Chapter 22 - Receiver Fundamentals

A radio receiver is the heart of any amateur radio installation, whether it is a stand-alone receiver or combined with a transmitter as a transceiver. It is relatively easy to build a good transmitter – all you really need is good frequency stability, adequate power and a clean output signal (no harmonics, key clicks or inter-modulation distortion). It is much harder to build a good receiver, and consequently there is more variation in receiver capability amongst both commercial and homebuilt designs.

When conditions are good (i.e. radio signals are propagating long distances) the amateur bands can be a very crowded place. If you listen during any CW contest, for instance, you will hear signals spaced 200 to 300 Hz apart over the entire CW section of a band. So the first attribute a good receiver must have is **selectivity**, the ability to distinguish between close spaced signals and receive only the one that the listener is interested in. Many of the signals on amateur bands are very weak, having come from low-powered transmitters a long distance away, so the second attribute an amateur receiver needs is **sensitivity**, the ability to "hear" very weak signals. And since these weak signals may be adjacent to strong signals, perhaps from other amateurs in your town, amateur receivers need another attribute: **dynamic range**.

Dynamic range is the ability of the receiver to receive signals of widely different signal strengths. This is usually achieved by the receiver self-adjusting sensitivity for widely differing signal strengths. This is usually done through the AGC mechanism. Where the AGC has insufficient range, this can be supplemented by an input attenuator or by manual adjustment of the RF Gain.

To get an idea of the challenges faced by receiver designers, a typical weak signal on an amateur band might deliver a power of −120 dBm from the antenna – that's one billionth of a microWatt. A strong signal might deliver −30 dBm, or 1 μW. So a strong signal could be 90 dB (one billion times) as strong as a weak signal – and yet the receiver might need to select and amplify the weak signal to a usable level, without being affected by the strong signal a few kilohertz away!

This module introduces two simple receiver designs – the tuned radio frequency receiver and the direct-conversion receiver, and considers how well they meet these requirements. It also introduces many of the concepts that you will need for the next module, which covers the **super-heterodyne receiver**.

The Tuned Radio Frequency (TRF) Receiver

One of the simplest receiver designs, which has been with us almost since the dawn of radio, is the tuned radio frequency receiver. The principle is simple: you use a band-pass filter to select the signal you want, amplify the weak radio signal, demodulate the signal (to recover the audio modulating frequency) and then amplify the recovered audio sufficiently to make it audible in headphones or a loudspeaker. The block diagram below shows the layout of a TRF receiver. The block labelled "detector" is a half-wave rectifier to demodulate AM signals.

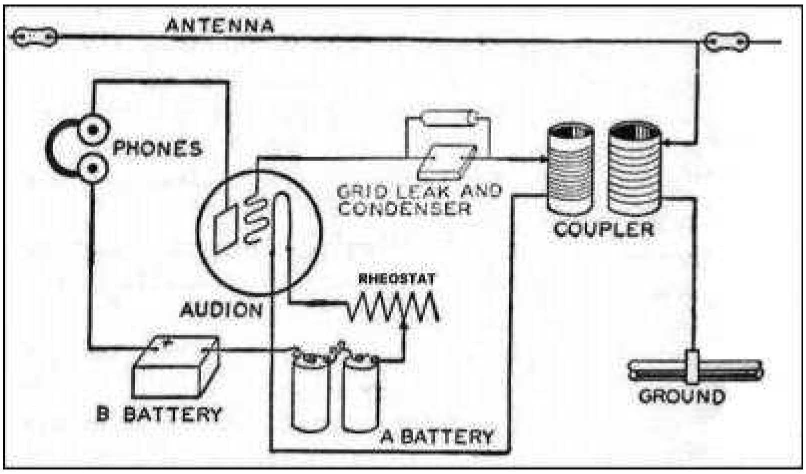

A Tuned Radio Frequency Receiver with Regeneration

The arrows through the bandpass filter indicate that they are tunable, so they can be used to select the desired signal. The dotted line joining the arrows on the two bandpass filters mean that they tune together, so a single control will change the tuning of both filters together.

