

Performance Characteristics of AM Ferrite Antennas and AM Broadcast Band radios, for the Mediumwave DX'er and Antenna Builder.

Rob Conboy, August 2022

Ferrite loopstick antennas are magnetic loop antennas. They started being becoming common in AM radios in the late 1950's and had become the norm by the 1960's. These antennas in AM radios have design tradeoffs that limit their performance. Luckily, this kind of antenna does not have high performance requirements. A loopstick is a magnetic loop antenna. It's performance characteristics are sensitivity, selectivity, directivity, and noise rejection. (rejection of near e-field radiation)

Sensitivity is the measure of how strong a signal with a standard modulation needs to be to get a standard quality at the loudspeaker. By definition, "sensitivity" is the input power of an RF signal with a 50 % AM modulated 1 kHz tone, needed to get a 12 dB SINAD output from the speaker. SINAD is a measure of audio quality. It's the same as signal to noise ratio, except it includes distortion byproducts in addition to noise. Sensitivity is measured with a signal generator connected to a standard loop coupled to the radio, and an audio analyzer connected to the speaker. This is done in a lab in a shielded booth or chamber.

Worse sensitivities require a stronger output from the generator to obtain a standard quality output from the speaker. Sensitivity is quoted in microvolts or in a dB scale of signal power. Lower numbers are better. When shopping for a radio, different manufacturers may list sensitivity differently from each other (for marketing purposes), so it can be hard to compare radios by info printed in their brochures. More reputable manufacturers include a reference in their specifications.

In practice, sensitivity of the radio only needs to be good enough for received background noise to be stronger than the internal noise generated within the radio. This performance is super easy to attain because the noise level of the AM band is very high, even in isolated rural areas, and can easily mask noise generated within the radio. Sensitivity correlates with the size of the ferrite rod. An 8 to 10 mm diameter rod which is 150 to 200 mm long has enough sensitivity for most conditions. However, most AM radios have smaller ferrite rods and will see a performance boost if upgraded. Once the received background noise from the antenna becomes higher than the radio's noise floor, any additional signal level boost from the antenna does not improve sensitivity. In fact, it degrades the dynamic range of the radio by making it easier to overload the radio's front-end. Therefore higher "S" meter readings are not necessarily better. Once amplitude from the antenna is high enough to make background (dead air) noise louder than the radio's internally generated noise, the only way to improve sensitivity is to reduce the background noise without reducing desired signals. The first way is to use a directional antenna that attenuates noise (and stations) coming from the antenna's disfavored directions. This reduction in noise helps unmask weak stations coming from the antenna's favored directions. A second way to increase sensitivity is to decrease noise, by decreasing the receive bandwidth through improved selectivity. Selectivity is usually obtained by a filter located much further down the signal path from the antenna.

Some radios have lower internally generated noise than others. These radios require

less signal to overcome this noise and therefore are inherently more sensitive and can take advantage of a smaller antenna. If they don't more easily overload on strong signals, then they have better dynamic range too, otherwise, increasing the size of the antenna might cause problems with strong signals.

Selectivity is a measure of rejection of signals on frequencies other than the desired frequency. For generally acceptable selectivity, a tuned ferrite antenna only needs to be selective enough so that signals received at the mixer's image frequency are attenuated enough to make their amplitude fall below the noise generated within the radio. This is an easy performance requirement to meet because the resonant frequency of the antenna is far enough from the image frequency. The loopstick antenna functions as a preselection filter. More selectivity is provided by a filter elsewhere in the radio.

Selectivity is often specified as the center frequency of the filter, the filter's insertion loss in dB, the -3 dB bandwidth of the filter, the -20 dB bandwidth, and the stop-band attenuation of frequencies far from the filter's passband. Most reputable manufacturers will list all the figures and provide a graph of the filter's selectivity curve with these figures labeled.

Practical aspects of selectivity are explained in more detail later, and Image frequencies are covered as well.

By definition, magnetic loop antennas are small in proportion to a wavelength. Since AM wavelengths are hundreds of meters, and ferrite antennas are centimeters long, even the biggest ferrite antenna is electrically tiny. Magnetic loop antennas are inherently bidirectional. They are most sensitive to signals in the plane of the individual wire turns. In a loopstick, that's broadside to the ferrite rod. In an air-core loop, that's in the plane of the coil. Loop antennas are insensitive to signals coming from off the ends of the ferrite rod. In a round air core loop, looking at the loop antenna as a circle, the antenna is insensitive to signals passing straight through the circle. Looking at the end of a ferrite rod antenna, the coil winding also is seen as a circle. Therefore, the antenna is insensitive to signals passing straight through it. This directivity lets you steer the antenna for best reception. A significant, but less obvious benefit is that directivity results a lower level of received background noise coming from directions the antenna disfavors. That noise tends to mask very weak signals from any direction. A directional antenna removes some of this noise, making the antenna more sensitive to weak signals coming from directions favored by the antenna.

A good loopstick antenna with the resonant circuit wired in a balanced configuration is excellent at rejecting locally generated e-field interference. Loopsticks, being magnetic loop antennas, respond to the magnetic field component of electromagnetic (radio) waves, not the electric field component. The electric field impinges on opposite ends of the coil equally, so it produces no difference in potential across the coil. The most common source of e-field interference is the electric field surrounding digital and power circuitry in small electronics near the radio. Note that e-field radiation is a local phenomenon. With increasing distance, a magnetic field component develops orthogonally to the e-field. It becomes a radio wave and the antenna will be sensitive to it. Fortunately, amplitude drops rapidly with

distance, and for small electronics, the amplitude of their interference field usually drops below ambient before any significant magnetic component develops. However, when a small electronic device is plugged into an outlet, interference can be carried some distance from its source through the wires. Home AC power wiring, even the safety ground, usually distributes interference throughout the house.

When speaking of loop antennas, the directivity is predictable only when the loop is electrically well balanced. If you ground one end of the loop (unbalancing it), the directional reception pattern is spoiled. When one end of the coil is grounded, e-field signals impinge on the opposite end of the coil and produce a potential across the coil. This potential is transferred into the radio. This e-field sensitivity is likely to be omnidirectional. This, when summed with the primary pickup of the antenna, causes the antenna to lose its predictable directional characteristic. Unfortunately, many radios ground one end of the antenna coil, ruining the antenna's ability to reject locally generated e-field interference.

It is possible to spoil the bidirectional pattern in a controlled way, changing it into a cardioid pattern, but it's difficult to get it right, and even harder to keep it right when changing frequencies or moving the antenna. Since the magnetic and electric components of radio waves are orthogonal to each other, all it takes is another 90 degrees phase rotation between an antenna element that is omnidirectionally sensitive to the electric field, and an antenna element that is bidirectionally sensitive to the magnetic field, to make the two field components, when combined, either enhance each other, or cancel each other. What happens is that signals impinging on the omnidirectionally sensitive antenna element will be in-phase with signals impinging on the bidirectionally sensitive element when coming from one direction, so that its component is added to the directionally sensitive antenna component, enhancing it. But, from the other direction, the directionally sensitive component is out of phase with the omnidirectionally sensitive component, canceling it at least partially. If the amplitudes of the two components are the same, but the phase between them is diametrically opposite, they completely cancel each other out, resulting in the antenna having no sensitivity to signals coming from the direction that fulfills this condition. The difficulty in creating a cardioid pattern lies in equalizing the sensitivities while making the phase between them exactly opposite. Even if you were to get a good, stable cardioid pattern, it will be very susceptible to locally generated interference since half the sensitivity is to the e-field. The environment needs to be free of locally generated interference. Improved and more stable directivity can more easily be attained by combining two antennas of the same type, separated by enough distance that propagation delay rotates the phase. There are schematics on the internet for building phasers for combining two antennas. Combining antennas can greatly enhance DX capability. More on this later.

