

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

Chapter 7

A CODE PRACTICE RECEIVER

This chapter describes how to build a simple *direct conversion* ham receiver for the 40-meter CW band. This receiver can be used to listen to CW ham stations for code practice. It will also receive ham single sideband (SSB, voice) stations and foreign shortwave broadcasts. When combined with the CW QRP transmitter described in chapter 6, you could even use it to talk to other hams.

The main virtue of this receiver is its simplicity. The basic design has only 5 transistors and is an excellent first project for a new ham. It can be powered with a 9 volt battery. However a 12 volt battery made from small AA flashlight batteries will last much longer and provide more volume in your headphones.

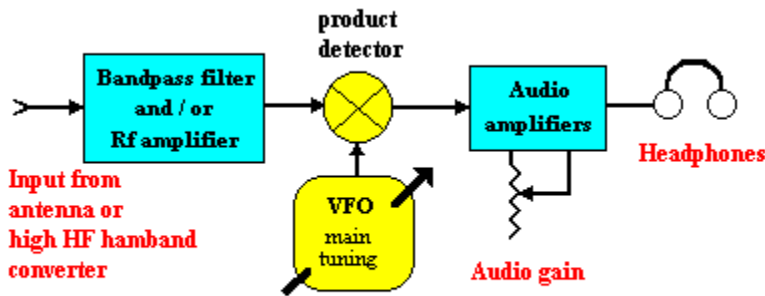


This receiver has good sensitivity and good stability. Unfortunately, by modern standards the selectivity of the basic receiver is poor for receiving CW. On many evenings the 40 meter CW band will be too crowded to listen to a single station easily. Adding the 700 Hz audio filter described later in the chapter will improve it greatly. Also, if you have a powerful AM broadcast station in your neighborhood, you may need to build a high pass filter to get rid of that interference. The broadcast filter is the little box shown to the right of the receiver.

When you're ready for a first rate receiver with selectivity as well as sensitivity, you may want to buy a commercial receiver or begin work on a superhetrodyne receiver as described in chapter 13. It is possible to get the same performance out of a direct conversion receiver, but that will require super-selective audio filters, image-canceling, and other R&D that I've never attempted. If you're interested, go for it!

Direct conversion receivers

A direct conversion receiver (DCR) has 4 basic circuit blocks. They are a band-pass filter, a variable frequency oscillator (VFO), a mixer or "product detector," and an audio amplifier.



Direct Conversion Receiver


The frequency tuning of a DCR is accomplished with an adjustable frequency oscillator called a VFO (variable frequency oscillator). This sinewave signal is mixed with the incoming radio signals in a special amplifier called a *product detector*. The signals heard in the headphones are those that differ from the VFO frequency by a difference within the audio range, say 20 Hz to 2.5 KHz. For example, suppose you are listening to a Morse code station on 7.100,000 MHz. You might tune the VFO to 7.100,700 MHz. That is, you tune 700 Hz above the actual transmitting frequency. The frequency difference will be heard as an audible, musical, 700 Hz Morse code tone.

The difference frequency is filtered, amplified, and passed on the earphones - pretty simple radio! Considering how few parts it has, it's amazing how sensitive it is. I measured the performance of mine and found it could easily detect a 0.5 microvolt signal on 40 meters.

The input filter

The product detector has a weak diode-like characteristic. That means it is prone to act like a crystal set and detect strong signals present on the antenna that are not near the VFO frequency. Consequently, the signals coming in from the antenna must first be filtered to restrict signals to approximately 7 MHz. As you may recall from your experience with crystal sets, the tuning selectivity of a simple input LC filter is usually extremely limited. However, a simple filter can get rid of background AM broadcast signals or strong foreign AM shortwave broadcast stations operating near 7 MHz.

40 Meter Hamband Passband Filter



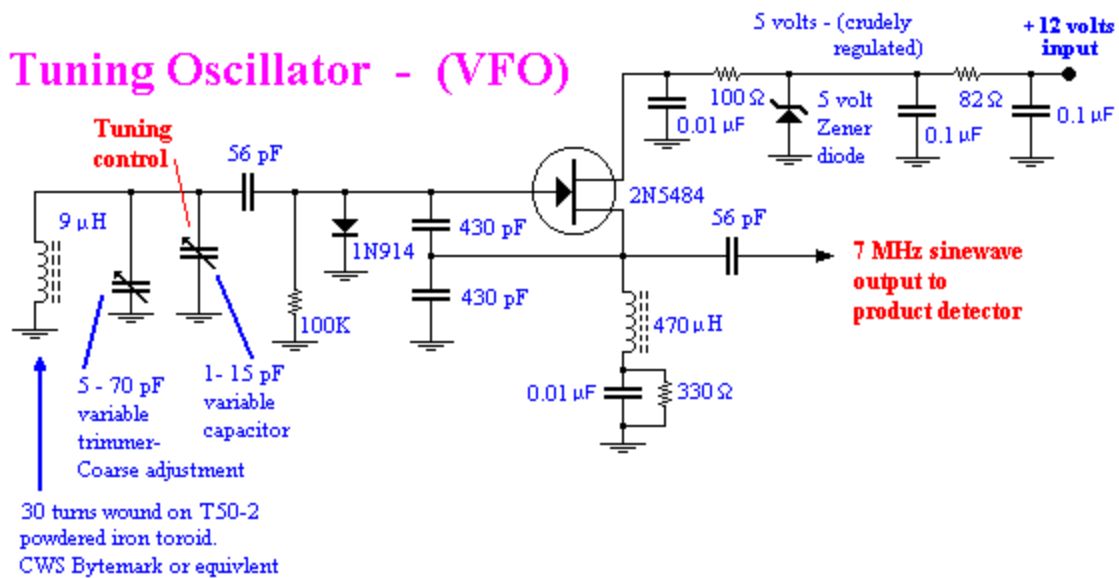
The input filter consists of two tuned LC circuits. The input from the antenna is wired as an RF transformer that matches the high impedance filter and high impedance product detector with a low impedance antenna. In other words, the antenna is assumed to be approximately 50 ohms, like a typical ham dipole with a coax lead wire. The transformer steps up the RF voltage about ten times using the 1:10 turns ratio. If the receiver will be used with a high impedance long wire or whip antenna, the input lead should simply connect to the top of the trimmer capacitor through a small blocking capacitor.

