiers in the filter train. This mixer is just a class C amplifier stage with a 500 ohm emitter resistor. We can use class C because the input signals are much larger than 0.6 volts. One input, usually the VFO, feeds into the transistor base in the usual fashion. However, this drive level is adjustable with the input pot.

The local oscillator input is applied across the emitter resistor. I usually inject the higher frequency input across the resistor, but I've done it both ways. Unlike the input to the base, the input to the resistor has no amplification gain whatsoever. In order for the emitter signal to produce a big signal on the collector, all of the signal amplitude must be impressed onto the 500 ohm resistor.

What I learned the hard way is that, **the local oscillator drive must be strong enough to turn the mixer stage full on and off like a switch, cycle by cycle**. I use an input sinewave of at least 20 volts peak-to-peak. A wimpy little 2-volt crystal oscillator signal will produce little difference-frequency component in the output and it will take many stages of filtering to extract the desired frequency. To get 20 volts pk-to-pk signal drive, I had to amplify the output of the crystal through an amplifier stage before it went into the mixer. I threw out two boards before I figured this out. (I'm not too bright.)

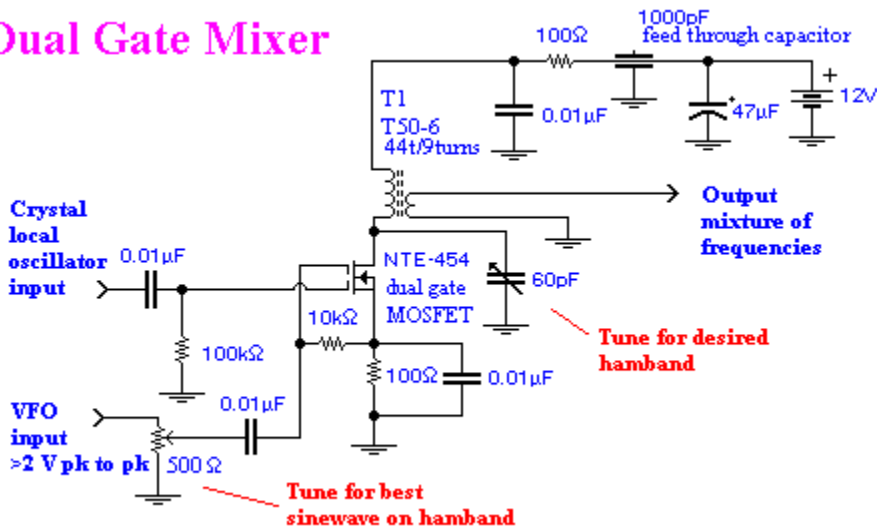
On the other hand, the second input signal, the VFO, can be small because it is amplified

by the transistor. Later, when you're tuning up the whole filter/ amplifier string for best output, you'll find that the maximum output and purity occurs at a specific setting of the input pot. The optimum VFO input level is not simply maximum input.

The LC filter/ tank circuit on the mixer collector is tuned to the desired sum or difference frequency. Using the formulas in your CWS (Amidon) T50-6 core literature, calculate the inductance needed to go with your trimmer cap to resonate at the desired band, just like we did back in chapter 6. I found that T37 cores were too small and didn't produce the gain per stage that I got out of the T50's. In contrast, the T68 cores were unnecessarily large.

