BFO Theory

The beat-frequency oscillator, or BFO, is generally the simplest type of metal detector and is a good starting point for learning how metal detectors work. The dominant technology throughout the 60’s and much of the 70’s, the BFO has now practically vanished from the market and can only be found in the poorest quality toy detectors. Although it has run its course commercially, the BFO can still make a pretty good detector and there are a few tricks that can make it an excellent detector under some conditions. We will not investigate these “tricks” in this paper but will focus on the fundamentals.

Basic System

The basic BFO employs two radio frequency (RF) oscillators which are tuned to very nearly the same frequency. The frequency of each oscillator is usually determined by an L-C resonant circuit, that is, an inductor and a capacitor. One is called the *search* oscillator and uses the search coil as its inductor; the other is called the *reference* oscillator and uses an internal inductor. The outputs of the two oscillators are fed into a mixer which produces a signal that contains the sum and difference frequency components (plus other intermodulation components) of the two input signals. As we will see shortly, the difference frequency is what we are interested in, and a low-pass filter at the output of the mixer will remove the other components. Figure 1 shows a block diagram of the oscillator/mixer setup.

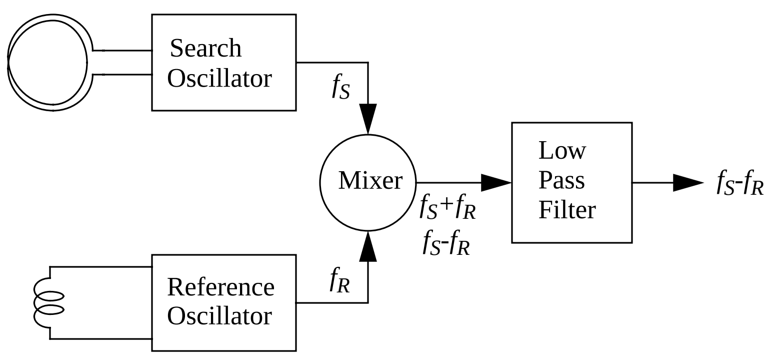

FIGURE 1. Basic Oscillator/Mixer

As long as the two oscillators are exactly the same frequency, the mixer-filter output will have no difference signal, just DC. If the frequency of the search oscillator shifts slightly, then a frequency difference signal will appear at the mixer-filter output. The frequencies of the two oscillators are normally chosen such that a shift in the search oscillator frequency will produce a mixer output signal that is in the audio frequency range. For example, if the reference oscillator is set to 100KHz and the search oscillator is at 100.5KHz, then the difference will be a 500Hz signal which is on the moderately low side of the audible range. The sum frequency will be 200.5KHz which is of no use, so the low-pass filter removes this component (plus the intermods). Figure 2 shows the mixer signals, and henceforth we will consider the mixer signal as being only the low-pass (difference) component.

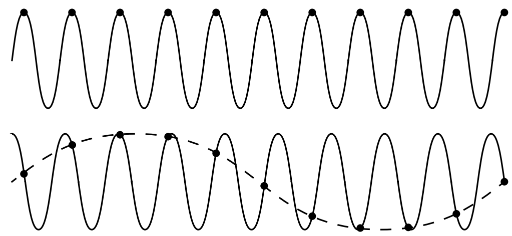

FIGURE 2. **Mixer signals**

Operation

Because the search oscillator uses the search coil as one of its components, its frequency will be sensitive to the inductance of the coil. The inductance of the search coil is, in turn, dependent on nearby conductive objects. The objects proximity, size, and type of metal determines how much inductance shift the search coil experiences. Two different mechanisms result in ferrous[1] metals causing the inductance to increase and non-ferrous metals causing a decrease, thereby producing a discrimination effect. Before we consider these two mechanisms we should first take a brief look at inductors.

