

by Carl Moreland
Revision 0.5 — 11 Mar. 2003

The complete circuit is shown in Fig. 1 and, at first glance, appears to be complicated because of all the options that were designed in. However, the design options really don't add much complexity, and allow for different configurations and performance comparisons. Some of the options include:

- Passive or active MOSFET turn-off
- Inverting or non-inverting preamp configuration
- Preamp compensation support for NE5534, LM318, and others
- Single-ended or differential integrator
- VCO or non-VCO audio

- Transmit pulse rate
- Transmit pulse width
- Integrator sample delay
- Integrator sample pulse width
- Secondary integrator sample delay
- Threshold
- Sensitivity
- Autotrack/autotrack speed
- Volume

Circuit

Power supply

Power for PI-1 is provided by a single 12-volt battery pack, preferably an 8-pack of AA's. A 9-volt pack can be used with a reduction in battery life, and even a 9-volt transistor-radio-type battery can be used, with a significant reduction in battery life, providing a proper coil tank is used (I'll discuss this later).

The power supply circuitry (Fig. 3) consists of a voltage doubler, and 3 regulators. The voltage doubler, built around a 7660 supply chip (IC1), generates a voltage that is roughly twice the battery voltage. At this point, the easiest way to visualize the supply voltages is to consider the positive side of the





Fig. 1: PI-1 Schematic




battery to be “ground”, the negative side of the battery to be the “minus rail”, and the voltage doubler output to be the “plus rail”. The schematic is labeled as such, with the battery connected between ground and “-VB”.

A 7805 regulator (IC2) produces +5 volts from the output of the voltage doubler, and is used as the positive supply for the analog opamps. There are two 7905 regulators, one for the -5 volt opamp (analog) supply (IC3), and one for a -5 volt digital supply (IC4). A sep-

arate digital regulator is used to help isolate digital noise from the analog path.

Clocking

The master clock generator is a 555 timer (IC5) connected as an astable



oscillator. Four different clocks are potentially needed, depending in the design options used. The main clock drives the coil switch Q3 which pulses the search coil. A second clock, delayed from the trailing edge of the main clock, controls the main sampling switch (IC7a) after the preamp (IC6). A third clock, delayed from the second one, controls an optional secondary sampling switch (IC7b), used for differential integration. Finally, a fourth clock can be used to gate the speaker for non-VCO audio. The delay of this pulse should be such that it gates the speaker sometime after the receiver sampling is done. Otherwise, current pulses created by the speaker can disrupt the supplies and distort the coil decay. Fig. 4 shows a timing diagram.

The 555 oscillator has two controls. R2 varies the frequency of the main pulse (and all derived pulses), and R3 varies the pulse width of the main pulse only (TP1). There is a slight amount of frequency variation with pulse width adjustment. The frequency of the main pulse determines how many times per second the coil is pulsed, and also how many samples per second the integrator (IC8a) is integrating. Generally, a higher pulse frequency allows for a faster coil motion and perhaps a little better noise performance, while consuming more power. Varying the main pulse width determines how long the coil is turned on, which can affect depth of detection and power consumption.

The frequency range for this design is 900Hz to 2300Hz, and the pulse width can be adjusted from 10 μ s to 75 μ s. R1-R4 can be modified for other frequency and pulse width ranges. Note that these

ranges, and subsequent ones, are subject to variation due to component tolerances. The numbers presented here are measured from a single build and may not be representative of nominal ranges.

The remaining pulses are generated by dual monostable multivibrators IC11 and IC12. IC11 creates the first delayed pulse for the main sampling switch (TP6), and IC12 creates an additional delay for the secondary sampling pulse (TP7). In each case, the first half of the multivibrator chip is used to delay the input pulse, and the second half sets the output pulse width. The delays, and the pulse widths, are set by RC time constants. The R-component consists of a potentiometer for adjustability, plus a fixed resistor to set a minimum delay & pulse width.

