Matt Behrend
http://home.earthlink.net/~electronxlc/

How a Tesla Coil works

On this page, I will explain the basic theory of how a Tesla coil works. Below is the Table of Contents for this page.

Description of Components

A Tesla coil is a high-voltage air-core resonant transformer. A Tesla coil has 6 basic components. The first is the primary transformer, which is a high-voltage iron-core transformer. The second is the tank capacitor, which is a high-voltage capacitor that is usually homemade, but can be purchased for a high price from commercial suppliers. The third is the spark gap, basically two wires separated by a small gap of air. The fourth is the primary coil consisting of about 10 to 15 turns of thick heavy gauge wire wound around the base of the secondary coil. The fifth is the secondary coil, and it consists of many hundreds of turns of relatively thin, small gauge enameled wire. The primary and secondary coils make up an air-core transformer. That means that there is no iron core inside of the coils. The sixth basic component is the toroid. It is usually an aluminum doughnut-shaped object, and placed on top of the secondary coil. The high-voltage sparks radiate in all directions from the toroid out into the air.

How the Components Operate

The primary transformer converts the AC line voltage (120/240 volts AC) to over 10,000 volts. This energy is used to charge the high-voltage capacitor. The capacitor is wired in series with primary coil to the output of the transformer. When the voltage builds up enough to break down the spark gap, a spark will bridge the spark gap and complete the circuit between the capacitor and the primary coil by shorting out the transformer. A spark ionizes the air, causing it to become much more conductive. All the energy stored in the capacitor will be forced through the primary coil. The process of charging the capacitor and firing the spark gap occurs very rapidly. The spark gap may fire from 120 to over 1000 hertz, but 120 hertz delivers the largest energy bursts. When the energy is transferred to the primary coil, an electromagnetic field is generated and surrounds the secondary coil. The secondary coil absorbs this energy and magnifies the voltage further. The resulting voltage could be a couple hundred thousand volts for small coils, or millions for very large ones. It is very important to efficiently ground the coil. Grounding the coil connects it to the second plate (the earth) of the toroid capacitor. The bottom wire of the secondary coil is grounded via a heavy gauge wire that leads directly to a ground rod. The top wire of the secondary is connected to the toroid. There are usually a couple inches between the top winding of the secondary and the toroid.

The tank capacitor is wired in series with the primary coil. This does not affect the charging of the capacitor since the primary has a very low reactance at 60 Hz. The transformer is better protected from RF noise by this configuration than if the capacitor and spark gap positions were reversed.

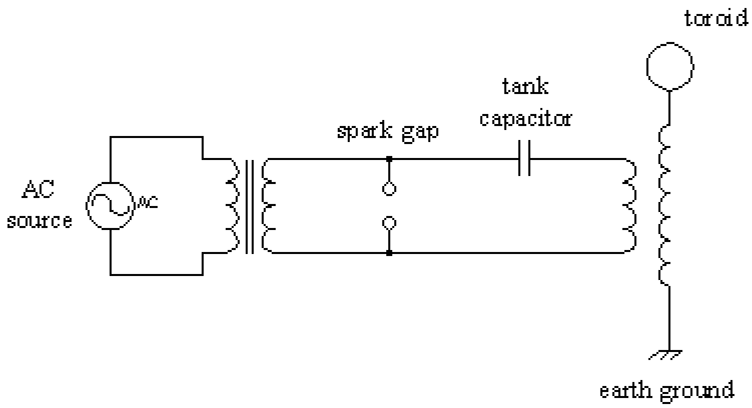

A schematic diagram of a Tesla coil.

LC Circuits

Tesla coils are composed of two LC circuits. Both of those circuits must be tuned to the exact same frequency for optimum efficiency. The first LC circuit is the tank circuit. The tank circuit is an oscillator composed of the capacitor, spark gap, and primary coil. The second LC circuit, which is a series resonant circuit, is the secondary coil and the toroid.

An LC circuit contains an inductor (a coil) and a capacitor respectively. An LC circuit is at resonance where the inductive reactance of the inductor equals the capacitive reactance of the capacitor at the same frequency. Reactance is the resistance to the AC waveform. Another way to define reactance, is the resistance to any change in voltage or current flow from the current instantaneous value. Alternating current is constantly changing because it flows in the form of a sine wave. Inductors resist changes in current flow and capacitors resist changes in voltage. Reactance is not the same as simple DC resistance. DC resistance is the resistance to the flow of electrons through a material, not the resistance to changes in voltage or current. The reactance of a capacitor or an inductor is based upon its electrical value and the frequency at which it operates. Finding the reactance of a component is explained on the calculations page. The graphic below shows the relationship between frequency, reactance, and resonant frequency. Capacitive reactance is denoted X_C and inductive reactance is denoted X_L.

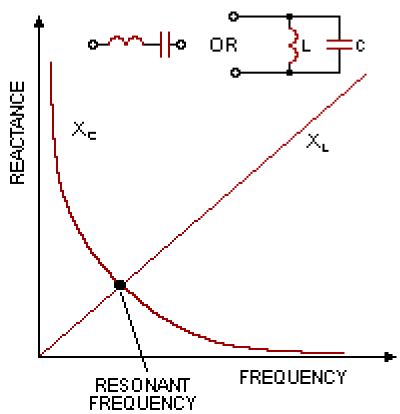

For any two given values of L and C their reactances will only be equal at one frequency. In a shock excited oscillator such as the tank circuit, the circuit will oscillate at that frequency where the reactances are equal. The circuit oscillates at that frequency because the circuit is balanced at that frequency. When the reactances are equal, both components have an equal charge and discharge time delay. It takes the inductor the same amount of time to discharge into the capacitor as it did for the capacitor to discharge into the inductor. The amount of time it takes for the discharge cycles to complete determines the number of cycles per second, which is the frequency of the oscillator circuit.

The toroid is actually one plate of a virtual capacitor. The second plate of the virtual capacitor is the earth itself, because the bottom wire of the secondary coil is grounded. The larger the toroid, the more capacitance it has. The toroid must have the right amount of capacitance so that its reactance equals the secondary coil's inductive reactance at the Tesla coil's designed resonant frequency. The secondary coil itself also has capacitance. The wire of the coil has capacitance to ground, and there is capacitance between each turn of the coil. Although these capacitances are extremely minute, at high frequencies they become quite significant. These effects on the circuit are discussed in more detail on the calculations page.

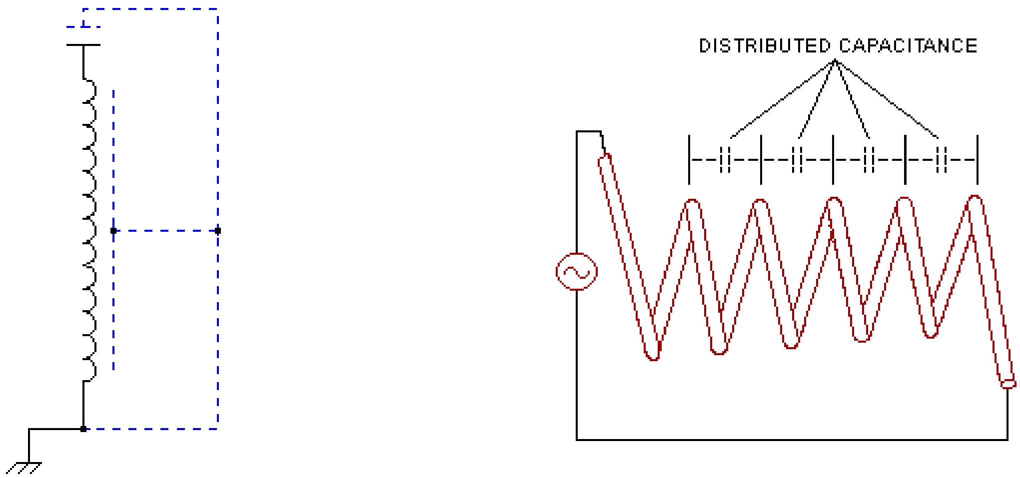

These graphics show the distribution of capacitance throughout the secondary circuit.

Voltage Gain

The Tesla coil achieves a great gain in voltage in a very different way than a conventional transformer. A transformer's change in voltage is dependent upon the turns ratio of the primary and secondary coils. If the primary coil of a transformer has 5 turns and the secondary coil has 100 turns, then the secondary voltage will be 20 times that of the primary. This does not fully apply to the interaction of the primary and secondary coils of a Tesla coil. Instead, a Tesla coil's voltage gain is based upon the different impedances of the primary and secondary circuit components. For example, a 0.06 µF tank capacitor has an impedance of 27 ohms at 100 kHz. The toroid has a capacitance of 30 pF and an impedance of 53,000 ohms at 100kHz. The greater impedance causes a greater voltage potential and less current. This behavior is better explained by the conservation of energy. The energy you put into a system is the energy you get out if losses are ignored.

Capacitance	Inductance
$J = 0.5\ V^2\ C$	$J = 0.5\ I^2\ L$

J = joules of energy stored
V = peak charge voltage
I = peak current
C = capacitance in farads
L = inductance in henries

The joule is a unit used for measuring the amount of energy stored in an inductor or capacitor. The formulas are derived by integrating the power (in watts) flowing through the component over time. The approximate output voltage of a Tesla coil can be found by this formula. The formula assumes that all energy is transferred from the primary circuit to the secondary circuit without any losses. That can't happen, so the actual voltage will be somewhat less than the calculated value. Don't use an RMS value in these equations, the peak values are required.

The secondary circuit voltage is found by setting the energy in the primary equal to the energy in the secondary, and solving for secondary voltage. The secondary capacitance (C_S) includes both the topload and the coil's self capacitance.

$$V_S = V_P \sqrt{\frac{C_P}{C_S}}$$

V_S = peak secondary voltage

V_P = peak charge voltage of tank capacitor
C_S = secondary capacitance in farads
C_P = tank capacitance in farads

Note that energy in equals energy out. I want to emphasize the fact that Tesla coils do not produce "free energy". Some people may argue that energy was produced since the high voltage transformer charged the tank capacitor to 10 kV at 100 mA and their secondary coil runs at 100 kV and several amps. This is a difference in power, not energy. Power is the rate at which energy is moving, and energy is power multiplied by time. It took 1/120 of a second to charge the tank capacitor but it took a mere 1/400000 of a second for the tank capacitor to move the same energy into the primary coil.

There are certain limitations to voltage gain. Using an extremely small toroid will not give a greater spark length or peak voltage. If there were no topload the output voltage would theoretically be at a maximum. However, when there is little capacitance, the coil will "break out" (ionize the air and create a spark) at a lower voltage. If an arc to ground occurs when the secondary is charged halfway, then most of that energy is spent in the arc and dissipated as heat. The secondary continues to be charged by the primary, but that energy is used to maintain the short arc rather than to continue building up the circuit's stored energy. This condition does not allow the secondary circuit to reach its peak potential, and spark length is severely decreased. Also, spark length is not merely based upon voltage, but is a function of both voltage and current. The toroid should be much greater than the self capacitance of the coil. I have noticed that a topload any less than the self capacitance of the secondary coil greatly decreases spark length. Also, if the topload is too great, then it may not be able to break out at the given power level.

The primary and secondary circuits must resonate at the same frequency, because all of the primary circuit's energy is not transferred to the secondary in a single 1/4 cycle of the AC waveform. The low coefficient of coupling makes it necessary for each cycle of energy coupled to the secondary to build upon the existing energy already transferred. This condition is called resonant rise, and it requires both circuits to have the same resonant frequency.

In an LC circuit that is being continuously supplied with power, the voltage will be equal to the product of Q and E, where Q is the quality factor of the inductor, and E is the applied voltage. In a Tesla coil, the power supplied to the secondary circuit is limited to the joules stored in the tank circuit. When all energy has been transferred, the peak voltage is found by the energy formula.

Resonant Rise

Resonant rise occurs in a series LC circuit when the frequency of the AC waveform that feeds the circuit is equal to the resonant frequency of the LC circuit.

The inductive and capacitive components in a circuit cause a phase shift between the waveform of the current flowing through the component and the voltage drop across the component. In a purely resistive circuit, the current is in phase with the voltage. As the input voltage to the resistor increases, so does the current flowing through it (Ohm's law).

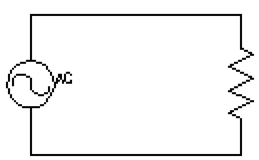

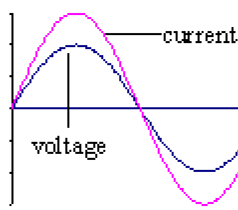

In a purely inductive circuit, the current waveform lags the voltage waveform by 90°, 1/4 of the sine wave. In other words, the voltage waveform leads the current waveform by 90°.

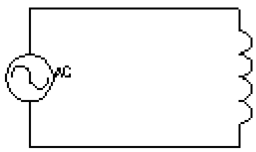

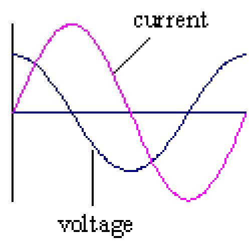

In a purely capacitive circuit, the current leads the voltage, or the voltage wave lags the current wave by 90°.

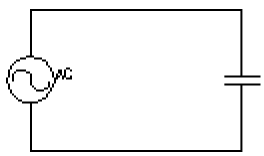

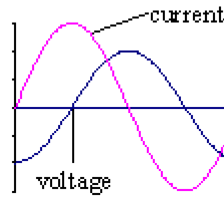

In a circuit containing a combination of resistance and capacitance or inductance, the phase angle between the voltage and current waveforms of the entire circuit is some number less than 90°, depending on the component values. The total impedance and the phase angle of an LCR circuit is found by vector addition, but that will not be covered on this page. In a tuned resonant circuit, the reactive components have a canceling effect, and the phase angle is zero (the waveforms are in phase).

In any parallel LC circuit, the voltage drops of all components are in phase and equal, while the current waveform of each component is out of phase (by 90°), and the magnitude of those waves may differ. In any series LC circuit, the current waveform flowing through the entire circuit remains constant, and the voltage drops of the L and C components are out of phase with the current. The magnitude of the voltage drops may vary. The secondary circuit is a series resonant circuit with the induced primary energy as the source.

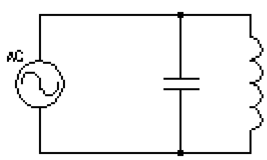

A Parallel LC circuit.

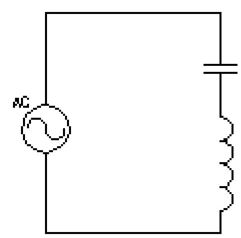

A series LC circuit.

The voltage drop of a component is the voltage potential that exists on the two leads of the component. The constant RMS voltage drop of an L or C component is found by multiplying its reactance by the current flowing through it (Ohm's law). In a series circuit at resonance, the reactances of the components, and therefore the voltage drops are equal. The current waveform is the same for each component. The capacitor voltage lags it by 90°, and the inductor voltage drop leads it by 90°. Therefore, the voltage drops are exactly opposite (180° out of phase), and equal in magnitude. This causes the reactive loads in the series resonant circuit to exactly cancel each other out. This allows the maximum amount of current possible to flow through the circuit unrestricted by the reactive components. The only impedance that limits current flow remaining in the series circuit is the DC resistance of the wire the inductor is wound from. An inductor always contains some value of resistance in its wire, and some value of inductive reactance (based on the inductance value), and these two impedances act in a series configuration.

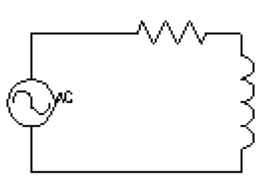

Since the opposite voltage drops exactly cancel each other out, and only the resistance is left, the energy being supplied by the primary sees the secondary circuit as a low-resistance load. This causes the current flow to increase, and thus the voltage drops also. It is in this manner that the primary builds upon what it has already transferred to the secondary, until all energy has been coupled to the secondary circuit.

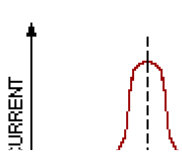

This graphic shows the current (voltage/reactance) in a series LCR circuit in relation to the frequency of the voltage waveform applied to it. At the resonant frequency, the circuit's current and compopnent voltage drops are at maximums.

Because of the way the primary and secondary coils are loosely coupled, they do not behave in the same way as an iron-core transformer. The induced energy in the secondary coil acts as a voltage source placed in series with the secondary coil and toroid capacitor. The primary coil (tank circuit) must be tuned to resonate at exactly the resonant frequency of the secondary circuit for maximum efficiency.

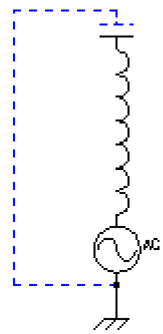

This graphic shows the induced voltage as an AC source to the secondary series resonant circuit.

Resonant rise is commonly misunderstood to occur over multiple "bangs" (a single firing of the spark gap), however this is not true. All of the examples and explanations of resonant rise in this paper describe the accumulation of energy in a single bang, where each cycle of the oscillation in one circuit builds upon the energy in the other circuit. The energy that is transferred to the secondary circuit in one bang is quickly dissipated and is not stored long enough for the next bang to build upon it. Once the spark gap quenches (stops conducting), the tank capacitor begins to charge again and any energy left in the secondary circuit is converted to RF radiation and heat (from arcs and wire resistance). The duration of the oscillation generated is very short compared to the charge time of the tank capacitor. When running a Tesla coil, you will probably notice that the maximum spark length grows in the first couple of minutes when it is first turned on. Although this would appear to be caused by the next bang building upon the previous, it is actually the result of the ionization of the air. As the coil runs, it fills the room with gasses that are more conductive. The arcs also create ionized channels that are more prone to break down with a spark on the next cycle. For this reason, a greater number of firings per second may increase a Tesla coil's maximum spark length.

Oscillation and Tuning

The output waveform of a Tesla coil that is commonly referred to as "ringing" is caused by the oscillatory response of the tank circuit and secondary resonant circuit. Oscillation is the process where an interchange of energy takes place between the inductor and the capacitor in a tuned resonant circuit. This oscillation occurs in the tank circuit and in the secondary circuit.

The firing of the spark gap is the electrical equivalent to striking the side of a bell with a hammer. The LC circuit is like the bell. A bell, when struck, will ring at a particular musical tone. The physical properties of the bell determine the frequency of its tone, just like the capacitance and inductance of a circuit determine the frequency of the circuit. The tank circuit will continue to "ring" for as long as the arc across the spark gap lasts.