Many TRF receivers use **regeneration**, which means feeding some of the signal from the output of the RF amplifier back to its input, in such a way as to reinforce the signal at the input of the RF amplifier. This is a form of **positive feedback**. It has the benefit of increasing the amplification of the RF amplifier (because some of the signal "circulates" through it many times, being amplified each time) and also increasing the selectivity, since the signal also passes through the band-pass filter at the output of the RF amplifier many times. Of course an amplifier with positive feedback is an oscillator, so if too much regeneration is applied then the circuit will oscillate. Regenerative receivers (a name for TRF receivers that use regeneration) usually have a control to adjust the amount of regeneration, which is adjusted to get the maximum possible sensitivity and selectivity without oscillation.

The advantage of TRF receivers is that they are simple to construct and require relatively few components – typically just two or three valves or transistors and a handful of other parts.

This made them attractive in the days before transistors, when thermionic valves were used for amplification in radio receivers, as valves were relatively expensive so the fewer the better!

Their big disadvantage is that they have very poor selectivity and dynamic range. Tunable bandpass filters just aren't capable of rejecting an unwanted signal that is only a couple of kilohertz away from the signal you are listening to, so unwanted signals will also get through to the detector and be recovered as audio or cause **inter-modulation distortion**. TRF receivers are also best suited for receiving AM signals. Although **regenerative receivers** can be used with CW and SSB signals, by adjusting the regeneration control so the circuit just oscillates, adjustment is tricky and the quality of reception poor. For these reasons TRF receivers are not widely used any more.

The Direct-Conversion Receiver

A design that is used in quite a few homebuilt receivers **[and nowadays commercial designs]** is the Direct Conversion receiver. In a Direct Conversion receiver, the radio-frequency signal from the antenna is mixed with a locally generated oscillator signal, producing the usual sum and difference mixing products.

The frequency of the oscillator that generates this local mixing signal – it is known as the local oscillator (LO) or beat frequency oscillator (BFO) – is set so the difference mixing product is at audio frequency. In this way the Direct Conversion receiver "directly converts" the desired radio-frequency signal to audio, where it can be filtered and amplified. Let's look at the circuit in a little more detail.

A Direct-Conversion Receiver

The signal from the antenna first passes through a bandpass filter. Unlike in the Tuned Radio Frequency receiver, this bandpass filter is not responsible for the overall selectivity of the receiver – its role is simply to reject interference from strong local commercial broadcast stations and the like. It does not have to be tunable – usually a fixed-tuned filter covering an entire amateur band will suffice.

The signal is then amplified by an RF amplifier and fed into the product detector, which We have represented on the diagram using the symbol for a mixer – the circle with a cross in it. ("Mixer", "Modulator" and "Product Detector" are different names for essentially the same circuit, depending on the exact role it plays.) The product detector mixes the amplified RF signal with a signal generated by the tunable local oscillator, generating the usual sum and difference mixing product.

Suppose we want to receive an upper-sideband signal on 14,200 MHz. By convention, we refer to the frequency of a single-sideband signal as the frequency where the carrier would have been if it had not been suppressed. So the upper sideband of this USB signal (i.e. all that is left of it after the carrier and lower sideband were removed) will range in frequency from 14,200 3 MHz to 14,203 0 MHz, 300 Hz to 3 kHz above the (suppressed) carrier. If the local oscillator is set to exactly 14,200

MHz – the frequency where the carrier would have been – then the difference mixing products will range in frequency between 300 Hz and 3 kHz. What we have done is to translate the USB signal from its frequency of 14,200 MHz back to the audio frequency range.

This graph shows how mixing the 14,200 MHz USB signal with a 14,200 MHz signal from the local oscillator generated a difference mixing product (signal frequency – local oscillator frequency) in the audio range and a sum product (signal frequency + local oscillator frequency) up above 28,400 MHz.

Although the example used an upper sideband signal, the same process would work equally well using a lower sideband signal, and the local oscillator frequency would still be 14,200 MHz, the frequency where the carrier would have been. The following graph shows the same process with a lower-sideband signal.

Once again the difference product is back at audio frequency, while the sum product is at around twice the signal frequency, 28,400 MHz. Also note how for the lower side-band signal, the mixing process has inverted the sideband (so the recovered audio is the mirror image of the sideband), which makes up for the sideband inversion that would have occurred when the LSB signal was generated.

So whether the signal is USB or LSB, mixing it with a local oscillator with the same frequency that the carrier would have had will demodulate it and recover the audio.

To complete the hat trick, suppose we have a CW signal at 14.200 MHz. All we need to do is set the local oscillator just below it – say at 14.1994 MHz, which is 600 Hz below the CW signal – and the difference mixing product will be a 600 Hz tone, just right for listening to CW. So we can also use the product detector to receive a CW signal. (Setting the local oscillator 600 Hz above the CW signal would work just as well.)