For the differential integrator to work properly, both sampling pulses need to have the same pulse width, so both pulse width monostables are controlled by the same RC time constant, R46/R46a/C26, by using isolation diodes D6


and D7. This scheme works nicely, as long as the two pulses have a sufficient amount of delay time between them, which is easily met in this application.



Thus, potentiometer R42 sets the main sampling delay, which is variable from 10–11 μ s to 75 μ s. The secondary sampling delay tracks the delay of the main sample, with an additional delay set by R45, variable from 75 μ s to 350 μ s. R46 adjusts both pulse widths from 3 μ s to 50 μ s. Note that extreme limits of various settings may not be compatible. For example, at the highest frequency (2300Hz), a combination of high pulse width, large sample delay, and maximum secondary delay might place the secondary sample pulse into the next transmit pulse. There is no danger in this, but it may result in bizarre behavior.

Transistor level shifters Q8 and Q9 convert the 0/-5 volt logic levels of the monostable outputs, into +5/-5 volt levels needed to drive the sampling switches IC7a and IC7b. Transistor Q10 is also a level shifter, and drives the sampling switch IC7c, which is used for the non-VCO audio option. If the VCO option is used, Q10, R49, and R50 are not needed.

Transmit/Receiver Front-end

The main device in a PI detector is the coil switch, usually a MOSFET or bipolar transistor with a high current capacity and a high breakdown voltage. The coil switch serves two functions: to short the coil across the battery voltage, which creates a large coil current, and to turn the coil current off as quickly as





possible. During the “on” time, current flowing through the coil creates a magnetic field around the coil which, when static, does nothing. But when the coil is switched off, the current ceases, and so does the magnetic field. It is the collapse of the magnetic field that induces eddy currents in a nearby metal target which, in turn, generates a counter-magnetic field. The faster the magnetic field is collapsed, the greater the target field that is generated.

The output of the 555 oscillator drives buffer/level-shifter transistor Q1, which primarily converts the 0/-5 volt levels of the 555 into 0/-VB levels. See Fig. 5. Transistor Q2 is simply a buffer, with low-value emitter resistor (R9) used as a fast pull-down for the main coil switch transistor Q3. R8 is an optional resistor for limiting the high voltage or peak current applied to Q3,

and might be useful if Q3 is a bipolar NPN. For the default setup, R8 should be shorted and the gate of Q3 connected to R9.

R7 provides an option in which Q2 can be used as an active pulldown device for Q3, instead of using R9. Fig. 6 shows this hookup. In a PI detector, the turn-off of the coil switch should be done quickly, and an active device pulling down the gate of Q3 will achieve this better than a resistor. Although R9 is probably sufficient, especially if its value is low enough, Q2 offers the experimenter the option of investigating active pulldown. If R9 is used, then different resistor values can be tried. Because the signal applied to R9 is a low duty cycle pulse, even a low value like 100Ω will not cause excessive power dissipation.

In the case of active pull-down, the control pulse to the gate of Q3 is inverted, so the 555 oscillator will have to be modified (via changes to R1-R4) to produce a “positive” pulse. I have not done this, so I do not currently have values for R1-R4. The value for R7 is a bit of a guess, and may need adjusting depending on the gate capacitance of Q3. Also, Q2 should really have a base resistor (like Q1) of about 10kΩ, which was not included in the PCB layout and will have to be kludged in.

Coil switch Q3 must have sufficient current capacity during the time it is shorting the coil, and it should also have a high breakdown voltage to allow for a high flyback voltage. NMOS device IRF740 is rated for 10 amps and 400 volts. Other devices can be substituted, including NPN’s and IG-BJT’s, with slight differences in performance.