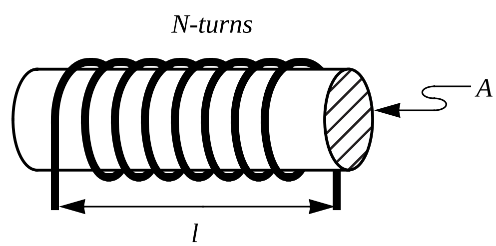

FIGURE 3. **Solenoidal Inductor**

The inductance of a simple solenoidal coil (Figure 3) is generally given as

$$L = \frac{\mu N^2 A}{l} \tag{EQ 1}$$

where N is the number of turns, A is the cross-sectional area of the coil, l is the length of the coil, and μ is the permeability of the core, or center, of the coil. Permeability is the characteristic of a material that describes its ability to "conduct" magnetic flux. In the case of the search coil, the core material is just air which has a permeability value of $400\pi \times 10^{-9}$henries/meter (or 400π nH/m). A material with a higher permeability, such as iron, will cause the inductance to be higher. Air (actually, vacuum) is often used as a reference permeability and is given the symbol μ_0; other materials are then given a relative permeability term called μ_r. Iron, for example, has a relative permeability of about 5000. Thus the permeability can be stated as

1. While there is a difference between a ferrous and a ferric target, we will use the term *ferrous* to include both.

$$\mu = \mu_0 \mu_r \quad \text{(EQ 2)}$$

An air core inductor can have its value substantially increased by adding an iron core. We shall see how this affects the operation of the BFO in the coming paragraphs.

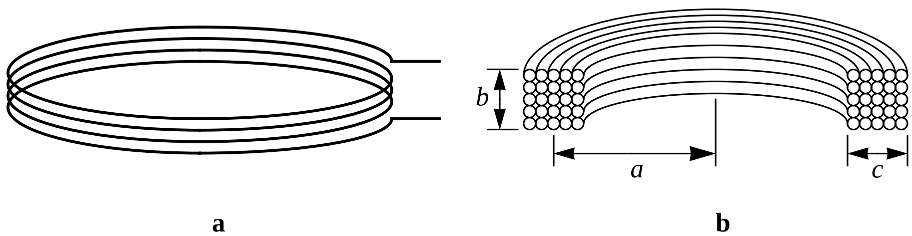

FIGURE 4. Multi-turn Coil

Metal detectors rarely use a solenoidal coil. Rather, they typically use some sort of multi-turn (scramble-wound) loop coil as shown in Figure 4a. The analysis for this type of coil is not very straightforward because the proximity capacitances of adjacent wires causes non-uniform current distributions. If the coil has the dimensions as shown in Figure 4b then the inductance can be approximated by

$$L = \frac{0.8(NA)^2}{6a + 9b + 10c} \quad \text{(EQ 3)}$$

When a current flows through an inductor a magnetic field is developed; if the current is DC the magnetic field is static, if the current is AC the magnetic field is dynamic and is coupled to a dynamic electric field which together comprise an electromagnetic field. Magnetic fields are usually represented graphically by their lines of magnetic flux and the field for a solenoidal inductor is shown in Figure 5. The field for a typical

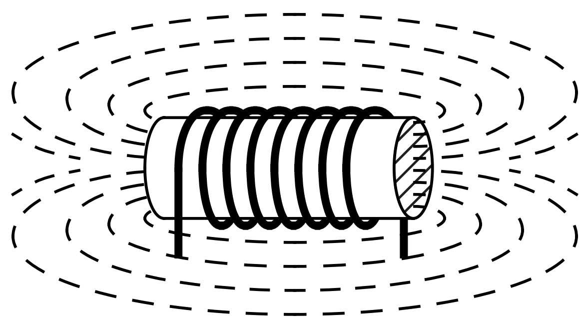

FIGURE 5. Solenoidal magnetic field

loop coil is similar and is shown in Figure 6. These diagrams are not to scale. The flux density - represented by the spacing of the flux lines - tends to be much weaker outside of the windings. The total magnetic flux is related to the coil's inductance and the current flow though it; the magnetic flux density at a given point in space is related to the total flux, the geometry of the coil (circular vs. rectangular, for example), and the position at which it is measured.

