

CRYSTAL SETS TO SIDEBAND

© Frank W. Harris 2021, REV 15

Chapter 7A

CODE PRACTICE RECEIVERS

Building a Direct Conversion Receiver

This three part chapter describes how to build simple ham receivers for one or more CW bands. They can be used for code practice and will also receive ham single sideband (SSB) voice stations. The first example is a Direct Conversion Receiver (DCR). 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 direct conversion 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. Unfortunately, a sharp audio filter and even an advanced image-canceling version of the DCR, will not make it competitive with a superheterodyne equipped with sharp crystal filters in the IF stage. 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 superheterodyne receiver as described in chapters 13A and 13B. In my opinion, it is not possible to get the same performance out of a direct conversion receiver. However, I don't know everything and I hope I'm wrong. If you're interested, build an image-canceling DCR and experiment!

Other time-honored ways to build a simple ham receiver are the regenerative designs discussed in Chapter 7B. Modern ARRL handbooks do not recommend building regenerative receivers. They are concerned that the oscillating detector-amplifier will broadcast noise and interference audible to other hams living in your neighborhood. If you are lucky, a regenerative receiver can be equal to the DCR described here in Chapter 7A. Building them is like the lottery. You can't be sure it will work well enough to bother. It took three tries, but eventually I built a good one that covered 40 and 30 meters. My main complaint about regenerative receivers is that they often assaulted my ears with explosions of noise.

In Chapter 7C I experimented with building a "simple" superheterodyne. Its main virtue was that it could cover the entire HF spectrum, not just one or two ham bands. By the time it was superior to DCR in more than one way, it was no longer "simple." No design is perfect. Explore!

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.

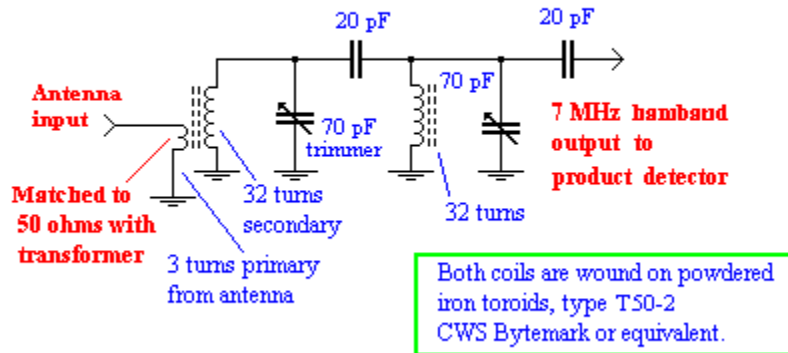
The difference frequency is filtered, amplified and passed on to the earphones - pretty simple radio! Considering how few parts it has, it's amazing how sensitive it is. I measured the performance of mine and found it could easily detect a 0.5 microvolt signal on 40 meters.

The input filter

The product detector (mixer) 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 should 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 limited. However, a simple filter can get rid of background low frequency 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 a whip antenna, the input lead should simply connect to the top of the trimmer capacitor on the left side through a small blocking capacitor.

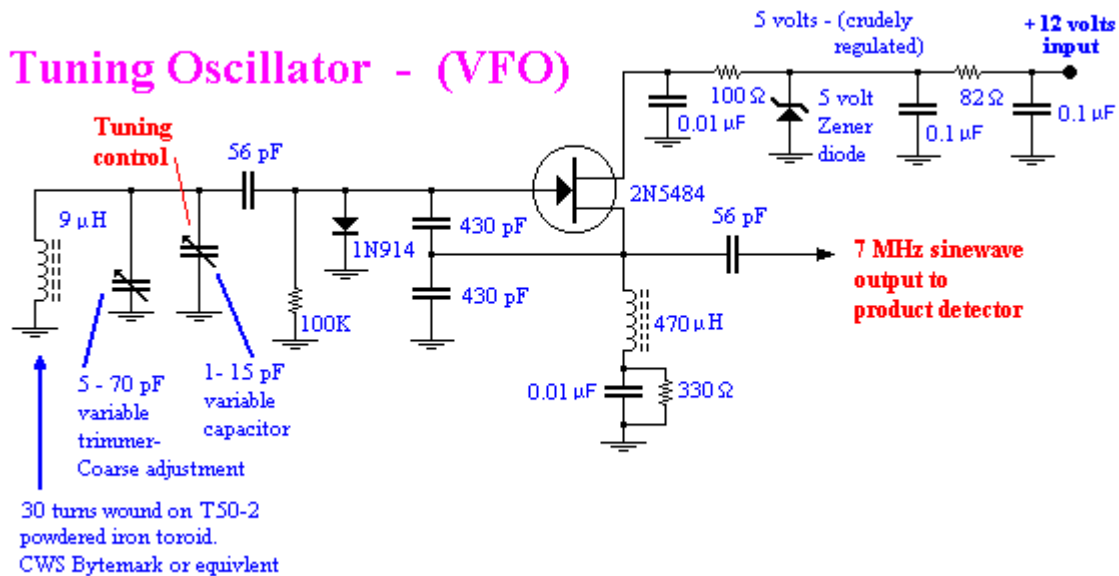
While listening to CW signals on 40 meters, tune the trimmer capacitors until the signals are as loud as possible and any background voice stations disappear. If this filter is insufficient to get rid of local AM broadcast signals, a 1.8 MHz high pass filter is described later in this chapter. Also, you might try reducing the 20 pF coupling capacitors between the LC stages to 5 pF or less.