While listening to live CW signals on 40 meters, tune the trimmer capacitors until the CW signals are as loud as possible and any background voice stations disappear. If this filter is insufficient to get rid of local AM broadcast signals, a 1.8 MHz high pass filter is described later in this chapter.

This input filter section could also include an RF amplifier stage. However, this would probably not be useful on 40 meters because atmospheric noise is relatively high below about 20 meters. On the other hand, if you would rather build this receiver for 20 meters or above, an RF amplifier will be a great help to sensitivity. RF input amplifiers are discussed in chapter 13.

The RF tuning oscillator (VFO)

The circuit below shows the VFO tuning oscillator. As shown, the inductor and capacitor values will tune the 40 meter ham band and the shortwave broadcast bands just above and below 7 MHz. The circuit is a tunable oscillator comparable in function to the crystal-controlled oscillator used in the QRP transmitter in chapter 6. Instead of a quartz crystal, the frequency control is a tunable LC circuit that has a range of hundreds of KHz. This VFO is too crude to control the frequency of a transmitter. It would drift too much and the guys you are talking to would complain about your unstable signal. However, because you will be listening to a KHz of audio at once, a drift of a 100 Hz per minute will be scarcely noticeable. Building a first rate VFO for a transmitter is described in chapter 10.



JFET TRANSISTORS

This tuning oscillator uses a Junction Field Effect Transistor, (JFET). They are ideal for building VFOs. Unlike bipolar transistors, the main current from the JFET drain to the source does not pass through any PN junctions. PN junctions change their characteristics with temperature. Therefore, VFOs made from bipolar transistors tend to drift more than JFETs. JFETs work on the same principle as a MOSFET transistor, but the control gate is a P-N junction diode rather than a tiny capacitor. Power MOSFETs were introduced in chapter 6 as a way to key the QRP power supply on and off.

P-N junctions


As you learned in chapter 4, *a semiconductor P-N junction conducts when positive voltage is applied to the P-type semiconductor*. Conversely, it does not conduct when positive is applied to the N-type semiconductor side. In the crystal set, we were dealing with such tiny currents (microamperes) that the transition at zero volts was plainly visible. Commercial diodes handle much larger currents, milliamperes or even amperes. With commercial diodes the forward conduction doesn't fully start until a forward offset is reached. For an ordinary P-N junction silicon diode, this is around 0.6 volts. For a commercial Schottky diode, this offset is around 0.2 volts. This threshold means that if we apply a positive voltage to the P-type semiconductor side below the threshold, say, 0.1 volts, the current will be far less than a milliampere. Since very little current flows, the diode is still essentially "OFF."

P-N junctions can work like capacitors


In a MOSFET, gate current is prevented from flowing into the main channel by means of a layer of glass insulation. In a JFET, the gate appears to be a forward biased diode but doesn't conduct much current unless the drive exceeds the forward breakdown voltage, about 0.6 volts DC. *With zero volts input on the gate, the JFET is already partly turned on*. In other words, with zero volts on the gate the transistor is already operating as a Class A amplifier and, unlike a bipolar transistor, no bias resistor is needed.

JFETs operate at input voltages less than 0.6 volts


The JFET turns full on, (that is, it saturates) when the gate voltage approaches or exceeds 0.6 volts. Notice the diode clamp that prevents the gate voltage from ever exceeding 0.6 volts. To turn the JFET completely off, reverse voltage must be applied to the gate so that the charge carriers, holes and electrons, are completely depleted. For example, in the N-channel JFET shown below, the transistor turns completely OFF when the gate goes negative with respect to the source lead of the transistor.



As you would expect, P-channel JFETs are mirror images of N-channel JFETs with respect to construction and polarity. To turn a P-channel fully off, its diode gate is back biased with positive voltage as shown below.