R10 and C10 form an RC tank for providing the transient current of Q3. These are not absolutely necessary, but can reduce supply noise, especially when using a 9-volt transistor-radio-type battery. If not used, omit C10 and short R10. In use, C10 should be large enough to supply the needed charge for the time Q3 is on. This depends on three parameters: the gate pulse width, the average coil current, and the supply voltage. As an example, assume that the pulse width is 25μs, the coil current is 2A, and the supply is 10v. Furthermore, assume we want the capacitor voltage to ripple no more than 100mV. Capacitor current is given by the equation $i=C*dv/dt$, so $C=i*dt/dv$, or $2A*25μs/0.1v = 500μF$. To find the value of R10, we assume the remainder of the clock time is available to charge C10 back up, so we can set the $R10*C10$ time constant appropriately. But this is a circuit that need not be operated optimally, so we can just as easily oversize C10 to, say, 1000μF and set R10 equal to 1Ω.

Resistor R11 is the coil damping resistor and prevents ringing when Q3 turns off. Ringing will occur because of the L-C tank created by the inductance of the coil, and the total parasitic capacitance, including the self-capacitance of the coil, capacitance of the cable, and Q3. The value of R11 is dependent on




Fig. 8: Differential Integrator & Single-Ended Integrator

the coil, and is best determined experimentally. If the resistor value is too low, it will overdamp the coil and reduce sensitivity. In general, R11 will be in the range of 200-800 ohms, with 680Ω used for the coils described in this article.

The coil signal is applied to preamp IC6, which can be connected as an inverting opamp, or a non-inverting opamp. In inverting mode (Fig. 7, top), gain is set to 1000 by R13/R12. R14 is not important, and can even be shorted. Diodes D3 and D4 provide clamping to protect the opamp. Without them, the

coil pulse of 200-400 volts will absolutely destroy the opamp, so do not operate the circuit without them! IC6 can be configured as a non-inverting amp (Fig. 7, bottom), by applying the signal to pin 3 of IC6, and using R13/R14 for the gain path. In both configurations, C11 can be added to experiment with bandwidth, but otherwise should be left out.

IC6 needs to be a high-speed opamp, with good transient recovery. The NE5534 and LM318 have been popular choices, but there are many more to choose from now. Different opamps

Important note: R11 absorbs the energy that was stored in the coil just prior to flyback, and dissipates the energy as heat. The amount of time-averaged energy is dependent on several factors: coil resistance, transmit pulse width (turn-on time), and transmit frequency. Low coil resistance, large pulse width, and high frequency combine to give the highest energy, and R11 can get *very hot*! It is recommended that R11 be no less than a 1/2-watt resistor, preferably metal film. If you intend to run it hot, consider a 1-watt resistor, and/or heatsinking it. Note that Q3 should not run hot, unless it is breaking down during flyback.

Parts List

Resistors (5% 1/4-watt)

R1	56k
R4, R42a	1.5k
R5, R6, R43, R44, R45a, R46, R47, R48	10k
R40	3.3k
R9	220
R10	1
R11	See Text
R12, R14, R24, R25, R28	1k
R13	1M
R16	22k
R20, R21	27
R26, R27, R29, R33, R38	100k
R30	39k
R35	See Text
R39, R46a	330

Potentiometers

R2, R18, R31, R34	100k
R3, R42, R46	10k
R23	100
R41	5k Audio Taper
R45	50k

Capacitors

C1	100μF elect.
C2	10μF elect.
C3	22μF elect.
C4, C5, C6, C7,	
C14, C15	47μF elect.
C8, C24, C25, C26	10nF poly.
C9, C17, C18, C23	0.1μF poly.
C10	1000μF elect.
C19, C20, C21, C22	0.47μF poly.

Diodes

D1-D4, D6, D7	1N4148
---------------	--------

Transistors

Q1, Q5, Q6, Q7	2N3906
Q2, Q8, Q9	2N3904
Q3	IRF740

Integrated Circuits

IC1	ICL7660(Voltage inv)
IC2	78L05 (+5v ref)
IC3, IC4	79L05(-5v ref)
IC5, IC9	NE555 (Timer)
IC6	NE5534(Opamp)
IC7	4066 (Bilateral switch)
IC8	TL072 (Opamp)
IC11, IC12	74HC221 (Multivibrator)

Misc

J1	Audio jack
SP1	Speaker

have different schemes for nulling off-sets and some use external components for frequency compensation, so this design provides several options. These options are discussed later. For all opamps, it is desirable to have the cleanest supply voltages possible, so R20-C14 and R21-C15 provide low-pass noise filtering for both rails.