In the case of small ferrous objects, the inductance of the search coil will increase because of the higher permeability presented by the target. An increase in inductance will result in a decrease in the frequency on

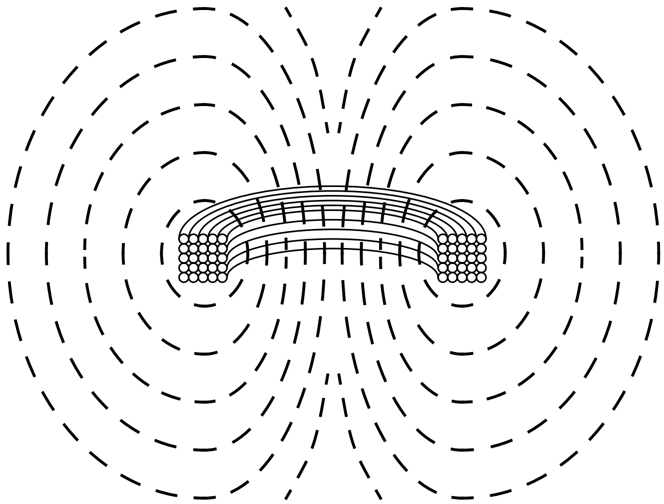

FIGURE 6. **Loop coil magnetic field**

the search oscillator and a corresponding decrease in the frequency output of the mixer. A non-ferrous object does not affect the permeability seen by the coil, so how is the inductance affected? The answer lies in a phenomena called *eddy currents*. When a metal object is subjected to an electromagnetic field a current is induced. This current will, in turn, produce its own magnetic field which is in opposition to the main field. This small, opposing field affects the overall field of the detector's coil and produces an apparent decrease in the inductance. This will increase the frequency of oscillation as well as the mixer output. Figure 7 attempts to illustrate this concept.

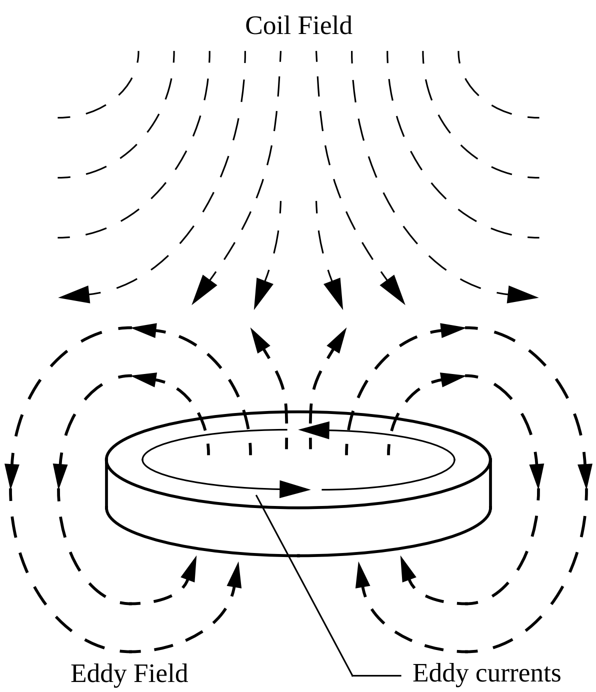

FIGURE 7. **Eddy currents**

It was noted that *small* ferrous objects offer a higher permeability which increases the inductance. Large ferrous objects also present a higher permeability, but they additionally generate a substantial eddy current effect which can override the permeability effect. Thus a large ferrous object can (and often does) give the same response as a non-ferrous object.