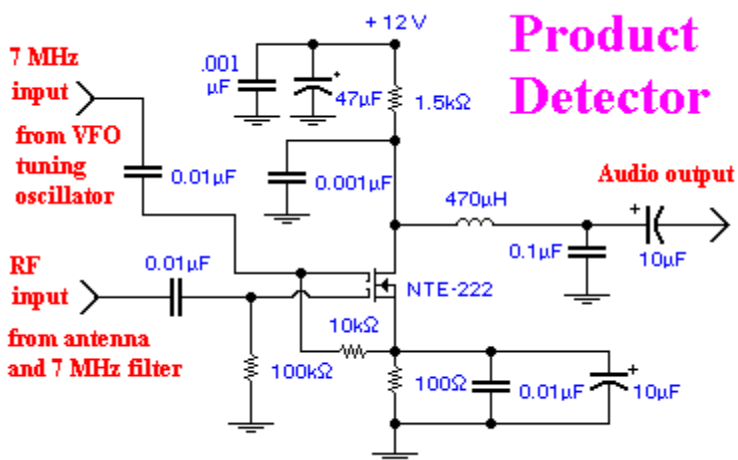
COMMON JFET TRANSISTOR CASES



The product detector

This receiver project uses a *product detector*. The product detector amplifies and mixes the VFO signal with the signal coming in from the antenna. The resultant mixture of radio signals is amplified and sent to an audio amplifier. Since the only audio component present is the *DIFFERENCE* between the two radio signals, that's what the audio amplifier passes on the earphones.

The product detector circuit described here uses a dual gate N-channel MOSFET transistor. A dual gate MOSFET works like a single gate MOSFET (or JFET). But as the name implies, it has two control gates modulating the drain-to-source current instead of just one gate. With two gates this MOSFET is convenient for mixing two separate signals because the output of the amplifier is an amplified mixture of the two inputs. Mixers and product detectors can be built with single gate MOSFETs and JFETs, but with a dual gate MOSFET, both inputs are amplified in one stage.




The gates of the MOSFET are tiny capacitors and therefore have extremely high impedance. The amplifier design is a *common source* design, which implies that the source is essentially at ground. This also means the load resistance, 1.5 K ohms, is on the drain side (+12 volt side) of the transistor. Because the desired output is an audio frequency signal, the load impedance is a resistor, not an inductor. In order to work at audio frequencies, an inductor would

have to be huge and for this application would have little advantage over a resistor. The 470 microhenry RF choke (inductor) and the 0.001 and 0.1 microfarad capacitors serve as a filter to remove the RF from the audio output. Without the filter, the audio amplifier would be more likely to rectify, (i.e. detect), strong RF signals that are able to get into the product detector and are not within the desired audio frequency difference from the VFO.


Modern designs often use integrated circuit mixer modules for product detectors. Personally, I prefer the dual gate MOSFET design simply because I know what's in there. It isn't just another mysterious integrated circuit. In this application any of the following dual gate MOSFET transistor types will work fine. I have successfully used the NTE221, NTE222, NTE454, or NTE 455. These are all "replacement" transistor types. So far as I know, any of the original 3N-series dual gates such as the 3N140 will also work. The 40763 by RCA is available from RF Parts Co. or from Mouser Electronics and should also work. I haven't personally tried it but I've seen it used in similar circuits.

MOSFET TRANSISTOR CASES

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

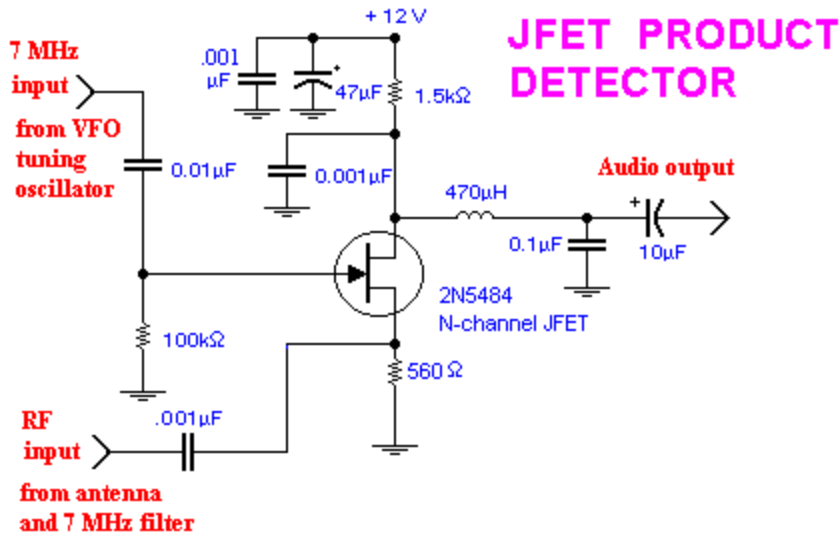


MOSFET POWER SWITCH
TO - 220 CASE
e.g., IRF9140, SMP20P10,
MTH25P05, etc.




A JFET product detector

Since this book was first written, dual gate MOSFETs have become increasingly expensive and hard to find. In the event that you need a substitute, here is a JFET circuit that works but in, my opinion, isn't quite 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.



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. Once again, most designs seen in ham magazines use integrated circuits marked “audio amplifier.” For example, the LM386 is a typical one-chip audio amplifier. I’ve used these and they usually work well. Of course I didn’t learn anything from the experience.