We now pass the recovered audio through a low-pass filter. The main purpose of the filter is to remove the difference mixing product from signals near to the one that we are listening to.

For example, suppose there is a CW signal at 14.205 MHz while we are listening to our 14.200 MHz USB signal. The difference mixing product of the 14.205 MHz CW signal and the 14.200 MHz local oscillator is 5 kHz – in other words, we have translated the (unwanted) CW signal downwards in frequency to the audio range just as we have translated the (wanted) USB signal to audio. However a low-pass filter with a cutoff frequency of around 3 kHz or so should be able to remove the unwanted CW signal without affecting the desired USB signal.

Because it is quite easy for a strong signal to overload a mixer, causing inter-modulation distortion, the gain ahead of the mixer (i.e. the gain of the RF amplifier) is usually kept quite low so as not to amplify unwanted strong signals and overload the mixer. This means that most of the gain in a Direct Conversion receiver is at audio frequencies, in the amplifiers following the low-pass filter.

The only remaining part of the circuit is the **Automatic Gain Control (AGC)** system. Because there is such a wide range of signal strengths on the amateur (and other) bands, it is useful to have some way of automatically controlling the gain of the receiver, so it can have a lot of gain to amplify weak signals, but reduce this gain to avoid overload when amplifying strong signals. While this could be achieved with a manually operated gain control, this is not very operator friendly because when tuning from a weak signal (with the gain set on full) to a strong signal, the strong signal can be painfully loud. And when tuning from a strong signal (with the gain turned right down) to a weak signal, you might miss the weak signal altogether unless you remembered to turn the gain up.

The solution is automatic gain control. The AGC detector samples the audio signal after the first audio amplifier, and automatically adjusts the gain of the RF amplifier and the audio amplifier to keep the output signal level fairly constant. The output signal is then amplified by a final audio power amplifier and used to drive headphones or a speaker. The AGC control voltage is often also used to drive a signal strength meter, known as an “S meter”, that indicates the strength of the received signal using a fairly arbitrary scale calibrated from S1 (a very weak signal) to S9 (a very strong signal).

The Direct Conversion receiver has several advantages over a TRF receiver. Most importantly, its selectivity is very good, because unwanted nearby signals are easily filtered out by the audio low-pass filter that follows the product detector. It is more stable, having no tendency to oscillate like regenerative TRF receivers do. And it is easy to receive single sideband and CW signals with a Direct Conversion receiver – you just tune the signal in, without having to fiddle with the regeneration control.

However the Direct Conversion receiver does have one significant disadvantage. Since the same local oscillator frequency can be used to tune either a upper sideband or a lower sideband signal, if you are listening to say an upper sideband signal and there is a different signal occupying the frequencies on the other side of the local oscillator where the lower sideband would have been, then the other signal will also be shifted to audio frequencies and will interfere with the station you are trying to listen to.

For example, suppose you are listening to an USB signal at 14,200 MHz as before, but there is also a CW signal at a frequency of 14,199 MHz. Mixing the 14,200 MHz local oscillator signal with the 14,199 MHz CW signal will generate a 1 kHz audio tone. Since this falls within the same 300 Hz – 3 kHz audio range as the desired USB signal, you cannot filter it out using the low-pass filter. And because the unwanted signal is so close in frequency to the desired signal, you can't use the RF bandpass filter to reject it either.

The unwanted signal on the other side of the local oscillator signal is called an "image", so the principal disadvantage of the Direct Conversion receiver can be described as its inability to reject images, or lack of "image rejection". There are more sophisticated variations of the basic Direct Conversion design that are able to reject images, but these are quite complex and fall outside the scope of this course.

Summary

The key attributes of a receiver are **sensitivity, selectivity and dynamic range**.