More on selectivity: A loopstick antenna, consisting of a ferrite core coil and variable capacitor, collects radio waves centered around its resonant frequency and converts them into electricity. The antenna attenuates signals that are different from its resonant frequency. There will be a frequency below resonance and another frequency above resonance where this attenuation reduces the signal power by half. Bandwidth is defined as the upper half-power frequency minus the lower half-power frequency. Quality Factor (Q) is defined as the resonant frequency divided by the

bandwidth. Signals from the antenna are attenuated more and more as the frequency departs further from resonance, but this attenuation can be quite gradual. Lower values of Q yield wider bandwidths and more gradual attenuations.

If the ferrite antenna has a Q of 100, then the half-power bandwidth, when tuned to 600 kHz, is 6 kHz. (600,000 divided by 100) This contributes little or no audible selectivity to the radio because the attenuation is too gradual. This value of Q is a good choice for an AM antenna. It will provide good image rejection and reduce the chance of strong signal overload from non-adjacent channels. The same Q of 100, at a station located at 1600 KHz, will yield a 16 kHz bandwidth. This overlaps the sidebands belonging to the adjacent channels. With this bandwidth, antenna tuning near the top of the band is very broad. This broad antenna bandwidth can make it so that fine tuning the antenna has no noticeable effect near the top of the band. That helps us, by masking an undesirable characteristic of many tuners: they crowd stations more and more tightly together toward the top of the band. Broad tuning of the antenna also helps reduce the effect of any mistracking between the local oscillator's frequency and antenna's resonant frequency.

In order to understand image rejection and the relationship between the local oscillator frequency and the antenna resonant frequency, we have to know the function of the mixer and what it does. The mixer transposes the tuned frequency from the antenna to an intermediate frequency (IF) for further processing.

This works by mixing the signal from the antenna with a signal from a local oscillator of a different frequency.

In most (AM) radios, the local oscillator is a set number of kHz (often 455 kHz) higher than the antenna's tuned frequency.

It is possible for a radio to be designed so that its local oscillator frequency is lower instead of higher, but doing so for AM creates problems to be explained later.

It is even possible to make the the local oscillator be the exact same frequency as the RF. In that case, the mixer theoretically outputs a 0 Hz IF (actually a DC offset depending on the phase relationship between the carrier and oscillator). The sidebands will be the baseband audio frequency. This is called a "product detector" and can decode SSB and AM transmissions. We'll touch upon that a little more, later.

The mixer has two input signals. They are the local oscillator just discussed, and the RF signal from the antenna.

The mixer's output contains two signals of interest: the sum of, and difference between, the oscillator frequency and the frequency of the signal from the antenna. This means that the mixer transposes two different frequencies from the antenna to the same IF frequency. We only desire to receive one of these frequencies. The undesired frequency is called the "image" and comes from the sum. The desired frequency comes from the difference. The resonance of the antenna passes the desired radio frequency, and filters out the undesired image frequency.

In a superhetrodyne radio, the local oscillator's frequency tracks the tuned frequency but is offset by a set number of kHz. This set number will determine the transposed intermediate frequency (IF). Because of the LO is offset to a higher

frequency, the ratio of highest to lowest oscillator frequency is smaller than the ratio of highest to lowest antenna frequency. This difference in ratios causes an issue to be discussed later.

The IF filter is downstream from the mixer in the signal path through the radio. The IF filter provides the selectivity that separates adjacent frequencies from the tuned frequency. The IF filter does not remove the image because the image is transposed to the same IF as the desired. Instead, the antenna's resonance removes the image before it is transposed to IF.

Assuming the IF filter center frequency is 455 kHz, then when the radio is tuned to:

Receive Frequency = 530 kHz - Bottom channel of the AM band, then

Local Oscillator (LO) frequency will be 985 kHz - This is the LO frequency for the bottom channel of the AM band.

With LO = 985 kHz and IF Filter = 455 kHz, the IF contains stations from two frequencies: the difference $985 - 455 = 530$ kHz (desired) and the sum $985 + 455 = 1440$ kHz (image)

Note that the image frequency is within the AM band, if there is a station there, there's potential for interference, but this Image frequency is well outside of the antenna's bandwidth. Therefore, the image is greatly attenuated.

Receive Frequency = 1700 kHz - Top channel of the AM Band

LO Frequency is now 2155 kHz. This is the LO frequency for the top channel of the AM band.

With LO = 2155 kHz and IF filter = 455 kHz, the IF contains stations from two RF frequencies: Difference $2155 - 455 = 1700$ kHz (desired) and sum $2155 + 455 = 2605$ kHz (image)

The loopstick is tuned to the desired frequency. Any signals at the Image frequency are greatly attenuated.

Tuning the antenna from highest frequency (1700 kHz) to Lowest frequency (530 kHz) yields a 3.2 to 1 tuning ratio. ($1700 / 530 = 3.2$)

Tuning the Local Oscillator from highest frequency (2155 kHz) to lowest frequency (985 kHz) yields a 2.2 to 1 tuning ratio ($2155 / 985 = 2.2$)

The ratio between the (span of RF frequencies) to the (span of LO frequencies) is approximately a 3:2 ratio. (actually 3.2 to 2.2)

Because these tuning ratios are different, and the antenna tuning capacitor and the local oscillator tuning capacitor share the same knob shaft, the two frequencies will not track perfectly in proportion to each other. There will be multiple places on the dial where RF and IF are not perfectly aligned, even when the radio is aligned correctly. Some tuner's variable capacitors have rotor plates with slots in them. The purpose of these slots is to enable plates to be bent partway through the capacitor's tuning range, therefore changing the capacitance slightly at some places on the dial to correct any mistracking. Doing this is rarely necessary and difficult to do without making matters worse or even short circuiting the capacitor..

Here is what happens if the radio were designed to have local oscillator be 455 kHz below the tuned frequency of the antenna instead of above:

Receive Frequency = 530 kHz - Bottom channel of the AM band, then
Local Oscillator (LO) frequency derived from 530 kHz minus 455 kHz equals 75 kHz - this is the LO frequency for the bottom channel of the AM band.

Receive frequency = 1700 kHz - Top channel of the AM band
Local Oscillator will be 1700 kHz minus 455 kHz equals 1245 kHz - this is the LO frequency for the top channel of the AM band.

Now let's look at the tuning ratio span from highest to lowest oscillator frequency: 16.6 to 1 (1245 kHz divided by 75 kHz equals 16.6)
Yet, the tuning ratio of the antenna is still 3.2 to 1
But now, span of LO vs. span of RF is a 16.6 to 3.2 ratio. (roughly 5:1)

This is the problem referred to earlier. Firstly, in the analog world, it's nearly impossible to make a variable oscillator that works continuously across an almost 17:1 frequency ratio, and even more difficult to have it share a tuning knob with the antenna's capacitor. Next, the roughly 5:1 ratio between LO and RF spans makes mistracking much more severe than it would be at 3:2. Furthermore, lower image frequencies are intrinsically harder to suppress, while static, noise, interference tend to have a higher amplitude at lower frequencies. These reasons are why we choose a local oscillator that has a higher frequency than the desired frequency instead of lower. In mixer parlance, this is called "high side injection".

To recap, because some mistracking is unavoidable, the RF (antenna) tuned selectivity needs to be broad enough to allow the signal to be received at points on the dial where mistracking is greatest. If the antenna selectivity is too "good", then performance may suffer because the antenna tuning is slightly inaccurate for some frequencies. Also, selectivity must not be too "poor", because the antenna needs to reject the image frequency. This is because the mixer converts two frequencies to IF, making the radio sensitive to two frequencies at once. The antenna selectivity needs to be narrow enough to filter out the undesired "Image" frequency. In other words, the antenna functions as a preselection filter. A radio alignment procedure may make a lower frequency and an upper frequency have LO and RF perfectly aligned, while other points on the dial may be slightly misaligned. The LO determines what frequency will be transposed to IF. The antenna will not be tuned to exactly this frequency at some points on the dial. The antenna's selectivity needs to be broad enough so this misalignment won't matter.