This input filter section could also include an RF amplifier stage. However, this would probably not improve reception on 40 meters because atmospheric noise is relatively high below about 30 meters. On the other hand, if you would rather build this receiver for 30 meters or above, an RF amplifier will be a great help to sensitivity. RF input amplifiers for 30 meters and higher bands are discussed in Chapter 13B. The relatively crude VFO design described here works well on 40 meters and it will also work well on 80 or 160 meters. At frequencies higher than 40, it will be increasingly unstable, meaning that the VFO drift will be more obvious. On the other hand, I haven't built a higher frequency version, so I don't really know that higher bands are impractical.

An input RF amplifier stage will also help prevent your local oscillator (VFO) from broadcasting out into the world from your receiving antenna. The direct coupled VFO generates a pure sine wave. Consequently, if your oscillator is audible to a nearby ham, it will only be one discrete, weak "whistle" on the band. In contrast, the noise generated by a regenerative receiver is a wide band of noise and signals. The oscillating amplifier-detector of a regen is re-broadcasting whatever signals are coming in from the antenna.

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 6A. 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 several KHz of bandwidth at once, a drift of 50 Hz per minute will be scarcely noticeable. Building first rate VFOs for receivers and transmitters is described in chapters 10 and 11.



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 VFOs made from 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 6B 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 homebuilt crystal set, we were dealing with such tiny currents (microamperes) that the transition at zero volts was plainly observable. 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, current will flow but will be far less than a milliampere. Since 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, the transistor is already operating as a Class A amplifier and, unlike a bipolar transistor, no bias resistor from the drain supply is needed. JFETs are examples of *depletion-type FETs*.

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 - the 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.



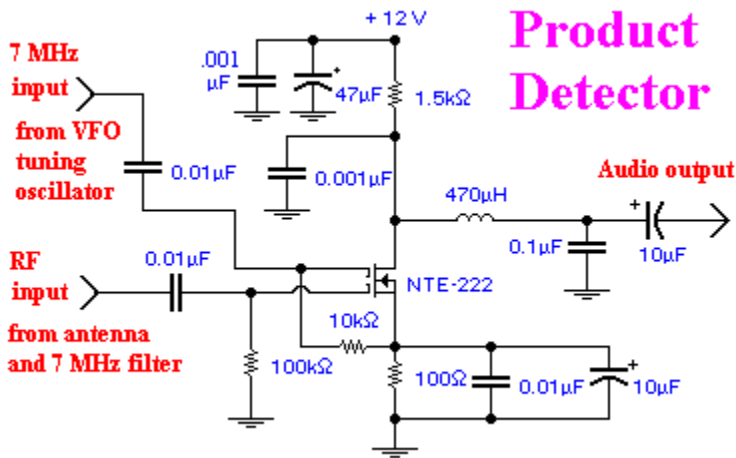
COMMON JFET TRANSISTOR CASES



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 to the earphones. A product detector is essentially the same as a "*mixer*." The only difference is that the product detector output is an audio signal whereas the output of a mixer is a constant frequency radio signal, an "Intermediate Frequency" (IF). In the superheterodyne design much of the signal amplification and frequency filtering occurs in the IF amplifiers, not in the audio stages. In the DCR design, most of the amplification occurs in audio amplifier stages. Unfortunately RF crystal filters can't easily be used in the DCR design.

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.



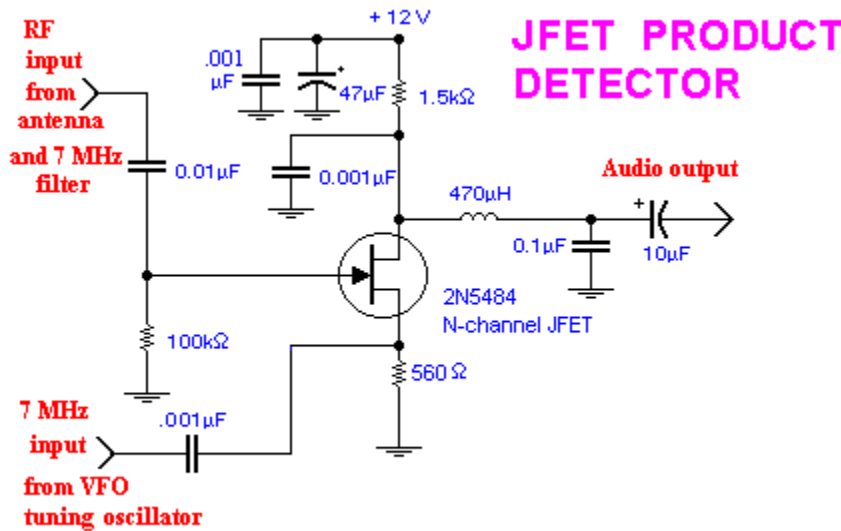
The gates of the MOSFET are tiny capacitors and therefore have extremely high impedance. In other words, DC current doesn't flow into the gates. The amplifier design is a *common source* design, which implies that the source is essentially at ground. This also means the load resistance, 1.5 K ohms, is 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 40673 by RCA should also work. It's available from RF Parts Co. and Mouser Electronics. I haven't personally tried it but I've seen it used in similar circuits. Using the new BF994S or similar (cheap) surface mount dual-gate MOSFET is discussed in Chapter 13A.

JFET product detector

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

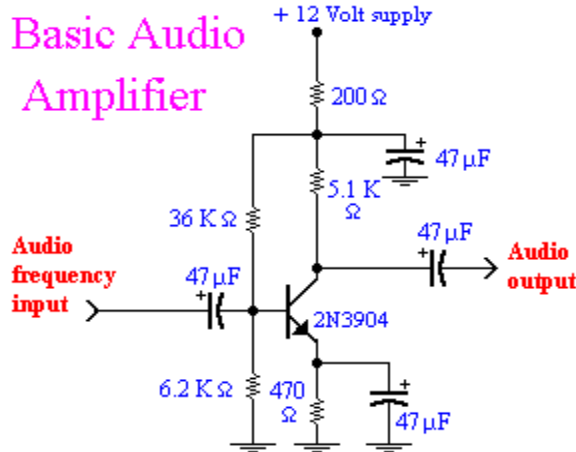
large VFO signal into a separate gate, it is introduced across the source resistor.