- **Sensitivity** is the ability to receive weak signals; **[Specific to BANDWIDTH of signal]**
- **Selectivity** is the ability to distinguish between nearby signals; **[Rejection of unwanted signals]**
- **Dynamic range** is the ability of the receiver to receive signals of widely different signal strengths. **[Rejection of strong signals at the same time as receiving a weak signal]**

In the tuned radio frequency receiver all signal filtering is done at radio frequencies. As a result they have poor selectivity. Regeneration, which consists of feeding some of the output signal back to the input of the RF amplifier, can increase both the sensitivity and selectivity of the TRF receiver, but makes it prone to oscillation. **[Interferes with nearby receivers – emc. Chapter 27]**

In the direct-conversion receiver, the incoming RF signal is mixed down to audio frequency using a product detector **[mixer]** and local oscillator. Most of the selectivity of a direct conversion receiver is contributed by audio filters following the product detector. Direct conversion receivers have much better selectivity than TRF receivers, but they suffer from an **image response** to the opposite sideband that can only be eliminated with complex designs. **[NOTE – In years gone by, complex designs meant expensive. Nowadays it can be easily and cheaply fabricated using I.C.'s and multiple transistors. Or in software digital signal processing. D.S.P.]**

http://en.wikipedia.org/wiki/Radio_receiver_design

Regenerative and Reflex Receivers - This file contributed by Kim Smith and The Radio Electronique.
http://pe2bz.philpem.me.uk/Comm/-%20Receivers/-%20Regenerative/Info-900-Regen-Misc/regenrx.htm

Chapter 23 - The Super-heterodyne Receiver

The Single-Conversion Superhet

The super-heterodyne receiver or “Superhet” as it is commonly known is the most widely used receiver design in amateur radio. It overcomes the lack of **image rejection** of the Direct Conversion receiver by converting the incoming RF signal to one or more **intermediate frequencies** **[I.F.]** before demodulating it. The block diagram of a typical single-conversion superhet (one with only a single intermediate frequency) is shown below.

A Single-Conversion Superhet Receiver

The RF signal from the antenna is first filtered by a band-pass filter. As in the Direct Conversion receiver this can be a fixed-tuned filter covering an entire amateur band, since the receiver does not rely on this filter (known as the pre-selector) for its selectivity. As we shall see, the main purpose of the pre-selector is to reject the image frequency. The signal is then amplified in an RF amplifier – once again, not too much amplification, to avoid overloading the mixer that follows (in some designs the RF amplifier may be omitted entirely).

In the first mixer, the RF signal is mixed with the signal from the tunable local oscillator. But instead of mixing it down to audio, this converts it to an intermediate frequency (IF).

Common intermediate frequencies for single-conversion superhets are 455 kHz, 9 MHz and 10,7 MHz.

Suppose for example we want to receive a signal on 14,200 MHz again, and the intermediate frequency is 9 MHz. Then we could use a local oscillator frequency of either 5,200 MHz (because the difference between 5,200 MHz and 14,200 MHz gives the IF frequency of 9 MHz) or 23,200 MHz (because the difference between 14,200 MHz and 23,200 MHz is also the IF frequency of 9 MHz). For this example, we will assume that we chose a local oscillator frequency of 5,200 MHz, since this is within the range that can easily be generated by a VFO.

The resulting 9 MHz IF signal is then filtered by the IF filter, which is a very narrowband bandpass filter. Modern designs typically use crystal filters, so for this example we shall assume a crystal filter with a pass-band of 9.0003 MHz (300 Hz above 9 MHz) to 9.0030 MHz (3 kHz above 9 MHz). Signals within the pass-band will be passed with little attenuation, while signals that fall outside the pass-band will be blocked. So what components of our original RF signal will fall within the filter pass-band? Well an RF signal at 14.2003 MHz would be mixed down to 9.0003 MHz by the 5.2 MHz local oscillator signal; and a signal at 14.2030 MHz would be mixed down to 9.0030 MHz. So the signals that originated at these frequencies – from 14.200 3 to 14.203 MHz – will make it through the IF filter. This corresponds to an USB signal at a frequency of 14.200 MHz.

What about signals on the “other side” of 14.200 MHz, from 14.1970 to 14.1997 MHz, i.e. the frequencies that would have caused an image in a Direct Conversion receiver? Well, they will be mixed down to between 8.9970 MHz and 8.9997 MHz, and will be rejected by the IF filter, so they do not cause a problem.

There is still an image, but in this case it is from 3.8003 MHz to 3.8030 MHz. A 3.8003 MHz signal mixed with our 5,2 MHz local oscillator will generate an additive (sum) product at 9.0003 MHz, and a 3.8030 MHz signal will generate a mixing product at 9.0030 MHz. So signals within the frequency range 3.8003 MHz to 3.8030 MHz when combined with the 5.2 MHz local oscillator signal will also generate products in the IF range from 9.0003 to 9.0030 MHz that will be passed by our IF filter.