The output signal of the preamp can be optionally filtered by R22-C16, which can be installed as a low-pass filter as shown in Fig. 1, or transposed to provide high-pass filtering. Both filter types have shown up in some commercial designs, and values for R22 and C16 will depend on the application. If not used, omit C16 and short R22.

The final part of the receiver front-end is the sampling switches. IC7a and IC7b are 4066 gates that sample the same output signal of the preamp, but at different times. Potentiometer R23 provides a method of fine-tuning the signal balance between the two switches, which will be explained in the **Calibration** section. IC7a is the main sampling switch, and will sample the signal during the decay of the target signal. The sample delay, and sample pulse width are set by R42 and R46 as described in the Clocking section.

IC7b samples the signal at a later time, determined by R45, and is used for interference cancellation in the differential integrator. This scheme is used in many commercial designs, but is often lacking in other PI projects. It is optional, and if not used, R23 can be replaced with a short from IC7a, pin 1 to R22.

Receiver Back-end

The remaining circuitry consists of the differential integrator, amplifier, and audio. IC8a is the integrator, and is set to a gain of 100 by R26/R24. C17 is the integration capacitor, and its value should be chosen to provide a sufficient, but not excessive, integration time constant. The purpose if the integrator is to average the sampled signals from the preamp. Noise, which is random in nature, averages to zero, or so we hope, while even weak target signals will correlate and remain. A short time constant will result in an excessively fast decay

of the target signal, while a long time constant will smear adjacent targets and reduce selectivity.

When used as a differential integrator (Fig. 8, top), the secondary signal from IC7b should be applied to a divider that matches the inverting gain network, such that the non-inverting gain is the same as the inverting gain. Usually, this will call for either precision-valued resistors, or at least matched resistors. However, standard 5% resistors are sufficient if R23 is used to calibrate the differential gain. If the integrator is used in a single-ended configuration (Fig. 8, bottom), then R25, R27, and C18 may be omitted, with IC8a pin 3 shorted to ground. Alternatively, placing a shorting switch across R27-C18 will allow a quick comparison between single-ended and differential modes.

The output of the integrator is AC coupled through C19 to the final gain stage, opamp IC8b. The gain of this amp is set to 100 by R29/R28. C28 provides some additional low-pass filtering. With the integrator signal AC coupled, resistor R34 is used to provide a discharge path for the non-inverting input of the opamp. Thus, static signals are rejected and only dynamic signals will be amplified. This means that when the coil is swept over a target, the signal will rise and fall and get amplified by IC8b. But if the coil is moved over a target and held steady, the signal will rise and then quickly fall as C19 is discharged. The rate of discharge, or the allowable motion speed, can be controlled by varying R34.

This type of motion circuit is common in PI detectors and helps reject slowly changing ground conditions. However, the need to keep the coil in motion makes pinpointing targets difficult. If R34 is removed, and the inputs of IC8b have a very low bias current, then the capacitor will not discharge and even static signals will pass. Therefore, IC8b (and thus IC8a) needs to be either a JFET-input opamp, or a CMOS opamp. Even a JFET-input opamp has a slight bias current, perhaps a few picoamps, which will charge C19 in a short time (perhaps 15 seconds or so), so this mode is likely not useful for continuous

hunting, unless a true zero-bias current CMOS opamp is used.

R34 is shown in the schematic as a switched potentiometer. It can be board-mounted and set for a preferred retune rate, or panel-mounted for easy field changes. This is a control that is not typically panel-mounted, and if it is, should have a fixed resistor (of about 20k) connected in series with the pot, to avoid adjusting all the way down to zero ohms, which will completely kill the signal, although with no risk of circuit damage. If there is no desire for a zero-motion mode, then the switch can be replaced with a short on the PC board. If a mode switch is used, it should either be a momentary-contact pushbutton switch located for easy thumb access, or a momentary toggle switch for index finger control. If a CMOS opamp is used for IC8, then a SPDT switch with one side momentary (On-Off-Mom.) can be used, similar to the mode trigger used on White's detectors.