The diagram above shows a basic audio frequency *RC coupled* linear amplifier. If you look closely at the component values, you'll see it is similar to the class A RF amplifier we used as a buffer amplifier in the QRP transmitter in chapter 6. The differences are the large sizes of the capacitors and the absence of inductors. A 5.1 K Ω resistor is used as a load where you would expect to find an inductor in an RF amplifier. You could use an inductor here, but at audio frequencies it would have to be huge and expensive. Notice that the 36 K Ω resistor turns the amplifier partly on so that an audio sinewave will "fit" between zero volts and 12 volts.

its sensitivity. Conversely, when the signals are strong, the collector has a big current flowing but a low average voltage from the collector to ground. This lower voltage biases the transistor more “off.”


Protecting your ears from strong signals

This 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 one 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” with opposite polarities across the headphones. This limits the positive and negative sound peaks to just 0.6 volts and my ears have been adequately protected.

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 worried about having enough gain, 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. That's the 10 μ F 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 didn't want to waste any gain, I put in the capacitor and my gain jumped up noticeably. This bypass has no disadvantage that I could detect.



Mechanical construction

I built my receiver in a box made from soldering pieces of two-sided PC board together. This is the same construction described for the QRP transmitter described in chapter 6. The black coax RG-174 cables running around the inside of the box connect the optional 700 Hz

filter module described later in this chapter. The coax shields the wires from stray RF that might be present. Only one end of the coax should be grounded.

The tuning knob ideally should be a mechanical, vernier, planetary-gear device to spread out the 40 meter CW band and make the signals easier to tune. Alternatively, one can use two capacitors in parallel with a smaller variable capacitor, say 3 to 15 picofarads (pF), for fine tuning. In this way, out of the 180 degrees of useful knob rotation, more degrees can be devoted to the CW band and less to the nearby broadcast stations. A 70 pF trimmer capacitor tunes the whole frequency spectrum around 40 meters to align the small capacitor to the ham band. The small variable capacitor is called an *electrical bandspread*. In my receiver the 15 pF fine-tuning capacitor has the big black knob, while the tiny 70 pF ceramic screw adjust trimmer is directly behind it.

The receiver power supply

The receiver will work well on 9 volts at about 10 milliamperes. However, when the voltage drops below about 8 volts, the volume and sensitivity will fall off dramatically. So, if you use a small alkaline 9 volt "transistor radio" battery, it will only work well for about 1/3 of the energy stored in the battery. *To get the most from alkaline batteries, you need to be able to run them down to 2/3 of their original voltage.* In this case the receiver would have to work well down to 6 volts. I recommend making up a 12 volt battery out of AA cells. Radio Shack and other companies sell battery holders that carry 6 or 8 cells to give you 9 or 12 volts. Or if you are using the receiver with the QRP transmitter described in the last chapter, the receiver can share the QRP power supply.

Hearing the VFO

When I was trying to use the DCR with my QRP transmitter, I found I couldn't hear my own transmitter VFO (crystal oscillator) over the stations I wanted to call. To fix that, I ran a coax cable over to the QRP and imported a weakly-coupled transmitter crystal oscillator signal into the receiver box. The center connector on the right side of the receiver is for that purpose. The receiver has is no direct connection to the transmitter oscillator. The connector merely projects a small piece of wire into the receiver box like an antenna to broadcast the weak oscillator signal.

FIXING THE LIMITATIONS OF DIRECT CONVERSION

One of the DCR limitations I had read about didn't need to be fixed. Because a DCR has nearly all its gain in the audio amplifiers, these amplifiers are supposed to be prone to *microphonics*. In other words, DC receivers are supposed to be sensitive to vibration. When you touch a knob or bump the table, you are likely to hear that sound amplified in the earphones. I observed no microphonics at all. I suspect microphonics may have been a common problem using vacuum tube amplifiers.

A DCR detects both sidebands at once

Unfortunately, the product detector mixer detects both the upper and lower sidebands simultaneously. This is OK for listening to AM radio stations, but has poor selectivity on a crowded CW band. A good CW receiver can select a bandwidth as small as 500 Hz or less. A

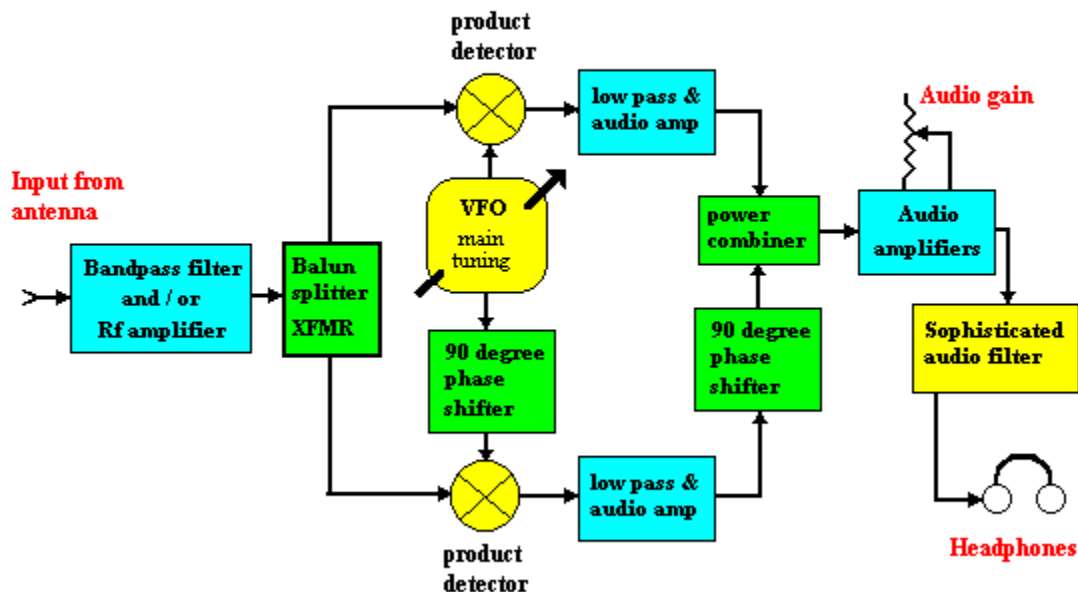
simple direct conversion receiver will have a bandwidth of 10,000 Hz or more, depending on the audio filtering.