However this time the image is far away from the desired signal at 14.200 MHz, so it can easily be filtered out before the mixer, and this is the main purpose of the pre-selector. It must pass the desired frequencies, around 14.2 MHz, while rejecting the image frequencies, around 3.8 MHz.

Fortunately because these frequencies are so far apart, it is fairly easy to get good “image rejection” from a simple bandpass filter made of inductors and capacitors.

To find the image frequency, just find the sum of, and difference between, twice the IF frequency and the desired receive frequency. So for the example above, with an IF of 9 MHz, twice the IF is 18 MHz. The sum of 18 MHz and the desired receive frequency of 14.2 MHz is 32.2 MHz. This is where the image would be if the design used a local oscillator with a frequency higher than the desired signal. The difference between twice the LO frequency, 18 MHz, and the desired receive frequency, 14.2 MHz, is 3.8 MHz, and this is where the image frequency will be with the local oscillator running at a lower frequency than the desired receive frequency, as it is in the example above.

Note that by varying the frequency of the local oscillator we can change what frequency RF signal will be mixed down to the 9 MHz IF. For example, a local oscillator frequency of 5.3 MHz would mix an RF signal of 14.300 MHz down to the 9 MHz IF, while our original reception frequency of 14.200 MHz would now be mixed down to 8.900 MHz and would be blocked by the IF filter. So can you tune a superhet receiver by varying the frequency of its local oscillator (the same as for a Direct Conversion receiver).

The circuitry after the IF filter is virtually identical to that of a Direct Conversion receiver. The IF signal is amplified, and then mixed with another locally generated oscillator signal – this time called the "Beat Frequency Oscillator" or BFO – to recover the audio signal, which is then amplified by an audio amplifier. Since the IF signal is at a fixed frequency – 9 MHz – the BFO does not have to be tunable so we can use a stable fixed-frequency 9 MHz crystal oscillator for the BFO.

The Automatic Gain Control (AGC) also works similarly to that of a direct conversion receiver, although in this case the AGC control voltage is derived from the intermediate frequency, rather than the audio frequency output. This gives us "IF-derived AGC" as opposed to the "audio-derived AGC" that we had in the direct-conversion design. IF-derived AGC is superior to audio-derived AGC as it is able to respond more rapidly to sudden changes in signal strength.

The same design can be used to receive CW signals as well. For example, to receive a CW signal with a frequency of 14.200 MHz, the local oscillator would be set to 5.1994 MHz, generating an IF signal at the difference between these frequencies, 9.0006 MHz, which is within the pass-band of the crystal filter. After being amplified it will be mixed with the 9.000 MHz BFO signal in the product detector, generating an audio tone of 600 Hz.

So how about lower sideband signals? Well the simplest approach would be to have a second IF filter with a pass-band from 8.9970 (3 kHz below 9 MHz) to 8.9977 MHz (300 Hz below 9 MHz) that can be selected in place of the 9.0003 to 9.0030 MHz filter when we want to receive an LSB signal. Then when switching from USB to LSB all you have to do is switch filters, the local oscillator and BFO frequencies remain the same. Since crystal filters are quite expensive, an alternative approach is to use the same IF filter for LSB and USB reception, and just change the frequencies of the local oscillator and BFO. For example, to receive a LSB signal at 14.200 MHz using the 9.0003 – 9.0030 MHz IF filter we could set the local oscillator to 5.1967 MHz and the BFO to 9.0033 MHz. We leave it to the reader to fill in the details.

Since we can receive USB, LSB and CW signals using this design, how about AM signals? Well there are two options. The simplest is just to leave the receiver design exactly as it is, and receive AM signals as though they were single-sideband signals, ignoring the carrier and the other sideband, which will be filtered out by the IF filter. A better approach would be to provide another selectable IF filter, this time with a pass-band from 8.997 to 9.003 MHz to accommodate the 6 kHz bandwidth of an AM signal. The product detector would then be designed so that in the absence of any signal from the BFO, it would act as a half-wave rectifier and would detect AM by rectifying the IF signal (an "envelope detector"). This would give us the benefits of "proper" AM demodulation, notably accurate reproduction of the frequencies of the original audio signal even if the receiver is not perfectly tuned.