An article in the May 2008 issue of World Radio News describes a DCR similar to mine designed by Wayne McFee, NB6M. The interesting difference is that the JFET product detector was implemented with two JFETs in series. In that way both inputs, the VFO and the antenna as well, both go to gates and the output signals are added together, literally one on top of the other. The antenna amplifier JFET takes the place of the 560 ohm source resistor above. Since the McFee design has often been copied, I gather it works well, but I haven't tried it as a product detector. I did try to use it as a mixer in a superheterodyne, but its performance was not significantly different than a single JFET mixer.

A generic R-C coupled audio 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.

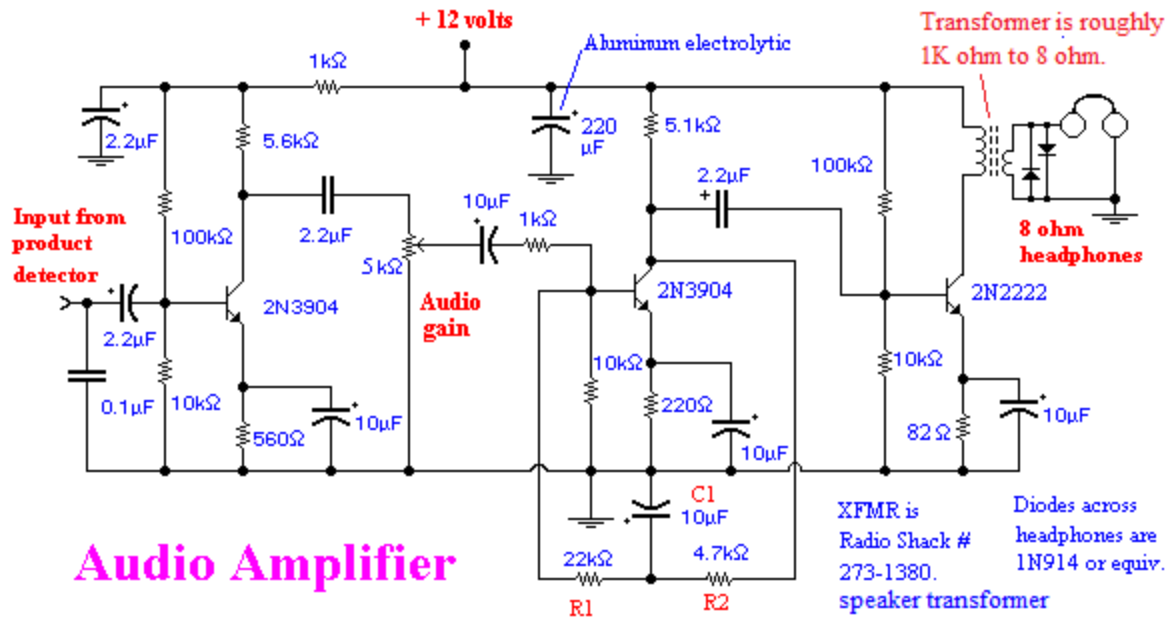


The diagram above shows a basic audio frequency *RC coupled* linear amplifier. If you look closely at the component values, you'll see it is similar to the class A RF amplifier we used as a buffer amplifier in the QRP transmitter in Chapter 6A. The differences are the large sizes of the capacitors and the absence of inductors. A 5.1 KΩ 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 KΩ resistor turns the amplifier partly on so that the audio sinewave will "fit" between zero volts and 12 volts. Otherwise the amplifier would only amplify positive signals larger than about 0.6 volts. Without the forward bias, the negative half of the sinewaves would be clipped off and the sound would be distorted. Because of the 0.6 volt base voltage offset, small signals wouldn't be heard at all.

The 470 ohm resistor stabilizes the transistor operating point, giving it a small amount of DC negative feedback. When the transistor conducts DC voltage, the voltage across this resistor in effect decreases the voltage on the base, thereby turning the transistor more off. This reduces the amplifier gain slightly, but makes the amplifier less likely to oscillate or run hot. The adjacent 47 μF capacitor allows the AC audio signal to go to ground without passing through the resistor. This allows us to have it both ways: The amplifier is more stable and the audio amplification is essentially unchanged.

The DCR audio amplifier

My complete audio amplifier is an extensively modified version of an example in the ARRL 1986 handbook. It is the same circuit used in the receiver in Chapter 13A. It looks like three straight-forward R-C coupled amplifiers in series. But the original design had extra filter components I didn't quite understand. Every part that I didn't understand, I left out. That's how I learn how circuits work. 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 I turned it on, the amplifier seemed completely dead. I put those mysterious components back in the circuit and ... 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 kind of audio Automatic Gain Control (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 small DC current is injected into the base. This increases the base voltage above 0.6 volts so that it's always turned on. The low frequency feedback adjusts the bias for weak and strong signals. When the signals are weak, the middle transistor is turned off, so it's collector voltage is high (+12 volts) and unchanging. This big collector voltage is leaked into C1 through R2. The voltage on C1 then biases the base through R1 to provide a forward bias for the transistor, turning it on and raising 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."

Headphone "impedance matching"

First, why does the total AC resistance of a load, such as a headphone, have to be "matched" to the impedance of the source driving it? It turns out that, in order to transfer power from a voltage source - a battery to a resistor, a transmitter to an antenna, or audio amplifier to a loudspeaker - *the internal impedance of the voltage source must match the impedance of the load.*

Draw a closed loop schematic of a battery. All voltage sources have a limit in how much current they can provide. This limit can be thought of as "the effective internal resistance." Big batteries have a lower internal resistance than big ones. If you have a little battery with a high

11. Chapter 7A, Harris

resistance, you need a load with the same internal resistance to suck the maximum power out of the battery. If the load resistance is too low, big currents will flow, but most of the power will be burned up inside the battery's internal resistance.