Although audio filtering is helpful for a DCR, the big problem is that you'll still be listening to audio from both "sidebands" at once. CW by definition is extremely narrowband. It covers a few Hz at most. Therefore, CW doesn't really have sidebands. However, to hear the signal as a musical tone, a receiver must be tuned either above or below the actual frequency by an amount equal to the frequency of the desired musical tone. Most modern hams are using transceivers that automatically listen about 700 Hz above or below their transmitting frequency. This offset frequency is adjustable, but is nearly always set between 500 to 1000 Hz. By convention, ***On 40 meters and below, they are listening BELOW the transmitting frequency.*** On HF bands on or above 30 meters, today's hams listen ABOVE the transmitting frequency. This convention has its origin in the method for generating single sideband phone that used to be common 40 years ago. The old SSB design isn't used anymore, but the convention lives on. The new 60 meter band is specified for upper sideband only and is an exception. The result is that if you answer a station on the upper sideband side that is listening on the lower sideband, he will never hear you because modern equipment only hears one sideband at a time.

Sharp audio filters aren't enough

A sharp audio filter on the DCR output does not correct this problem. For example, suppose you build a sharp audio filter to select 700 Hz audio tones only. Yes, you'll hear your 700 Hz offset CW station just fine. But you're likely to be confused by an ***image*** signal that you can also hear 1.4 KHz away on the opposite sideband. That is, $2 \times 700 \text{ Hz} = 1.4 \text{ KHz}$. When you answer a CQ, this confusion will make it hard for you to know where your contact will be listening. Remember, when he switched to that band, he selected upper or lower sideband and didn't have to worry any more about this effect.

Modern Image-Rejecting Direct Conversion Receiver



A sophisticated direct conversion (DC) receiver eliminates one sideband by essentially building two DC receivers. One sideband is cancelled out by phase shifting the input VFO RF signal and also phase shifting the resulting audio signal to remove one sideband. So, although a DCR is “simple,” you might need to build two of them, plus a sophisticated audio filter. Personally, I decided that a dual conversion superhetrodyne was the best way to go.

AM broadcast interference

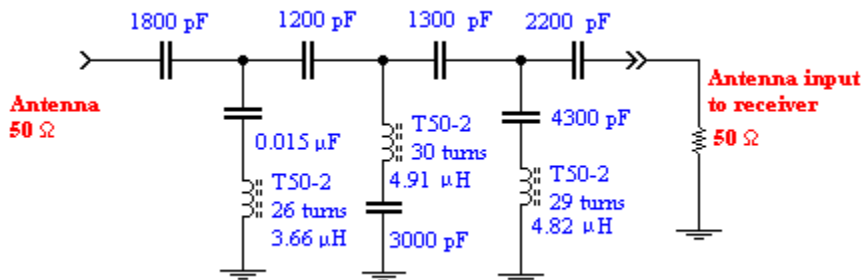

In my city there is a powerful AM radio station that broadcasts at high power during the day. When I put an oscilloscope on my 40 meter dipole lead wire, I see a 1.75 peak volts RF signal on the wire even with a 50 ohm resistor connected to the antenna. This represents 31 milliwatts of power! No wonder my crystal set was so loud.

$$\text{Power} = \text{Voltage (RMS)}^2 \div \text{load resistance} = V_{\text{RMS}}^2 / R$$

$$\text{RMS voltage} = \text{Voltage peak} \times (0.707)$$

$$\text{Power} = [(1.75 \text{ volts peak}) \times (0.707)]^2 \div 50 \text{ ohms} = 31 \text{ milliwatts}$$

The ham signals I was trying to listen to were buried in the voltage storms from the local AM rap-music station. The input filter for the DC receiver described earlier in this chapter was apparently inadequate to get rid of such a large signal. So, although I could hear ham stations, in the background I could always hear the rap music leaking through like a crystal set. I needed a better filter. The good old '86 ARRL handbook had such a filter design. It worked the first time without tweaking and reduced the RF voltage coming in on my antenna down to about 0.15 volts peak without affecting the strength of the 40 meter signals.



An AM broadcast filter

A 700 Hz audio filter – why you need one

As described earlier, if you intend to answer a fellow calling CQ, you must be sure you are listening on the same side of his transmitter frequency that he is. Because your direct conversion receiver has no image canceling, you will hear his Morse code both above and below his actual transmitting frequency. If you zero in your transmitter on the wrong side of his signal, your transmitter will be 1.4 KHz (2 X 700 Hz) away from where he is listening and he'll never hear you.