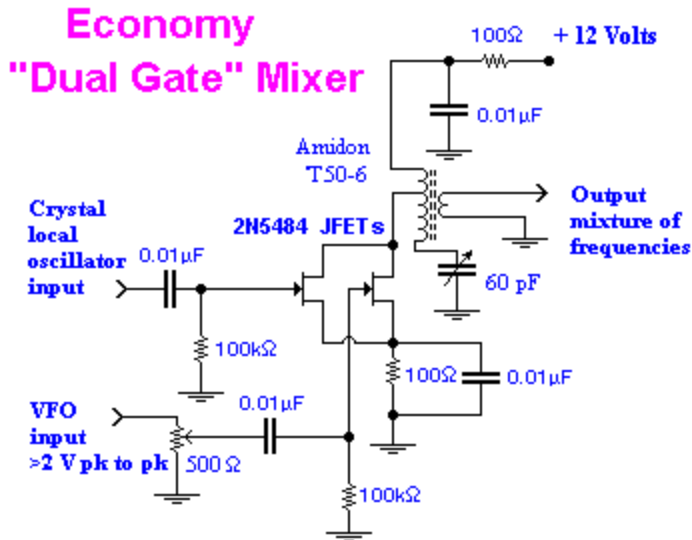
Dual gate MOSFET mixers

Dual Gate Mixer

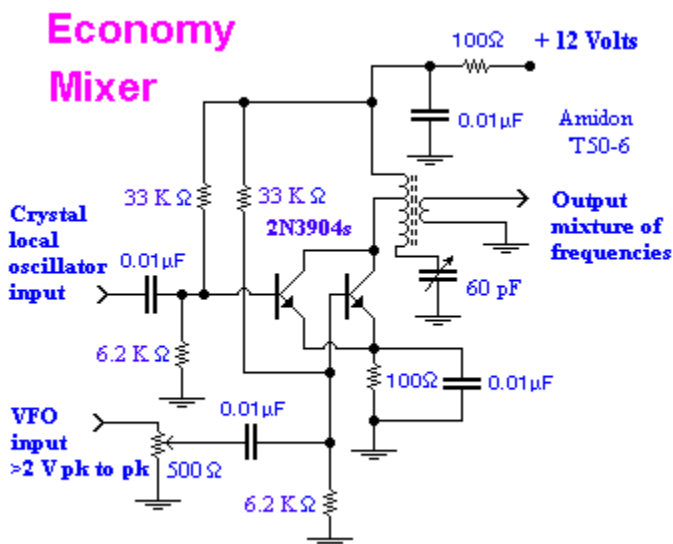


The principal advantage of the bipolar transistor mixer shown earlier is that *it's cheap*. I began using dual-gate MOSFET transistor mixer stages in a receiver project and found them superior in several ways. A dual gate MOSFET is a small RF transistor with *TWO* input gates. Otherwise, in principle a dual gate works just like the power MOSFETs described in chapter 6. Since both gates have lots of voltage gain, small signals may be used on both inputs. I found that each gate only needs 2 volts peak-to-peak and the output is much easier to tune and filter. Unfortunately dual-gate MOSFETs cost \$5 or more each compared with 20 cents. This saves me some complexity and my later converters have usually used the pricey mixer. I have used NTE-221, NTE-222 and NTE454. This application is quite uncritical and I believe any dual gate will work well. You will find that mixers in superhetrodyne receivers are not so uncritical.

The economy dual gate mixer



The above “dual gate mixer” is made from two paralleled JFETs. Two JFETs are about a tenth the price of the MOSFET and solve the cost problem. Like the dual gate MOSFET, this circuit has the advantage that both inputs have gain. The two circuits are interchangeable for this PMO application. Be advised that when I tried to use this as a receiver mixer, this dual JFET circuit was too insensitive.



If you are really cheap, you can use the same trick with parallel bipolar transistors. If both input signals are small, both transistors will need forward bias, the 33K resistors. If one of the inputs is large enough, say 5 volts peak or larger, you won't need the forward bias for that input. This circuit should be plenty sensitive for a receiver mixer. However, because it has PN junctions, it will be noisier than the dual-gate MOSFET mixer and I don't recommend it for receivers.

Tuning the mixer

When you first apply the two input frequencies to the mixer stage, the scope will show a messy, complicated waveform on the collector (or drain). It will be impossible to see what setting of the trimmer capacitor on the drain is best. The first step is to turn off the input frequency that is farthest from the desired output frequency. Now it will be easy to tune the trimmer for maximum gain. For example, in the 20 meter filter train above, tune the first mixer/ amplifier to the crystal frequency of 18 MHz. When it's peaked at 18 MHz, look to see if the trimmer is at or near maximum or minimum capacitance. If it is near minimum capacitance, but not quite there, the mixer stage should tune up well at the lower frequency, 14 MHz. If the gain at 18 MHz is maximum while the capacitance is tuned to either minimum or maximum, then you know that your toroid coil has too many or too few turns.

Later in the receiver chapter (13) and the sideband chapter (15) there are examples of un-tuned, broadband mixers that could also be used. Broadband mixers don't need tuning and they are less likely to oscillate. However, they have less gain and you may need more stages to reach the same power level.

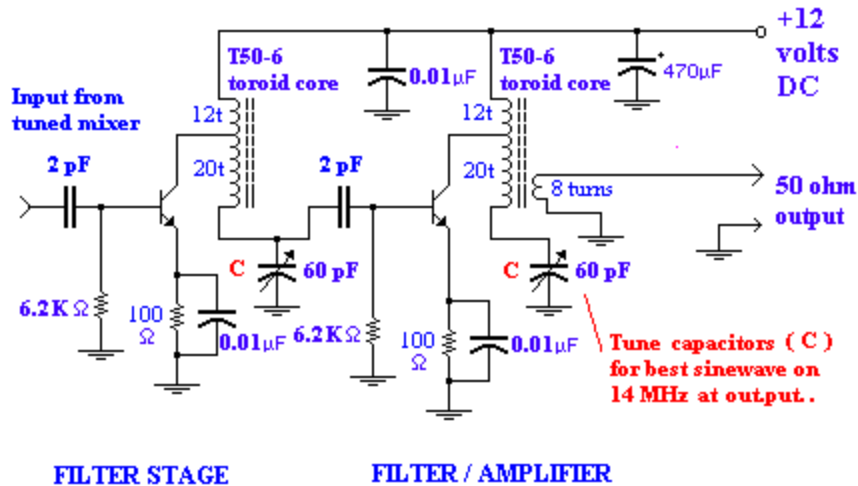
Filtering the desired frequency from the mixture