The rated impedance of headphones and loudspeakers is the internal resistance

As you know, headphones convert varying low frequency current into sound. They consist of relatively high inductance inductors wound around an iron core. The magnetized core is positioned next to a thin steel diaphragm that is attracted to the varying magnetic field generated by the headphone. I went through my collection of headphones, ear buds and loudspeakers and measured their resistance with a ohm-meter. I discovered that "8 ohm" labeled headphones have 8 ohms of plain old resistance. "4 ohm" labeled loudspeakers have 4 ohms resistance, ancient 1,000 ohm headphones have 1,000 ohms resistance and so on. When tested with a capacitance meter, they all have zero capacitance and look like short circuits. In other words, *they are all inductors*.

I had two misconceptions. First, I thought a loudspeaker labeled "impedance" number was the resultant AC resistance caused by the inductance plus any resistance, if any. Nope. It's just the everyday resistance. Second, I assumed that modern tiny headphones and ear buds were exotic piezo-electric devices - capacitors. No, again. At least all of the samples I tested were inductors.

Modern earphone impedance

I have some modern super light-weight headphones from Radio Shack and a couple pairs that came with a computer language course. They had no labeled "impedance." They measured 25 to 30 ohms resistance. And that explains why some headphones are *MUCH* louder than others in my radios. An audio amplifier with an 8 ohm or 4 ohm matching output transformer produces little sound in a 30 ohm headphone. I plugged in a tiny 4 ohm earbud and had loads of volume.

Protecting your ears from strong signals

This triple amplifier has enough gain to almost deafen you 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. In practice, with sensitive, modern 8 ohm headphones, I found that much less than one volt peak is plenty of volume for me. Eventually I put in two Schottky diodes across the headphones which "short" the output with opposite polarities. This limits the positive and negative sound peaks to just 0.2 volts and my ears are 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 bypass on the input back in and, as I expected, the audio sounded more "bass" and became somewhat weaker. However, getting rid of that piercing, hissing static was well worth the loss of gain. Also, once I began listening to voice stations, I became more aware of having too much treble to be pleasant. Eventually, I added a 0.3 μF between the base of the second 2N3904 and ground. Experiment!

12. Chapter 7A, Harris


The original design also had no emitter bypass capacitor. That's the 10 μF 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.

An auxillary stand-alone audio amplifier



Richard Fisher, KI6SN, found a related use for this same amplifier circuit. He noticed that many small portable QRP transceivers used by backpackers, don't have enough audio volume when used out in the woods with the wind blowing, the brook babbling and so on. He built his prototype using the punched-disk (a.k.a. Manhattan style) homebrew printed circuit method. He says it has more gain than he wanted, but it is easier to reduce gain than add it.

Mechanical construction



I built my receiver in a box made from soldering pieces of two-sided PC board together. This is the same construction (and size) described for the QRP transmitter described in Chapter 6A. 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 shortwave broadcast stations. I used a 70 pF trimmer capacitor to tune the entire frequency spectrum around 40 meters to align the small capacitor to tune the ham band. A low capacitance variable capacitor like this is called an *electrical bandspread*. In my receiver the 15 pF fine-tuning capacitor has the big black knob, while the physically tiny 70 pF ceramic screw adjust trimmer is directly behind it. It would be better to have a calibrated vernier dial of some sort to adjust the fine tuning capacitor.

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 extract the most energy from alkaline batteries, you must 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 transmitter VFO

I tried to use the DCR with my QRP transmitter which was placed some distance away on my desk. I found I couldn't hear my own transmitter VFO (crystal oscillator) "Spot signal" 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 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.

In 2015 I decided to dust off the DCR and combine it with the QRP to make a complete station. When I placed the QRP directly on top of the DCR, the "Spot signal" was much too loud, even without the antenna-like connection. I also had difficulties with feedback (squeals) when I ran it on 12 volts. Using a simple LM317L voltage regulator circuit, I made the 12 volt supply adjustable. (See Chapter 8) Also, when I switched the 700 Hz filter in and out, the audio volume was not the same. I put a 2K pot on the output of the filter so the volume could be equal.

FIXING THE LIMITATIONS OF DIRECT CONVERSION

One of the DCR limitations I had read about did not 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 when using vacuum tube amplifiers in which the grids and plates in the tubes are prone to mechanical vibration.