More on selectivity:

AM broadcast channel spacing is 9 or 10 kHz (depending on country). The signal is amplitude modulated, resulting in an RF bandwidth that is twice the demodulated audio bandwidth. Broadcast AM signals have the same audio content in both sidebands. The sidebands are identical mirror images of each other, centered around the carrier. Most of the energy is in the carrier. The audio is in the sidebands. Only one sideband is actually needed to be able to recover the audio, the other sideband is redundant. Unfortunately, unless a radio station limits its audio to 5 kHz, the

sidebands of their transmissions overlap the sidebands of adjacent channel transmissions. Most transmit audio to around 9 kHz or a little higher, to include treble in the transmission. This makes music sound better. It also means their transmissions occupy 18 kHz or more of bandwidth. Adjacent channel interference is unavoidable regardless of the radio's channel selectivity.

The receive channel bandwidth is determined by the radio's IF filter, and is usually 10 kHz or less. IF filters are often narrower than the broadcast bandwidth. If the IF filter is 10 kHz wide, the audio will be limited to 5 kHz when tuning is centered on-frequency. Audio frequencies above 5 kHz will be attenuated. If you detune the radio while keeping the carrier within the IF filter passband, the audio frequency response will extend to a higher frequency, as high as just below 10 kHz, because this particular IF filter's passband is 10 kHz wide. Detuning helps reject adjacent channel interference when it extends into the opposite sideband.

Some radios give you a switchable choice of IF filters of different bandwidths. Wider ones yield better audio fidelity on strong signals. Narrow ones are better for weak signals that are accompanied by noise and interference. Ceramic, crystal and mechanical IF filters have (relatively) flat frequency response within their passband, but attenuate very steeply outside their bandwidth. This combination of flat passband and sharp cutoff yields excellent selectivity.

If the IF filter is narrower than is broadcast, then strong signals will sound a little harsh while missing some treble. This harshness occurs because of the filter's abrupt transition from flat passband to steep rolloff. The reason the broadcast does not sound so harsh with weak signals through a narrow filter is because noise masks the harshness. The filter's attenuation of that noise makes the channel more audible. This is how narrow IF filters make the radio more sensitive for weak signals - by removing noise that would otherwise mask the broadcast audio content.

Beware that one vendor's "6 kHz AM filter" may be 12 kHz wide and limit audio to 6 kHz. Another's might be 6 kHz wide and limit audio to 3 kHz when tuning is centered on the carrier. It depends on how the vendor defines it. This can trip anyone who neglects to read the filter's datasheet before purchasing an upgrade.

If you install an improved replacement antenna with much higher Q, then you might want to realign your RF/LO relationship so that the LO/Rf mistracking is worst at the high end of the band and best at the low end. This is accomplished by using proportionally lower frequency alignment points instead of those called for in the alignment procedure, i.e. 1140 kHz and 570 kHz instead of 1400 kHz and 700 kHz. The antenna bandwidth is widest at the top end, where mistracking will have the least consequence. The mistracking may be slight enough to make no difference. The only way to know is to try a better antenna of the exact same inductance and see if the radio needs realignment.

Determining the exact resonant frequency of an antenna can be tricky unless you couple the antenna to a network analyzer and make an S11 measurement. The Nano VNA is a good network analyzer that costs below \$100. LO alignment to the antenna can be determined by measuring the LO frequency and subtracting the IF frequency. The

difference should equal the resonant frequency of the antenna. Any inequality is due to mistracking. The LO frequency can easily be measured using a frequency counter or an SDR's spectrum display. A good Inductance/Capacitance meter for under \$100 is the "KNACRO LC200A Handheld L-C Meter Inductance Capacitance Meter Range 1pF-100mF 1uH-100H". Frequency counter/generators can be found on Amazon for roughly \$100 too, but make sure the one you select works below 1 MHz, as some of them don't. The Airspy HF+ Discovery is an under \$200 SDR that works really well as a spectrum analyzer, even though it is a radio - a darned good one! It's smaller than a deck of cards with an antenna spigot on one end, and a USB jack on the other to connect to a computer.

Manually tunable ferrite sleeve loop (FSL) external antennas increase the sensitivity of the radio manyfold and also improve selectivity, but selectivity is not improved nearly as dramatically as it might with some additional amplification and feedback circuitry. Since they're manually tuned, automatic mistracking cannot be an issue. For those who are more familiar with large hand-made tuned loop antennas mounted on a stand - a 2 by 2 foot tuned box loop has roughly equivalent signal gathering ability to an 6 inch diameter FSL made with mix 61 ferrite rods that are 200 mm long and 10 mm thick. (Please forgive my mixture of measure units) Differences in performance between these two will most likely be due to choice of components and construction techniques. Signal amplitude from the FSL will be higher because its field is concentrated into a much smaller space than that of the box loop. If the amplitude is too high, move the radio a little further from the antenna. A 4 inch diameter FSL made from the same rods, outputs roughly equal amplitude to a FOUR by FOUR foot box loop, but I'm more interested in signal to noise ratio than I am signal amplitude.

Without additional circuitry, a tuned ferrite antenna cannot be made to have the same bandwidth at all frequencies because Q remains relatively constant for all frequencies. What happens is the bandwidth becomes wider and wider with frequency. ($Q=F/BW$) In practice, due to imperfections, especially in the ferrite, Q usually degrades somewhat as frequency rises, making bandwidth even wider at higher frequencies. The only way to make bandwidth stay nearly the same when tuned frequency changes is either by spoiling Q at lower frequencies, making tuning uniformly broad throughout the band, or preferably, by increasing Q as frequency goes up, by using an amplifier with a variable feedback network.

Again, the received bandwidth of the antenna increases as tuned frequency goes up. Near the top of the band, tuning is so broad that fine tuning does practically nothing.

Tuning any radio around the band would seem to contradict the above statement. This is because channel selectivity is normally provided by the IF filter, not the antenna. The selectivity is intentionally broad until the signal reaches the IF filter. Since IF is centered around a frequency that is always the same, the filter's selectivity is always the same for every station. "Selectivity" is the function of this filter - to separate stations.

It is possible to make a resonant antenna with selectivity that rivals or exceeds the selectivity of an IF filter. Doing so will also greatly increase the

sensitivity. Because of design tradeoffs in superhet radios, an antenna you design for improved selectivity should be external, manually tuned, and coupled through the air to the internal antenna instead of directly replacing it. (if you wired up an output jack, the antenna can be connected to a jack on the side of the radio if it has one). One advantage of a manually tuned enhanced selectivity antenna is the ability to shift tuning to favor either sideband, even with a digitally tuned radio that tunes in 10 kHz steps. If tuned to a strong signal that threatens to overload the radio, just move the outboard antenna further away or don't use it.

A Q value of 1000 for the antenna is usually attainable with good components and careful construction. Maximum receive sensitivity will be enhanced, and selectivity will noticeably be better at the bottom half the band (towards 530 kHz), likely narrower than the bandwidth of the IF filter. The IF filter may be 6 kHz wide with steep rolloff. When antenna Q is 1000, the bandwidth from the antenna will be 530 Hz (when tuned to 530 kHz) with more gradual rolloff. Strong signal reception will yield audio with somewhat reduced treble. Yet, at the high end of the band, the antenna's bandwidth will be roughly 3 times wider and you may not notice any reduction of treble sounds. RF/LO mistracking can become an issue when Q is 1000, so results might be worse than expected in the lower half of the AM band. Note: Most modern digitally tuned radios, including the more recent DSP based designs, don't mistrack!