When the capacitor is charged to its peak voltage, the spark gap fires. The capacitor then discharges its energy into the primary coil. The capacitor starts off at its peak voltage and is discharged to zero volts, forming 1/4 cycle of the tank circuit's resonant frequency. As the capacitor is being discharged, the voltage drop is decreasing and current flowing through the circuit is increasing, causing the field strength of the primary coil to rise. The primary coil stores this energy in the form of its electromagnetic field. When the current flowing through the primary coil reaches its peak and the capacitor voltage is zero (because the capacitor is drained), current flow begins to fall because the capacitor can no longer supply current. The primary coil resists this decrease in current flow by producing emf (electromotive force). This emf is caused by the collapsing of the coil's magnetic field as the current flowing through it drops. A coil's magnetic field is sustained by current flow. The emf voltage causes current to flow in the same direction and the primary coil charges the capacitor with its emf at the opposite of the capacitor's original polarity. Now that the capacitor is charged at the new polarity, it also discharges into the primary coil at that new polarity, causing the electromagnetic field of the primary coil to change polarity. As the emf energy stored in the capacitor drops to zero, the primary produces emf again, causing the capacitor to be charged at its original polarity, and the cycle repeats itself. The oscillations continue to repeat until the spark gap is quenched, or the circuit runs out of energy (in secondary circuit).

This back and forth recycling of energy can be compared to the swinging of a pendulum. The pendulum starts off in neutral position, pointing straight up and down. When it is pushed left it swings to the left until it peaks, and then swings back to the right. It goes past the neutral position, peaks, and swings back to the left again. It continues swinging, traveling a little less distance each cycle until it runs out of energy or is stopped.

Below is a graphic of how oscillation behaves. The diagram shows the polarity of charge, the direction of current flow, and pieces of the current flow wave generated by each part of the cycle. The sine wave shown is in relation to the flow of current when oscillation begins.

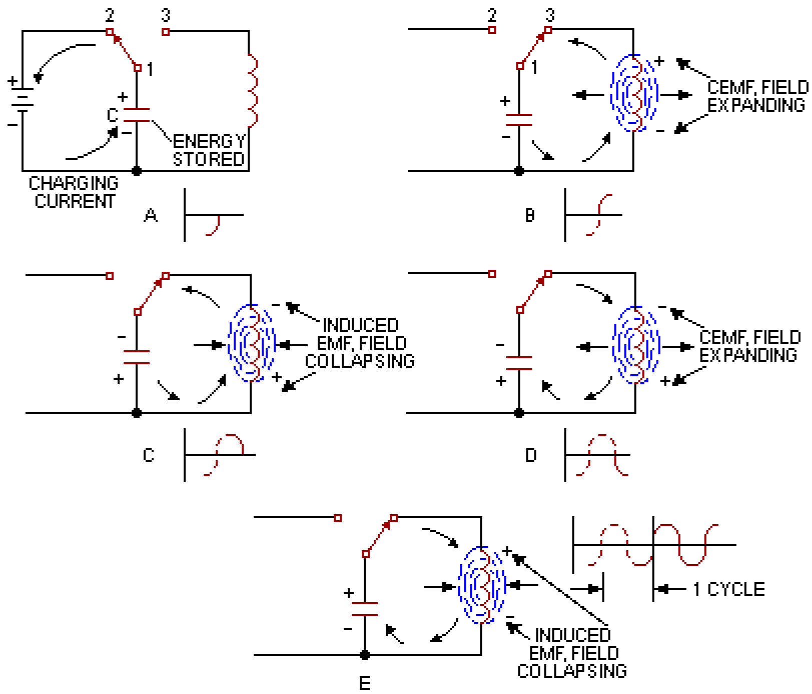

The result of the oscillations is a damped sine wave. Like the pendulum analogy, each cycle is weaker in magnitude than the previous. The decreasing of each cycle is caused by losses within the components of the LC circuit. The wave is also dampened by the transfer of energy to the secondary circuit. Losses occur because of the resistance of the wire in primary coil, not its impedance (inductive reactance), and the dissipation factor of the capacitor. For the tank capacitor, you should use a very low loss dielectric; "leaky" capacitors waste energy. The copper conductor of the primary coil offers significant resistance because of the skin effect. At high frequencies, electricity flows near the surface of a conductor more than through its center. The electrical current may only penetrate the conductor by a few mils, making the actual resistance much higher than the DC resistance of the wire. The solution to this problem is to use a conductor with maximum surface area. Below is a graphic of the damped sine wave produced in an LCR circuit. If there were no losses within the circuit, the oscillations would continue indefinitely and there would not be a dampening effect (excluding that of energy transfer), but such a condition does not exist.

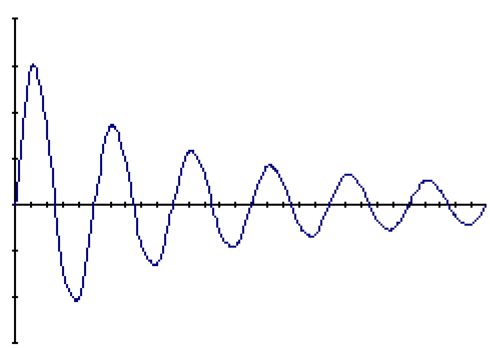

Damped sine wave graphic.

It is very important for the primary coil to resonate at the same frequency as the secondary coil so that each pulse from the ringing of the primary is in sync with the secondary coil. This situation is like pushing a swing set. You must give the swing a push at the same frequency that it is moving back and forth in order to maintain its motion. If the swing is flying toward you and you give it a push before it is ready to go back the other way, then you will cancel out the motion and the swing will stop.

The secondary resonant circuit of the secondary coil and toroid is an oscillator. Although the secondary circuit is magnetically induced with the primary circuit's energy, the secondary circuit can ring at a different frequency that is independent of the primary circuit. The oscillatory response of the secondary circuit operates in the same manner as in the primary circuit. The difference in the secondary circuit is that its excitation source is the primary coil, rather than the initial charge of a capacitor. Throughout the oscillations, the primary coil builds upon the oscillation of the secondary circuit by resonant rise. This is why it is absolutely critical that the primary circuit should resonate at the same frequency as the secondary circuit. Just as the circuits reinforce each other when they are in tune, they can cancel out the resonant rise in the secondary circuit if they are out of tune.

The graphics below show the voltage waveforms of the tank circuit (pink) and the secondary circuit (blue). These are close representations of the waveforms produced by a Tesla coil. I created these graphics with Excel, using a mathematical equivalent to the waveforms I viewed with my oscilloscope.

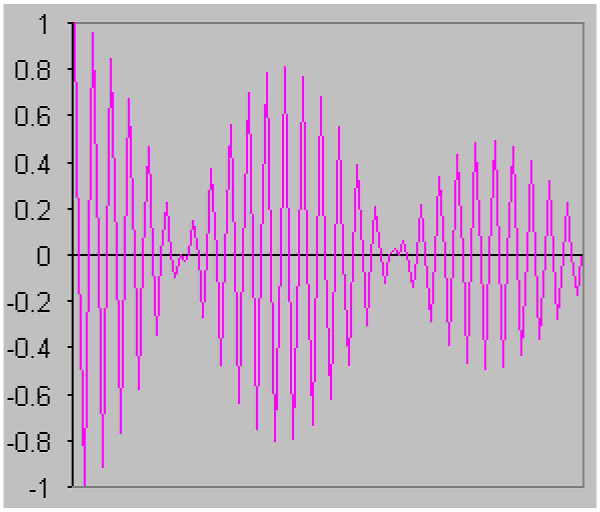

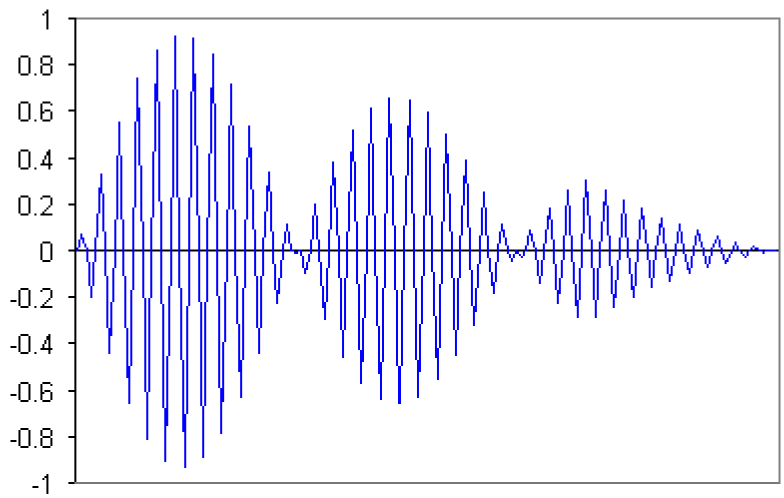

Notice how the energy in the secondary circuit gradually builds up to a peak. At that point, the secondary begins to transfer energy back to the primary circuit (oscillation between the two circuits). The spark gap quenches at the point where the graph of the primary circuit ends. The graph of the secondary shows the circuit ringing down after that point. There is no more interaction between the primary and secondary circuits, because the continuity of the primary has been broken. The secondary's remaining energy is converted to heat and radiated RF energy, creating a damped sine wave.

In these graphs, the spark gap has quenched just before the third notch. A notch is each point in time when all of the tank circuit's energy has been transferred to the secondary circuit. It is ideal, although difficult, to achieve first notch quenching. This is important, because the longer the energy stays in the circuit, the more time is allowed for heat losses to waste energy. The tank circuit is the greatest source of loss, especially the spark gap. The tank circuit operates on hundreds of amps, causing the resistance of the circuit to dissipate a great amount of energy as heat. The amount of power lost is equal to $I^2 R$, where I is current and R is resistance. The frequency of the oscillations between the primary and secondary is determined by the coefficient of coupling, the tuning, and less significantly, the losses.

In order to achieve the longest spark length (the goal of most coilers), it is necessary for the spark to break out at the first notch. This requires that the topload be of the correct size to only allow break out to occur at the peak voltage of the first notch. A given toroid will charge to a specific potential before it will allow the air surrounding it to break down and form a spark to ground. The optimum toroid size is usually found by experimentation, as it is affected by many factors.

It is important to note that these waveforms are produced by a Tesla coil that is not arcing to ground. If there were a spark to ground, the waveforms would be quite different, but still similar. At the point in time when the arc completed its path to ground, the energy level of the secondary would drop very quickly, and would not increase further. Any energy that the primary could continue to supply would be consumed in maintaining the arc. There would also not be any further oscillation between the circuits, since the secondary would never reach a peak energy level after the arc struck ground.

Quarter Wavelength Frequency

For a long time, it has been believed that Tesla coils are governed by the principles of the quarter wavelength theory. Much work has also been done in the recent years to model the secondary coil using transmission line theory. Both models seem to deviate from measurements others have taken. The lumped component model used for designing Tesla coils works well, but it is not capable of describing the internal behavior of the secondary coil. Modeling the operation of the secondary coil is not a simple task. Some complexities of the secondary coil are the magnetic coupling between adjacent turns, the distribution of capacitance between turns, capacitance to ground, capacitance from windings to toroid, and the added lumped capacitance of the toroid to ground. The secondary coil and toroid compose a system, which by itself deserves attention. Modeling of this system will explain the uneven voltage distribution along the coil and the phase angle between current at the base and top of the coil that has been observed by some people. Accurate models will provide greater insight into this part of the Tesla coil and aid in optimizing design. The following is a basic introduction to 1/4 wavelength theory for antennas and how it may relate to secondary coil operation. It is not necessary to design a Tesla coil to resonate at the 1/4 wavelength frequency of its windings.

A monopolar antenna most efficiently radiates its RF energy when it is operating at its 1/4 wavelength frequency. The output of the antenna is at its peak when the maximum voltage is at the top of the antenna while maximum current is on the bottom. Maximum current occurs at the point on the sine wave when voltage is at a minimum, and minimum current occurs at the point when voltage is at a maximum. The graphic below depicts the voltage and current waves on a monopolar antenna. These waveforms do not represent the actual instantaneous RF waves flowing through the antenna, but rather their peak values at each point along the length of the antenna. These are called standing waves.

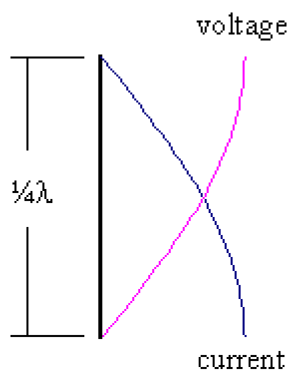

Monopolar antennas can also operate at odd multiple harmonics of the 1/4 wavelength frequency. Harmonic waveforms are shown on the two antennas below.

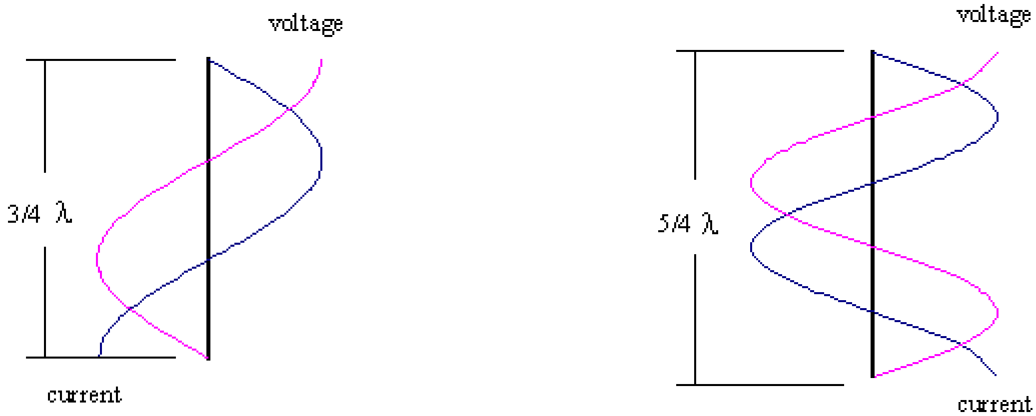

Antennas with 3/4 and 5/4 wavelengths.

An analogy to this condition would be the waves that can be induced into a rope. One end of a long rope is fastened to a pole (or anything), while a person holds the other end in his hand. If the person swings the rope from side to side very rapidly at an appropriate frequency, standing waves can be seen in the motion of the rope at harmonics to the frequency that the person is swinging it.

If the antenna was not operating at its 1/4 wavelength frequency or a harmonic, then when the point of maximum current was on the bottom of the antenna, the point of maximum voltage would either be before or past the end of the antenna. This results in a decrease in radiated energy.

Electricity travels at the speed of light. If a coil has a length of wire "X", then the quarter wavelength frequency of that coil will be the frequency at which 1/4 of its cycle is completed when the electricity travels the distance "X".

Below is a formula for calculating the 1/4 wavelength frequency of an antenna. Simply solve for frequency or length. The variable 'W' is multiplied by 4, because 4 times the 1/4 wavelength of wire is the distance spanned by a full cycle of the 1/4 wavelength frequency. The wavelength of a single cycle multiplied by the number of cycles per second (1/4 wavelength frequency) will be the distance traveled by electricity in one second.

$4WF = (5{,}280)(186{,}000)$

W = the secondary coil's wire length in feet
F = the quarter wavelength frequency of the secondary coil in hertz
186,000 = the speed of light in miles per second
5,280 = feet in a mile

It is critical to the efficiency of a Tesla coil to have all the components properly in tune with each other. Having just a little too much or too less of any electrical value can have a drastic effect on the output of your coil. The calculations page explains the steps involved in designing an efficient, tuned Tesla coil and finding the necessary values for each component.

Tesla Coil Calculations

On this page of my site, I will explain how to design a Tesla coil to achieve the highest voltage gain possible with maximum efficiency. Computing all the component values for optimum performance is the most critical part of building a Tesla coil. A Tesla coil is a finely tuned resonant circuit. Simply slapping all the components together with random values will not get you anywhere. Get out your calculator, because there is a lot of math involved in this.

Table of Contents

There are two main ways to design a Tesla coil. The first is to start with the transformer, and the second is to start with the secondary coil. Both methods involve some guess and check work if you are designing the Tesla coil around a specific transformer to power it. If you have already bought a transformer then you should use the first method. Find the maximum capacitance value for the tank capacitor and then design the primary and secondary coils. You will have to design a primary coil and calculate the resonant frequency of the tank circuit. Then design a properly coupled secondary coil to match the primary coil.

The second method is to start with the secondary coil and work down. If for some reason you want your Tesla coil to operate at a specific frequency, then you can use this method. First design the secondary coil to the desired frequency and design a primary coil to match it. Next, based on the inductance of the primary you designed, calculate the necessary capacitance to make the tank circuit resonate at the same frequency you have chosen for the secondary coil. Then, find the specifications for a primary transformer that is capable of charging the tank capacitor. If you want a certain power level for the primary transformer, then adjust your secondary coil design until the tank capacitor matches the transformer you want.

Since these calculations can be rather time-consuming, several computer programs have been developed by other coilers to help you in this process. They are available on my download page. I have also written Java applets for the necessary electrical formulas and they are on: http://home.earthlink.net/~electronxlc/java/javacalc.html.

Primary Transformer Calculations

I suggest that you actually obtain your transformer before you start building the other components, unless you are absolutely sure that you can get a specific type of transformer. If you build your Tesla coil to be powered by a certain transformer, and then when you try to buy it you can't find the right one, you could have quite a problem on your hands.

There are two important things you need to know about your transformer. The first is the impedance of the secondary coil in the transformer. The second is the amount of current your transformer will draw from the 120/240 volt line source. This is only an important factor if you are using a high-power transformer such as a pole pig.

Different types of transformers will tell you different electrical characteristics. Most transformers, such as pole pigs will have the voltage, and VA (volts multiplied by amps) ratings written on them. Neon sign transformers will often give the voltage and current values of the secondary winding (the high-voltage coil of the transformer).

If you have a power transformer with a VA and voltage rating, then use the formula below. First, divide the transformer's VA rating by V (voltage rating) to find the current rating I.

$$\frac{VA}{V} = I$$

Next, use Ohm's law to find the impedance of the transformer's secondary coil (the high-voltage windings). Ohm's law states that E = IZ, where E = volts, I = amps, and Z = impedance in ohms. Divide the transformer's output voltage by its current rating to find Z.

$$\frac{E}{I} = Z$$

You will use this impedance value for the capacitor calculations in the next section of this page. The impedance of a transformer's secondary coil limits the amount of current it is capable of supplying.
If your transformer gives you its voltage and current ratings, then use the same Ohm's law formula above to solve for impedance.

If you are using a pole pig, calculate the current draw of its primary coil by dividing the VA rating by the input voltage, 120 or 240. It is usually better to power a pole pig with 240 volts instead of 120. Most houses have a pair of 240 volt wires at 100 amps. The two legs of the 240, plus the neutral line make two sets of 120 at 100 amps each. For example, a 10,000VA pole pig would draw 41.7 amps from 240V, but would draw 83.3 amps at l20V. At 240 volts, the transformer draws half the current. Having less current decreases the thickness of wire you need to power it, and you won't have to use a huge variac. Note that these current calculations assume a maximum load across the output of the transformer.

After you know how much current your transformer will draw, you need to use AC cable that is rated for that much current. If you want to be able to plug your transformer into the wall, you may have to use a heavy-duty appliance outlet with a high enough current rating. Your local hardware store will carry them. If you have a load that exceeds the wire rating, you will start a fire.