Problems

We have looked at the effects of desirable targets, but there are also some undesirable issues that we must deal with. The first is that of ground mineralization. Most soils contain some trace amounts of conductive minerals including iron oxide and salts. This mineralization alters the inductance of the search coil and therefore shifts the search frequency. The end result is a mixer frequency which varies with search coil height. Ground mineralization tends to be pretty consistent over a small search area, so the immediate solution is to retune the detector after the search coil is lowered. Unfortunately, any height variation of the coil as it is being swept will result in false signals. There is no easy solution for this and it was a major problem of detectors in general until ground balanced IB designs came along.

The search coil windings are also sensitive to some undesirable effects. In Figure 1 we saw that the search oscillator frequency is set by the search coil inductor and a capacitor. The value of this capacitor is mostly determined by a physical device, but there is also a parasitic term to it. The parasitic portion comes from distributed capacitances in other circuit devices, the search coil cable, and the search coil windings. We are not so much concerned with the existence of these parasitics, but with their stability. The search coil windings usually present the greatest problem so we will focus on that.

In the BFO, the search coil is made of multiple windings of thin wire that has an even thinner coating of insulation. Two physically adjacent wires will therefore have a parasitic capacitance between them. In all, the coil will look like an extremely complex distributed LC network[2] for which Figure 8 does not even come close to an adequate approximation. Again, we are not too concerned with the existence of these parasitics as long as they do not change. Because the parasitic capacitances are determined by the physical proximity of wires, it is crucial that these wires **not move** during normal operation. Therefore, special care must be taken when making the search coil to ensure rigidity.

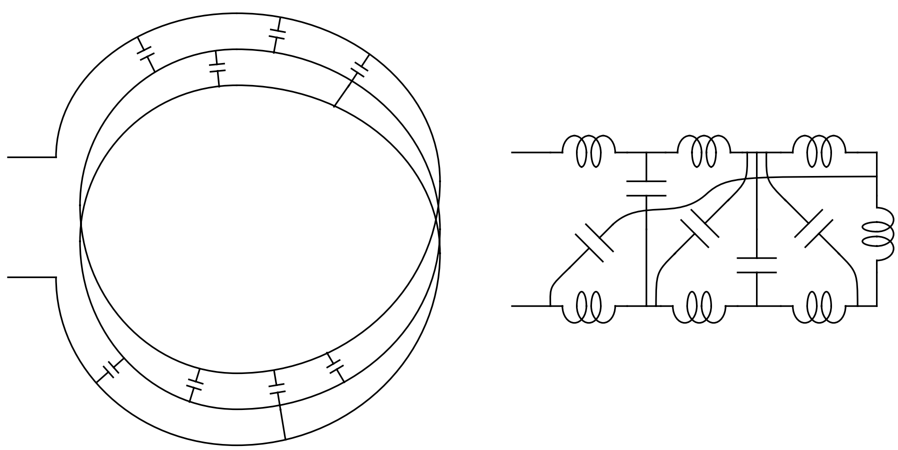

FIGURE 8. Coil parasitics

2. There is also a distributed resistance term that is not being considered at the moment.

Another related problem is ground capacity which is a proximity effect that also alters the total capacitance. It is mostly a nuisance in wet grass. As the search coil is lowered to the ground, the total capacitance “seen” by the search oscillator is slightly altered by the distributed parallel coil-to-ground capacitance as shown in Figure 9. This change in capacitance will alter the search frequency as well so that the mixer output frequency will again vary with the height of the coil above the ground. This effect can be minimized with a Faraday shield which is a conductive (aluminum) foil wrapped almost completely around the loop and connected to the circuit ground. It presents a uniform coil capacitance to the oscillator which the ground capacitance does not affect. It will not affect the magnetic field of the coil as long as the shield does not completely enclose the full circumference of the coil. If it does, then it will present a short circuit for eddy currents and severely attenuate the magnetic field; thus a gap is left in the shield.