For example, on 40 meters, as you tune *UP* the band you'll first hear his Morse code on the low side of his actual frequency. That is the lower sideband. If you continue tuning *VERY* carefully, the tone you hear will fall to a low pitch, then disappear, then rise back up to the same musical tone again. That is the upper sideband.

If you wish to answer the fellow's CQ, listen on the lower sideband on 40 meters. Now set your transmitter to "spot" and tune your transmitter's crystal-pulling capacitor *DOWN*-frequency until you first hear an equal tone. If you have a little musical talent, you should be able to match the tones pretty well. *In summary, when listening for CQs, tune UP the 40 meter band. When spotting your transmitter oscillator on top of the station you wish to call, tune the transmitter DOWN the band.* If you were on 20 meters, you would reverse these instructions.


Those darn modern hams are listening with exquisitely narrow audio filters. They will usually only hear you if your signal is making a tone within a couple hundred Hz above or below their listening frequency. When they switch in the really sharp digital filters, they only hear stations within a few Hz of their listening frequency.

An analog audio filter



If we had ideal, small cheap components, the easy way to build analog audio filters would be to use inductors and capacitors the same way we build radio frequency filters. We might even imagine that there would be audio frequency quartz crystals for building super precise audio filters.

Unfortunately, in the real world those parts don't exist. In order to work at such low frequencies, the inductors would have to be gigantic and expensive. Since that approach isn't practical, we use resistor/ capacitor networks. *These RC networks don't resonate, but rather they just attenuate some frequencies more than others.* A large capacitor charges more slowly than a small one. When combined with a resistor, the frequency of the sinewave across either the capacitor or the resistor will be attenuated at either the high frequencies or the low frequencies, respectively.


700 Hz Low Pass & High Pass Filters



Low pass filter



High pass filter



Low pass filter

As shown above, a simple resistor and capacitor makes a low pass filter when the output is taken off the capacitor. If the capacitor is large, relatively high frequency sinewaves will be "shorted out" to ground. Remember, *the voltage across a capacitor can't change instantly*. In contrast, smaller capacitors will be negligible to low frequency sinewaves that have plenty of time to charge the capacitor during each half cycle. The resistor loads down the circuit and draws current into the capacitor. Similarly, during the next half cycle, the resistor has sufficiently low resistance to fully discharge the charged capacitor and get ready for the next half cycle.

High pass filter


When the R and C are reversed and the output is taken off the resistor, then the RC filter becomes a high pass filter. High frequencies pass right through a capacitor to the output as if it weren't even there. Remember, *the voltage across a capacitor cannot change instantly*. Therefore, the capacitor looks like a wire with zero voltage drop across it. At high frequencies there isn't time to charge. At low frequencies, the capacitor charges up completely and the voltage drop across the capacitor approaches the entire voltage of the input sinewave. These *RC coupled filters* all have a particular frequency called the *break point* where the capacitor becomes relatively insignificant in the circuit. A simple RC filter attenuates the sinewave voltage ten times (20 decibels) for a tenfold change in frequency away from its break point.

Compensating for attenuation

A problem with RC filters is that the resistors attenuate ALL the frequencies to some degree, not just the ones we don't want. Fortunately if an RC network is combined with an audio frequency linear amplifier, we can compensate for the attenuation. By amplifying the output, we can restore the signal strength of the desired frequencies back to its original strength.

Bandpass filter = RC networks plus amplifier

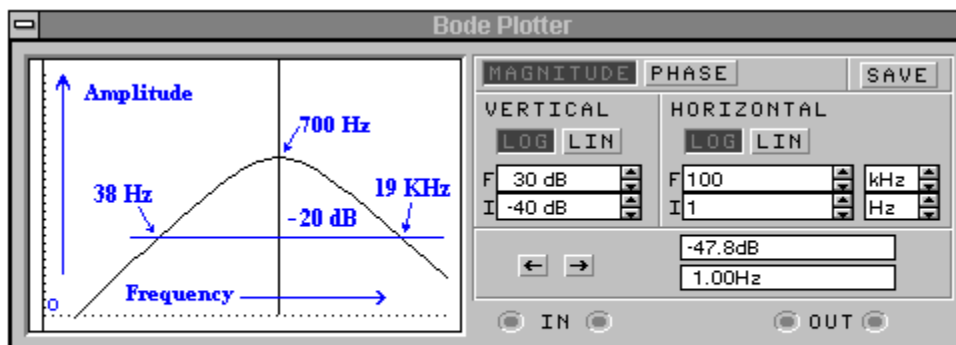
To make a bandpass filter we need to combine a low pass filter with a high pass filter and then amplify the result. The circuit below is a one-stage R-C (resistor/ capacitor) audio filter/ amplifier that accentuates 700 Hz tones and attenuates tones above and below roughly 600 and 800 Hz.



Single Stage 700 Hz Bandpass Filter

In this filter the Rs and Cs are not arranged as you might expect. However, if you carefully analyze how they affect signals passing through the circuit, they work the same as the separate circuits described earlier. That is, low frequencies and high frequencies are attenuated and the desired 700 Hz frequency range is relatively unaffected. The input RC, the 2K Ω and 0.15 μF capacitor, attenuate low frequencies because low frequencies, (for example, 60 Hz and direct current), can't get through the capacitor. Therefore the input network is a high pass filter.