AC coupling the integrator signal into IC8b also causes the output of IC8b to have a quiescent (average) value, plus a dynamic target signal. R30-R31-R32-R33 provide a way to adjust the quiescent output level of the amp, which is a useful way of setting an audio threshold. C21 simply filters the adjustment voltage applied to R33.

The last part of the backend, and the most complicated looking, is the audio circuitry. There are two major options for the audio: VCO, and non-VCO. With the VCO option, a 555 is used as a voltage-to-frequency converter, such that changes in signal level result in changes in the audio frequency. With the non-VCO option, the output signal of IC8b is chopped using one of the 4066 switches, so the frequency remains constant (and equal to the main clock frequency) and signal level only affects the apparent loudness of the audio. The circuitries for the two methods are different enough that they need to be explained separately.

One of the hookup options for the preamp (IC6) was as an inverting or non-inverting amplifier. Because the integrator stage also inverts the signal,

the inverting preamp option results in a non-inverted signal at the output of IC8b (i.e., a target signal creates an increase in the output voltage of IC8b). The VCO audio option needs a non-inverted signal, while the non-VCO option needs an inverted signal. Therefore, depending on how the preamp is connected and which audio option is used, the output of IC8b may need an additional inversion. Q4, R36, and R37 provide for this.

For VCO audio using an inverting preamp, D5, Q4, R36, and R37 may be omitted. As shown in Fig. 9, R35 drives the common-base transistor Q6, whose emitter is biased to 0v by Q5-R38. Since R35 has signal on one side and virtual ground on the other, it acts as a voltage-to-current converter, with a gain equal to the value of R35. Q6 passes this current to the 555 oscillator's RC timing, making the oscillator frequency increase with increasing current. The output of the 555 drives buffer transistor Q7, which drives a speaker or headphones.

Normally, the output audio frequency is adjusted to a certain quiescent value, and a target signal causes the frequency to increase. Setting the quiescent audio frequency is done with the threshold adjustment R31, described previously. For the inverting preamp option, R32 is not needed because R35-Q6 can only generate a positive current. The value of R35 determines the amount of VCO current that is generated for a given target signal, therefore it can be made variable for adjusting sensitivity.

For VCO audio when using a non-inverting preamp, Q4, R36, and R37 must be included to invert the signal. Fig. 10 shows the hookup for this. In this case, R35 no longer affects the voltage-to-current gain, but R37 does.

Non-VCO audio can also be achieved, either with an inverting or non-inverting preamp. Fig. 11 shows the non-VCO audio hookup. R35 sets the audio threshold, but the audio response is now loudness, not frequency. IC7c chops the output signal of IC8b, so the audio frequency is constant and equal to the master clock

frequency. Keep in mind that if the sample frequency is varied (via R2), the audio frequency will follow. They cannot be set independently.

Once the threshold loudness is adjusted, a target signal will cause an increase in loudness. As the coil swings past the target, the output of IC8b can go negative. In the VCO option, Q6 prevents negative signals from generating negative currents. In the non-VCO option, negative signals are clamped by diode D6.

The non-VCO option potentially offers better performance than the non-VCO option. Driving a speaker (or headphones, to a lesser extent) creates large current spikes that can feed through to the supply voltages, and cause interference with the receiver. With the non-VCO option, the audio frequency is synchronous with the sampling frequency, and if feedthrough occurs, it can be moved in time by delaying the chopping signal until it is placed at a non-sensitive point in the sampling period. With VCO audio, the interference in non-synchronous and cannot be controlled. No tests have been performed to see if this is really a problem.

Construction & Calibration

Building the PI-1 project is pretty straightforward, and can be done in discreet stages. As already mentioned, a DMM is needed, and while an oscilloscope is useful, it is not necessary. Waveforms will be shown so those with an o-scope can verify circuit operation, and those without can see what the circuitry is supposed to be doing. The remainder of this section will describe the construction of "Option 1", which has an inverting preamp and VCO audio. Other options will be discussed in the **Other Options** section.