The radiation resistance of a loopstick antenna is a fraction of an ohm. Efficiency and especially antenna Q is degraded by the resistance of the coil, resistance within the capacitor, and by any load placed upon the resonant circuit. Efficiency is pretty much unimportant at MW because it's so easy to get enough signal from the antenna. Q is more important. Why is the Q of your homebuilt antenna not as good as you expected? If you're winding your coil on ferrite, then the ferrite itself causes hysteresis and resistance losses due to its secondary permeability (μ''). This damps the resonance. Also, any load coupled to the resonant circuit, such as coupling to a radio, will damp the system. To a lesser degree, any nearby conductive objects also load the antenna, even the Earth itself loads the antenna, as does the antenna's coupling to signals on the air. One way to minimize the load presented by the radio is by critical or looser coupling. This is achieved by changing the distance between antenna and radio, or if direct connected, by keeping the load impedance as high as possible. So, how can 1 or 2 turns of wire used as a pickup coil present a high impedance? Remember that we're dealing with a parallel resonant circuit, and impedance is highest at resonance. A single turn wound in the middle of a 40 turn coil has a 40:1 turns ratio which presents a 1600 to 1 impedance step down. Then with an additional 5:1 turns ratio transformer, you can connect to an amplifier with a 50 Ohm input. This arrangement will transform 50 Ohms up to 1600 Ohms, and then transform 1600 Ohms up to be a 2,000,000 ohm load on the resonant circuit. This 2-Meg load is close to the equivalent load created by coupling the antenna to the environment. It will still damp the resonance, but far less than directly connecting to the coil.

Resistance within the coil: RF conducts on the outside surface of a wire. This is called "skin effect", and it makes wire have higher resistance for RF than it would for DC. This resistance increases as frequency rises. Litz wire addresses this problem by having multiple individually insulated strands bundled together. This increases the total surface area manyfold, reducing the resistance to radio

frequencies. Unfortunately, Litz wire introduces another problem: When conductors are running parallel to each other (such as within a Litz bundle), the current flow through each strand tries to induce current flow in the opposite direction in adjacent strands. This manifests as an increased RF resistance of the bundle. This is called the proximity effect. Proximity effect also increases resistance as frequency rises, but at a different rate than skin effect. Skin and proximity effects combined makes Litz wire better than solid wire below some frequency, and worse above it. When Litz wire has 46 AWG strands, this frequency is near 1400 kHz.

Many antenna builders opt to space wind their coils, leaving a gap between turns. This supposedly works to reduce the parasitic capacitance within the coil, but the real benefit to doing this may have more to do with transmission line theory (RF propagating within the coil) i.e. reducing proximity effect between turns.

You might see Litz wire sold using a specification like "46/600". This means 46 gauge individually insulated strands bundled together in a cable made of 600 strands.

Most MW loop antenna builders opt for 46 gauge Litz wire for maximum Q. This is actually a tradeoff. One limitation of this 46 gauge Litz wire is that, above 1400 kHz or so, solid wire of the same diameter as the bundled litz cable, performs as well as Litz, and even better than Litz at higher frequencies.

Litz with 48 gauge strands is also a good choice because it is optimized to extend beyond the top of the band. Because each strand is thinner, it makes a thinner cable for the same number of strands. Therefore the minimum resistance of 48/x may be slightly worse than 46/x Litz at low to mid frequencies - but with 48/x, RF resistance may have less variance across the band. A coil wound with 46/x may theoretically have a higher maximum Q for the same number of strands, but 48/x Litz degrades the Q less at higher frequencies than 46/x does.

Before you go and invest in the most exotic Litz wire, remember it won't improve the losses within the ferrite or nor will it improve losses from coupling to and from the antenna. Increasing the strand count beyond a certain point (500 to 1000 strands) may yield no useful additional improvement of Q, because attainable Q is also limited by other parameters, not just wire RF resistance. In fact, solid core wire can be a good choice for the coil because it can be better than Litz near the high end of the band, helping to reduce Q degradation with frequency.

Capacitor effect on Q

MW broadcast band air variable capacitors have two capacitances in parallel. (Ganged capacitors are discussed later)

We obviously have the variable capacitance between the rotor plates and stator plates.

We have additional capacitance between the stator plates and the metal frame. These two capacitances are in parallel and therefore add together. Some variable capacitors also have in compression a mica trimmer capacitor between stator and frame for radio alignment. One of the connections to the capacitor is to the frame (housing). From there, RF travels through sliding spring contacts that electrically

connect the frame to the rotor. The other connection to the capacitor is to the stator. The stator is suspended above the frame using an insulated mount. The Q of the capacitor is mainly impacted by resistance in the spring contacts between the rotor and frame. Q is affected to a lesser extent by the dielectric qualities of the insulation comprising the stator mount. Ceramic insulation is superior.

At resonance, current circulates back and forth between the coil and capacitor as a low impedance circuit.

The resistance of the spring contacts in this low impedance circuit degrades its Q . Spring contact resistance, though small, is also likely to be greater than the radiation resistance of the whole antenna.

When the capacitor is rotated to minimum capacitance, the aforementioned contact resistance is of less consequence.

This is because much of the remaining capacitance is directly between the stator and frame, bypassing the spring contacts. Therefore the Q of the capacitor becomes less and less degraded by contact resistance as the plates are opened.

This works to our advantage, but not enough.

The spring contacts in most capacitors are sliding wipers that use spring tension to push between frame and rotor as the rotor is turned. The more spring contacts built in to the capacitor, the better. The cleaner they are, the better. Some (hard to find) capacitors have a spring connection solidly attached to both rotor and frame. This spring winds and unwinds as the shaft is rotated. This type avoids the problem of contact resistance by not having any.

Other capacitors have a nonconductive ceramic frame. This allows a lower minimum capacitance by eliminating the secondary capacitance to the frame. This does not relieve the problem of contact resistance, but these capacitors are often constructed with better materials and may have higher Q overall.

Then there is the butterfly capacitor, which has the advantage of higher Q , and is often used in RF circuits where either Q or power handling is important. They are commonly used in magnetic loop antennas for transmitting. A butterfly capacitor is made with two opposing stators with a shared rotor between. The rotor has plates shaped like butterfly wings. One wing meshes with one stator. The other wing meshes with the opposite stator. This makes two identical variable capacitors in series sharing one rotor. There aren't any spring contacts in the RF path to degrade the Q .

Butterfly capacitors are less suitable for our use because they are much larger (for a given maximum capacitance and voltage rating) and their capacitance range is not tapered well for our use. Furthermore, their large bulk and surface area increases parasitic capacitance to the surrounding environment. They do not have the tapered plate shapes and offset rotor shafts that good MW broadcast band air variables have for avoiding crunching the frequencies together near the top end of the dial. This crunch happens because the effect of increasing the capacitance from 10 to 20 pF is many times greater than the effect of increasing it from 350 pF to 360 pF, even though the change remains 10 pF in both cases. (One doubles the capacitance, the other changes it by one thirtysixth, but both take the same number of degrees of shaft rotation)

One issue with tuning capacitors is that their conductive mass introduces parasitic capacitance between it and the environment. This parasitic capacitance has an antenna effect which is sensitive to electric fields. You do not want this but you're stuck with at least some. This not only makes the antenna system more sensitive to e-field noise, it also skews the figure-8 directional pattern of the antenna as a system. This can make large capacitors, especially butterfly and vacuum variables, even less suitable for our use, in spite of their better Q. The physically smaller the capacitor, the better, just as long as Q and capacitance range are not compromised by miniaturization, for example, if the capacitor is constructed with the plates sandwiched between polymer sheets instead of air as a dielectric. Materials with higher dielectric constants enable proportionally smaller plate surfaces but usually compromise Q. Also, when you make plates smaller, the capacitance they contribute becomes proportionally less in relation to the capacitance within capacitor's internal support structure.

If you want the higher Q of a butterfly capacitor but with a smaller size and better taper, you can use a 2-ganged broadcast band air-variable capacitor and connect your coil between the stators. Both sections of the capacitor need to have identical capacitance and range. Connecting between the stators divides the capacitance by 2, therefore you will need to double the inductance. This takes the spring contacts out of the RF path. Q is also higher because tuned impedance is higher when Inductance is higher and capacitance is lower. Again, both sections of the capacitor need to be identical. Some dual ganged capacitors have different size sections with different capacitance ranges. Those are for tuning an antenna and a local oscillator together. The oscillator section is a little smaller. Those with different ranges are not candidates for connecting between stators.

Doubling the inductance doubles its reactance, but because doubling the inductance only requires about 1.5 times as much Litz wire, the increase in RF resistance (which is detrimental to Q) is less in proportion to the increase in reactance (which is beneficial to Q).