Tank Capacitor Calculations

Now that you know the impedance of your transformer, you can calculate the amount of capacitance you need for the tank capacitor. Below is the capacitive reactance formula.

$$2\pi X_C F C = 1 \longrightarrow C = \frac{1}{2\pi F Z}$$

X_C = capacitive reactance in ohms

F = frequency in hertz

C = capacitance in farads

Z = impedance of the transformer in ohms

Capacitive reactance is the capacitor's resistance to the AC waveform. The dielectric in a capacitor isn't conductive, but a capacitor does resist the changes in the AC waveform, and thus it is a load. When the capacitor is being charged, and after it is fully charged and the AC voltage begins to drop below the capacitor's voltage, current flows. Whenever the input voltage across the capacitor changes from its last instantaneous value, current flows through it.

This formula converts the load that a capacitor imposes on an AC circuit into standard ohms. Capacitive reactance is based on the capacitance value, and the frequency of the waveform it is being charged with. The higher capacitance value a capacitor has, the less impedance it has, the more current it requires to charge it, and the more energy it can store. The constant 2pi in the formula is the circumference of the unit circle and is multiplied by the frequency F to find the angular velocity of the waveform. The angular velocity of a waveform is the ratio of the rate of change in angle (radians/sec), and is related to the rate of change in the magnitude of the waveform.

The idea here is to find the highest capacitance value possible that will not overload the transformer. By substituting the impedance of the transformer (Z) for the capacitive reactance (X_C) in the formula, the capacitor will impose a load of Z ohms, which is equal to the maximum current capabilities of the transformer. If the capacitance is too high, the load on the transformer will be too great which causes excess stress on the transformer, and the amount of power transferred to the capacitor will be decreased. However, the capacitance may be less than the calculated maximum value. If the capacitance is less, you won't get as much of a bang and you won't be utilizing the full capabilities of your transformer (the capacitor will not draw as much current as the transformer is capable of delivering).

For the variable F, plug in 60, because the AC line source operates at 60 hertz. Plug in your transformer's impedance for variable (X_C), and solve for C (this is shown in the formula on the right).

To convert farads to microfarads, multiply by 1,000,000, and divide microfarads by 1,000,000 to find farads. The adapted formula below uses microfarads instead of farads.

$$2\pi X_C FC = 1{,}000{,}000 \qquad C = \frac{1{,}000{,}000}{2\pi FZ}$$

If you have decided to use a lower capacitance value, then you can plug in your capacitance value, and solve for (X_C). You can use this impedance value to calculate how much current your capacitor will draw from the secondary coil of the transformer by using Ohm's law. Divide the transformer's voltage by the capacitive reactance to find the current flow. From here, you can also calculate approximately how much current the primary coil of the transformer will draw based on the load imposed on its secondary. This is not really necessary, unless you are just curious or need to know the current draw of the transformer. If a load across the output of a transformer consumes W watts of power, then the primary of the transformer consumes W watts of power. In other words, power in equals power out. Ohm's law also states that P = IE where P = power in watts, I = amps, and E = volts. Thus, the following formula is derived:

$E_P I_P = E_S I_S$

E_P = primary voltage
I_P = primary current in amps
E_S = secondary voltage
I_S = secondary current in amps

Primary Coil Calculations

The next calculations will be for designing the primary coil. The size of the primary is an important factor. You should generally have about 10 to 15 turns in your primary coil, and tapped at a slightly lower number. This usually gives the best output for most coils, but needs to be varied by experiment to suit your design. You should use an adjustable tap for your primary coil to fine tune it. For example, when calculating the inductance of your primary, plug in a number of 12 turns, but actually build your primary with 14 or 15 turns to allow for tuning. The primary coil should be wound from thick conductor with a large surface area, such as copper tubing. The spacing between turns in the primary coil must be large enough to prevent arcing between turns. Spacing should usually not exceed 1/2 inch between turns.

There are 4 basic types of primary coil designs.

1. The single turn primary. It is a single turn at the base of the secondary coil, and usually has a large diameter
2. The helix primary. It is a single-layer, cylinder-shaped coil, just like the secondary coil, but much shorter.
3. The inverse conical section, or saucer coil. It is shaped like a V, and usually rises at a 30° angle.
4. The Archimedes spiral, or flat pancake coil. It is a flat coil shaped in a spiral.

I do not recommend the first two coil designs, based on what I have read. They are not as efficient as the second two designs. The saucer coil is good for small to medium power Tesla coils. I would not use it on a high-power coil, because the sparks from the toroid are more prone to strike it since it is upraised. For high-power coils, the flat spiral is the best type of primary. The diameter of these two types of coils is usually about equal to the length of the secondary winding.

Here are the primary coil inductance formulas. You can use them to find the inductance of the primary you are designing, or to design the dimensions of a primary to have a specific inductance. You will need to know the inductance of the primary coil to calculate the resonant frequency of the tank circuit.

Archimedes Spiral

$$L = \frac{(NR)^2}{8R + 11W}$$

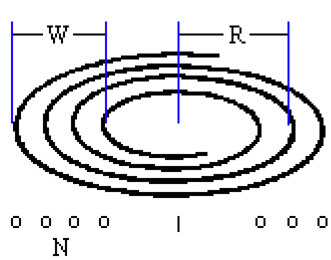

L = inductance of coil in microhenrys (μH)
R = average radius of the coil in inches
N = number of turns
W = width of the coil in inches

Helical Coil

$$L = \frac{(NR)^2}{9R + 10H}$$

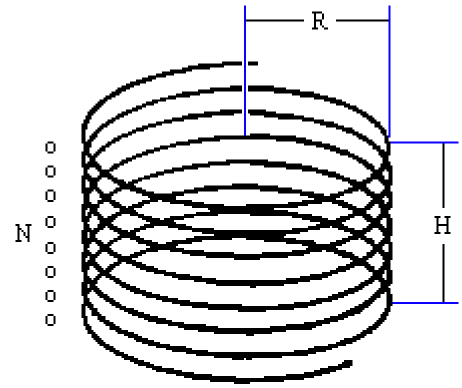

L = inductance of coil in microhenrys (μH)
N = number of turns
R = radius of coil in inches (Measure from the center of the coil to the middle of the wire.)
H = height of coil in inches

Inverse Conical Coil

$$L_1 = \frac{(NR)^2}{9R + 10H} \quad L_2 = \frac{(NR)^2}{8R + 11W}$$

$$L = \sqrt{(L_1 \sin(x))^2 + (L_2 \cos(x))^2}$$

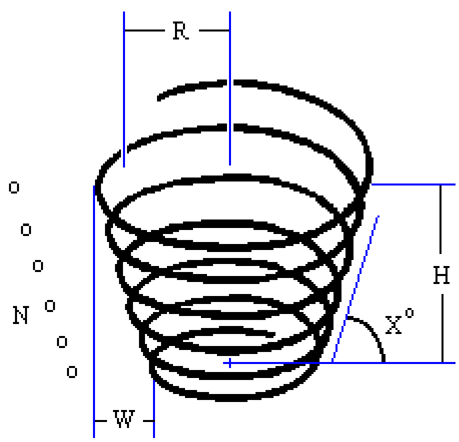

L = inductance of coil in microhenrys (μH)
L_1 = helix factor
L_2 = spiral factor
N = number of turns
R = average radius of coil in inches
H = effective height of the coil in inches
W = effective width of the coil in inches
X = rise angle of the coil in degrees

In some formulas, you will have to input the inductance value in henrys, not microhenrys. To convert microhenrys to henrys, divide by 1,000,000 or 10^6, and multiply henrys by 1,000,000 to find microhenrys.

As discussed on the *How a Tesla Coil Works* page, a resonant circuit is in tune when the inductive reactance (X_L) equals the capacitive reactance (X_C) at the same frequency. Here are the inductive and capacitive reactance formulas:

Inductive Reactance

$$X_L = 2\pi FL$$

X_L = inductive reactance in ohms
F = frequency in hertz
L = inductance in henrys

Capacitive Reactance

$$2\pi X_c FC = 1 \longrightarrow X_c = \frac{1}{2\pi FC}$$

X_C = Capacitive reactance in ohms
F = frequency in hertz

C = capacitance in farads

The circuit will resonate at the frequency F for which $X_L = X_C$. So, by the combination of these two formulas the resonant circuit formula is derived.

$$X_L = X_C \longrightarrow 2\pi FL = \frac{1}{2\pi FC} \longrightarrow 4\pi^2F^2LC = 1$$

F = frequency in hertz
L = inductance in henrys
C = capacitance in farads

If you want to solve for F, use the simplified formula below.

$$F = \frac{1}{2\pi\sqrt{LC}}$$

The resonant circuit formula can be used to find the frequency of a given LC circuit, or to find the necessary inductance and capacitance to make the circuit resonate at a frequency F. Depending on how you are designing your Tesla coil, you will use this formula to calculate the missing variable in the tank circuit.

Next, I will explain secondary coil design. As mentioned before, you may have to make these calculations several times to make everything fit.

Secondary Coil Design

The next step is to design the secondary coil. The secondary coil must resonate at the same frequency as the primary coil in the tank circuit. One of the factors to consider in designing the secondary coil is the aspect ratio, which is the ratio of the length of the secondary coil to the diameter of the secondary coil. The chart below shows the basic guidelines that typically produce the best output. You do not have to follow this exactly.

Secondary coil aspect ratios:

All dimensions are in inches.

Coil Diameter	Aspect Ratio	Length of Coil
3	6:1	18
4	5:1	20
5	4.5:1	22.5
6	4:1	24
7	3.5:1	24.5
8	3.1	24
8+	3:1	multiply by 3 to find length

The general guideline for the type of wire to use for the secondary coil, is 22 gauge or thicker. Really thin wire is a bad idea. It causes too much resistance and too much self capacitance. Use a high-quality enameled magnet wire. Don't use regular insulated hookup wire, unless you want to effectively space out the turns. Below is a list of wire types compiled by Richard Quick.

AWG SIZE	D.C. OHMS PER 1000 FT	WIRE DIAM INCHES	APPROX. TURNS PER INCH, SOLID ENAMEL COVERED	FEET PER POUND
1	.1264	.2893	X	3.947
2	.1593	.2576	X	4.977
3	.2009	.2294	X	6.276
4	.2533	.2043	X	7.914
5	.3195	.1819	X	9.980
6	.4028	.1620	X	12.58
7	.5080	.1443	X	15.87
8	.6405	.1286	7.6	20.01
9	.8077	.1144	8.6	25.23
10	1.018	.1019	9.6	31.82
11	1.284	.0907	10.7	40.12
12	1.619	.0808	12.0	50.59
13	2.042	.0720	13.5	63.80
14	2.524	.0641	15	80.44
15	3.181	.0571	16.8	101.40
16	4.018	.0508	18.9	127.90
17	5.054	.0453	21.2	161.3
18	6.386	.0403	23.6	203.4
19	8.046	.0359	26.4	256.5
20	10.13	.0320	29.4	323.4
21	12.77	.0285	33.1	407.8
22	16.20	.0253	37.0	514.2
23	20.30	.0226	41.3	648.4
24	25.67	.0201	46.3	817.7
25	32.37	.0179	51.7	1031
26	41.02	.0159	58.0	1300
27	51.44	.0142	64.9	1639
28	65.31	.0126	72.7	2067
29	81.21	.0113	81.6	2607
30	103.7	.0100	90.5	3287
31	130.9	.0089	101	4145
32	162.0	.0080	113	5227
33	205.7	.0071	127	6591
34	261.3	.0063	143	8310
35	330.7	.0056	158	10480
36	414.8	.0050	175	13210
37	512.1	.0045	198	16660
38	648.2	.0040	224	21010
39	846.6	.0036	248	26500
40	1079	.0031	282	33410
41	1323	.0028		
42	1659	.0025		
43	2143	.0022		
44	2593	.0020		
45	3348	.00176		
46	4207	.00157		
47	5291	.00140		

"For winding Tesla secondary coils the general consensus is to use number 22 AWG, or larger diameter, double Formvar Magnet wire. Often surplus partial spools of odd wire sizes are found, so I took time to post a more complete chart of AWG numbers and fractional diameters than is usually available. All information above is approximate despite the decimal places. Wire diameter, turns per inch, resistance, feet per pound, etc. all vary slightly from one manufacturer to another. Insulation thickness will vary depending upon type and the supplier. Turns per inch will also vary with the quality of the winding. In practice, Tesla coil windings are not perfect (nor do they need to be) so expect some slight variations from the specification table above."
Richard Quick

Choosing the coil form is the next step in designing the secondary coil. Your coil should usually not be much less than 4 inches in diameter. Large Tesla coils can have coil forms over 12 inches in diameter. The type of coil form used can be important, so try to use a low loss material. PVC is commonly used for Tesla coils because it is cheap, but it is a very high loss material. Materials such as polyethylene, polystyrene, polypropylene, lexan, and Plexiglas are more suitable. The coil form should be as thin as possible, such as 1/8 inch. Unless you really care about that extra inch or two of spark length, use whatever is cheapest.

Once you have chosen a wire gauge and coil diameter, you need to calculate the number of turns and length of the coil. Use the formula below to make these calculations. Secondary coils should generally have 800 to 1000 total turns.

$$L = \frac{\pi DAH}{12} \qquad T = AH$$

L = length of wire in feet
D = outer diameter of coil form
H = height of coil in inches
A = number of turns per inch
T = total number of turns

To find the approximate number of turns per inch, use the wire gauge chart to find the average thickness of the wire you chose.

$$A = \frac{1}{B}$$

A = turns per inch
B = thickness of wire in inches

These coil calculations are approximate and will have to be made again when the secondary coil construction is complete. Wire thickness for the same gauge wire will vary for each manufacturer. The same is also true for coil forms. Even though a piece of pipe may be labeled as being 5 inches, it may actually have an outer diameter of 5.25 inches.

After the dimensions of the secondary coil are designed, you need to find the inductance of the coil. Use the helical coil inductance formula below.

$$L = \frac{(NR)^2}{9R + 10H}$$

L = inductance of coil in microhenrys (µH)
N = number of turns
R = radius of coil in inches
H = height of coil in inches

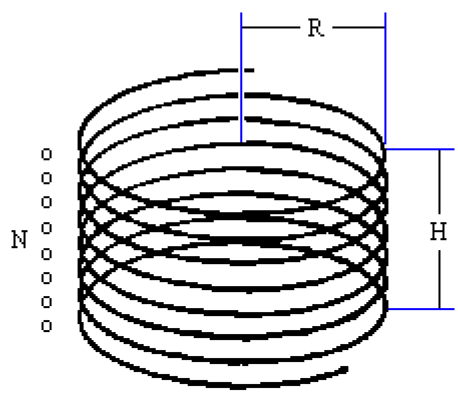

This inductance value will be used to find the necessary toroid capacitance.

Toroid Calculations

The toroid (discharge terminal) is just as critical as every other component in a Tesla coil. The toroid is what tunes the resonant frequency of the secondary circuit. The toroid is one plate of a virtual capacitor. The dielectric is air, and the second plate is the earth itself.

The inductive reactance of the secondary and the capacitive reactance of the toroid need to be near the designed operating frequency. The primary coil is then tapped at the appropriate number of turns to exactly match the secondary frequency. The larger the surface area of the toroid, the more capacitance it has. There is also another factor in determining how much toroid capacitance you need, and that is the virtual capacitance of the secondary coil itself. At these frequency levels just about everything becomes a significant capacitor, and the slightest amount of capacitance can affect the frequency of the secondary resonant circuit. The capacitance of the secondary coil needs to be considered in finding the required toroid capacitance value.

The capacitance of the secondary coil is found by the Medhurst formula. This formula is only a good approximation, and is fairly accurate for most coils. There may be a difference of 1 or 2 picofarads from the real capacitance, but this is easily overcome by adjusting the primary coil tap. This formula is inaccurate for very small coils with a large number of turns. The Medhurst formula is shown below.

$$C = 0.29\ L + 0.41\ R + 1.94\sqrt{\frac{R^3}{L}}$$

C = self capacitance in picofarads
R = radius of coil in inches
L = length of coil in inches

The capacitance of the secondary coil depends on many variables such as the dielectric constant of the wire enamel, wire thickness and number of turns, proximity to ground, and is difficult to accurately calculate. Terry Fritz has written an excellent program for accurately measuring the capacitance of a coil with or without a toroid. His ETesla5 is available on his site. You can also measure it. My oscilloscope has been a very helpful tool in designing and tuning my Tesla coil. If you do not have an oscilloscope, then I suggest that you get one. If you can't afford one, then borrow it. It is an extremely useful tool, and will save you a lot of time.

To measure the capacitance of the secondary coil, you will first have to measure its resonant frequency when it is operating without any toroid or extra capacitance. Connect a small wire to the probe of your oscilloscope as an antenna. Run your Tesla coil without any toroid or discharge terminal, and look at the graph on your scope to find the frequency it is resonating at. Simply putting the scope in the same room with a small Tesla coil will give you a good reading, but keep the scope's wire well out of spark range of the coil, or you'll fry it. Then, use your frequency measurement, and the inductance of the secondary coil to calculate its self-capacitance. Use the resonant circuit formula.

The total capacitance of the secondary circuit will be a little less than the sum of coil capacitance and toroid capacitance. Their close proximity causes the toroid to shield some of the secondary coil surface area from ground. Below is the capacitance formula for a toroid.

$$C = 1.4\left(1.2781 - \frac{D_2}{D_1}\right)\sqrt{\pi D_2(D_1 - D_2)}$$

C = capacitance in picofarads
D_1 = outside diameter of toroid in inches
D_2 = diameter of cross section of toroid in inches

Capacitance of a sphere:

I do not recommend using a sphere as the secondary capacitor, because it is not as space efficient as a toroid, and a toroid is also much more impressive-looking. A toroid works better than a sphere in reducing the intensity of the electric field near the top of the secondary, thus preventing sparks from breaking out of the top secondary coil turns. Below is the formula for a sphere anyway.

$$C = \frac{25.4\ R}{9}$$

C = capacitance in picofarads
R = radius in inches

Now that the component specifications have been calculated, a properly tuned Tesla coil can be built. When making these calculations, keep in mind that they are not completely exact. The secondary coil inductance will have to be recalculated once it is built due to tolerances in wire thickness and coil form diameter. When purchasing the material for the coil form, allow an extra 3 inches beyond the actual coil length. You can always cut off extra material after winding if necessary. Recalculate the inductance of the secondary coil after it has been wound by carefully counting the turns per inch and measuring its dimensions. The Building the Components pages explain the basic steps for constructing each component of a Tesla coil.

RotarySparkGap

My motor controller is designed to use a DC or AC motor to drive a synchronous rotary spark gap. A synchronous RSG only fires at specific points on the AC waveform that charges the tank capacitor. Mine is configured to allow maximum energy bursts by firing at every peak of the charging waveform which is 120 times per second (for 60Hz AC line).