In building circuits, I generally like to use IC sockets. It allows for the re-use of ICs if you decide later to scrap the circuit, and it enables the easy swapping of ICs if one is bad, or if you decide to experiment with a different IC. In this design, I would suggest socketing at least IC6, and perhaps IC8. These opamps are the two most likely to be experimented with.

Also, there are several resistors and capacitors that can be experimented with. To make these easy to swap out, pin sockets can be used, or any other component socket that can be split into single pins. Make sure to use sockets that are big enough for the component leads, as there are sockets available for several different lead sizes.

The schematic for Option 1 is shown in Fig. 12, and a complete list of components is given in the sidebar. Resistors are 1/4-watt 5%, except for R11 which should be 1/2-watt. Electrolytic caps should be rated for at least 25v on the VB lines, and 16v elsewhere. For the other caps, polypropylene and polystyrene offer the best stability, but monolithic is sufficient.

As in the circuit description, we'll start with the power supply section. Install IC1-IC4, C1-C7, and D1-D2. Watch the polarities on the electrolytic caps, as well as the diodes. Temporarily short SW1, apply 12v to the circuitry and use a DMM to measure the voltage outputs of IC2 (pin 1 = +5v), IC3 (pin 3 = -5v), and IC4 (pin 3 = -5v), with respect to ground. Remember that "ground" is the positive side of the battery, not the negative. Since IC1 is a voltage doubler, the input to IC2 (pin 3) should be roughly equal to the battery voltage, but could be as low as 80% due to the inherent nature of the chip. Battery voltages drop quickly from their rated voltage, so the output of IC1 could be as low as 8 or 9 volts. Other measured voltages should differ by no more than a few percent.

Next is the clocking circuitry. Install IC5, R1-R4, and C8-C9. Also install IC11-IC12, R42-R48, C24-C26, D6-D7, and Q8-Q9. Make sure all appropriate shorting wires are placed for this option. Power up the circuitry again, and measure the supply voltages on IC5 (pin 1 = -5v), IC11 (pin 8 = -5v), and IC12 (pin 8 = -5v). The output of IC5 (TestPoint 1) should measure slightly negative, and vary with the Frequency and Pulse Width adjustments. It should fall between -0.25 and -1.25 volts, and be roughly -0.5v with both adjustments centered.

The transmit circuitry and the receiver preamp are next: IC6, R5-R21, C10-C15, D3-D4, and Q1-Q3. Remember to install only the components needed, and jumpers where appropriate. As mentioned before, R9 normally sees short pulses of current, and should not overheat. However, if the output of IC1 should get stuck in a low state, R9 can destructively overheat (*burn up*), as well as MOSFET Q3. Therefore, make sure IC1 is always operating, or pull R9 if needed (this will also protect Q3).

Damping resistor R11 will depend on the coil used, but for the 500uH coils described in the **Coil** section, it should be 680Ω. R11 should also be rated at 1/2-watt, as it can get a bit warm, especially when the TX pulse width is set high. Leave the coil L1 off for now. Set R18 to its center position. Power up the circuitry and check the supply pins on IC6: pin 7 = +5v and pin 4 = -5v (roughly; R20 and R21 will slightly reduce these values). Adjusting R18 should vary output of IC6 (pin 6, TP4) from about -4v, to +4v.

Continue with the remainder of the receiver: IC7, R23-R34, and C17-C21, plus any needed jumpers. *Do not install IC8 at this time.* Temporarily short SW2 and set R34 to its maximum value. Again, apply power and measure the power pins on IC7 (pin 14 = +5, pin 7 = -5v) and the pads of IC8 (pin 8 = +5, pin 4 = -5v).