Another advantage to connecting your coil between the two stators is that the capacitor as connected will be electrically symmetrical. This helps make the parasitic capacitance to the environment cancel itself out. E-fields impinging on both stators are equal, resulting on no e-field potential between stators. This also makes the frame and rotor of the capacitor float at or near ground potential. The rotor now partially shields the stators. Grounding the frame won't affect tuning but will have an unpredictable effect on noise pickup. Putting your coil between the stators also makes the capacitor much less sensitive to the proximity of the hand while tuning. The disadvantage of connecting your inductance between the two stators is, with higher circuit impedance and lower capacitance, parasitic capacitance becomes more of an issue. A coil with more turns will also have a lower self-resonant frequency. It makes it more difficult to attain a frequency range that covers the whole band. Layout, shortest lead lengths possible, and lead dress become much more important.

Notice that capacitors sold for MW band tuning have construction differences from one brand to the next. Some are better than others. Some have rotor plates that are

half-round. Some have rotor plates that are tapered. Some have stator plates that are tapered. Some have rotors that are offset to one side of the stator plates. Some have rotor shafts that are centered over the stator. Some have only one set of spring contacts. Others have more.

The worst capacitors for our use have half-round rotor and stator plates on a shaft that is centered over the stator plates. These crowd stations together at the high end of the band the most. Offset shafts or tapered plates are better, and having both offset shaft and tapered plates are best. The holy grail is the straight line frequency capacitor with a coiled spring connection between rotor and frame (no contact resistance issue). These vintage capacitors have long, narrow crescent shaped rotor plates on shafts that are offset far to the side. The offsets and tapers make it so capacitance changes exponentially more slowly as you tune up in frequency. In a straight line frequency capacitor, turning the knob to double the capacitance should always involve about the same amount of shaft rotation. i.e. changing from 10 pF to 20 pF should involve a similar amount of shaft rotation as changing from 150 pF to 300 pF.

When trying for extremely high Q factors, you have to contend with losses arising from near field coupling of your antenna to things in the surrounding environment. Even if constructed with extreme care, your tuned circuit using exotic parts will not be able to attain the extremely high values of Q calculated using the parameters of these components. The highest Q you may be able to realize, with careful construction, might still be below 1000, even if the components yield Q values in the thousands. This is because coupling the resonant circuit to the environment places a load on it, introducing losses which damp the resonance. What this all means is there is a point beyond which, the environment will be the limiting factor in the selectivity performance of your antenna. The way to overcome these limiting factors is through regenerative feedback. This positive feedback is so effective it can even overcome losses incurred from using cheap materials. One can easily attain Q factors exceeding 5000. Then it will often be better to select a wide IF filter and let the antenna provide selectivity. With the right selectivity, you can tune the antenna to one sideband while rejecting the other - yet the radio, centered on the carrier, will still provide clean audio in AM mode. With regenerative feedback, the only difference between exotic components and cheap ones, is that cheap ones need more gain to overcome losses (and are therefore noisier), and it is more difficult to attain stability at all frequencies. With regenerative feedback, the road to the best possible performance does not require exotic components. Construction technique matters most. With extremely high values of Q, the antenna becomes extremely sensitive to the environment. If not carefully laid out, simply reaching toward the antenna can detune it enough to notice. Regenerative feedback requires the use of an RF amplifier. Amplifier output is fed back into the antenna to make up for the losses that damp the resonance and degrade Q. Gain must be slightly below 1 after making up for losses. Amplifier gain is controlled by a variable attenuator. A simple variable resistor will work. Feedback is in phase. The amplifier MUST be stable even if the load is not 50 Ohms. Some amplifiers aren't, and 50 Ohm tuned circuits are never 50 Ohms except at one frequency. Mini Circuits ZFL-500 and ZFL-1000 series low noise amps work well.

Where to physically locate the capacitor?

Most builders place the capacitor near either end of the coil, but doing so creates 4 detrimental conditions:

1. The coil ends are of highly sensitive. The proximity of the capacitor creates an imbalance in the antenna, skewing the reception pattern and coupling noise into the antenna.
2. The differing lengths of the two connections from the coil to capacitor create an imbalance that skews the reception pattern and increases sensitivity to noise.
3. If the capacitor is located off the end of the coil, its conductive bulk behaves as a shorted turn coupled to the coil, degrading Q.
4. The ends of the coil are very sensitive. Bringing your hand close (to tune the capacitor) detunes the frequency, degrades the Q and couples noise into the antenna. This is rarely an issue with a passive antenna, but with a regenerative antenna's enhanced sensitivity and needle sharp selectivity, it can ruin the performance. That being said, success is possible with your capacitor located here. Using an insulated shaft extender and a vernier dial helps immensely.

It's best to avoid grounding the coil. Worst case is if you ground one end of the coil. Mounting the capacitor can ground it, and therefore one end of the coil. The antenna will still work, but will become far more sensitive to noise pickup. This is how most old tabletop AM radios with internal loop antennas are wired. That's why they pick up so much locally generated interference. Those old radios were made in an era that had much less locally generated interference than there is today.

With one exception, the best place to place the capacitor is to the side of the coil, in the middle where it's relatively insensitive to e-field coupling into the antenna.

If you're using a dual ganged capacitor with identical sections, you can make the antenna perfectly symmetrical while increasing Q.

This makes the capacitor less sensitive to hand capacitance.

If you keep coil/capacitor circuitry, and any auxiliary windings as balanced circuits, then the middle of the coil is relatively insensitive to parasitic capacitance including that to your hand.

If you need to make a wire connection to a radio, the best way to do this is by winding a balanced pickup coil using as few turns as possible to keep the impedance ratio as high as possible. You may decrease the degradation of Q by inserting a step down transformer. Signal levels are likely to be very low, so may need a preamp. Place the pickup coil in the middle of the tuned coil for best noise rejection. The transformer will behave as a balun, preserving the balanced nature of the tuned circuit.

Ferrite sleeve antennas have an RF dead zone along the axial centerline mid way through the sleeve. If the sleeve is of sufficient diameter, the tuning capacitor may be located here, inside the sleeve. The sleeve inner diameter needs to be more than 4 times wider than the capacitor. In other words, the ferrite sleeve needs to be fairly large to locate the capacitor inside it without the capacitor ruining the performance. A 2x2x2 capacitor, when located inside the sleeve, noticeably degrades performance when the sleeve is 6 inches in diameter. An 8 inch diameter ferrite sleeve should perform well. You will need an insulated shaft coupler to connect an

extension shaft to the capacitor and some way to mount the capacitor and the vernier dial. In this configuration, it is best if the coil is centered amid the length of the sleeve. You may wish to leave a gap between two rods of the sleeve so that the coil ends can be passed through the ferrite, keeping the coil connections to the capacitor very short. This will reduce the parasitic capacitance to enable higher frequencies to be tuned. Another advantage is that an insulated shaft extension is needed on the capacitor to bring control outside of the ferrite sleeve. This keeps hands away from sensitive areas. Furthermore, the ferrite sleeve itself helps shield the capacitor. The gap between two rods has negligible effect on the sensitivity, but reduces coupling between any conductors running axially through the sleeve. It also relieves one from the problem of fitting imperfectly straight rods. It may be best for directional pattern symmetry to orient the antenna so this gap runs along the bottom of the sleeve.

The superhet radio design is older than the invention of the ferrite rod antenna. Many older AM radios, many from the 1930's, do not have any internal antennas. Instead, they have terminals to connect an external wire antenna and a ground connection. These radios have a tuned preselection filter inside for image rejection instead of using a tuned antenna. Others do have internal loop antennas, but they're not wound on ferrite. Instead, the antenna coil has a larger "air core", the windings are held in place with beeswax, and the loop coil is usually mounted inside, at the back of the radio.

Most high performance receivers with PLL oscillators and digital frequency display, are of dual conversion superhetrodyne design. This splits the functions of image rejection and channel selectivity into two separate IF stages, one after the other. A dual conversion design has 2 local oscillators, 2 mixers, 2 IF Frequencies, and 2 IF filters.