At the mixer drain (or collector) there are four frequency components and you must filter the one you want into a pure sinewave. Using the design tables from a recent handbook, it should be possible to design a Chebyshev or other design of bandpass filter to extract the ham band frequency from the three other frequency components. You'll find this takes multiple toroids and numerous specific values of fixed capacitors. My approach is to use two sharply tuned amplifiers as filters, just like an IF amplifier in a receiver.

The ease of filtering depends on how far the desired frequency is from local oscillator and other mixer products. For example on 14 MHz, the 4.0 MHz VFO is 28% of the desired frequency. 14 MHz compared with the 18 MHz crystal oscillator is 77% of the desired frequency. This is pretty close but not a problem. Now suppose we use a 32 MHz crystal on 10 meters (that is, 28 MHz). With a 4.0 MHz VFO, the desired frequency is 88% of the crystal frequency. You will find that tuning this up is much more "tweaky," but still practical. In general, having the crystal oscillator *BELOW* the desired frequency makes tuning the filters easier.

Bandpass "filter amplifiers"

Each amplifier stage is essentially like the bipolar transistor mixer shown earlier. However, the emitter resistor is bypassed with a capacitor so that, from the point of view of the RF, the emitter is connected to ground. The purpose of the RC in series with the emitter is to stabilize the gain and reduce the DC current drawn by the stage. You may use either class A or class C amplifiers. I often use class A amplifiers, meaning the stage is biased "on" at all times with a 33K resistor, just as we did back in chapter 6. Class A's draw more current than class C amplifiers that are basically the same circuit. However, they handle wave components of any amplitude. Or to say it another way, class A's work over a wider range of input amplitudes and don't introduce harmonics that must be filtered out.



FILTER FOR 20 METER CW (14 MHz)

An RF filter/ amplifier stage.

Two stages of filter amplifiers are shown above. Including the tuned mixer stage, three stages of tuned amplifier was enough for any HF band using an 80 meter VFO. However, as explained earlier, when you are trying to separate two frequencies that are only 10% different, using only 3 filter stages is just barely practical. If you are having trouble getting a pure sinewave, just add another tuned stage. However, I am assuming that these stages will only be used for CW and not for single sideband. (See chapter 15.) As shown above, the above stages are class C. If desired, you could forward bias these amplifiers on with 33K ohm resistors and convert them to class A amplifiers. In this way, they could handle smaller signal levels.

Use tiny interstage coupling capacitors

The **BIG SECRET** in making “filter amplifiers” work is to **use tiny coupling capacitors between amplifier stages**. Notice the 2 PF capacitors between stages in the diagram above. The purpose of these stages is **filtering, not power gain**. The LC circuit “rings like a bell” when the input contains a frequency that resonates with the circuit. This ringing exaggerates the desired frequency component. If you load the LC circuit by trying to couple significant power to the next stage, it’s like putting your hand on a ringing bell - the ringing will be damped and the filter effect dies. **To avoid damping the ringing, use tiny, 2 pF capacitors.** OK, on 80 Meters maybe 5 pF is acceptable. And on 10 meters 1 pF would be better. However, 2 pF works over the whole HF spectrum. Notice that if you use large coupling capacitors, say 50 pF, that 50 pF becomes part of the LC resonance and will dominate the tuning. Also remember that your scope probe contributes another 5 pF or so. To make a final adjustment of a filter stage, you must put the probe on the output of the stage **following** the one you are adjusting.

With one stage of filtering after the mixer, the waveform will still look “messy” on the scope. But after two stages of filtering it should be possible for your frequency counter to lock onto the correct frequency. As you tune the VFO, the reading on the counter should track solidly with no drifting and dancing digits. When properly tuned, the sinewave will look nearly perfect on the oscilloscope after two stages of filtering. When you first try to tune up all three stages at

once, you may be frustrated, but keep trying. When your counter “locks” onto the desired band, look at the ringing on the trimmer of the last stage with the scope while tweaking all the previous stages for the best sinewave. Notice that perfection occurs when you trim the input level of the VFO on R1. Now you see why the input is applied through a trim pot.