Usually, synchronous RSG's are built using a synchronous AC motor, which runs at a specific speed corresponding to the AC frequency. The phase of gap firing is adjusted by rotating the motor body so the shaft is at the correct angle to fire during the peaks of the sine wave. However, my digital motor controller allows easy and fine adjustment of phase angle from a safe distance while the Tesla coil is in operation.

Schematic Diagram for Logic

Schematic Diagram for Data Interface Board

This circuit is used to isolate the sensitive logic circuitry from RFI on data wires going into the enclosure. These data wires include the IR sensor module and the circuit test points. Note that there is a separate power supply for the isolation circuitry. The abbreviations for the test module data signals are: ROT = rotary feedback waveform; AC = AC reference; CK = clock ACC = accuracy; PHS = phase correction; FRQ = frequency (speed) correction; OUT = final output; GND = negative supply; VCC = positive supply. Use a 10 foot shielded computer serial cable with DB-9 connectors to link the controller to a test point module. This will be a valuable tool for trouble shooting problems such as construction errors and finding where RFI is affecting the circuit.

PCB Layout for Logic

This circuit uses a double-sided board and has thin traces. You will need a pre-sensitized positive acting photo PC board. You will also need a UV exposing lamp and developer. These are available at good electronics stores (not Radio Shack) and are made by GC Electronics. Print the circuit onto a clear transparency sheet with a good ink jet printer for exposing the PC board.

PCB Layout for Interface Board

I designed this PC board to fit into the remaining space in my project box. All resistors are vertically mounted. Note that some chips are in reverse orientation as indicated on the layout. On the PCB there is only a single pin for the 1.2k pull up and 220 ohm resistors for the optoisolators. The resistors are vertically mounted and the second pins are wired together by hand and connected to Vcc or GND. If this layout is difficult to understand, then just design your own board for this part.

Theory of Operation

The rotary spark gap controller continuously regulates the alignment of the rotary electrodes in relation to the peaks of the capacitor voltage waveform. The capacitor reaches a peak charge at each zero crossing of the AC line due to its reactive phase shift when used with inductive ballast. This allows only peak energy packets to be delivered to the rest of the system.

The regulation circuitry is divided into two main sections: the frequency (speed) comparator, and the phase comparator. The frequency comparator maintains a constant firing rate of 120 hertz, while the phase comparator ensures that the firings occur only when the tank capacitor is fully charged.

Monitoring of the electrode positions is achieved through an IR emitter / receiver module mounted near the side of the wheel, allowing the electrodes to pass through it. The module is positioned to be 90 degrees out of phase with the electrode alignment at the AC peaks, so that when the edge of a rotating electrode first beaks the light path in the IR module, the midpoint between two electrodes is aligned with the stationary electrodes. The regulator causes the electrodes to pass through the IR sensor at the zero crossing of the AC sine wave. When an electrode passes through the sensor the output goes low and then returns high after it passes.

A pull-down resistor is connected to the second lead of the IR receiver to drive the output low when light is blocked by a passing electrode. This lead is connected to the inverting input of OP-Amp U8A that is wired in a voltage comparator configuration. The OP-Amp is used to square up the waveform. The non-inverting input of U8A is connected to a voltage reference point between the high and low logic states. This is biased closer to VCC so that the OP-Amp changes states as soon as the sensor begins to change states. The OP-Amp is in an inverting configuration because the RS flip-flop it drives is positive edge triggered. The optoisolator block shown is part of the data interface circuit.

The output of U8A is buffered and connected to the clock input of U6A, an RS flip-flop in a divide by two configuration. The output of U6A changes states every time an electrode breaks the light path in the IR module. The position of the IR module is important to the phase comparator, which compares the state change to the zero crossing of the AC waveform. The use of both U5A and U5B is necessary to buffer and stabilize the waveform.

Note: A resistor and capacitor are placed on pin 3 of U6A to filter out RFI that tends to be picked up by U5. Use this same configuration on any other data signals that you may find to be malfunctioning due to RFI. Protecting sensitive CMOS devices in a Tesla coil environment is like holding a lightning rod in a storm and trying not to get zapped.

The feedback waveform the electrode sensor is compared to a square wave sample of the AC waveform. The AC is squared up by U8B, also a voltage comparator, which changes states on the zero crossing. The regulator circuit does not take into account polarity when comparing the waveforms, because it is not necessary for certain electrodes to only fire on the positive peak and others on the negative.

The Frequency Comparator

The frequency of firings is measured by counting the number of clock cycles that pass during the positive half of the rotary feedback waveform. A 555 timer, U7, provides the clock at a frequency of 11,040 hertz. Decade counters U14 and U15 count the number of clock cycles. When the rotary feedback goes low the counters are reset to zero. AND gate U1C has its input connected to Q9 of U15 and Q2 of U14, setting the count reference to 92. The frequency comparator will regulate the motor speed so the count reaches 92, causing a firing rate of 120 hertz. If the count does not reach 92, then the regulator will slow down the motor, and speed it up if the count reaches or passes 92. The output of AND gate U1C is connected to the clock input of RS flip-flop U13A wired in a latching configuration. If the count reaches 92, then Q will latch high; otherwise it will remain low. Q of U13A is connected to the data input of RS flip-flop U13B. The rotary feedback is inverted by U5C and connected to the clock input of U13B, which is positive edge triggered. The reset input of U13A is connected to the rotary feedback, which is negative edge triggered. When the rotary feedback goes low it simultaneously transfers the data input of U13B to its Q and resets latch U13A. Q of U13B (the frequency correction data) is connected to the input of a tristate buffer, U10.

Another section of the frequency comparator checks the accuracy of the current frequency. If the frequency is too high or low, then the phase comparator cannot give an accurate correction. This accuracy comparator is part of the circuit that switches the regulator's output between control of the frequency and phase comparators. The two comparators take turns controlling the speed of the motor. If the frequency is not accurate, then the phase comparator will not receive its turn and the frequency comparator will be given more time to correct the motor's speed.

The accuracy comparator checks if the count reaches exactly 92. Count numbers 92 and 93 are sampled by AND gates U1C and U1B respectively. The outputs of the AND gates are each connected to a latch. The 92 count latch serves both the frequency and accuracy comparators. The outputs of the two latches are connected to the inputs of U2C, an exclusive NOR gate. Its output is low only if count 92 is reached, but not 93. A low will allow the phase comparator to take its turn.

The turns are defined by RS flip-flop U6B that is a divide by two counter with the rotary feedback as its clock. It gives the two comparators turns of one complete cycle of the rotary feedback, which is regulated to 1/60 of a second. While the output of U6B is high, the frequency comparator is given control of the regulator's output and the phase comparator while it is low. U6B's output and the accuracy comparator's output are connected to the inputs of OR gate U3B. The output of U3B is connected to the control input of tristate buffer U10A which controls the frequency comparator. U3B's output is inverted and connected to the control input of U10B, which controls the phase comparator. The phase comparator is given control only if the outputs of U6B and the accuracy comparator are both low. The tristates allow data to pass through if the control inputs are high, and only one tristate is allowed to be on at a time. The two outputs of the tristates are connected together to produce the final output of the regulator.

The Phase Comparator

The phase comparator is divided into two main sections: the phase angle comparator and the lead comparator. The phase angle comparator determines whether the rotary and AC waveforms are out of phase by less than or greater than 90 degrees. The lead comparator determines which waveform leads the other, providing the polarity of the phase angle.

The phase angle comparator uses XNOR gate U2A to compare the rotary feedback and AC sample. The output of U2A is high when the polarities of the waveforms are the same. Decade counters U16 and U17 are used to count the number of clock cycles during the positive part of the cycle. If the count reaches 46, then the phase angle is less than 90 degrees; otherwise it is greater than 90 degrees.

The inputs of AND gate U11B are connected to Q6 and Q4 of U16 and U17 respectively. The output of U11B is latched by U12, just as in the frequency comparator. When the output of XNOR U2A goes low, data is transferred to the output of U12B and the counter and latch are reset.

The lead comparator detects if the AC sample or the rotary feedback goes high first. The rotary feedback is connected to the clock input of U9B and the AC sample to the data input. If the AC waveform goes high before the rotary waveform does, then the high AC data will be transferred to Q. If the rotary waveform leads, then the output will be low.

The outputs of the phase angle comparator and lead comparator are connected to the inputs of XNOR gate U2B. The output of U2B gives the phase comparator's correction of motor speed. The phase comparator does not care if the waveforms are 180 degrees out of phase, just so long as the zero crossings (state changes) of the waveforms line up.

The Four Phase Conditions

If the waveforms are less than 90 degrees out of phase, then the phase counter will pass 46 and the output of the phase angle comparator will be high. The phase angle is closer to zero degrees than 180, so the phase comparator will give a correction to maintain a phase angle of zero degrees. If the AC sample leads, then the output of the lead comparator will be high. This will cause the motor to speed up until the phase angle becomes zero.

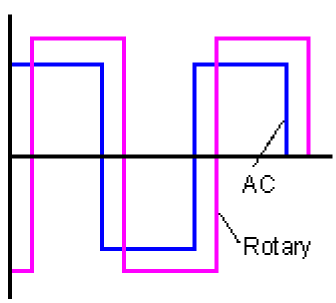

If the rotary feedback leads with an angle less than 90 degrees, then the lead comparator will give a low output and the output of U2B will be low. This will cause the motor to slow down until the phase angle becomes zero.

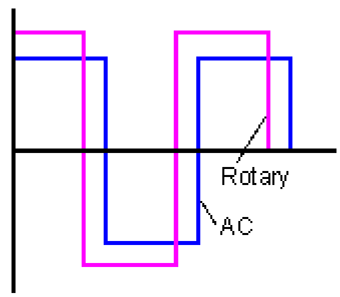

If the waveforms are greater than 90 degrees out of phase, then the output of the phase angle comparator will be low. The phase angle is closer to 180 degrees than zero, so the phase comparator will give a correction to maintain a 180 degree phase angle. If the rotary feedback leads, then the output of the lead comparator will be high, causing the motor to slow down until the phase angle becomes 180 degrees.

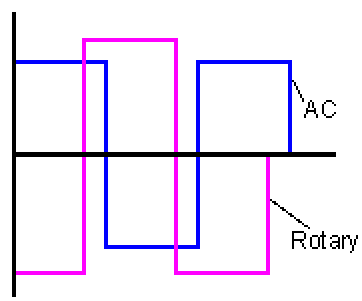

If the AC sample leads with an angle greater than 90 degrees, then the output of the lead comparator will be low and the output of U2B will be high. This will cause the motor to speed up until the phase angle becomes 180 degrees

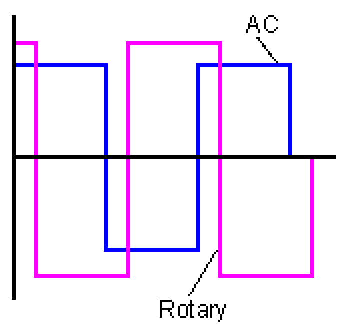

As the phase comparator makes adjustments to the motor speed, the firing frequency will exceed the error tolerance set by the accuracy comparator. This will cause the frequency comparator to take control of the motor and the phase comparator will continue its adjustments on its next turn. Although the phase comparator has strict operating guidelines, it does not take much time for it to correct the phase angle. These guidelines are also necessary to prevent overcorrecting of the phase angle, which would cause the electrode alignment to be erratic. The phase comparator's adjustments cause the rotary feedback waveform to hover around the ideal phase angle of zero or 180 degrees, and within a tolerance of ±10 degrees for my AC motor. Much better accuracy can be achieved with a DC motor.

The motor is controlled by rapidly turning it on and off with a solid state relay or MOSFET for DC motors. The two logic control terminals of the relay are connected to logic ground and to the output of the regulator. When the regulator's output is high, the motor is turned on. To prevent jerking of the motor's shaft position when it is turned on, an inductor is wired in series with the motor. I use a 1000 foot spool of 22 gauge insulated wire for my 1/6 HP AC motor. Since the impedance of the inductor decreases motor speed, the motor speed or number of electrodes must be great enough so that the firing rate is at least 120 hertz without regulation.

The AC reference waveform is supplied by a simple series RC circuit at 24 volts AC. Use a 100V capacitor in series with a potentiometer and connect to the pins of the pot or the capacitor to achieve the desired phase angle. Your high voltage transformer must be inductively ballasted for optimum efficiency. The ballast should be resonant with the tank capacitor at 60 hertz for peak charging voltage. Since the circuit will be series resonant (great resonant rise across tank cap and transformer is possible), do not turn on the coil until the RSG is running. If gap does not fire, then the next cycle will build upon the previous, causing over-voltage. This is why a safety spark gap is a must. The inductive reactance needed in series with the primary to be resonant with a capacitive reactance on the secondary is given by the following formula.

$$X_L = \frac{E_P^2 X_C}{E_S^2}$$

X_L = inductive reactance in ohms for ballast (60 hertz)
E_P = primary voltage
E_S = secondary voltage
X_C = tank cap capacitive reactance (60 hertz)

It is very important to consider that by using this method of charging, the capacitor voltage at full charge will be equal to pi/2 times the peak transformer voltage. This occurs because there is a resonant rise over 2 successive 1/4-cycles of the AC source. Make sure that this will not destroy the tank capacitor or the high voltage transformer. To prevent overdriving the transformer, do not bring your variac above 66%, or you can use a large inductance in series with the secondary of the transformer to ballast it.

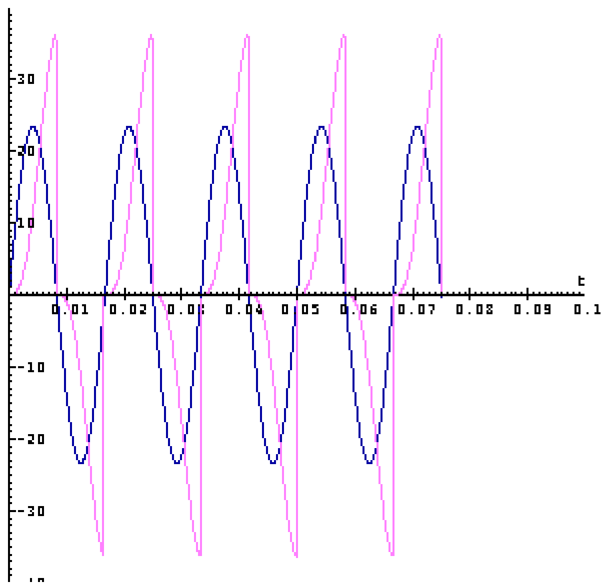

The blue curve is the AC source and the purple curve is the capacitor voltage.

My Tesla Coil Specifications

My Tesla coil project began in August of 1997. Since that time, most of my work has been related to this project. I had originally intended to build a small bipolar coil, but after experimenting with the traditional monopolar design, my project scope rapidly expanded. I designed and constructed each component of the Tesla coil to specifications that would make it suitable in a larger system than my current design at that point. The result of this process is my highly modular Tesla coil with an easily adjustable primary coil, a 75kV DC tank capacitor with selectable capacitance, and a digitally synchronized rotary spark gap. I engineered my Tesla coil to give myself more flexibility in the design of the next component and system capabilities, and to facilitate experimental adjustments.

The specifications of my Tesla coil are as follows.

The Power Transformer

1. The transformer is a 3025 VA instrumentation transformer with a secondary voltage rating of 16,500 volts and a primary of 120 volts. The transformer uses inductive ballast and has AC line filters on each leg of the primary. Input power is controlled by a 120V 45A Powerstat variac, Model D5500H. The variac is connected to the AC line with a 30A appliance plug. All 120 volt wiring uses 10 gauge wire. The transformer was given to me by a friend from the local power company. The 16.5 kV transformers were being phased out of the system and I was able to get one that was hardly used.

Frame

2. The frame is constructed from two circular pieces of 3/4 inch plywood. Its diameter is 33" and height is 22". The lower platform has four casters mounted beneath it for mobility. Four 2x4's support the upper platform. These boards are jointed to the platforms by wood screws and angle brackets.

The Tank Capacitor

3. The tank capacitor is a 0.0278 microfarad polyethylene vertically-stacked plate capacitor. A total of 99 packets of two 60 mil and one 30 mil poly sheets with aluminum foil plates are wired in parallel. The total of 150 mils of dielectric thickness is a conservative value and gives a high safety margin for my 16.5 kV transformer. The typical dielectric strength for polyethylene of 500V per mil rates my capacitor at 75kV DC. The capacitor is contained in a polypropylene plastic storage bin and is filled with Diala-AX transformer oil. The capacitor is actually two independent capacitors of 0.0092 and 0.0186 μF. This allows easy adjustment of power level and frequency for experimentation. There is a common connection and two terminals: one for each capacitance value. These are both used under normal operation at full power. The capacitor is contained in a 15 gallon Rubbermaid bin and filled with 5 gallons of oil.

The Rotary Spark Gap

4. The rotary spark gap uses a 1,000 RPM, 1/6 horsepower motor. It has eight rotary electrodes and four stationary electrodes, creating a series of four gaps. All electrodes are machined from 1/4 inch high-strength stainless steel rods. The stationary electrodes are supported by 1 inch diameter aluminum alloy rods. The stationary electrodes are mounted on Plexiglas platforms to isolate them from the base of the rotary spark gap. The long aluminum rods are excellent heat sinks. The rotary electrode wheel is a circle of 1/4 inch clear Plexiglas and the electrodes are held in place by Allen screws. The Plexiglas may have to be replaced by a more heat-tolerant material later on, but so far it is working fine. The base of the unit is a piece of 5/8 inch plywood with two coats of spar varnish. The rotary gap is coupled with a unique regulation device that I designed to maximize its firing efficiency. I am especially proud of the rotary spark gap because I put a great deal of my time into it, and it is the most elegant component of my Tesla coil.

The Primary Coil

5. The primary coil is a flat spiral coil of 1/4 inch copper tubing. The coil is supported every 90 degrees by four pieces of Plexiglas with holes drilled in them. Each Plexiglas support is cut lengthwise directly through the row of holes. This allows me to simply remove the top section of the support to place or remove the entire primary coil. The coil has 13 turns with an inner diameter of 8.5" and an outer diameter of 20.5"

The Secondary Coil

6. The secondary coil is a single-layer, close-wound coil on a piece of 4 inch PVC pipe. The coil is wound from 774 feet of 22 gauge wire. I used a high quality enameled motor wire, which wasn't cheap. There are 717 turns at about 35 TPI, and the windings are 20.5 inches long. The coil's inductance is 9,990 microhenrys and it has a self capacitance of 9.14 picofarads.

The Toroid

7. The toroid is made from an 8 foot length of 6" diameter flexible aluminum air conditioning ducting. The outer diameter is 32 inches. The total secondary circuit capacitance with the toroid is 44.0pf. The two ends are joined together by a thick coat of epoxy and the seam is hidden with silver paint.