In this type, the first IF is for image rejection. It works by converting RF to a much higher IF frequency. This enables designing a preselection filter in the front-end to cover the entire band instead of just one channel. This is made possible because the image frequency is dozens to hundreds of times higher than the desired frequency. For example, at 70 MHz IF, image rejection of the entire AM band at once is easy because the image frequency is many times higher than the band. In these radios, the preselection filter is usually a band pass filter made with discrete coils and capacitors. These filters have a flat passband and a moderately steep stop-band rolloff. Some preselection filters in communications (shortwave) receivers serve a dual purpose. The main function is image rejection. A secondary function is to keep strong, out of band signals away from any RF preamp and away from the mixer. For example, keeping products from the AM broadcast band from affecting the 160 meter band. Many radios also have switch-selectable preamps and attenuators after the preselection filter. Unfortunately some radios automatically attenuate the AM broadcast band while disabling any preamp, because signals are usually higher there. The first IF filter needs to be wider than the bandwidth of any kind of modulation the radio is capable of receiving. If the radio has a noise blanker that works at IF, it will need to be sensitive to transients that have wider bandwidth than any modulated signal. The first IF filter may need to let some of that transient energy through to be used to control the noise blanker. Noise blankers work by momentarily silencing the radio during transients. This silencing

can be so short that we don't notice it.

The second IF stage in a double conversion superhet is for selectivity. The second IF is usually a very low frequency, such as 50 kHz. Here it is easy to make high performance channel selection filters inexpensively; often made from discrete coils and capacitors.

Dual conversion superhet is a superior, higher performance design. Most tabletop communications receivers use a design like this. This kind of radio usually uses a broadband, untuned antenna, such as a long wire attached to a tree. Still, if one or more in-band stations are strong enough to desensitize the radio to other stations, or make phantom stations appear because of intermodulation between strong stations, then a tuned antenna will help to attenuate everything but the station you want to receive. If signals are too strong, the only other options are to use a smaller antenna, or reduce overall sensitivity by placing an attenuator between the antenna and the radio. Many radios come with a built in switchable attenuator for this purpose. They may also have a switchable RF amplifier as well.

Most serious AM DX work using communications receivers and SDR's is done in SSB mode, even though SSB mode is intended for decoding SSB transmissions, not AM broadcasts. Using SSB mode for broadcast AM offers distinct advantages in intellegability during DX conditions, yet suffers from fair to poor audio fidelity. SDR and communications receivers usually have a stable and accurate enough SSB mode for broadcast AM reception. In comparison to AM mode, using SSB mode for AM broadcasts greatly improves both sensitivity and selectivity performance, and it reduces fading of broadcast AM signals.

SSB transmissions suppress the carrier, and transmit only one sideband because the opposite sideband is redundant. This makes transmission far more efficient by only broadcasting the audio content. The carrier contributes nothing other than serving as a reference. This reference can be supplied by an oscillator in the receiver. The carrier, without contributing, contains most of the signal power. That power is wasted, therefore, it is not broadcast during SSB transmission. SSB voice transmissions sound like quacking ducks when in AM mode because the carrier, which is the reference, is left out of the transmission. SSB mode uses a BFO (beat frequency oscillator) in the receiver to restore this reference. AM band radio transmission is different. The carrier and both sidebands are broadcast. When SSB mode is chosen for AM broadcast signals, audio fidelity may suffer a distinct shift in pitch. This improves with improved IF filter attenuation of the carrier, and improves when the BFO and carrier frequencies are within a few Hz of each other.

If the BFO and carrier aren't close enough in frequency, and the IF filter lets the carrier pass, the beat frequency between BFO and carrier will be an audible tone, and the broadcast audio pitch will be shifted by the frequency of the tone. Very annoying but intellegible.

If the IF filter sufficiently attenuates the carrier, but the BFO and carrier aren't close enough, the IF filter will remove the tone, but the pitch will be shifted by the frequency difference between the BFO and the carrier. It can sound dreadfully annoying.

If BFO and Carrier are within a few Hz, and the IF filter lets the carrier pass, the audio's phase rotates at the rate of the beat frequency. The accompanying shift of audio pitch may be imperceptible, but the audio will sound "swishy" or have a vibrato. Less annoying.

If the BFO and Carrier are within a few Hz of each other, and the IF filter removes the carrier, then audio fidelity will be just about normal, but missing some low bass, because frequencies near the carrier will be in the IF filter's stop band. This is the holy grail.

In practice, Since the IF is shifted to place one sideband of the broadcast within the passband of the IF filter, it enables using a filter with less than half of the bandwidth of a filter designed for broadcast AM. This increased selectivity also increases sensitivity by reducing noise. SSB mode also allows you to choose which sideband has less interference. SSB mode also attenuates the carrier (where most of the signal power is), preventing it from governing AGC (automatic gain control), therefore gain will be much higher. Since the BFO replaces the carrier, audio remains intelligible during common instances of selective fading that obliterate the carrier. This selective fading comes and goes. When the radio is in normal AM mode, during selective fades, audio becomes distorted beyond intelligibility. In SSB mode, audio remains intelligible during selective fades. Furthermore, since selective fading is always changing in frequency, depth, and exact physical location, a wire antenna whose length reaches outside a selectively faded location has an advantage over an antenna wound on ferrite. SSB mode removes that particular advantage by replacing the carrier with an oscillator inside the radio.

Most serious dx work is done with with disabled AGC (automatic gain control), sometimes called AVC (automatic volume control), with RF or IF gain set manually. This implies that signals are weak, often weakened by having the IF filter reject the strong carrier when in SSB mode. When IF gain is too high, audio will become badly distorted. The problem with AGC is that it has time constants that are tuned for casual reception and will make the audio bounce in volume during dx conditions, or silence the radio too long during static crashes or strong transients. AGC tries to keep the audio volume the same for all signal strengths. It does this by automatically changing the IF gain to maintain a constant level. The S-meter is actually a measure of a DC voltage derived from the IF output. This is the AGC control voltage. The IF amplifiers have voltage controlled gain. The higher the control voltage, the lower the IF gain. AGC control voltage is often derived by a rectifier and filter capacitor tapped into the IF output next to the detector. This circuit is very similar to the detector, except that it has much longer time constants. The rectifier changes the tapped IF signal to DC. This DC charges a capacitor. The voltage across the capacitor is distributed back to the IF amplifiers to control the IF gain. AGC has attack and release time constants. A series resistor between the rectifier and capacitor provides an "attack" time constant that slows the AGC response to sudden increases of signal strength. Another resistor across the capacitor bleeds off its charge, setting the AGC's "release" time constant. It responds to decreases in signal strength. If you stick a voltmeter across the capacitor, that voltmeter is the same thing as an S-meter, except, the S-meter scale is calibrated in dB.

Phasers on stun:

Sometimes you will come across a carrier but the sideband amplitude is just below that of the atmospheric noise, so there will be no audio other than noise. If there is a carrier, but received noise completely masks the sidebands, you might be able to reduce the noise enough to hear the station, by using a directional array of 2 antennas combined through a "phaser". This example is most likely daytime conditions, but that's the best way to learn how to use a phaser.

A phase combiner, or "phaser" does two things. It lets you adjust the relative amplitude between two antennas, and lets you adjust the phase of one antenna in relation to the other. You want to adjust the phaser to make the two antenna outputs from one direction be 180 degrees out of phase, so when combined, they cancel each other out. If the amplitudes are the same, the cancellation is complete. The "phaser" lets you adjust to make it exactly 180 degrees phase difference and lets you adjust the amplitudes until they're exactly the same. SDRplay's RSPduo is a radio that can do this natively with two antennas without needing a phaser. I prefer using a phaser.

Adjusting a phaser is both easy and hard. It is hard to know what the directional pattern actually as you make adjustments. But at the same time, it's easy to hear reception improve (or get worse) while adjusting.