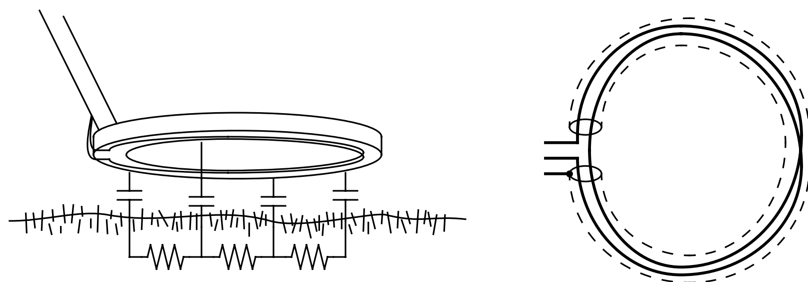

FIGURE 9. Ground capacitance & Faraday shielding

Sensitivities

The response of the BFO to targets and parasitics depends on one other parameter: sensitivity. Sensitivity is basically what it takes to get the search oscillator frequency to shift by a certain amount. From the preceding discussions we can see that it is desirable to have the oscillator very sensitive to a change in inductance but insensitive to changes in capacitance[3]. The ideal L-C oscillator has a frequency that is

$$f_0 = \frac{1}{2\pi\sqrt{L \cdot C}} \quad \text{(EQ 4)}$$

where C is the total capacitance, including parasitics. The frequency sensitivity due to *small* changes in L and C will be

$$\frac{\Delta f}{f_0} \cong \sqrt{\frac{\Delta L}{L} \cdot \frac{\Delta C}{C}} \quad \text{(EQ 5)}$$

What we see from this is that a small L has more sensitivity to ΔL and a large C has less sensitivity to ΔC. This will give us the best response to a target and the least response to proximity effects. As always, there are trade-offs in other areas.

A small inductor value will produce a weaker magnetic field than a large inductor for a given current. In fact, for a simple solenoidal coil the magnetic flux density (at the center of the coil) is

3. A proximity detector, on the other hand, would need a the opposite sensitivities.

$$B = \frac{\mu N I}{l} \qquad \text{(EQ 6)}$$

To compensate for this a larger current must be used with small inductance coils which will affect battery drain. Another effect is low value inductors with fewer turns also tend to have a lower quality factor, or Q. Q is equal to the resonance frequency divided by the resonance “bandwidth” and basically describes how “tight” the oscillation is (See Figure 10). High Q circuits have more sensitivity, not only to targets but also to other effects such as temperature, coil wire movement, and ground effect. Low Q circuits have less sensitivity but more stability. Therefore, we should try to attain the highest possible Q within the manageability of the parasitic effects.