Miscellaneous

8. All tank circuit connections are made with 1 gauge fine stranded welding cable. The base of the secondary coil has a flat copper ground strap to connect it to the RF ground block. The RF ground block is connected to an eight foot radio ground rod with an 8 gauge cable. The lower level of the frame holds the tank capacitor and rotary spark gap. On the upper level of the frame, the primary and secondary coils are mounted. The transformer is mounted by itself on a piece of plywood, also with four casters on the bottom so I can easily roll it around. The complete system with toroid stands 4'5" tall.

Formula Pages

On these pages I have posted Tesla coil formulas for a reference to those who prefer to make calculations on paper.

Ohm's Law

$E = IZ$ $P = IE$

E = volts
I = current in amps
Z = impedance or resistance in ohms
P = power in watts

Transformer Input and Output

$E_P I_P = E_S I_S$

E_P = primary voltage
I_P = primary current in amps
E_S = secondary voltage
I_S = secondary current in amps

Capacitive Reactance

$$X_C = \frac{1}{2\pi FC}$$

X_C = capacitive reactance in ohms
F = frequency in hertz
C = capacitance in farads

Inductive Reactance

$$X_L = 2\pi FL$$

X_L = inductive reactance in ohms
F = frequency in hertz
L = inductance in henrys

Resonant Circuit Formula

$$4\pi^2F^2LC = 1$$

$$F = \frac{1}{2\pi\sqrt{LC}}$$

F = frequency in hertz
L = inductance in henrys
C = capacitance in farads

Spiral Coil Inductance

$$L = \frac{(NR)^2}{8R + 11W}$$

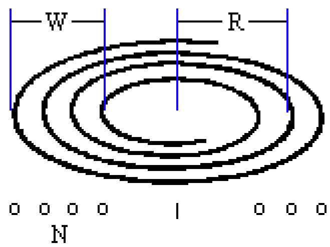

L = inductance of coil in microhenrys (µH)
R = average radius of the coil in inches
N = number of turns
W = width of the coil in inches

Helical Coil Inductance

$$L = \frac{(NR)^2}{9R + 10H}$$

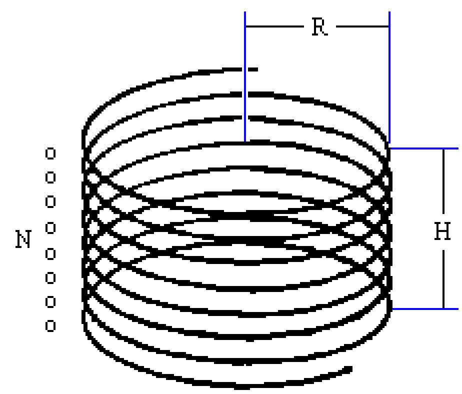

L = inductance of coil in microhenrys (µH)
N = number of turns
R = radius of coil in inches (Measure from the center of the coil to the middle of the wire.)
H = height of coil in inches

Inverse Conical Coil Inductance

$$L_1 = \frac{(NR)^2}{9R + 10H} \qquad L_2 = \frac{(NR)^2}{8R + 11W}$$

$$L = \sqrt{(L_1 \sin(x))^2 + (L_2 \cos(x))^2}$$

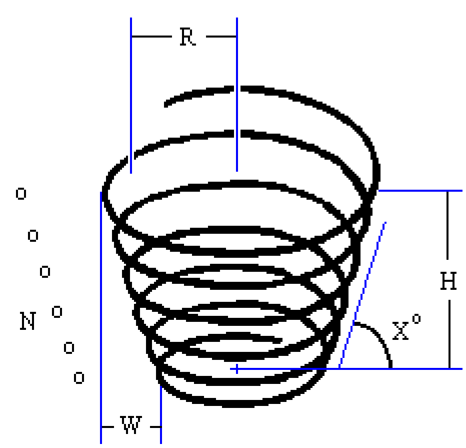

L = inductance of coil in microhenrys (µH)
L_1 = helix factor
L_2 = spiral factor
N = number of turns
R = average radius of coil in inches
H = effective height of the coil in inches
W = effective width of the coil in inches
X = rise angle of the coil in degrees

Secondary Coil Dimensions

$$L = \frac{\pi DAH}{12} \qquad T = AH \qquad A = \frac{1}{B}$$

L = length of wire in feet
D = outer diameter of coil form in inches
H = height of windings in inches
A = number of turns per inch
T = total number of turns
B = thickness of wire in inches

Medhurst

$$C = 0.29\ L + 0.41\ R + 1.94\sqrt{\frac{R^3}{L}}$$

C = self capacitance in picofarads
R = radius of secondary coil in inches
L = length of secondary coil in inches

Toroid Capacitance

$$C = 1.4\left(1.2781 - \frac{D_2}{D_1}\right)\sqrt{\pi D_2(D_1 - D_2)}$$

C = capacitance in picofarads
D_1 = outside diameter of toroid in inches
D_2 = diameter of cross section of toroid in inches
This equation courtesy Bert Pool.

Sphere Capacitance

$$C = \frac{25.4\,R}{9}$$

C = capacitance in picofarads
R = radius in inches

Plate Capacitors

$$C = \frac{0.224KA(N - 1)}{1,000,000D}$$

C = capacitance in microfarads
K = dielectric constant
A = area of each plate in square inches
N = number of plates
D = distance between plates in inches (thickness of dielectric)

Leyden Jar Capacitors

$$C = \frac{0.224\pi KD(H + 0.25D)}{1,000,000T}$$

C = capacitance in microfarads
K = dielectric constant
D = diameter of jar in inches
H = height of jar in inches
T = thickness of jar in inches

AC RMS and Peak Voltage

$E_{RMS} = 0.7071 \cdot E_P$

E_{RMS} = RMS voltage
E_P = peak voltage

Rotary Spark Gap Firings per Second

$$F = \frac{RE}{60}$$

F = firings per second (hertz)
R = motor RPM rating
E = number of rotary electrodes

Rotary Spark Gap Electrode Speed

$$S = \frac{\pi RD}{1056}$$

S = electrode speed (MPH)
R = motor RPM rating
D = diameter of electrode placement circle (inches)

Energy for L and C

Capacitance	Inductance
$J = 0.5\ V^2\ C$	$J = 0.5\ I^2\ L$

J = joules of energy stored
V = peak charge voltage
I = peak current
C = capacitance in farads
L = inductance in henries

I stated peak values of V and I because I want to emphasize not to use RMS values. The energy stored at any given time is of course: $J(t) = 0.5\ [V(t)]^2\ C$ and $J(t) = 0.5\ [I(t)]^2\ L$.

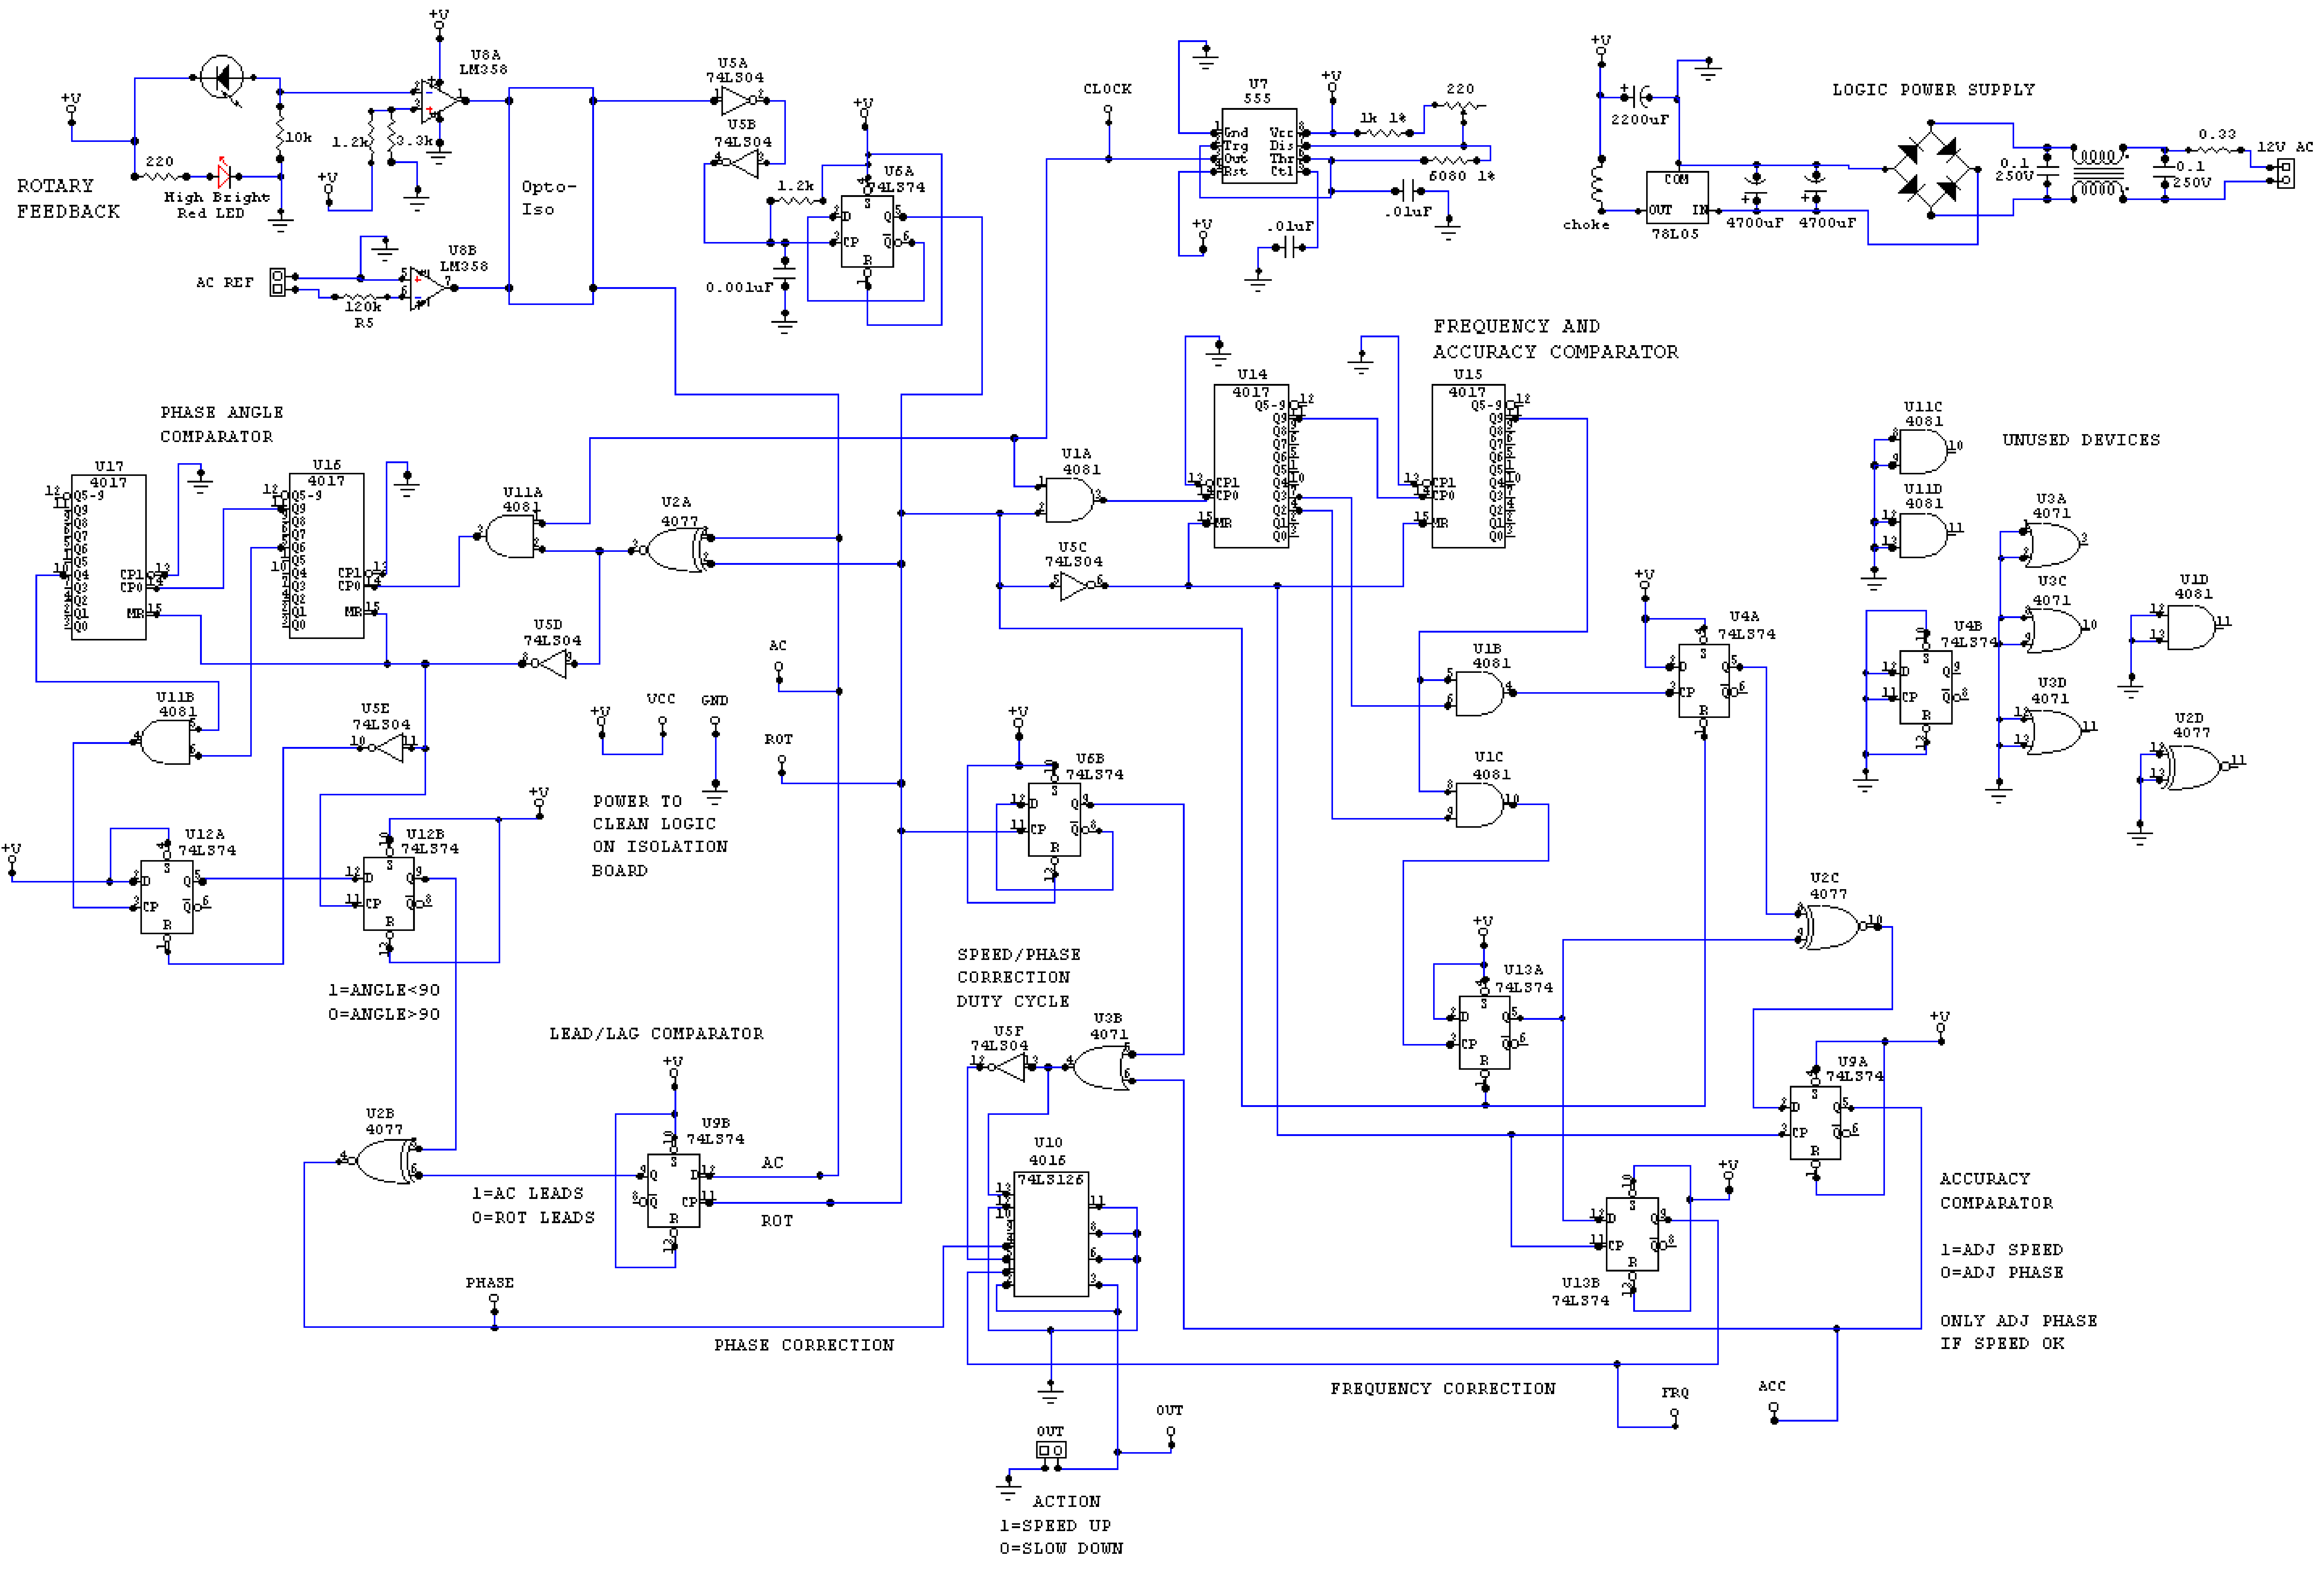

ROTARY
FEEDBACK
220
High Bright
Red LED
10k
1.2k
3.3k
+V
U8A
LM358
U8B
LM358
AC REF
120k
R5
Opto-
Iso
U5A
74LS04
U5B
74LS04
1.2k
0.001uF
U6A
74LS74
CLOCK
U7
555
Gnd
Trg
Out
Rst
Vcc
Dis
Thr
Ctl
1k 1*
220
6080 1*
.01uF
.01uF
2200uF
choke
COM
OUT
IN
78L05
4700uF
4700uF
LOGIC POWER SUPPLY
0.1
250V
0.1
250V
0.33
12V AC
FREQUENCY AND
ACCURACY COMPARATOR
U14
4017
U15
4017
PHASE ANGLE
COMPARATOR
U17
4017
U16
4017
U11A
4081
U2A
4077
U1A
4081
U5C
74LS04
U5D
74LS04
UNUSED DEVICES
U11C
4081
U11D
4081
U3A
4071
U3C
4071
U3D
4071
U1D
4081
U2D
4077
U4B
74LS74
U11B
4081
U5E
74LS04
U12A
74LS74
U12B
74LS74
1=ANGLE<90
0=ANGLE>90
POWER TO
CLEAN LOGIC
ON ISOLATION
BOARD
+V
VCC
GND
AC
ROT
U6B
74LS74
U1B
4081
U1C
4081
U4A
74LS74
U2C
4077
U13A
74LS74
U9A
74LS74
U13B
74LS74
SPEED/PHASE
CORRECTION
DUTY CYCLE
U5F
74LS04
U3B
4071
LEAD/LAG COMPARATOR
U9B
74LS74
AC
ROT
1=AC LEADS
0=ROT LEADS
U2B
4077
PHASE
PHASE CORRECTION
U10
4016
74LS126
FREQUENCY CORRECTION
FRQ
ACC
OUT
OUT
ACTION
1=SPEED UP
0=SLOW DOWN
ACCURACY
COMPARATOR
1=ADJ SPEED
0=ADJ PHASE
ONLY ADJ PHASE
IF SPEED OK