For proper differential integration, the differential gain paths of IC8 should be reasonably well matched. The inverting

path gain is $A_{inv} = \frac{R26}{R24}$. The non-inverting path gain is

$$A_{noninv} = \frac{R27}{R25 + R27} \times \frac{R24 + R26}{R24}$$

which simplifies to $R26/R24$, when $R25=R24$ and $R27=R26$. One way to match resistors is to use 1% or 0.1% precision resistors. Another way is to calibrate the gain. In the spirit of making this project easy to build, without specialty parts, we will calibrate.

Short pad 1 of IC8 to a convenient ground (simple to do when using IC sockets — IC9 pin 4 is ground). ***This should never be done with IC8***

installed. R24-R26 is now a simple voltage divider, the same as R25-R27. A DC voltage can be applied to R23, and if there is a mismatch in the two dividers, R23 can be adjusted to null it out. This should be done with the analog gates turned on, so short pins 1 & 2, and pins 3 & 4, of IC7. If IC7 was not socketed, shorting these pins will not harm the chip.

Use a voltmeter to measure the voltage at TP4. Adjust R18 until this voltage is about 2 volts. Now measure the voltage across pins 2 and 3 of IC8, and adjust R23 until it is zero. Readjust R18 until the output of IC6 (TP4) is 0 volts. Remove power to the board, remove the short to pad 1, the shorts on IC7, and install IC8. Set R31 to its midpoint.

Finally, install the audio circuitry: R35-R41, Q5-Q7, C22-C23, IC9, and necessary jumpers. Set R35 to its midpoint and R41 to maximum. At this point, the circuitry is ready to be packaged and completed. Using a suitable enclosure (remember to allow for a battery pack) add SW1, SW2, the speaker and/or headphone jack, and mount whichever pots will be user-accessible.

Coil

This design, like most other PI designs, is set up to use a simple mono coil. A mono coil has only one winding, as opposed to coils with multiple windings, which often require delicate alignment. Important coil parameters are coil diameter, number of windings, and wire gauge.

For all-purpose coin hunting, a good coil size is 8-10 inches (20-25cm) in diameter. A smaller coil will be more sensitive to smaller targets, but will achieve less depth. Conversely, a larger coil will gain depth, but lose sensitivity to small targets. One popular setup for PI detectors is to use a very large coil, such as 1 meter, for deep cache hunting.

The number of windings affects the strength of the magnetic field that is generated. Theoretically, field strength is proportional to the square of the number of turns, so doubling the number of turns results in four times the field strength. This seems to be a fantastic

return on wire usage, except for two drawbacks. More windings results in more coil weight, and also results in more coil resistance and capacitance.

Coil weight must be kept reasonable to allow for long usage, and resistance can easily limit the amount of “on” current, which is a limitation on the field strength. Another effect of additional windings is an increase in interwinding capacitance, which tends to limit dI/dt during turn-off, again reducing the transient field strength. Wire gauge pretty much trades off resistance and weight.

The standard coil for this design has a diameter of 8-inches (20cm), with 30 turns of 26 AWG enameled wire. A simple way to wind a coil is to drive a circle of nails of the correct diameter into a piece of plywood, and wind the wire around the nails. Then, pull out a few nails and slide the coil off the jig.

Before removing the coil, ensure that both ends have about 6-inches (15cm) of overlap. Also slide some short pieces of tape under the wire, between nails, and wrap tightly to keep them bundled after the coil is removed from the jig.

This is really all that’s needed to wind a mono coil, and the wire ends (pig-tails) can be soldered to a length of coaxial cable that feeds the main circuitry. However, a critical addition is a ground shield which, for PI detectors, is mostly needed to shield from outside electrical interference.

There are a number of ways to shield coils, including foil, wire wrap, and metallic tape. A simple method is to wrap the coil with aluminum duct tape, the kind easily found at the local home improvement center. This is the tape that is true aluminum, not the silvery metallic-looking fabric duct tape. Standard duct tape is conductive, but not enough for shielding.

Begin by cutting a length of tape equal to the coil circumference, minus about 1/4” (6-7mm). Duct tape is usually wide, so cut lengthways to get a piece about 1/2” (12.5mm) wide. Now cut notches every inch or so along both edges of the tape, such that the notches are about 1/3 the width of the tape, and

are staggered with the other side. See Fig. X for details.