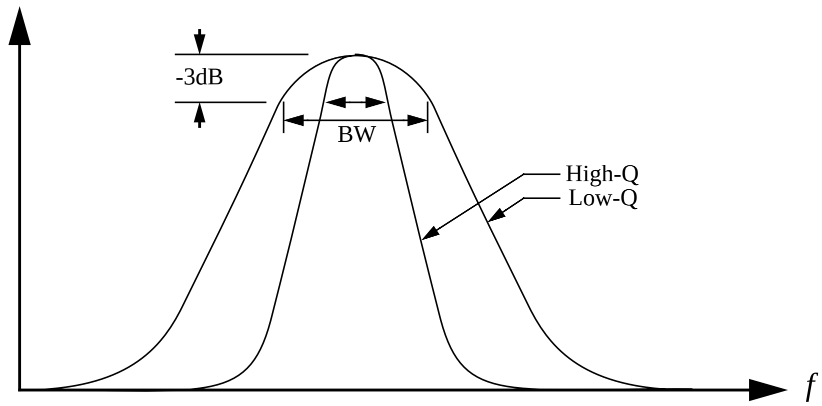

FIGURE 10. **Resonance Quality Factor**

As we have seen, BFO detectors use a very simple coil design unlike most induction balance units. Thus BFO’s can easily be outfitted with loops of almost any size, from 3 or 4 inches to a massive 36 inches or more. This is one of the reasons the BFO became known as an all-purpose detector. The smallest coils are very sensitive to small near-surface objects like gold nuggets, the mid-sized coils are normally used for coin hunting, and the largest coils were effective on large, deep objects like caches and relics. The only consideration in making different coil sizes is to match the inductance of each; otherwise, you will have to swap out the oscillator capacitor every time you change coils. An alternative is to place a parallel capacitor inside each coil with a proper value such that the total oscillator L-C remains constant.

Frequency Considerations

There are two frequencies of concern with the BFO: that of the search oscillator and that of the mixer output. We have already briefly looked at a couple of issues with the mixer frequency. To use this frequency as an audible indicator it must clearly fall in the audio range. But what is an optimum audio frequency, and why?

From Equation 5 we saw that the relative change in the search oscillator frequency depends on the relative change in the inductance of the coil. Our prior examples dealt with a frequency that was around 100KHz, and whether you tune the search oscillator to 100.5KHz or 100.05KHz a given target will still produce about the same amount of absolute frequency shift for each case. Let’s again assume that the reference frequency is 100KHz and that the target creates a +10Hz shift. For a search frequency of 100.5KHz, the nominal mixer frequency is 500Hz and the 10 Hz target shift will increase this to 510Hz. If the search frequency is 100.05KHz then the nominal mixer output will be at 50Hz and the 10 Hz target shift will increase this to 60Hz. The latter case results in a larger *relative* frequency shift at the mixer which is much easier for the ear to discern. Thus a lower mixer frequency is better, but keep in mind that the response of the ear drops off rapidly below 50Hz. However, by using a square wave as the audio output we can actually set the mixer output frequency *below* the audio range, 10Hz for example, and still hear the harmonic content. This has the additional benefit of producing larger relative frequency shifts at the harmonics as shown in Figure 11. Thus the higher harmonics that fall into the audible range will exhibit greater frequency shifts which has the effect of increasing sensitivity.

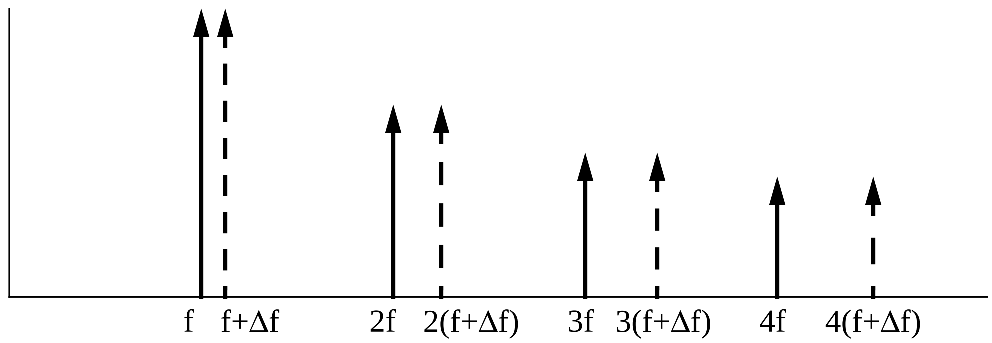

FIGURE 11. Harmonic Frequency Shifts

There is one drawback to setting the mixer output to such a low frequency. Suppose the mixer is set to 500Hz and a small ferrous object causes a 100Hz shift. The mixer output will then decrease from 500Hz to 400Hz. If the mixer is set to 100Hz then the target will cause a decrease to 0Hz, or DC, which makes the audio silent. At a lower mixer frequency (say 50Hz), the same ferrous object can actually cause the search oscillator frequency to drop below the reference oscillator frequency. When this happens, the mixer output frequency will decrease to DC, then roll over by 180 degrees and begin increasing. So when the nominal mixer operating frequency is too low, it is possible for even a small ferrous target to appear non-ferrous. However, you can usually audibly recognize when this situation occurs.