A DCR detects both sidebands at once

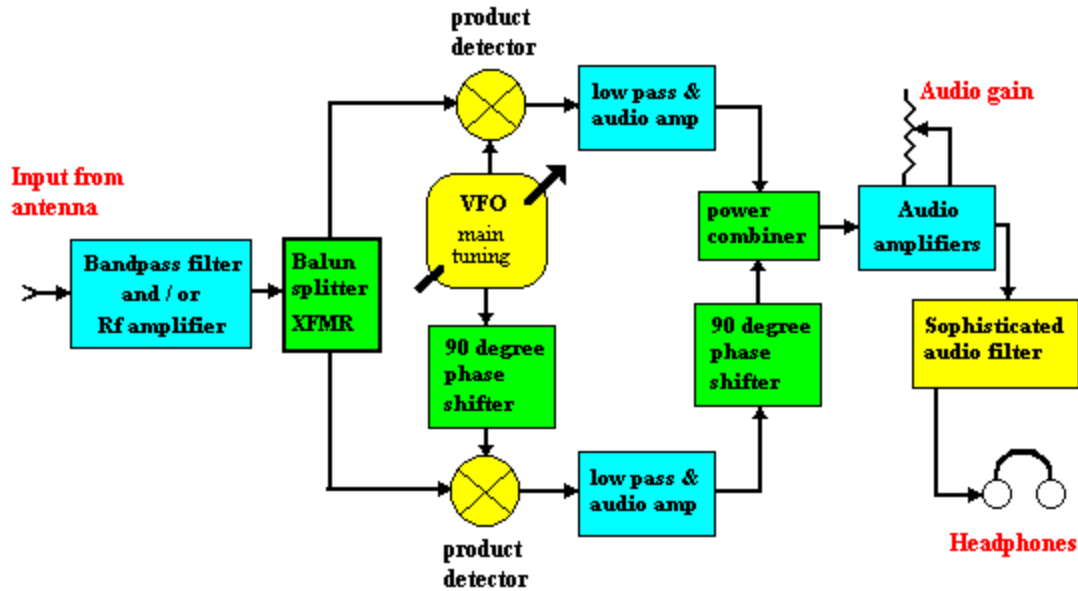
Unfortunately, the product detector mixer detects both the upper and lower sidebands simultaneously. This is good for listening to AM radio stations that broadcast both sidebands, but it has poor selectivity on a crowded CW band. A good CW receiver can select a bandwidth as small as 500 Hz or less. A 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 usually set between 500 to 1000 Hz. By convention, ***On 40 meters and below, they are listening BELOW the transmitting frequency. On HF bands 30 meters and above hams listen ABOVE the transmitting frequency.*** This convention has its origin in the method for generating single sideband phone that was formerly common in the 1950s. The old SSB design isn't used anymore, but the convention lives on. The relatively 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 while the station is listening on the lower sideband, he won't hear you because modern equipment only receives 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 his commercial transceiver to that band, he selected upper or lower sideband and didn't have to worry 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.

Building a DC receiver for multiple HF bands and narrowing the bandwidth

Once I used my DC receiver on the air, I realized that even canceling the unwanted sideband would not fix the reality that the signal bandwidth that reaches the detector is too wide and always will be. Even with the audio filter, any signal that reaches the detector(s) that contains 700 Hz will be detected and will reach your ears. If the band is crowded, I will always hear several stations at once. The DC receiver bandwidth will be somewhat narrower if the antenna signal first passes through the antenna tuner you will probably be using with your transmitter. Antenna tuners are described in Chapter 9.

The only real solution is to sharply limit the bandwidth *before* the signals reach the product detector. This can be done by using a superheterodyne receiver design (Chapters 13A and 13B). OR you could use a string of tuned bandpass filters ahead of the product detector. In other words, using the TRF design. TRF receivers use a string of simultaneously tuned, matched RF amplifiers ahead of the detector. This is quite difficult to do, unless the receiver is built for some special purpose that only needs to receive a single frequency.

I saw an elegant TRF AM radio from the late 1920s for sale at a ham fest. It had an ornate, polished wooden cabinet and was twice the size of a typical kitchen microwave. Inside were huge tubes and complex gears and shafts to tune and adjust each stage synchronously. Following the last stage of the RF amplifier chain was a simple rectifier AM detector and the audio amplifiers. In this way each sharply tuned amplifier could track with the others together

across the entire AM band.

Personally, I decided that a dual conversion superheterodyne with crystal filters is a far better solution to the bandwidth problem. Also, a superheterodyne receiver can be used on multiple bands. A superheterodyne doesn't require simultaneous active tuning of multiple RF amplifiers. And best of all, you can use crystal filters to narrow the bandwidth. An all band HF superheterodyne is described in Chapters 7C, 13A and 13B.

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.


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

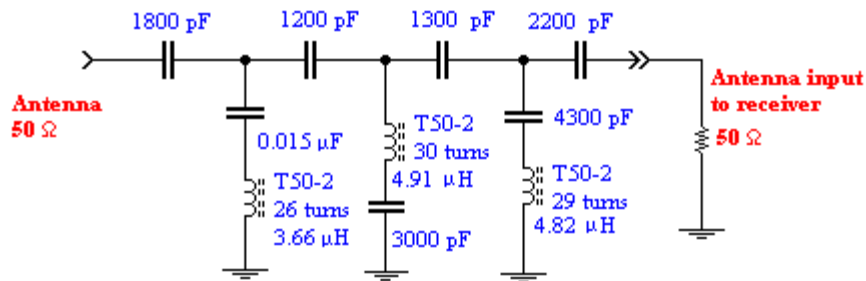
$$\text{RMS voltage} = \text{Voltage peak} \times (0.707)$$

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

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. When I used the DCR in conjunction with a transmitter and shared the transmitter "transmatch" antenna tuner, the problem mostly disappeared. However, if you are *really* close to an AM broadcast station, the transmatch may not remove all of it.

If you're not using a transmatch, the good old 1986 ARRL handbook has an AM broadcast rejection 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

If you continue to experiment with homebuilt receivers, this DCR project may not be the last time you need to suppress your local AM radio interference. Building this filter as a separate little unit enclosed with a metal lid allows it to be used in future radios. In Chapter 13A I describe a complex superheterodyne in which each ham band has its own preselector filter. You'd think that a separate AM suppressor would not be needed. Well, maybe not in your neighborhood, but in mine the darn rap music gets into everything. The superheterodyne described in Chapter 13 is assembled from individually packaged modules - VFO, power supplies, band converters, mixer, etc. Because the modules are coupled together with the usual RCA connectors and coax, I was immediately able to suppress the rap music by inserting the above module between the 80 meter preselector and the mixer modules. The small size allowed it to be slipped inconspicuously inside the chassis without dangling outside the enclosure. Placing the filter in series with the antenna or other locations was not nearly as effective.

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. Unless of course, his receiver is as primitive as yours.

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 fellow's 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 30 meters or above, you would reverse these instructions, because upper sideband is the convention up there.