Another use of a phaser is to use it to reduce interference from noise generated by a nearby source. One antenna is used to pick up as much of the offending noise as possible. The other antenna is set up in a less noisy, stronger signal environment. The phaser inverts the noise antenna's phase and combines it with the main antenna. The idea is to cancel out the noise, leaving only the desired signals. This use of a phaser is best when there is only one source of noise you're concerned about. For example, you might place a small antenna (for sensing noise) next to your breaker panel in the basement. Here, house noise is strong, and radio stations are weak due to the sense antenna being below ground level. You might then place the main antenna in the attic on the opposite side of the house, and located away from any power outlet wiring or ceiling lamp fixture on the floor below. While this strategy can improve reception immensely, it often won't completely eliminate local noise.

The phaser has two variable gain preamps or variable attenuators to connect two antennas. One antenna then goes through a circuit that lets you adjust the phase of its signal on a continuously variable basis. Then a hybrid signal combiner joins the two signals and outputs them to a radio.

A use case for two antennas connected through a phaser:

If you tune into a graveyard channel at night in AM mode, you are likely to hear just a rumble. There will be many carriers that have their sidebands below the noise floor. The frequency of each carrier is slightly different. Each carrier is subject to selective fading, so its amplitude is always changing. This selective fading is different for each and every carrier. Therefore the dominant carrier is always changing. A radio sees this continuously changing dominant carrier frequency as random low frequency modulation, and therefore it creates an audio rumble. When the

AGC responds to the dominant carrier the moment it is dominant, it pushes down the gain. This pumping the gain (amplitude) up and down adds to the modulation effect of the combined carriers. This gain pumping up and down IS amplitude modulation.

Due to selective fading, sometimes the sidebands of one or more stations rise above the noise floor and you will hear the broadcast audio from one or more stations for a moment. A moment later audio may be from another dominant station.

Sometimes a newly dominant carrier will push down the one you're trying to listen to at exactly the moment they're broadcasting the station ID.

Two antennas combined through a phaser will concentrate reception from one direction while lowering down the amplitude from other directions. This can clear up reception enough to make the difference between a hopeless jumble, and logging new stations.

Hybrids:

Do not confuse the "phaser" phase combiner with a different component sometimes also called a phase combiner. That other component is a signal splitter/combiner, also sometimes known as a "hybrid". These hybrids are used to connect two sources to a load while isolating the two sources from each other. They're also used to connect two loads to one source while isolating the two loads from each other. A "phaser" likely contains a "hybrid" at its output. One might also use a hybrid to connect one antenna to two radios.

A rudimentary hybrid can be made out of a center tapped autotransformer with a resistor tied across the two outside poles. A signal between one pole and the center tap will induce an inverted signal between the center tap and the opposite pole. The resistor (ohms equal to twice the impedance) isolates the two poles from each other because the current from an opposite pole has equal amplitude and opposite polarity from that of the current through the resistor.

The center tap is actually unnecessary because another winding provides a third port to carry the combined signal to or from both ports. This third port is called "common". The other ports are called 1 and 2. Some hybrids keep the center tap and place the resistor between the center tap and ground instead between opposite ports. It does the same thing, except the resistor value is now half that of the system impedance instead of double.

Hybrids have 3 dB of loss while isolating ports 1 and 2 from each other. These hybrids work best if all connections are impedance matched. Before you go thinking that cutting the power in half is inefficient, consider that as a divider, half of the power input into the common port comes out port 1. The other half comes out port 2.

This is the same device as a cable tv splitter, but a cable tv splitter will not work (at all) for us because it's designed for VHF and higher.

FSL's

The reason why a longer FSL can provide better selectivity is because of the length

to diameter ratio, and because a longer sleeve permits greater spacing of wire turns for lower proximity effect between turns. The larger length to diameter ratio provides a greater effective permeability, therefore requiring fewer wire turns to attain the required inductance. This helps offset the greater number of turns needed on a smaller diameter sleeve to attain the same inductance.

Effective permeability is a fraction of the actual magnetic permeability of ferrite material. When the initial permeability of the ferrite is above 100, a ballpark guess for effective permeability for an FSL may realistically be 4 or 5. The reason why the effective permeability of the FSL core is so much less than the ferrite itself, is because the return path for magnetic flux, from one antenna end to the other, is through the air. Much of the effective permeability depends on the dimensions of the ferrite material. Effective permeability improves with ferrite length. This also means that a large increase in initial permeability of the ferrite material itself (by switching to higher permeability ferrite) results in a small increase in effective permeability. In other words, the initial permeability (μ') of the ferrite, which contributes to inductance, is less important than someone might think. The same is not true for the secondary permeability (μ''), as secondary permeability is resistive, and therefore lossy. You want ferrite with a high ratio between μ' and μ'' . Ferrite mix 61 is the best material for ferrite antennas for this reason, even though the initial permeability is only 125. Mix 61 is also good for winding low loss high Q inductors.

Ferrite with high secondary permeability μ'' is not useful for antennas. These ferrites are most useful for things like choking off RF noise on power wiring or suppressing common mode noise on coaxial cables. Most high μ'' ferrites are toroids, beads, or sleeves, not rods. A coil wound on material with high secondary permeability behaves somewhat like an RF resistor (or more accurately, an RF transformer coupled to a resistor). In this case, power is dissipated when eddy currents are induced into the ferrite material. An antenna wound on ferrite with a μ'' being a significant proportion of total permeability will have a very low Q and most of the RF picked up will be dissipated in the ferrite. This secondary permeability damps resonance much like how draping a piece of cloth over a bell will dull its ring.

In a coil wound around ferrite, most of the magnetic flux in the ferrite is concentrated in the first few millimeters below the surface. What this means is that a ferrite sleeve is about as effective as a solid slug of ferrite but weighs (and costs) only a fraction as much. In other words, a solid cylinder of ferrite has about the same effective permeability as a hollow cylinder.

Turn-to-turn proximity effect raises the amount of RF resistance, therefore degrading Q and selectivity

To reduce proximity effect, any gap between turns is better than none; the optimum turns spacing to reduce proximity effect is about 6 wire diameters per turn. This is one reason to make a long ferrite sleeve instead of a fat one. It provides more space to spread out your coil turns.

Increasing the ferrite diameter has about the same effect on signal amplitude as increasing the length because doing either increases the amount of ferrite. However,

with either length or diameter increase, further increase has a diminishing contribution. There will be a point beyond which no additional increase in dimensions will yield a notable increase in amplitude contributed by ferrite. What you're really doing when changing the length to diameter ratio of the ferrite, is manipulating the effective permeability.

At large diameters with a very small ferrite length to diameter ratio, the antenna starts behaving almost as if it were air wound without ferrite. In other words, with decreasing length to diameter ratio, the effective permeability approaches 1 - the permeability of air. With increasing length, an incremental increase becomes less and less in proportion to the total length. Eventually, each incremental increase in length becomes insignificant because the increase in effective permeability approach 0.

Remember that received signal amplitude and sensitivity are related but not the same thing. Once the received background noise is raised above the noise floor of the radio, further amplitude increases become detrimental to the dynamic range of the radio. Therefore, increasing the size of the ferrite sleeve antenna can be detrimental.

Ferrite rods and bars are a sintered ceramic material. Therefore, hard and brittle. If one breaks, it can be superglued back together without affecting its performance. All rods are imperfect - most bend and shrink a little during the sinter process, causing variance from rod to rod. What that means is if you designed your antenna to exact dimensions using an exact number of rods, you will likely be unable to fit the last rod into place. Therefore, designs need to accomodate this issue.

Hollow construction, such as a self-supporting ferrite sleeve {brittle but least heavy), or sleeve built on a rigid nonconductive cylinder, provides more versitility in design choices. Good materials include sturdy cardboard tubes of sufficient diameter (absorbs moisture), big PVC pipe (can be heavy), and rigid styrofoam such as that used in exercise rollers. Rigid styrofoam can be hollowed through the use of a hole saw and a tool to hack out the foam between borings of the hole saw. Styrofoam is much lighter in weight, but still offers protection to the brittle ferrite. When ferrite is mounted on a rigid backing, the backing provides a great deal of structural strength to the brittle ferrite. Ferrite rods or bars can also be rubber banded onto a temporary form and glued together using epoxy or MMA adhesive, to make a rigid, self-supporting ferrite sleeve. Self supporting ferrite sleeves can be built leaving a small lengthwise gap between two rods. This relieves the problem of fitting an exact number of imperfectly straight rods. It also makes securing the coil ends much easier. Since the rods will have varying lengths, make sure they sit flat on one end of the self supporting sleeve to provide a mounting surface. If of sufficient diameter as discussed above, the capacitor can be mounted inside of the sleeve or onto a wooden disk attached to one end.