Note that the 6.2K Ω ohm and 0.02 μF capacitors are wired between the input and output of the linear amplifier. This feedback network is the low pass filter. This combination of components tends to "short out" the amplifier and attenuate everything above about 700 Hz. Note that the small 0.02 μF capacitor has little effect on low frequencies because it charges essentially immediately. Because the capacitor charges almost instantly, its voltage follows the low frequency sinewave curve without significantly attenuating it. However, at high frequencies a 0.02 μF capacitor becomes charged more slowly than the high frequency sinewave is changing. Between half cycles, the 6.2K Ω ohm resistor has low enough resistance to be able to discharge the 0.02 μF capacitor significantly so that the network is restored to block the next half cycle.



Single Stage 700 Hz Audio Filter

The frequency plot above shows the response of the single stage filter. As you can see, this filter is pretty crude. In order for the voltage of a sinewave to be attenuated more than 10 times (20 decibels), the frequency must be below 38 Hz or above 19 KHz. To be blunt, this won't help much. It will decrease some of the unpleasant high frequency noise from static, but won't be particularly helpful in relieving the QRM (interference from other stations).


More precise filtering can be achieved by putting multiple audio filters in series. The plot below shows four of the same filters in series. This is a big improvement and makes it easier to listen to just one station at a time. It also eliminates the high frequency noise and makes listening less tiring.



Four Stage 700 Hz Audio Filter

Implementing the filter

We can build filters that work just like the graphs above using transistor linear AF amplifiers made from discrete components. Here is what a single stage 700 Hz filter would look like:



Operational amplifiers – "perfect" amplifiers

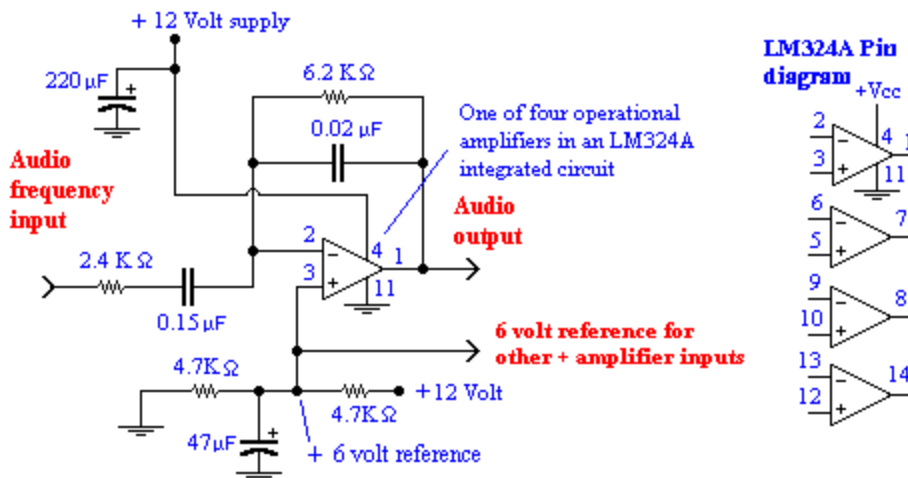
Notice the big blocking capacitor needed to prevent the DC bias for the class A amplifier from being changed by the 6.2 K ohm resistor. In any case, normal non-eccentric hams do not

make audio filters using linear amplifiers made from discrete parts. Instead, they use integrated circuit operational amplifiers. Being obsessive-compulsive, I thought about wiring all those bias resistors and 47 μF capacitors. Then I decided, what the heck! I'm checked out on audio amplifiers. I'm going to use a quadruple op-amp integrated circuit.

Op-amps are integrated circuits containing 12 to 30 transistors and resistors in a network that produces a "perfect" or idealized amplifier. How much gain should a "perfect" amplifier have? The best answer is that the gain should be infinite and then feedback can be used to produce any gain you like. Therefore real operational amplifiers have *near infinite gain*, and *near infinite input impedance*, and a *voltage source output*.

Typically there are two or four operational amplifiers on a single chip. The symbol for an operational amplifier is a triangle. The typical op-amp has an output pin and two input pins. For example, look at pins #1, #2, and #3 below. A **positive** voltage on the **positive input** pin makes the output pin voltage rise, that is, go positive. In contrast a **positive** voltage on the **negative input** pin makes the output pin go down. **Because of the feedback, the op-amp will "attempt" to drive its output pin up or down until feedback from the external circuit components cause the two inputs to be equal in voltage.** If the feedback is insufficient to produce equilibrium, the output will bump into either ground or the voltage supply, whichever happens first.

To bias an operational amplifier so that a large sinewave will "fit" between zero volts and 12 volts output, the amplifier must have a baseline about halfway between the supply voltage and zero. In this case the resting level will be about 6 volts. In the filter below, the positive input is connected to a 6 volt reference level made from a voltage divider connected between 12 volts and ground.



Single Stage 700 Hz Bandpass Filter


Homebrewing with integrated circuits

Homemade "gouge boards" described in chapter 6 are usually much too crude to use with integrated circuits. Therefore when I build with ICs, I use "perf-boards." These are plain, 1/16 inch thick fiberglass boards. They have no copper layers but have a grid of tiny holes drilled on 1/10 inch centers. The older ICs have their pins spaced 1/10 inch apart and the two rows of pins

are spaced apart in even multiples of one-tenth inch. Modern ICs are often available either as old-style, through-hole or modern surface-mount. To use through-hole, you just push them through the perf-board and then solder your components and wiring on the pins on the other side.