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 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. In the real world audio crystal filters don't exist. However, if you don't mind inductors on the order of 80 millihenries, you can combine them with electrolytic capacitors to build LC audio filters. To my surprise the 2005 ARRL Handbook has plans for a 5 stage LC audio filter. It is kinda big, but not impractical.

The designs illustrated below use RC networks and operational amplifiers to accomplish the same goal. *The simplest RC networks don't resonate, 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" to ground. Remember, *the voltage across a capacitor can't change instantly*. In contrast, the effect of smaller capacitors will be negligible to low frequency sinewaves that have 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, to high frequencies the capacitor looks like a wire with zero AC voltage drop across it. At high frequencies there isn't time to charge. At low frequencies, the capacitor charges up completely and the AC voltage drop across the capacitor approaches the entire voltage of the input sine wave. 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 sine wave 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 level.

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. It accentuates 700 Hz tones and attenuates tones above roughly 800 Hz and below roughly 600 Hz.



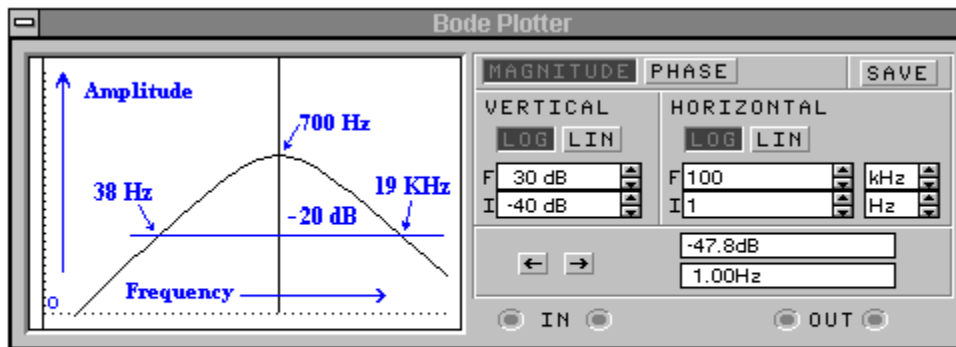
Single Stage 700 Hz Bandpass Filter

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 negative 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.

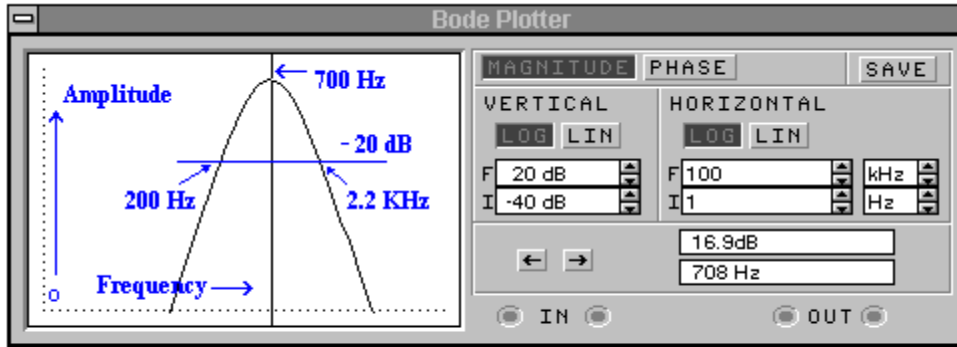
The two resistors, the 2.4 K Ω and 6.2 K Ω , also control the gain in this circuit. The triangle amplifier symbol is assumed to have infinite gain. It doesn't of course, but this symbol generally means an operational amplifier, which is supposed to act like a perfect transistor with unlimited gain. The gain is set by the external resistors. There's no need to go into the details, but, the ratio of the feedback resistor to the input resistor is 6.2 K / 2.4 K = 2.5 times. That means, the gain is approximately 2.5. To say it another way, the voltage lost across the 2.4 K resistor while entering the amplifier is compensated by the much larger voltage gain across the 6.2 K. This more than makes up for the signal voltage lost in the filtering process.



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).

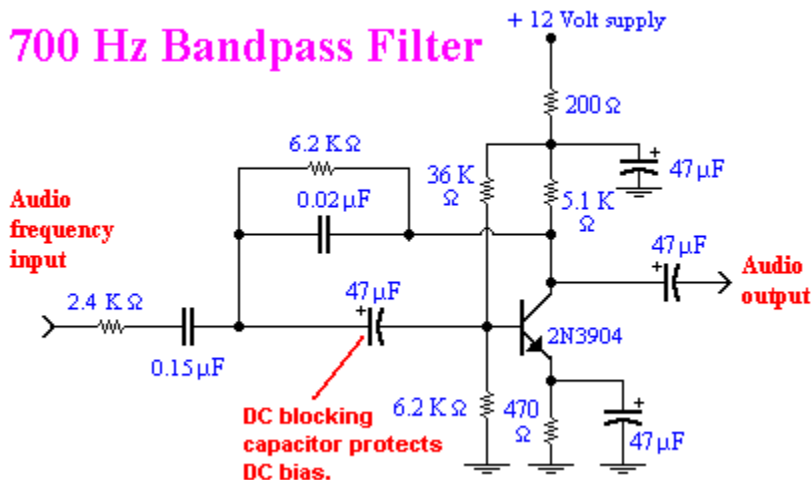
More precise filtering can be achieved by putting multiple audio filters in series. The plot below shows four of the same filters in series. This is a big improvement and makes it easier to listen to just one station at a time. It also eliminates the high frequency noise and makes listening less tiring.