Multiple-Conversion Superhet Receivers

When choosing the IF frequency for a single-conversion superhet, there is a trade-off between image rejection and selectivity. It is easier to make highly selective filters at a low IF – say 455 kHz. However a low I.F. means that the **image frequency** is close to the desired frequency, making it difficult to effectively reject the image. Conversely, a high IF makes a large separation between the image frequency and the desired signal, making it easy to reject the image while passing the desired signal. However a high IF makes it harder to achieve the desired selectivity.

The classical solution to this dilemma has been to use a superhet design with two intermediate frequencies – a high first IF for good image rejection, followed by a low second IF for good selectivity. However modern crystal filters generally make this unnecessary in HF receivers, since very good selectivity is available from crystal filters at intermediate frequencies in the 9 MHz region, which is a high enough IF to attain good image rejection as well. Of course in VHF and UHF receivers, a higher first IF may be required to prevent unwanted image responses.

Despite this, the multiple-conversion superhet is still the most common approach for commercial HF receivers, but for a slightly different reason. Most commercial receivers and transceivers today offer "general coverage receive", meaning that they can receive on any frequency in the MF and HF bands, typically from 500 kHz to 30 MHz. Unfortunately this gives them a problem with IF leak-through, which occurs when the first mixer is not exactly balanced, allowing some of the original RF signal to appear at the IF output. If the RF signal is at the same frequency as the IF, then it will be passed by the IF filter, causing the radio to respond to a frequency that it shouldn't, a phenomenon known as a "spurious response". This would not be a big problem for an amateur-bands-only receiver, because an IF frequency like 8,5 MHz could be chosen that is not close to any amateur band. Then the pre-selector, possibly assisted by a dedicated notch filter at the IF frequency, will be able to reject incoming RF signals at the IF frequency, so there are no signals in the RF input that could "leak through" into the IF stages.

However the designer of a general-coverage receiver is not so fortunate. If the chosen IF frequency is anywhere in the receiver's frequency range, then it will be impossible to reject RF signals at the IF frequency, since these might include the frequency the receiver is tuned to! The solution is to choose an IF frequency that is either above or below the receiver's frequency range. However now the selectivity versus image rejection tradeoff comes back with a vengeance because a filter that is above the frequency range of a typical general coverage HF receiver – that is, above 30 MHz – will not have the necessary selectivity; while a filter at an IF that is below the receiver's coverage – say 455 kHz – will not allow adequate image rejection.

The usual solution is a multiple-conversion superhet where the first IF is above the receiver coverage range, allowing good image rejection and IF leak-through rejection, while the second IF is at a lower frequency where better selectivity can be obtained. This is known as an "up-conversion" design, since the incoming signal is first converted up to a higher frequency. The IF filter at the high first IF is often referred to as a roofing filter and is generally wide enough to permit signals of all modes through, up to 12 or 15 kHz in the case of a receiver that supports FM as well as other modes. Much narrower filters are provided for the different modes (e.g. a 6 kHz filter for AM and a 2,4 kHz filter for SSB) at the lower second IF. The block diagram below shows the "front end" (the circuitry from the antenna to the IF filter) of a typical general-coverage dual-conversion superhet.

Front-End of a General Coverage Dual-Conversion Superhet

The design includes a bank of switched bandpass filters in the pre-selector, to allow coverage of the range 0.5 – 30 MHz with good image and IF leak-through rejection. The first local oscillator is a frequency synthesizer running from 60.5 to 90 MHz, which up-converts the RF signal to the first IF of 60 MHz. Here it is filtered by the roofing filter, which would typically have a bandwidth of 12 kHz or so. The purpose of the roofing filter is to reject signals which are close enough to the desired frequency to be passed by the pre-selector, but which might cause either inter-modulation distortion or an image response in the second mixer. The IF signal is then amplified and converted back down to the second IF frequency of 9 MHz. From here on the circuitry would be similar to the single-conversion design featured earlier.

Noise Limiters and Blankers

Many common sources of amplitude-modulated noise generate amplitude "spikes" of short duration but high amplitude, which extend over a wide range of frequencies. These may contain substantial energy due to their large amplitude, even though their duration is short. Such noise is generated both by natural sources, such as thunderstorms, and by man-made ones, like inadequately suppressed ignition systems. Interference from these noise sources can be reduced by noise limiters and blankers, which are available on almost all modern amateur transceivers.

A noise limiter is a very simple circuit that limits the maximum amplitude of the received signal.