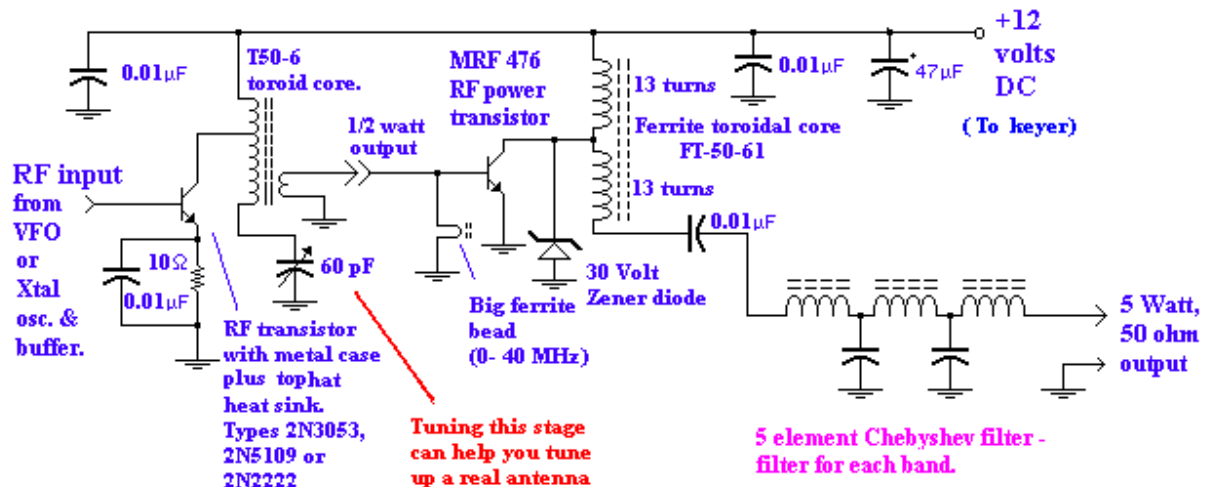
Another approach – untuned amplifiers and passive filter networks

Another practical circuit scheme uses un-tuned, broadband amplifiers and passive LC filters. An example is shown in chapter 15, page 13, where it was used in a single sideband transmitter. By condensing the tuning components into separate blocks from the amplifiers, it becomes practical to switch different filters into the converter and thereby cover multiple bands with one mixer/ amplifier assembly. Each filter block is connected to the switch by lengths of shielded RG-174 coax.

Where to get crystals for your local oscillators

Yes, you need a separate crystal for each band. Fortunately, standard microprocessor frequencies can cover the major ham bands. (e.g., 11 MHz, 18 MHz, 25 MHz, and 32 MHz cover 40, 20, 15 and 10 meters.) Mouser Electronics and Digi-Key sell these for about \$1 each. For WARC bands and 160 Meters you may have to spend some money or be creative. As explained earlier, don’t use those TTL oscillator blocks. The cheap ceramic resonator “crystals” are also a poor idea. Extra drift is not worth saving a few pennies.

The QRP power amplifier stages



QRP 5 watt amplifier

Your VFO now tunes the HF band of your choice. To increase the gain of this signal up to 3 to 5 watts, you will need two or three stages of power gain as described earlier in chapter 6. My transmitter chassis has screw holes that accept up to three QRP boards designed for different bands. To change bands, I move the input and output plugs to another board.

My QRP boards use two power amplifier stages. The first is a tuned stage. The second is

a broadband amplifier followed by a Chebyshev low pass filter designed for 50 ohms. This scheme seems to combine the advantages of both systems. For example, suppose I connect the QRP output to a non-inductive 50 ohm dummy load. All seven QRP boards I have built have no trouble delivering a clean sinewave into a dummy load. That is, tuning is easy until you have to connect it to a real antenna or to a final amplifier.

Tuning the QRP output to an antenna or amplifier

Suppose after tuning it up with a dummy load, I connect the QRP to the final or an antenna tuner. Suddenly I discover that the QRP output is badly distorted. The output stage broadband amplifier may even go into “noise mode.” If I had designed both stages as broadband, I wouldn’t have anything to adjust. Strange as it seems, tweaking a tuned stage will usually match up the broadband output stage to my final amplifier. In general, the lower the frequency, the easier it is to match stages and the antenna. Matching 80 and 40 meters is like falling off a log. 10 meters is tough and I still haven’t gotten my linear final (described in chapter 12) to put out more than about 20 watts on that band. Needless to say, I am in awe of guys who homebrew transistorized UHF rigs.

In summary, when you build an amplifier input that is nominally designed for “50 ohms resistive,” you may find that it has lots of reactance (inadvertent capacitance and or inductance) and is quite different than planned. Notice that Chebyshev filters are designed for specific input and output impedances. In other words, filters don’t filter when they are mismatched.