DATA I/O INTERFACE TO LOGIC

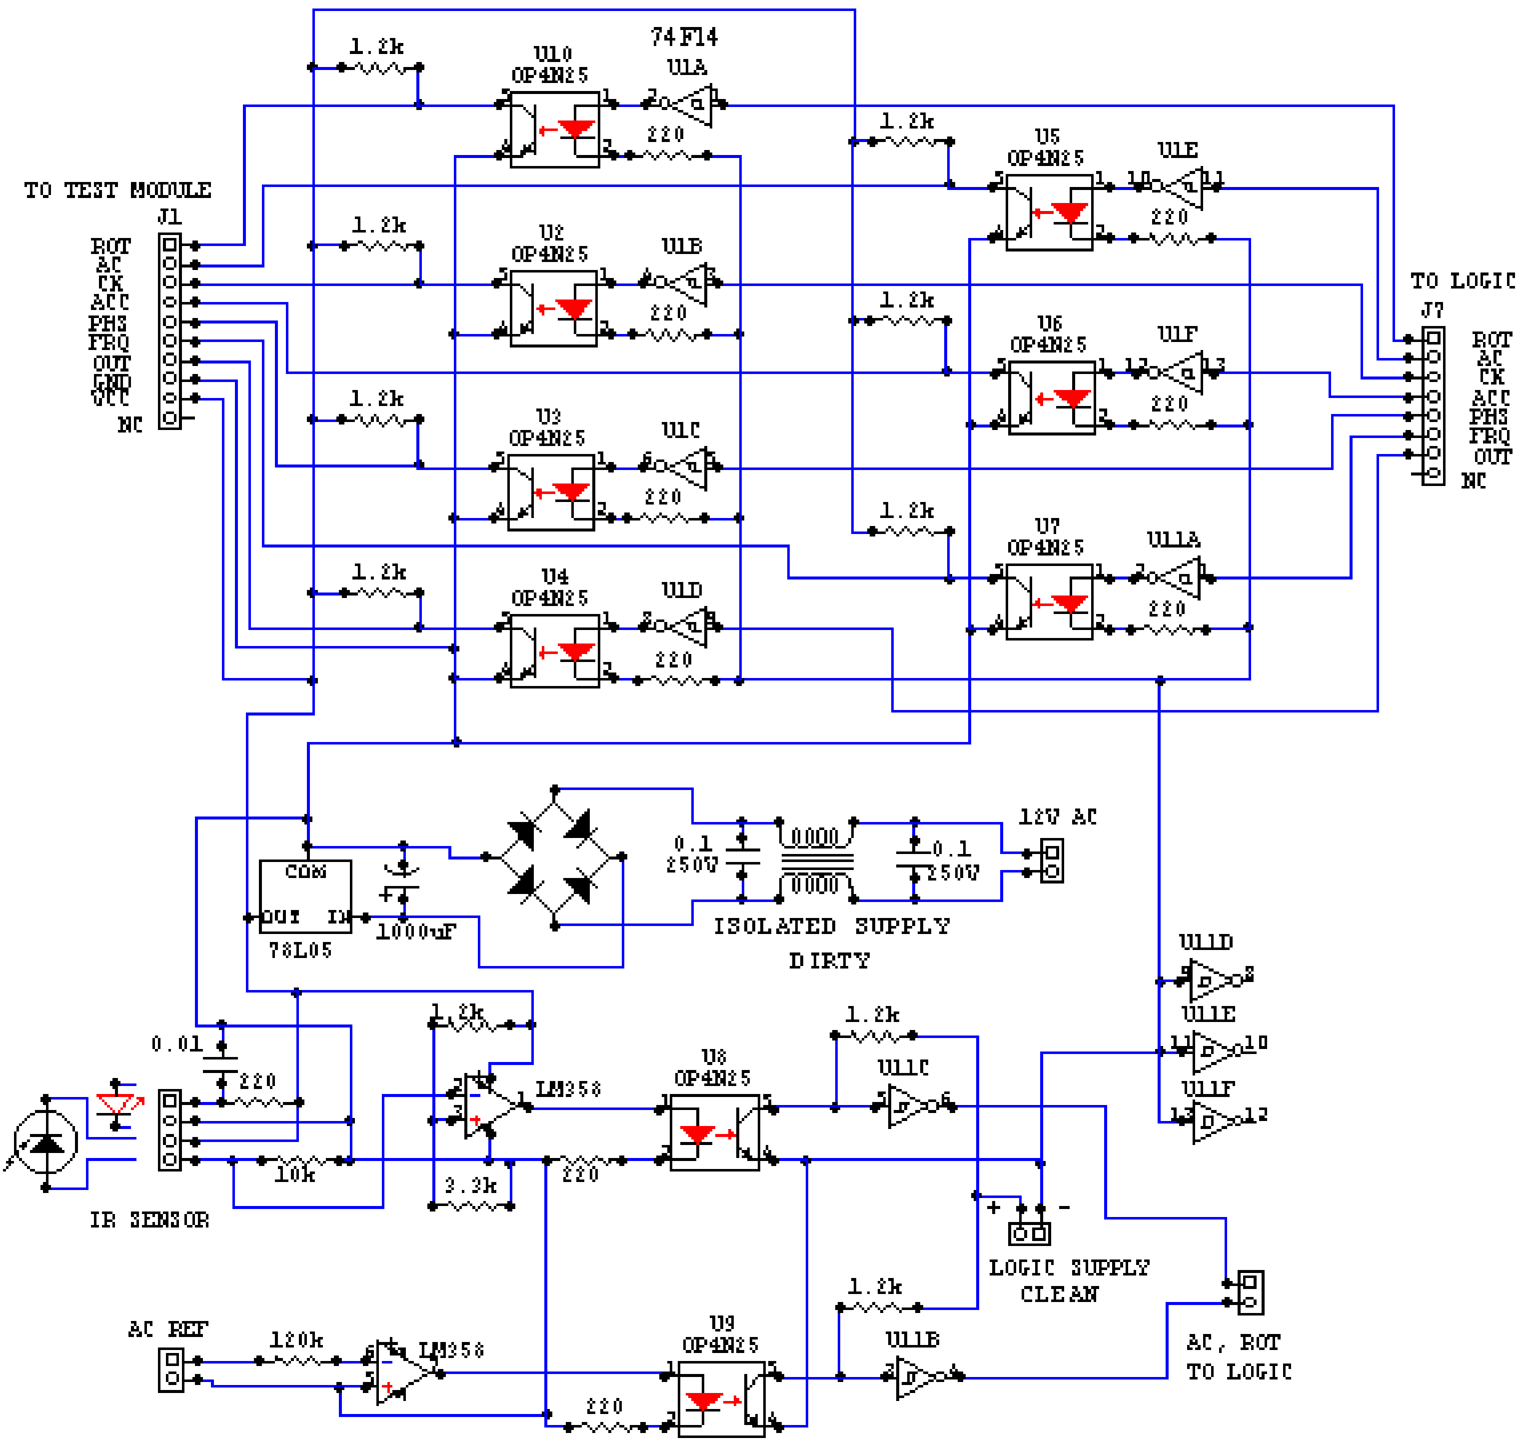

Rotary Spark Gap

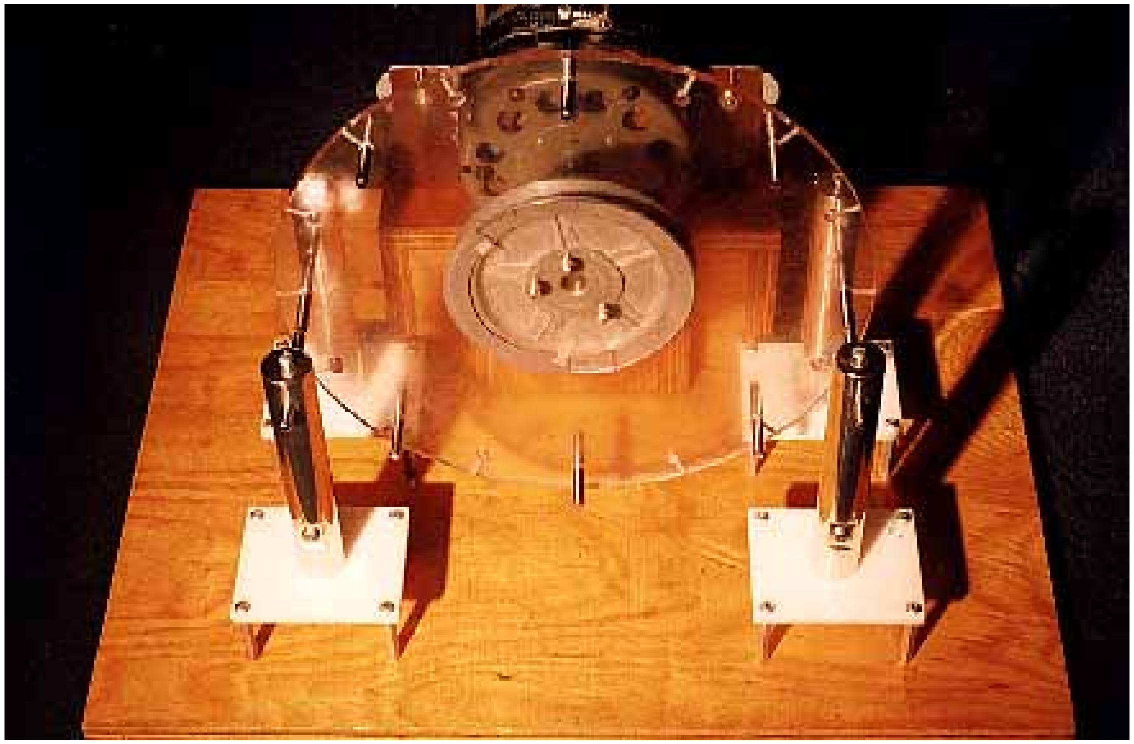

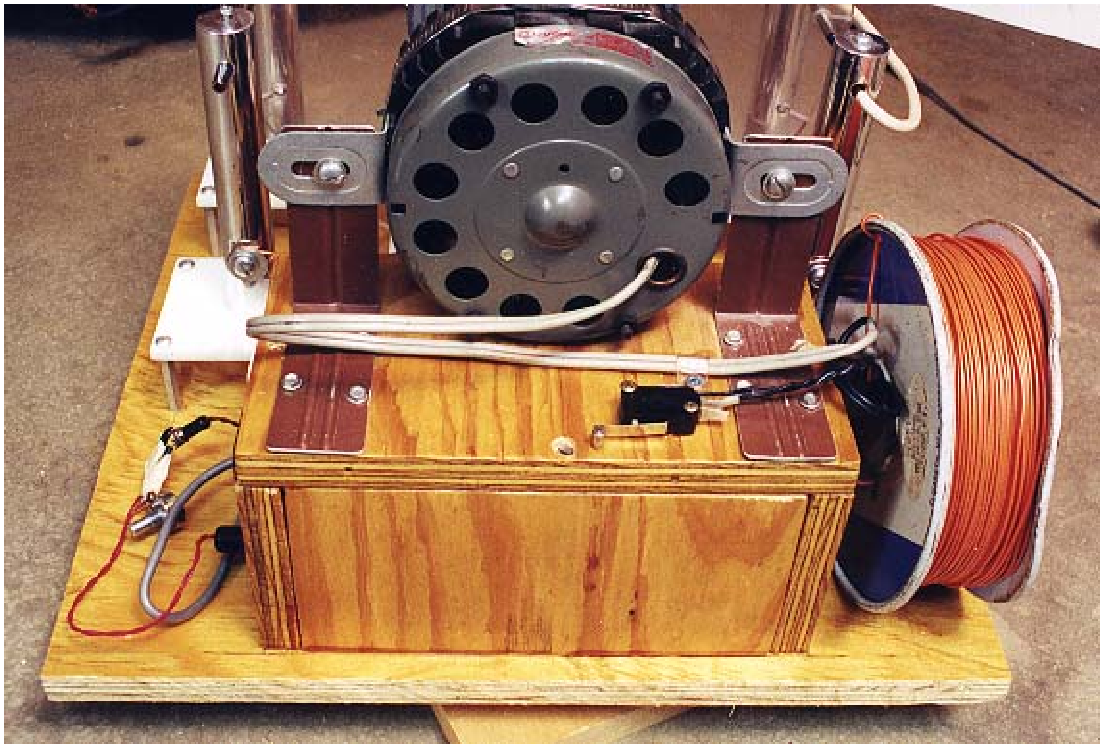

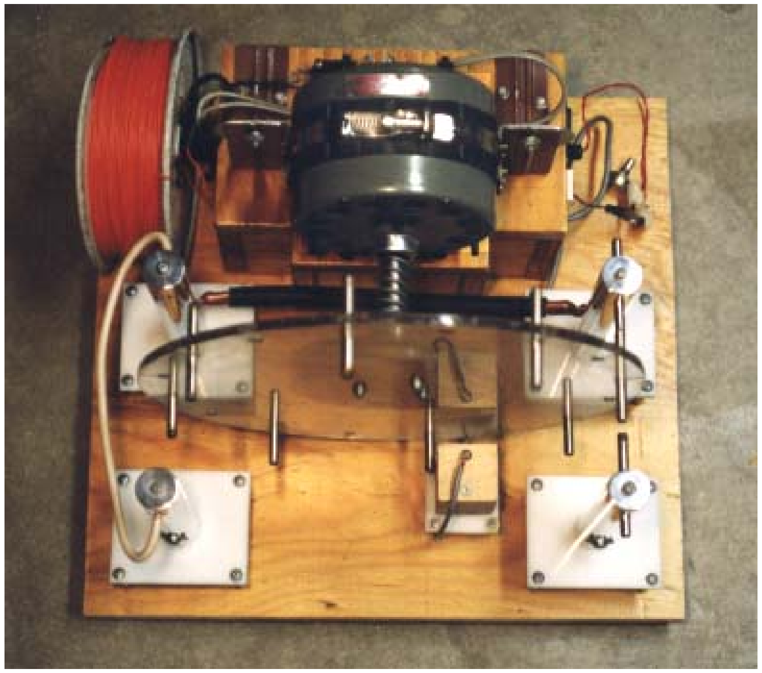

Rotary Spark Gap Motor Controller

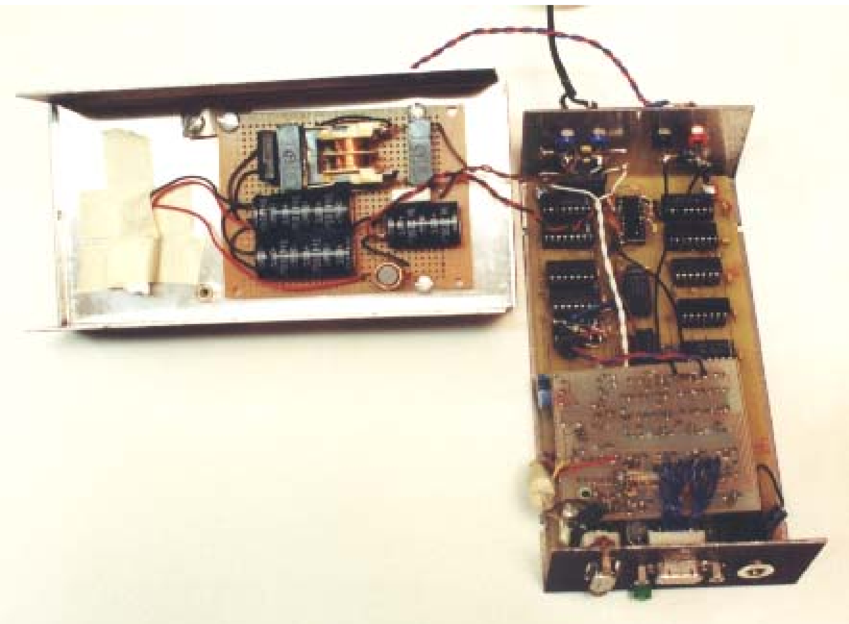

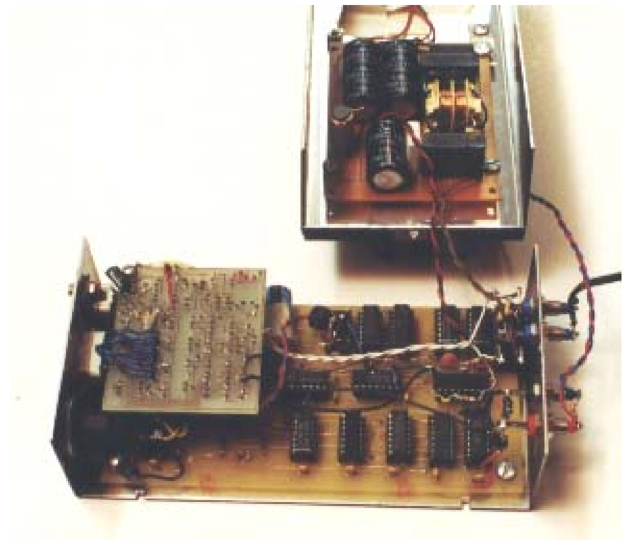

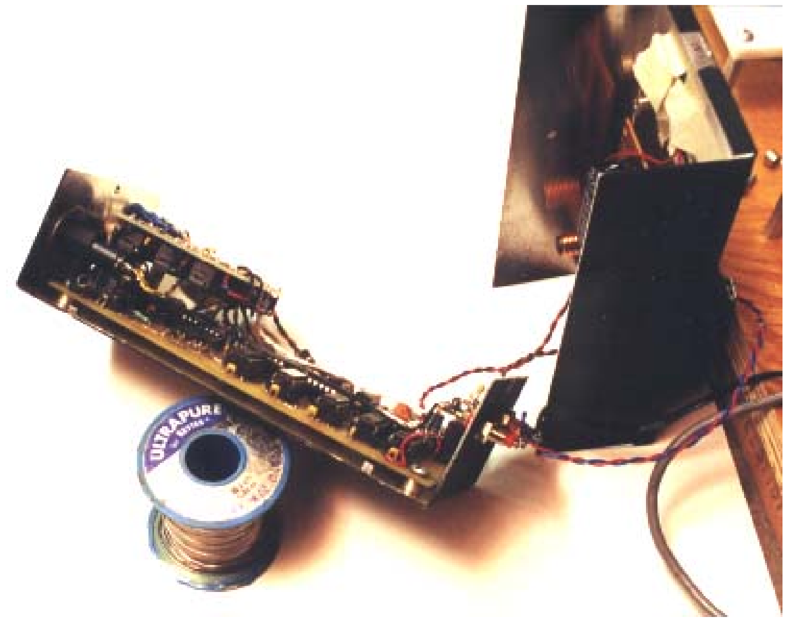