Circuit Diagram of a Noise Limiter

Assume the input signal has a maximum amplitude of 0.5 V peak under normal circumstances. This is less than the 0.6 V forward bias voltage of the diodes, so they do not conduct, and the input signal will be passed to the output unchanged. Then suppose a noise pulse generates a signal amplitude of 5 V. As soon as the amplitude exceeds 0.6 V, the diodes will conduct, effectively limiting the maximum output to 0.6 V peak and substantially reducing the energy of the noise signal.

The noise blanker is a more sophisticated variation on this idea. It detects the large amplitude of the incoming noise signal, and then immediately mutes (turns off) the audio output of the receiver completely for a predetermined time, typically a few milliseconds. Although this blocks the desired signal as well as the noise, this usually goes unnoticed by the listener as the human ear is quite insensitive to very short gaps in sounds, and the resulting signal degradation is much less than would have been caused by the high amplitude noise spike.

Frequency Modulation (FM) Reception

The basic superhet design can also be used to receive frequency modulated (FM) signals. However in this case, the product detector is replaced by a **Foster-Seeley discriminator** or a **ratio detector**. **[OK, This is no longer the case. Most modern FM discriminators use a simple 90 degree phase shifting circuit.]** These are circuits that convert frequency variations into a varying output voltage, so recovering the modulation from an FM signal.

The **discriminator** works by positioning the FM signal on the slope of a selective filter, so that variations in the frequency of the FM signal will result in variations in its amplitude. This converts the frequency modulation into a combined amplitude and frequency modulation, and a diode detector is used to recover the modulation from the AM component.

The graph shows how the slope of a high-pass filter could be used to convert frequency modulation into amplitude modulation. As the signal frequency increases from FC, the centre frequency, to FHIGH, the amplitude of the output increases from AC to AHIGH. If the frequency decreases from FC to FLOW, then the amplitude of the output will also decrease, from AC to ALOW.

Because the discriminator is also sensitive to changes in the amplitude of the incoming signal, it should be preceded by a limiter. This is a circuit that limits the amplitude of the signal, so that amplitude variations are not passed on to the discriminator or ratio detector that follows. The limiter circuit is identical to the noise limiter discussed earlier, except that in an FM receiver the circuit would be driven at a much higher input level, causing the diodes to conduct and clamp the output signal to 0,6 V peak. In this way the output of the limiter will always be at the same level (0,6 V peak), irrespective of the amplitude of the input signal. The block diagram below shows the final IF stage of a typical FM receiver

Final IF Stage of an FM Receiver

When the received signal is very weak the limiter is ineffective and the discriminator will respond to amplitude variations, which cause hiss in the audio. As the signal gets stronger and the limiter takes effect, the hiss decreases, a process called “**quieting**”. In order to prevent the hiss from bothering the listener when there is no received signal, most FM receivers incorporate a **squelch** feature, which mutes (turns off) the audio output when the received signal is below a minimum level known as the squelch threshold. The squelch threshold may be fixed or it may be adjustable using a squelch control.

Summary

The 'superhet' receiver converts the incoming RF signal to one or more **intermediate frequencies** before demodulating it. Superhet receivers have an **image frequency** that when mixed with the local oscillator will also generate the same **I.F.** as the desired receive signal.

The **image frequency** will be either the sum of, or the difference between, twice the IF frequency and the desired receive frequency. The role of the **pre-selector** is to reject incoming RF signals at the image frequency, preventing them from causing a spurious (unwanted) response in the receiver. The choice of intermediate frequency is a trade-off between selectivity (better at low frequencies) and image rejection (better with a higher frequency I.F.).

If a single IF cannot give adequate selectivity and image rejection, then a dual conversion design may be employed, with a higher first I.F. to give good image rejection, and a lower second I.F. to give good selectivity.

Noise limiters limit the amplitude of pulse noise, reducing the effect on the receiver. Noise blankers mute the audio output for a short time (a few milliseconds) when the higher amplitude associated with pulse noise is detected.

FM signals are detected using a **Foster-Seeley discriminator or ratio detector**. The discriminator should be preceded by a limiter to prevent it from being affected by variations in the amplitude of the signal. Weak FM signals have a characteristic hiss on them, and as the signal strength increases and the limiter becomes effective the hiss goes away, a process known as quieting. Most FM receivers incorporate a **squelch function, which mutes the audio output** when there is no received signal to avoid the annoying hiss.

http://users.tpg.com.au/users/ldbutler/Superhet.htm

http://zpostbox.narod.ru/tuned_radio_frequency_receiver_e.html