The final issue to consider is that of the search coil frequency. So far our examples have used 100KHz for two reasons: it's an easy number to use for the examples, and it is also a frequency that is often used. But is it the best frequency, and what factors determine a good frequency?

There are two primary factors that determine a good search coil frequency and they are, unfortunately, opposing. Lower frequencies penetrate the soil better than high frequencies and are better able to deal with mineralization which is why modern IB detectors typically use frequencies in the VLF (3-30 KHz) range. Moderately higher frequencies generate a better eddy current response up to the point where skin effect begins to degrade it. The choice of a moderately high frequency offers a balance between good depth penetration and good response. But with BFO's, the real reason for a higher oscillator frequency is that it gives better frequency sensitivity at the mixer output. For example, the same target that gave a 10Hz shift at 100KHz in the example above might only give a 1Hz shift if the search oscillator is operated at 10KHz. Thus if you ran the reference oscillator at 10KHz and the search oscillator at 10.05KHz, then the resulting mixer shift will be from 50Hz to 51Hz which is difficult to discern audibly. As before, there are ways to design BFO circuits such that VLF search frequencies can be used with a good audio response.

Potential Pitfalls

So far we have looked at operational theory and design considerations, mostly under ideal conditions. When you actually build a BFO circuit there are additional concerns that deal with non-idealities. The BFO relies on two non-identical oscillators to generate a stable (in the absence of a target) beat frequency, therefore it is important that both oscillators be stable. We have already looked at the effect of ground capacitance as an example of a source of instability.

Another important design factor with BFOs is that of thermal stability, i.e., the effects of temperature changes on the circuit. Practically all electronic components have a thermal sensitivity so as they heat up or cool down, their value or some other parameter changes. For example, transistor V_{BE} decreases about 2mV for each degree (Celsius) of temperature increase. Capacitors and resistors are also usually specified with a temperature coefficient. In the BFO, the main issue with thermal changes is that of the reference and search oscillator frequencies. As long as they both exhibit the same frequency variation with temperature then the mixer output will be fairly stable. But the oscillators are not entirely identical as one has a local inductor and the other a remote search coil, and one has some sort of frequency adjustment component. Thus you should take considerable steps towards making the detector thermally stable. This primarily involves matching the oscillators as closely as possible and avoiding dark, heat absorbing colors for the case and search coil.

It is possible for two oscillators that are in close proximity and running at very close to the same frequency to exhibit frequency lock. When this occurs both oscillators will run at exactly the same frequency and the mixer output will be at DC. For the BFO detector, it can take a fair amount of target signal to shift the search oscillator off of the locked frequency, therefore it is always advised to never run a BFO with a zero-beat nominal frequency as it can reduce target sensitivity. A low frequency beat, sometimes called “motor-boating,” will prevent frequency lock and give good audible sensitivity to small delta frequency changes.

There are some design practices that can help minimize frequency lock, or at least the difference frequency where it occurs (i.e., 1Hz vs. 10Hz). First, do not place the oscillator circuits in immediate proximity on the circuit board. The reference oscillator has a local inductor which generates an electromagnetic field, and this field can couple into the search oscillator and alter its frequency. Capacitors generate these fields as well. Oscillators also produce AC currents in the power supply and/or ground which can produce voltage fluctuations on these lines if their impedances are not zero (and they are not). Good supply isolation or low-pass filtering will help.

Conclusion

There seems to be a lot of issues to consider in a BFO design, but compared to the induction-balance detector it is very simple. The simplicity of the coil makes it easy to try out different shapes, sizes, and windings to compare sensitivities. The oscillator frequencies can be varied for even more comparisons. Such information will often carry over to more advanced designs. This makes the BFO an excellent first project that provides a solid foundation for learning.