I prefer to *use IC sockets*. If you accidentally damage an IC and it is installed with an IC socket, it can be instantly replaced. Without a socket, you'll seriously damage your wiring when you tear it apart to solder in a new one. The IC socket also has pins lined up just like an IC, so you push it through the board and solder your wires to the socket pins on the bottom.

For homebrew work you need the big, old time ICs with 1/10 inch spacing. Modern surface mount ICs have pin spacing as tiny as 20 thousands of an inch. They are designed to be installed with mass production robots and are extremely tricky to solder in your basement. One practical (but extremely time-consuming) home-brew approach is to make little etched circuit boards that accept the surface mount ICs. The little adapter boards have traces that serve to expand the pin lead spacing radially out to large pads that human fingers can work with directly. Using hemostat clamps, tweezers, and a microscope, you solder the IC onto the miniature board. After you have inspected your pin connections under a microscope, you can then solder these adapter boards onto a normal size perf-board.



A four-stage 700 Hz bandpass filter built on a perf-board.

Using the diagram for the single stage filter shown earlier, simply feed one output into the next input until you have four stages in series. All four op-amps can share the same 6 volt resting level reference.

Label your pins!


Keeping track of the pin numbers isn't easy unless you label them. I like to cut out a narrow strip of sticky white label paper and paste it between the pins. I write the pin numbers on the paper with a fine felt-tip pen adjacent to the respective pins. ***Pin # 1 is at the end of the integrated circuit that has the notch.*** Sometimes pin # 1 is indicated with a little number or a round indentation in the black plastic. The notch and indentation may be seen in the above photo. In any case, the pin numbers seen from the top are numbered counterclockwise. However, as seen from the bottom of the IC socket where you do your soldering, ***the pin numbers go around clockwise starting from the end with the notch in the plastic case.***

There are two kinds of IC sockets suitable for ham work - ***solder tail*** and ***wire wrap***. Wire wrap sockets have big tall pins that allow thin wires to be wrapped around them with a special tool. I prefer the solder-tail type. They are harder to work with, but aren't so tall, are less messy and don't take up so much room. Just to cover the subject of IC sockets, large complex

computer ICs use "grid sockets" and "solderball sockets" with hundreds of contact points. Unfortunately, working with these is nearly impossible at home because that would involve extreme precision and building multi-layer PC boards.

Wiring the 700 Hz filter into the circuit

I designed the op-amp filter to have a slight amount of amplification. Unfortunately, by the time I put four in series, I had as much gain as one of the audio amplifier stages. Therefore, the output of my filter goes to the input of the second AF amplifier stage and bypasses the first stage. When I returned the signal to the input of the 1st stage, it was much too loud and unstable. You could bypass the 2nd stage, but you would lose your AGC. A block diagram of my receiver is shown below. A double-pole, double-throw DPDT switch substitutes the filter for the 1st stage AF amplifier.



The digital solution to audio filtering

Actually, op-amp filters are also old-fashioned. The modern way to build audio filters is to use digital filtering. Basically, a microcomputer program measures the width of individual sinewaves with a counter-like algorithm. For example, a 700 Hz sinewave cycle is 1.43 milliseconds long. The digital filter measures the height and widths of the sinewave humps. If the widths are the desired frequency, then it just recreates waveforms of that length and amplitude. Sinewaves that are longer or shorter are ignored. Averaging and integrating math algorithms are used to massage the data and arrive at the most accurate and useful representation of the original spectral content of the signal. In practice, a digitally filtered passband of just one Hz is easy to build but is very hard to tune in with your VFO tuning knob. Therefore the passband of a digital filter is usually adjustable. When the passband is fairly wide, say 200 Hz, tuning isn't so difficult.

In conclusion,

The block diagram above shows the receiver I built. I packaged the AM broadcast band filter outside the receiver so I could use it with other receivers. A switch allows me to put in the 700 Hz filter when I need it. If you use it with single sideband phone, it will get rid of much of the noise, but voice quality will be poor. To tune in single sideband, you have to turn off the filter and tune your VFO to exactly the right spot to make the modulation intelligible. AM voice stations must also be tuned in perfectly, otherwise they make an obnoxious whistle. The whistle occurs because the AM broadcast includes a fixed carrier wave that is comparable to a Morse

21. Chapter 7, Harris

code signal with the key held down. To get rid of the whistle, you must "zero beat" the VFO so that there is no audio frequency difference between the carrier wave and your local VFO. When you do that, the whistle vanishes and you will just hear the audio sidebands.

So? How well does the direct conversion receiver work? If this were the year 1935, all your ham friends would probably be insanely jealous of its terrific performance. Unfortunately, by modern standards it is a bit of a toy. It's adequate for shortwave listening and an experienced CW operator can use it to talk to people. I managed to use it in conjunction with the QRP transmitter in chapter 6 to talk to hams in other states. I guess that makes me an experienced operator. Novice hams could easily use it for communications whenever the band isn't crowded. It has plenty of sensitivity, just not much selectivity.