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 the single stage 700 Hz filter would look like:



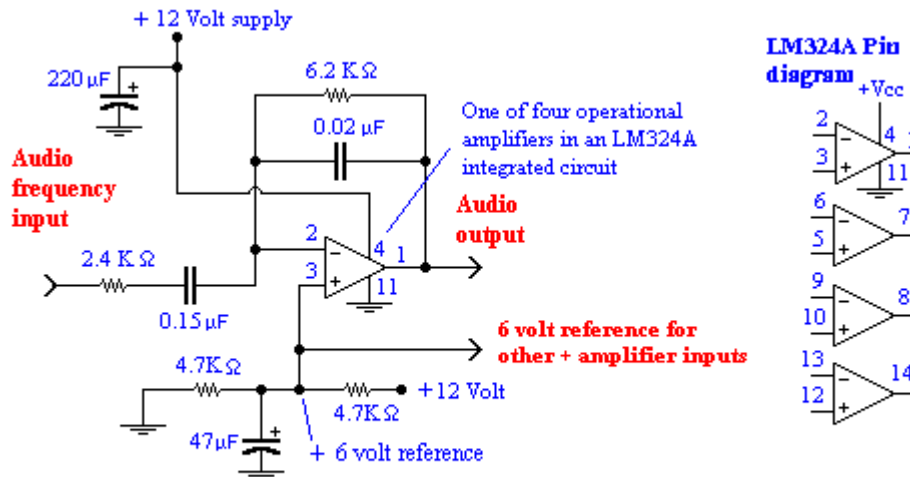
Operational amplifiers – "perfect" amplifiers

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 thought, 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, refer to 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 causes 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.

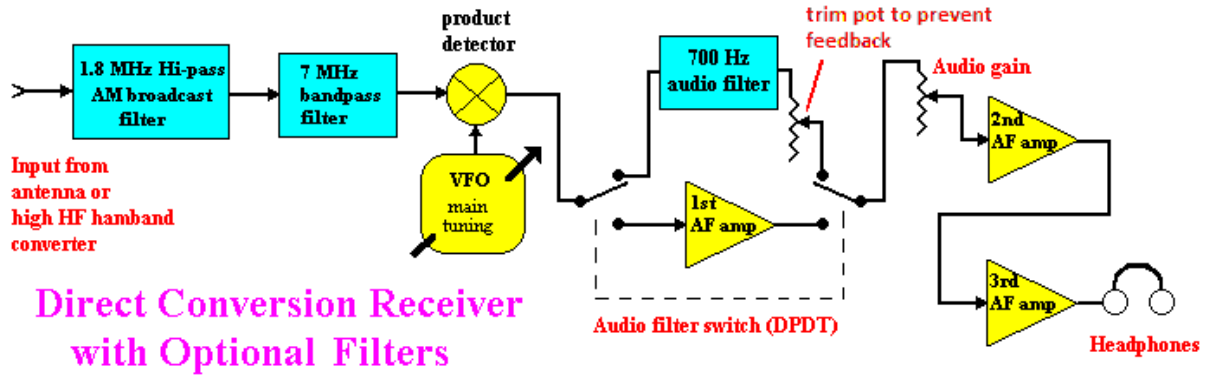
To bias an operational amplifier so that a large sinewave will "fit" between zero volts and 12 volts output, the amplifier must have a baseline about halfway between the supply voltage and zero. In this case the resting level will be about 6 volts. In the filter below, the positive input is connected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



Single Stage 700 Hz Bandpass Filter

Wiring the 700 Hz filter into the DCR circuit

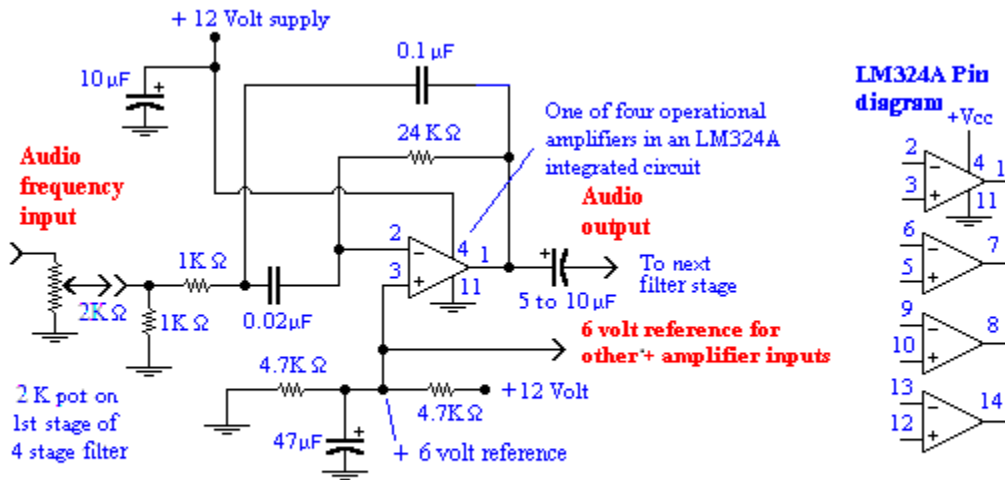
The op-amp filters have a significant amount of amplification. Unfortunately, by the time I put four in series, I had as much gain as one of the unfiltered audio amplifier stages. Therefore, the output of my filter goes to the input of the second AF amplifier stage and bypasses the first stage. When 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. I added a 2KΩ pot on the output of the filter to balance the volume with the unfiltered volume level. With higher power supply voltages, the filter may oscillate with an ear splitting squeal. This pot can trim the output to prevent this feedback.



In 2015 I revisited this old project and combined it with the crystal controlled QRP in Chapter 6A to make a (relatively) small portable transceiver. I used an antenna tuner and some D-cell batteries as shown in Chapter 9 to complete the little station. When I got on the air with this rig, I remembered why I prefer the complex receivers shown in chapters 13A and 13B. I was talking to a fellow in a nearby town but he was difficult to copy because I could hear two other CW stations simultaneously. Without the 700 Hz filter, I could hear 5 or 6 stations, plus a lot of noise. In my big receiver the same guy was all alone in a void of soft, distant, purring static. Improved selectivity made the 40 meter CW band seem 10 times bigger. Because of the lack of RF selectivity, the *non-resonant* 700 Hz filter described above is not as helpful as I expected.