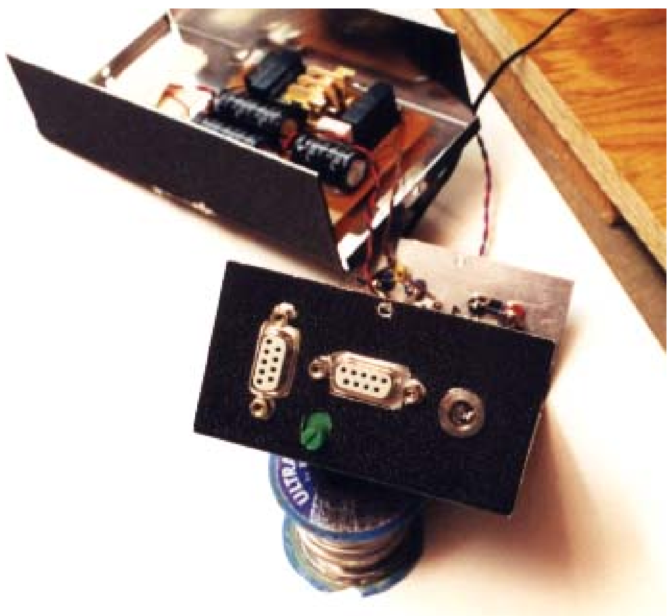

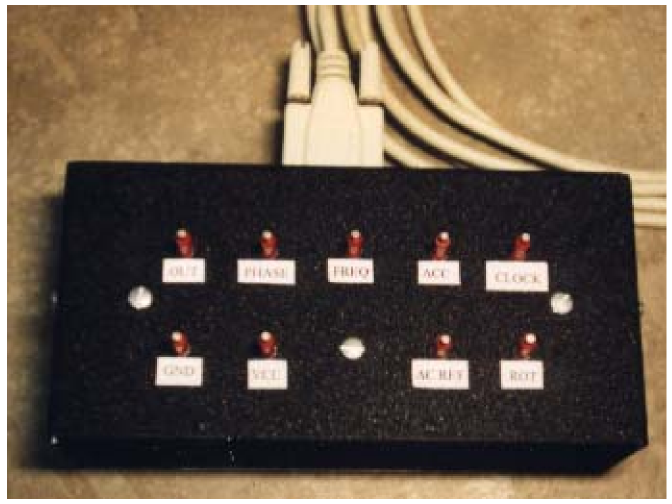

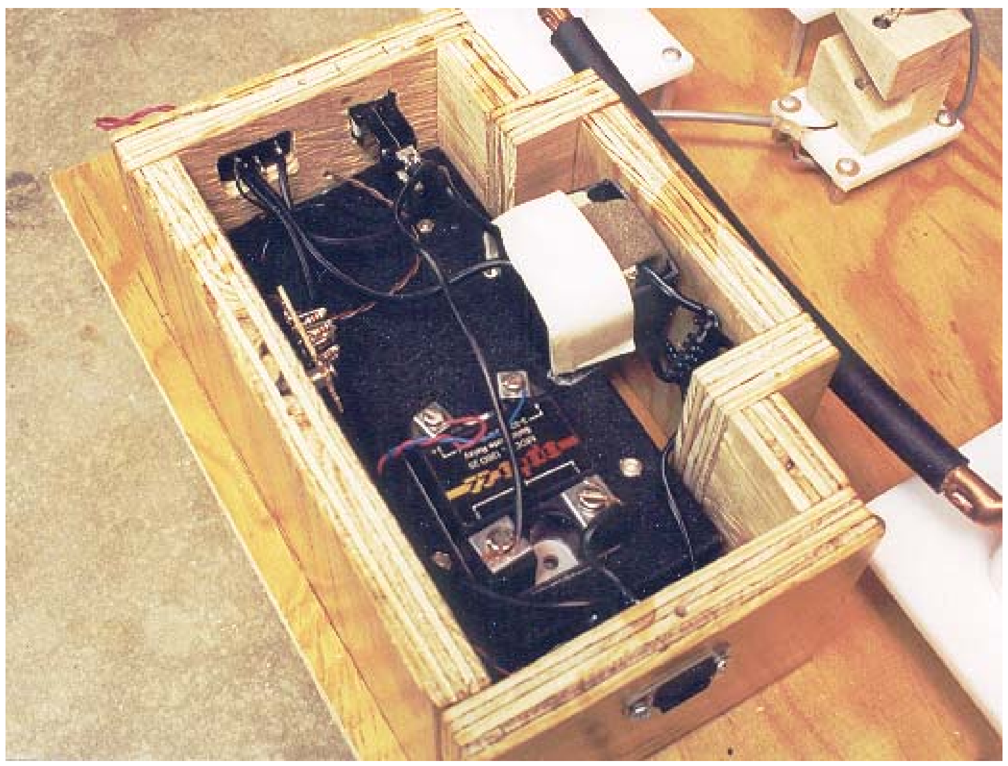

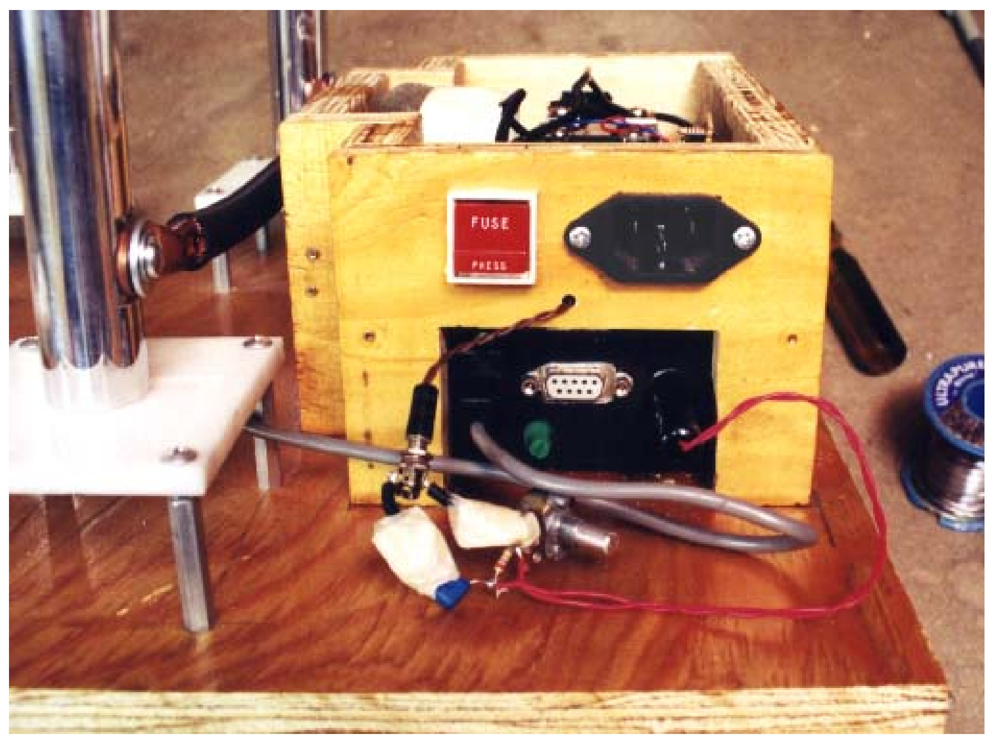

Transformer

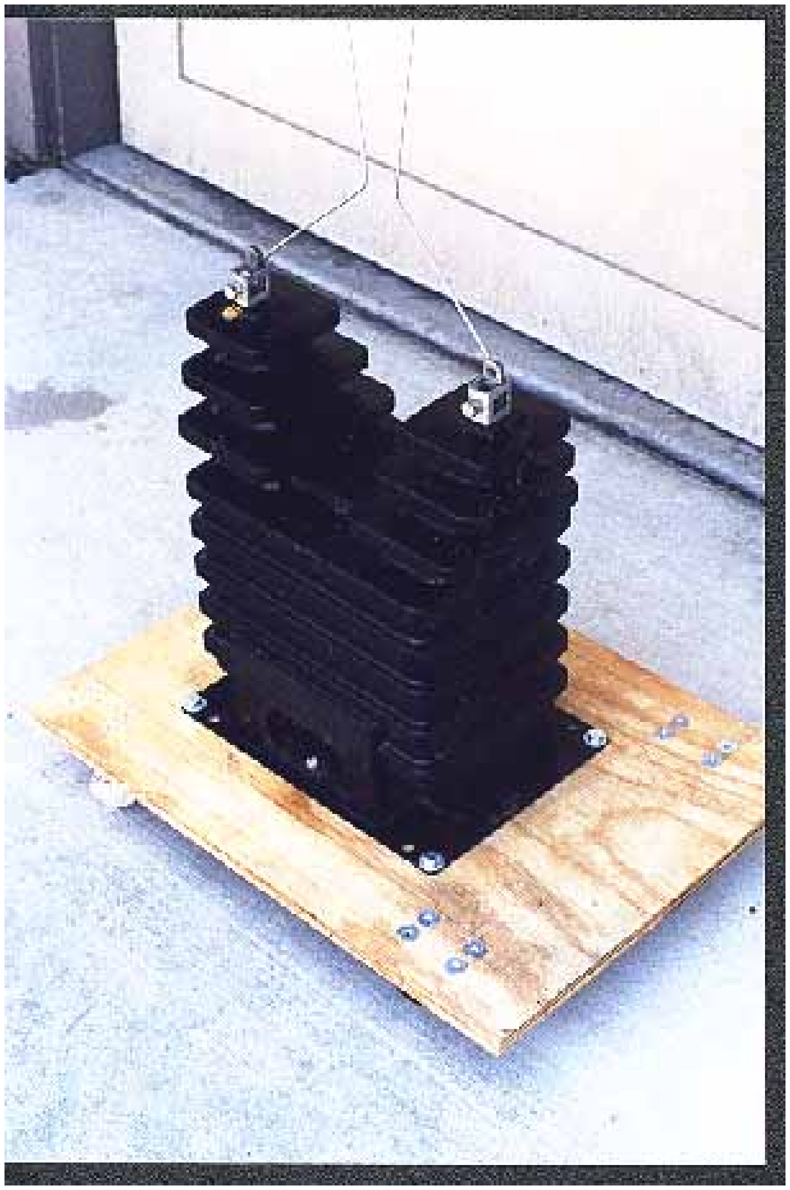

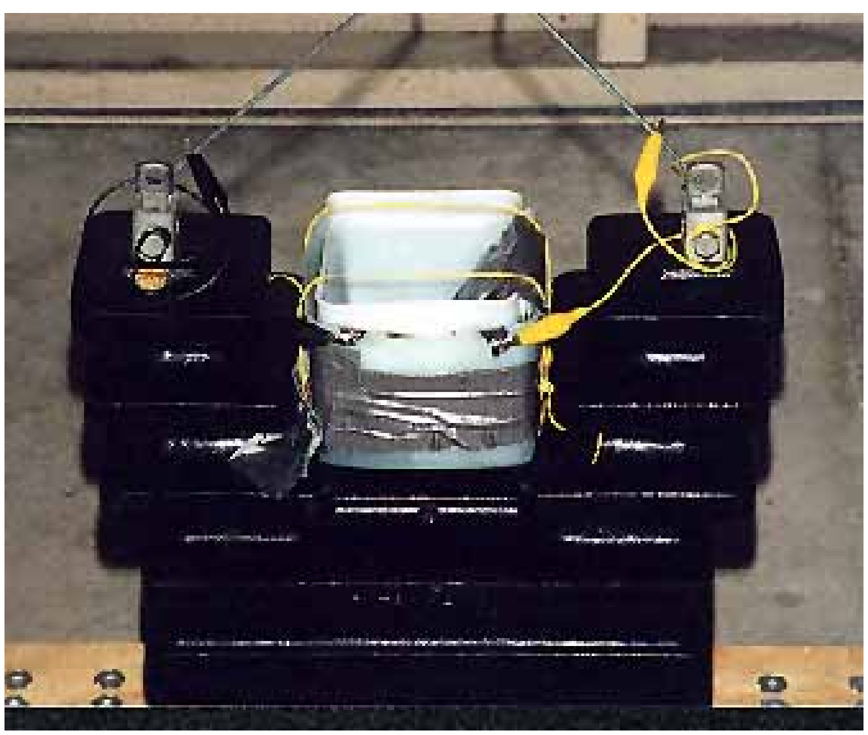

Tank Capacitor

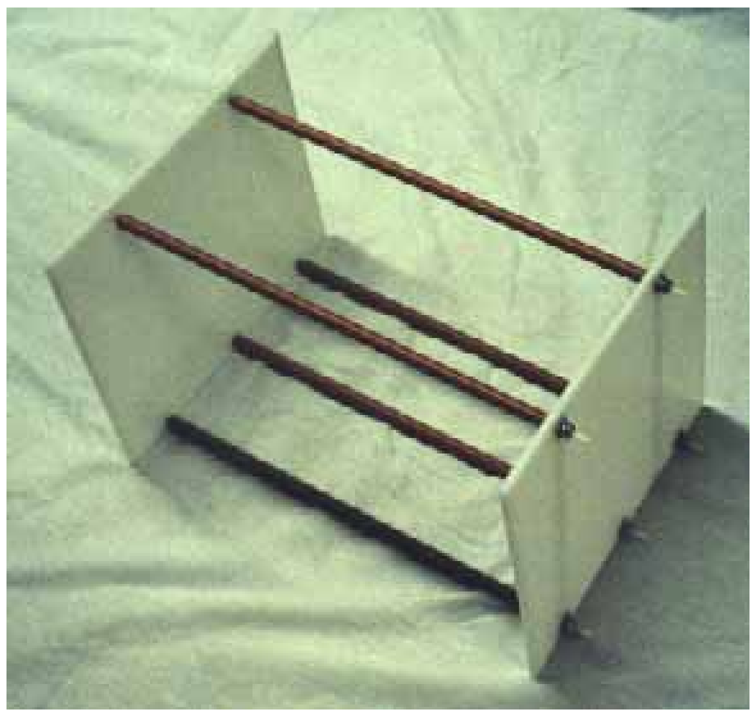
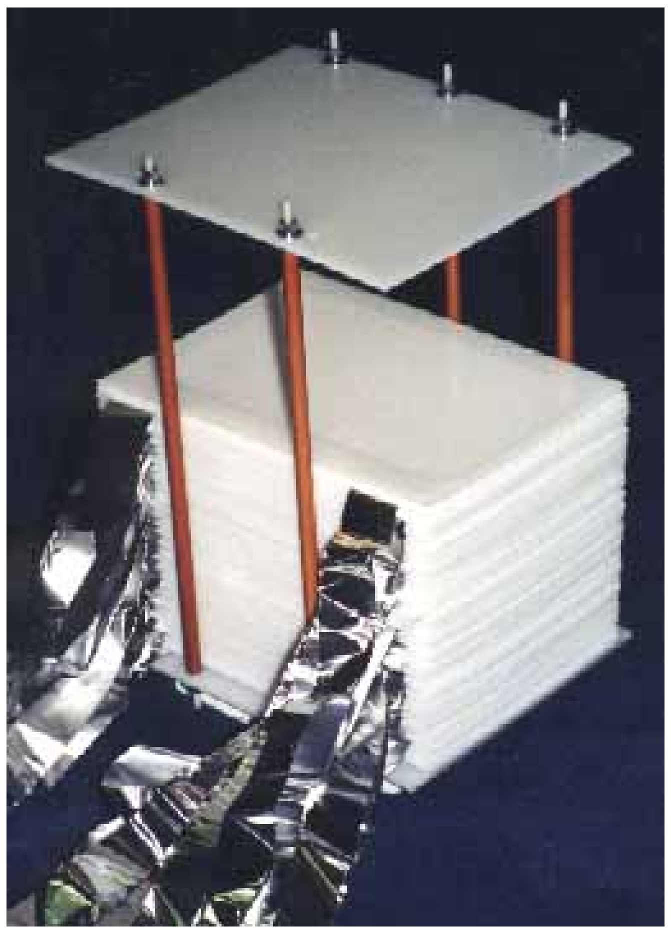