Another housing option is to glue two lids from plastic jugs together, lid to lid, using MMA low surface energy adhesive formulated specifically for PET polyester (recycle number 1). Drill some holes through the the two joined lids. Two of the holes will be to pass wires through. The rest of the holes will serve as reinforcement when filled with adhesive. Using MMA adhesive, mount the ferrite

inside one lid to be inserted into one jug. The ferrite jug will need some means to support the opposite end of the ferrite, such as foam tape around the inner sides near the bottom. Mount the capacitor inside the other lid using MMA adhesive. The drilled holes in the lid will reinforce the mounting, as will any holes in the back of the capacitor frame. You might need to drill a 1/2 inch relief hole in the center of the lids for the back end of the capacitor's rotor shaft. Mount a vernier dial to the bottom of the capacitor's jar, connecting the vernier to the capacitor using a shaft extension. The plastic jars provide substantial protection to the fragile antenna. If the FSL is going to be regenerative, a third jar can be affixed to the opposite end to house the electronics.

Large sugar-free Metamucil powder jars make excellent lid-to-lid housings for small diameter FSL's. These Metamucil jars share the same diameter and use the same wide lid as several shorter dietary supplement jars. These jars can provide excellent shock and water protection.

Some 2 Lb protein powder jugs mate perfectly with six inch diameter exercise rollers that are 1 foot long. The jugs can house the capacitor and vernier, while the roller supports the ferrite and coil. The coil ends can be brought through holes in the styrofoam to enter the plastic jug. Closed cell foam tape can bridge the gap between the roller and the jug if the jug is the same diameter as the roller. Cleanly cut the bottom off of the jar, coat the cut edge with contact adhesive and press the jar into place into gap between roller and tape. Then put a bead of silicone RTV around the inside seam between the jug and the roller. If a regenerative FSL, another jar can mounted on the opposite side of the roller to house the electronics and batteries. The foam tape also serves another purpose. It serves as a bumper to help protect the exposed Litz wire coil.

Choose your construction adhesives according to the duty they must do, as some glues don't work well, or at all, for some applications. Good construction adhesives include 5-minute epoxy (strong, brittle, but less brittle than ferrite, does not adhere to some plastics), Barge or Master contact cement (strong, very flexible, shear strength is better than tensile strength, which is fair to poor on some plastics, not gap filling, strong chemical smell until cured). The best for constructing FSL's is an expensive 2-part adhesive called MMA. Some "low surface energy" MMA adhesives are the only thing that work well on polyester, polypropylene, PVC or ABS plastics. It has an extremely strong chemical smell until cured. Some low surface energy MMA formulations take several hours to fully cure. Others take several minutes. MMA bonds need little surface prep. Bond strength may be greater than the bonded substrates, and bond flexibility is between that of epoxy and contact cement. MMA is commonly used to bond ferrite magnets in loudspeaker manufacture because of its strong, but compliant properties. Cyanoacrylate "super" glue of the thick, gap filling variety is excellent for tacking things together to hold them in place until a better adhesive cures. It, too, does not adhere to some of the the plastics mentioned above. Cyanocrylate accelerator fluid is a must for quick assembly.

Cyanoacrylate applied with accelerator works well on wood and wood composites, and are usually stronger than the wood, but can sometimes be brittle, so joints should be reinforced by nonmagnetic screws. Wood as a construction material contains

moisture, making it slightly conductive, and can introduce a little bit of parasitic capacitance when building a wooden box to house a self-supporting ferrite sleeve. This is usually inconsequential, but will likely be noticed when closing the lid on a regenerative FSL - the tuned frequency may shift 50 Hz or so. Beware that some black plastics and black neoprene rubbers and foam tapes contain carbon pigment which can introduce a surprising amount of parasitic capacitance, enough to completely ruin an antenna, so don't wind your coil on it. Some materials, to a lesser degree, can introduce some parasitic capacitance by having a higher dielectric constant than air. Therefore any layers of material used to separate the coil from the ferrite should be (clear or white) closed-cell foam. The only thing better than this foam might be bamboo or thin wooden strips laid in, and protruding above, the gaps between rods, and epoxied in place. This will suspend the coil above the ferrite with an air gap between the coil and the ferrite. Bamboo strips may not need to be laid between every rod; every second rod may suffice. "Stripes" of 5-minute epoxy laid perpendicular to the turns works well to rigidly anchor the coil to the bamboo strips.

Software defined radios.

A word to those of us who avoid SDR's thinking they take the challenge away: They don't. The challenge is the same. Switching to use an SDR is like getting a better antenna, in the sense that a better antenna has capabilities that a lesser antenna lacks. Tuned loop antennas, including FSL's, negate the wideband spectrum analysis features of SDR's by limiting the bandwidth passed to the radio. On the other hand, the wideband spectrum display makes FSL tuning visually obvious, and the channel spectrum display lets you see fine tuning and feedback Q adjustments in addition to the carrier and modulation.

A good SDR's raw performance and configurable filters and noise reduction allow you to pull signals out from deeper down in the difficult fringe - but that difficult fringe is always there to challenge you regardless of your radio or antenna. The raw performance of most SDR's is comparable to some of the best tabletop communications receivers ever made. Some SDR's are even better. Plus, SDR's have features that no tabletop radio has, and these features are actually useful to the DX'er - unlike the disutility of gimmic features many radios have, or that squelch button you keep accidentally activating.

In my opinion, feeling that using an SDR is cheating, is like feeling that sticking a tiny portable radio next to an FSL is cheating. While I don't actually consider either cheating, the FSL effectively pushes the front-end of the radio out of the radio and into the FSL. Your dx'ing takes advantage of capabilities of the FSL that the radio lacks. In a similar semantic way, an SDR has capabilities that a traditional radio lacks.

This brings us back to how important a good antenna is. If you have to choose between investing a better radio or a better antenna, then in most cases improving the antenna is the better investment.

I have been a MW DX enthusiast since childhood, focusing more on the technical side of the hobby, and have been experimenting with radios and antennas for over 50

years. In 2003 I came into possession of several boxes of leftover prototype mix 61 ferrite rods from an engineering project. Tasked with looking for ways to improve IBOC "HD" reception, I experimented with trying to improve sensitivity and selectivity by creating overcoupled "Batman mask" double peaked high selectivity antennas contoured to the IBOC broadcast bandwidth. These experiments were part of my job. I was able to get the selectivity I was aiming for, but it hurt instead of helped HD reception, so that avenue was abandoned. Although the antenna experiments were a failure, I was astounded by their dx capability when critically coupled. I wound my coils on cylinders of ferrite rods laid side by side, and then wrapped around cardboard concrete formers, sturdy shipping tubes, and PVC pipes. I discovered these ferrite cylinders performed just like box loop antennas many times larger. I also experimented with length vs. width configurations of the ferrite cylinders and confirmed that sensitivity correlates with the quantity of ferrite regardless of length or width. However, increasing length enabled slightly sharper selectivity, while a larger diameter yielded slightly narrower and deeper directional nulls. These differences were almost insignificant and it was difficult to determine if differences were because of sample-to-sample variances or due to configuration. Unfortunately none of the data and technical details were allowed to leave the lab. It wasn't until 2020 when I joined several MW dx groups online, that I learned that others had discovered and popularized this same antenna concept for dx around the same time I was experimenting. Aside from one report that was pseudoscience, their experiments independently corroborated many of the findings from my own, and vice versa. This kind of antenna became known as the Ferrite Sleeve Loop or "FSL" antenna. With pandemic time on my hands, I recreated some of my old experiments at home. Much of this essay has been an effort share things I have learned from all this.