A resonant audio filter

This resonant 700 Hz filter is a definite improvement. I replaced the old filter with this improved design which I first used in the new superheterodyne receiver described in Chapter 13B. The basic diagram of the improved filter is below:



This filter can be thought of as an audio frequency, series LC circuit. As you know, when signals pass through a series LC circuit, only one frequency is passed efficiently and all others are attenuated. The 0.02 µF capacitor serves as "the Capacitor" in this circuit. An inductor would

ordinarily be a large chunk of iron and wire. However, the $0.1\mu\text{F}$ feedback capacitor together with the operational amplifier work together to simulate "the Inductor."

The operational amplifier constantly adjusts its output voltage in order to hold the two input pins at the same voltage. Whenever the output changes, the feedback capacitor must charge before the voltage across it can change. As you know, the *voltage* across a capacitor can not change instantly. Similarly, the *current* through an inductor can never change instantly. Because the other end of the capacitor is connected to the input, the output can't keep up with the changing input voltage, so it must charge the $0.1\mu\text{F}$ capacitor continuously. This continuously charging current is analogous to an inductor charging. Because all the resistances and capacitors are fixed in size, there is a certain optimum frequency at which the sinewaves present at the input will appear at the output. All others are transferred inefficiently. If you change the value of *ANY* of these components, the resonant frequency will be different than 700 Hz.

In my new superheterodyne receiver, I built a double 4 stage resonant filter - a total of 8 stages. Well, it sure does limit the audio to 700 Hz! Even the static whistles at 700 Hz. Occasionally the QRM (interference from other stations) is so severe that 8 stages is able to select the desired station. But for most purposes, 8 stages is overkill. I predict you'll find that a 4 stage filter will be enough for your DCR. Rather than building another 4 stage filter, I believe the image canceling circuit for DCRs outlined earlier would be much more useful.

Homebrewing with integrated circuits

Homemade "gouge boards" described in Chapter 6A are usually too crude to use with integrated circuits. However, if you're careful, the gouge technique can carve fine enough traces to accommodate older ICs with $1/10$ inch spaced pins. In Chapter 13B the resonant filter just described was carved onto a circuit board.

When I built the DCR receiver filter I used a "perf-board." This is a plain, $1/16$ inch thick fiberglass board (available at Radio Shack). They have no copper layers but have a grid of tiny holes drilled on $1/10$ inch centers. The older ICs have two rows of pins spaced apart in even multiples of one-tenth inch. Modern ICs are sometimes available either as old-style, through-hole or modern surface-mount. To use through-hole, you could just push them through the perf-board and then solder your components and wiring to the pins on the other side.


I prefer to *use IC sockets* rather than hard-solder them on. 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 has pins lined up just like an IC, so you wire it the same way. Push it through the board and solder your wires to the socket pins on the bottom.



Push-in "micro-klip" pins are available for perf-boards which make it easier to connect the board to external power and RF connections. The pins are specifically designed for the holes in various sized perf-boards and grip the holes tightly. Push-in pins are inserted with a special tool as shown above.

Modern surface mount ICs have pin spacing of 50 thousands of an inch maximum. These days they are usually spaced even more closely than that. They are designed to be installed with mass production "pick and place" robots and are extremely tricky to solder in your basement. One practical (but 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 the perf-board.

Little adapter boards for surface mount ICs can sometimes be found for sale. Unfortunately, there are so many varieties of surface mount ICs that it is hard to buy the kind you want without buying big expensive sheets of various adapter boards, most of which are for ICs you'll never use.



A four-stage 700 Hz bandpass filter built on a perf-board. Push-pin connections can be seen on the periphery of the board. These are used to connect to external wiring.

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. The resonant version of the filter needs an input 10K Ω trim pot to adjust the gain of the 4-stage filter so that, when you switch it in and out, it doesn't oscillate.

Label your IC pins!

Keeping track of the IC 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 at the right hand end of the chip 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 "wire wrap" 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. I do use fine "wire wrap," Kynar insulated hook up wire for this purpose. It is easier to solder the fine wire to the tiny pins than thick, stiff, conventional, e.g., 26 gauge, hook up wire. 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 they require extreme precision and building multi-layer PC boards.

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 would be very hard to tune 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.

Measuring frequency with a counter

One advantage of a direct coupled receiver is that the VFO frequency and the frequency you're listening too are virtually the same. Consequently, if you wish to use a digital frequency counter to monitor the frequency of your receiver, it is easy to do. You just connect the counter to the oscillating JFET and read it directly. The same applies to regenerative receivers, like the one discussed next. However, the counter will NOT work properly with SUPER-regenerative receivers in which the oscillation occurs in short bursts. A digital counter is also difficult to use with a superheterodyne receiver like the ones discussed in Chapter 13. A single conversion superheterodyne has no oscillators on the same frequency you are receiving. A double-conversion superheterodyne is even more difficult. A frequency sample would have to be generated from the

VFO, BFO and converter crystal oscillator frequencies. It will be complicated to engineer.

In conclusion,

The block diagrams shown earlier show 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. It is too sharp to use with single sideband phone. It cuts off most of frequency range of speech and voices become unintelligible. 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 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 this direct conversion receiver work? If this were the year 1935, all your ham friends would be insanely jealous of its terrific performance. Unfortunately, by modern standards it is a toy. It's adequate for shortwave listening and an experienced CW operator can use it to talk to people. With difficulty I managed to use it in conjunction with the QRP transmitter in Chapter 6 to talk to hams in other states. Its best use is probably for practicing copying CW signals while you are studying for your license. Happy soldering!