Remove the backing and, starting at one of the pigtails, lay the tape around the outside of the coil, sticking the coil in the center of the tape. When the shield is placed all the way around the coil, stop just short of the pigtails, such that the two ends of the shield wrap do not touch. Allowing the shield to complete a short all the way around the coil, will allow significant eddy currents to flow, which will severely degrade performance. Now begin folding down the 1-inch tabs along one side of the tape, and then fold the tabs on the other side, such that they overlap slightly. One end of the coil shield will be soldered to the cable shield, so leave one of the tabs at the end unfolded.

Aluminum forms a thin oxide very quickly, so soldering it is not exactly easy, but with some effort can be done. Use a fairly hot iron, and work it around a small spot on an end tab, feeding solder to the spot. Use excess flux if you have it. In 15-20 seconds you should have a good solder contact. Solder down a short piece of wire for connecting the shield to the coaxial cable.

The final challenge of coil-building is to mount it to something rigid. This is actually the biggest challenge, as there is a limited selection of decent coil housings available. Metal detector manufacturers use injection-molded plastic shells. The coil is placed in the shell, and liquid epoxy added to secure it and add rigidity. For homebrewers, there are currently two sources of coil shells, and contact information is given at the end of the article.

Besides using an epoxy-filled plastic shell, there are some other alternatives. One of the simplest solutions is to epoxy the coil winding to a round piece of plywood of roughly the same diameter. Although this is an easy solution, it is not especially attractive or durable, and it would not be advisable to submerge it in water. But it is sufficient for experimentation.

Another solution is to simply buy a PI coil from a detector manufacturer. Some PI detectors use plug-in coils, and have accessory coils available. Other

models have hard-wired coils, but a replacement coil should still be available. In most cases, the value of the damping resistor (R11) will have to be tailored to the coil used.

Once the coil construction is complete, solder the pigtails to a length of coaxial cable, with the coil shield connected to the cable shield. Make sure the cable is securely clamped to the coil in some manner, so that there is no stress on the solder connections. Do not use any metal fittings in the construction of the coil, as they will desensitize the detector. Use plastic and nylon fittings instead. Connect the other end of the coax to the circuit at L1, with the shield on the ground side.

Adjustment & Usage

With the detector completely assembled and packaged, make sure the volume is turned down and apply power. SW2 should be closed. You should be able to adjust the Threshold and vary the audio tone from dead silence to a high pitch. Set the Threshold to get a very low frequency pattering audio. Wave a metal target close to the coil and the audio frequency should increase. If the metal is held absolutely still over the coil, the audio should retune itself back to the threshold tone.

With all the other pots set to the positions prescribed in the **Construction & Calibration** section, you should be able to detect a medium-sized coin several inches from the coil. Increasing the Sensitivity (by decreasing R35) might slightly improve depth, but the detector quickly becomes unstable, especially in noisy environments. Increasing the Pulse Width also improves depth slightly, at the expense of power consumption.

The most interesting controls are the Sample Delay (R42) and Sample Pulse Width (R46). Varying these can change sensitivity, even to different types of metal. In general, a short delay (around 15–20 μ s) gives the best sensitivity to low-conductance metal like gold.

Once the detector is operating and stable, you can turn SW2 off and test the non-autotrack mode. Now when a

metal target is held absolutely still over the coil, the audio should not retune to the threshold tone. This is normally used for pinpointing targets, not for general searching. In fact, unless a CMOS opamp is used for IC8, it is likely that this mode will become unstable after 15 seconds or so.

This design is intended to be a learning platform for how a PI detector works, so it has a lot of knobs to turn. Thus, it is quite easy to lose track of how adjustments interact, so when testing settings it is best to keep a good record. As you find an optimal setting for a certain knob, mark the position on the panel.

Other Options

The non-VCO option will be described here, as well as other configurations.

Substitutions

Component substitution options will be described here.