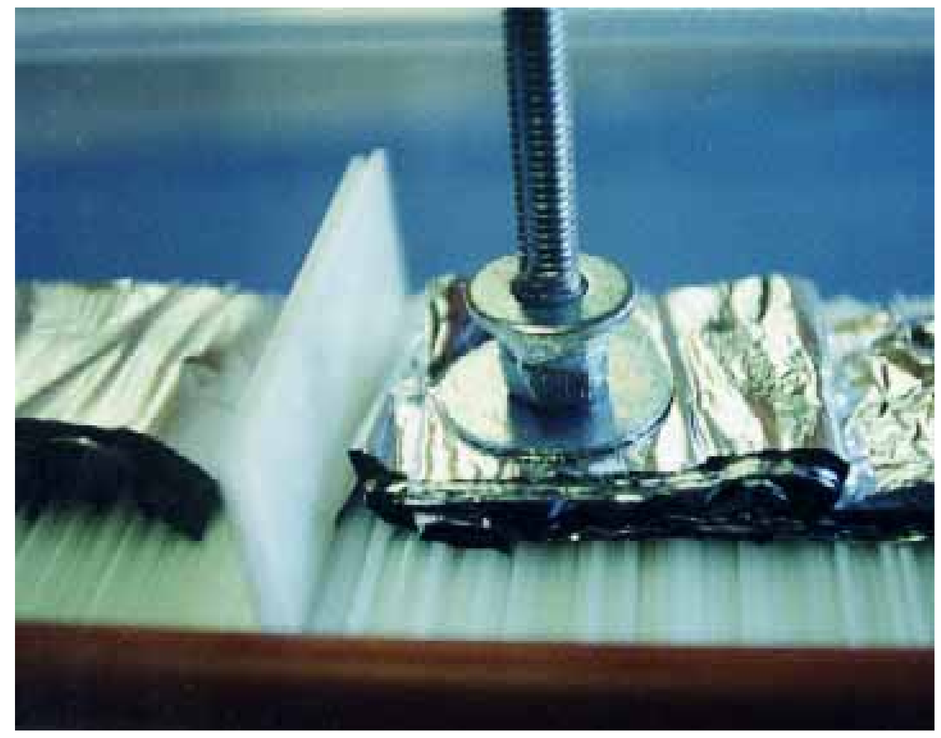
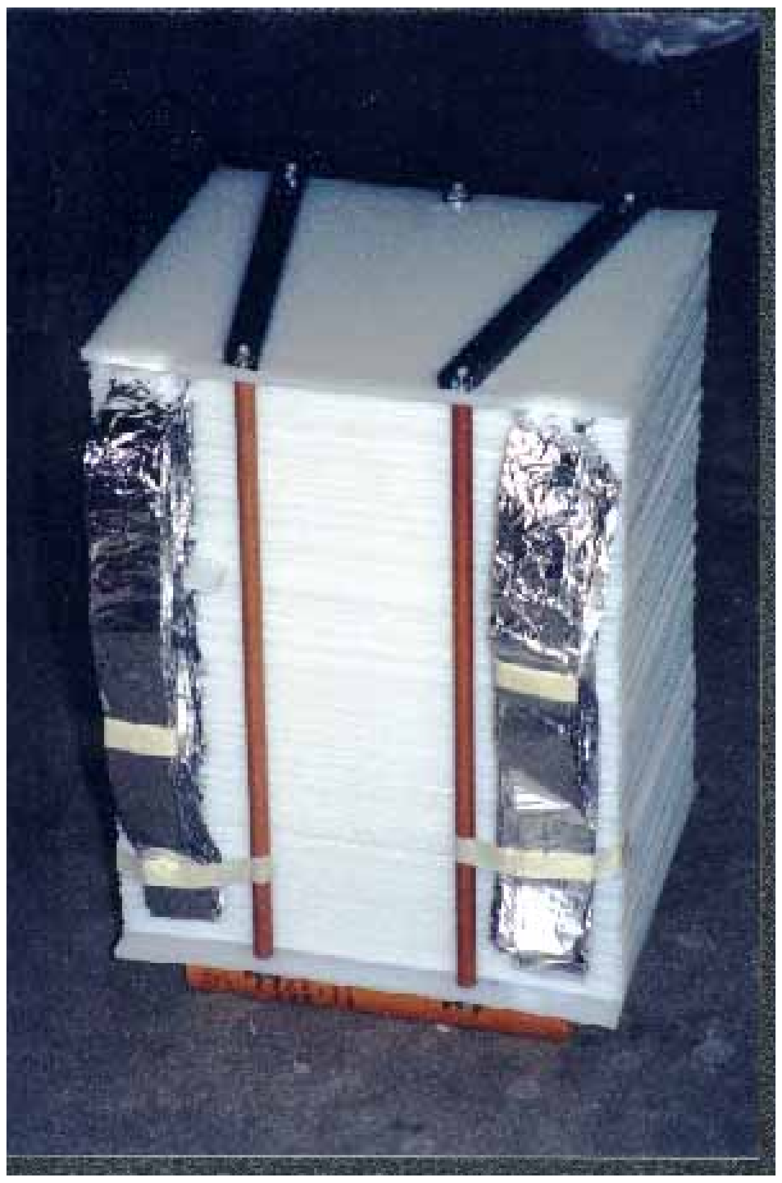

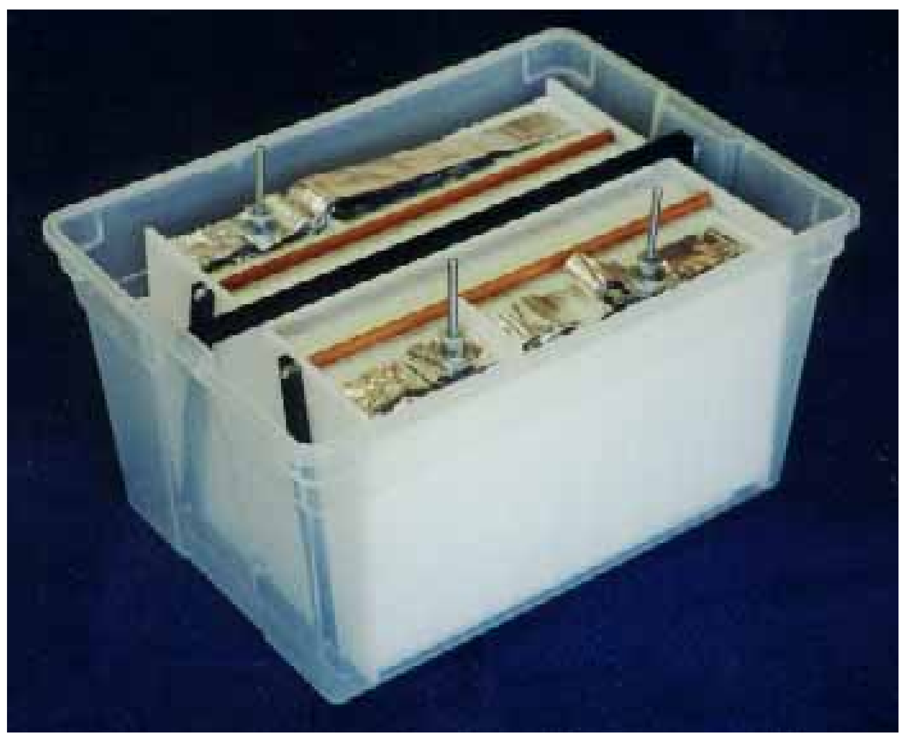




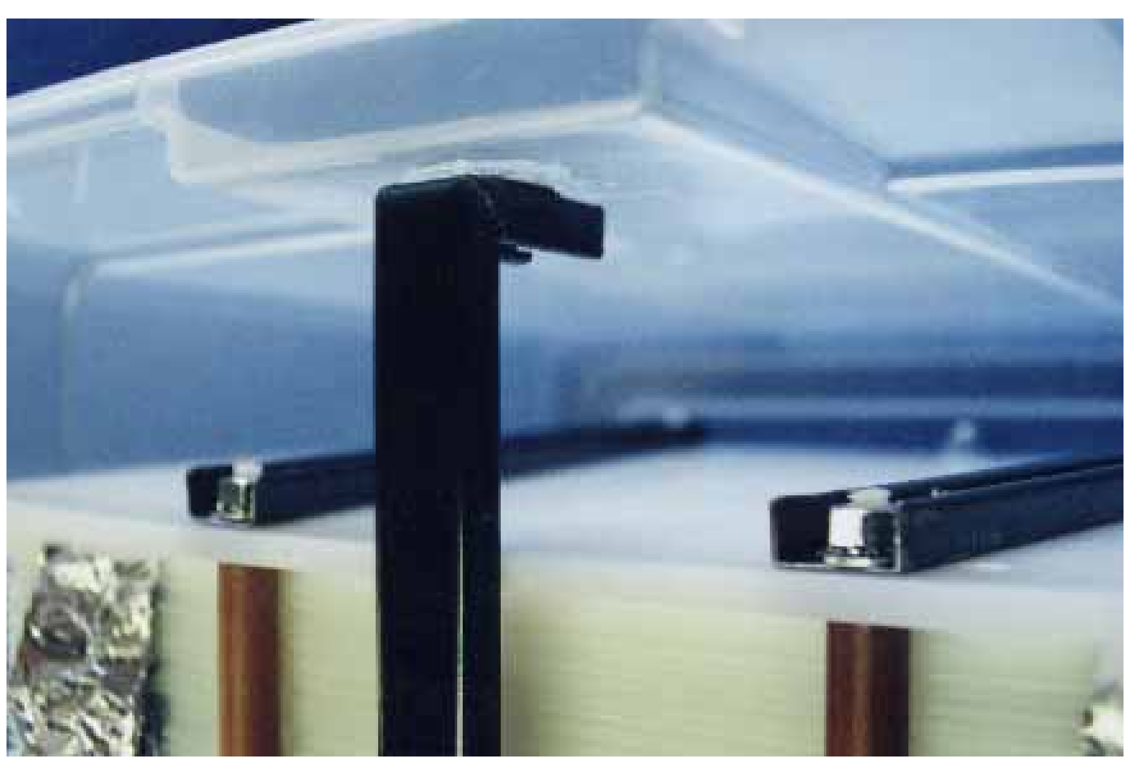

Primary Coil

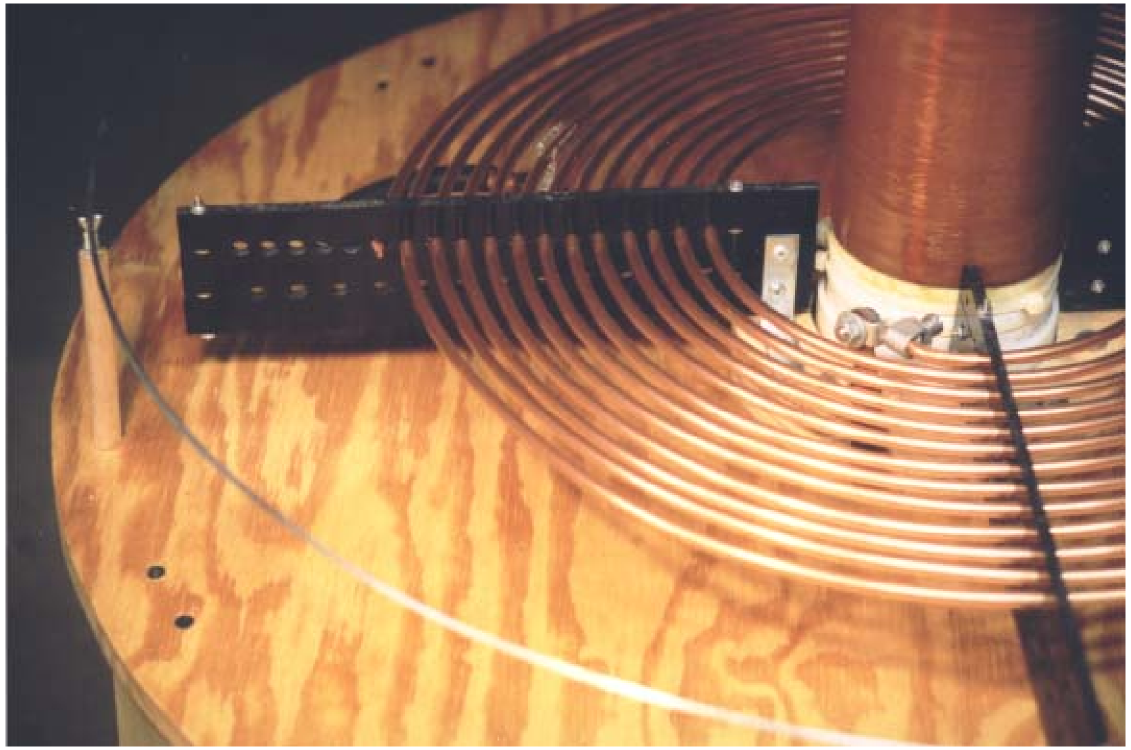

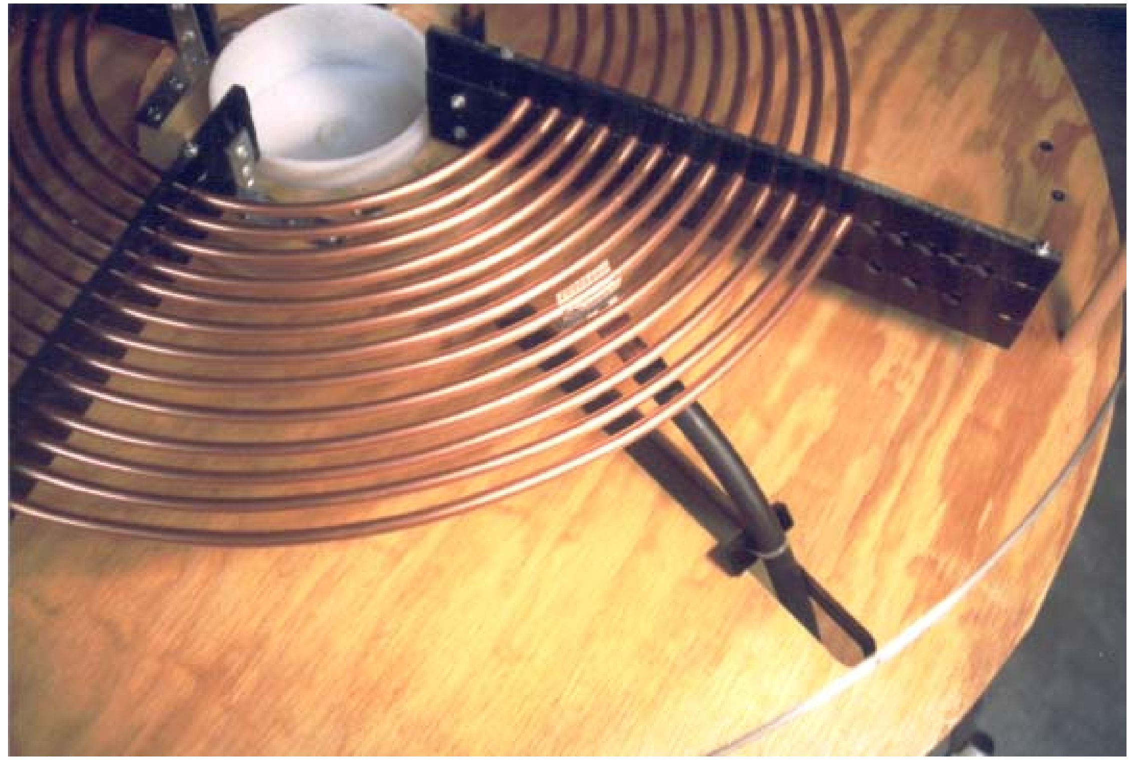

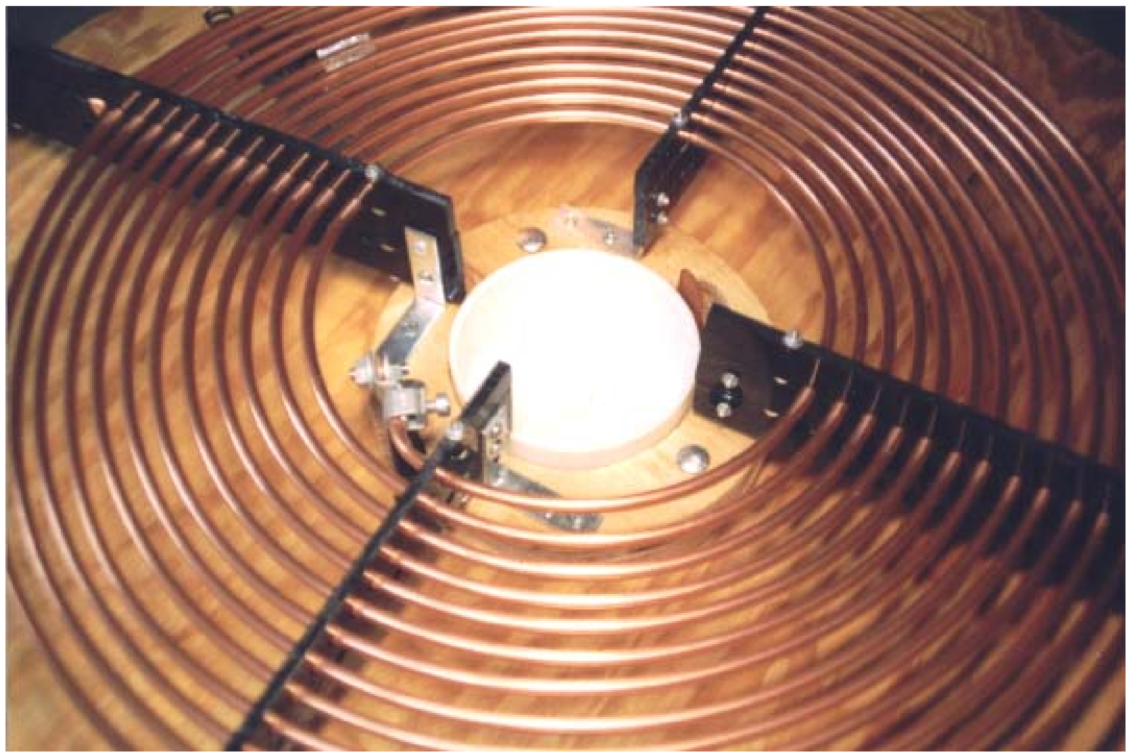

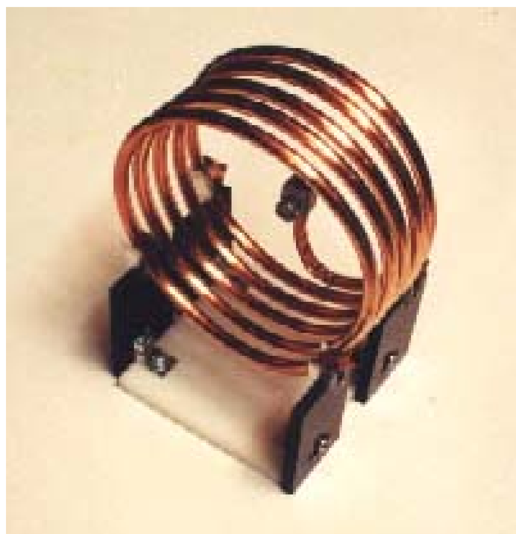


Secondary Coil

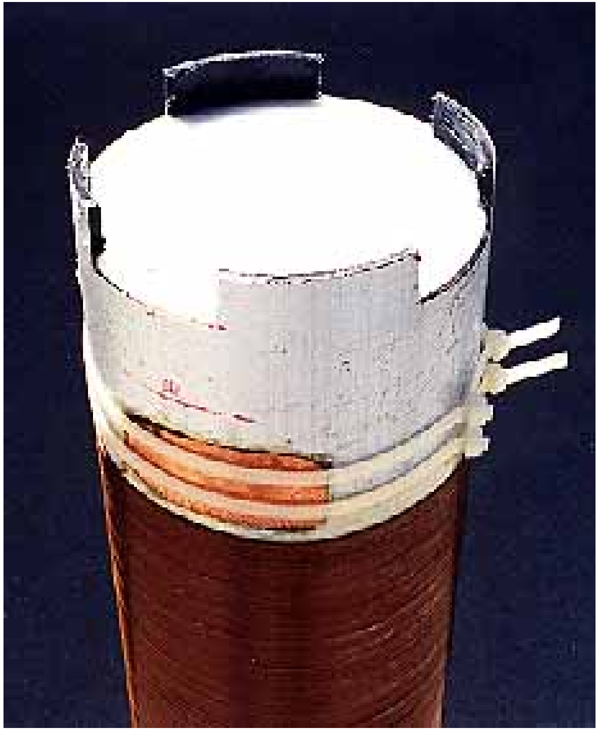

Torroidial Discharge Cap

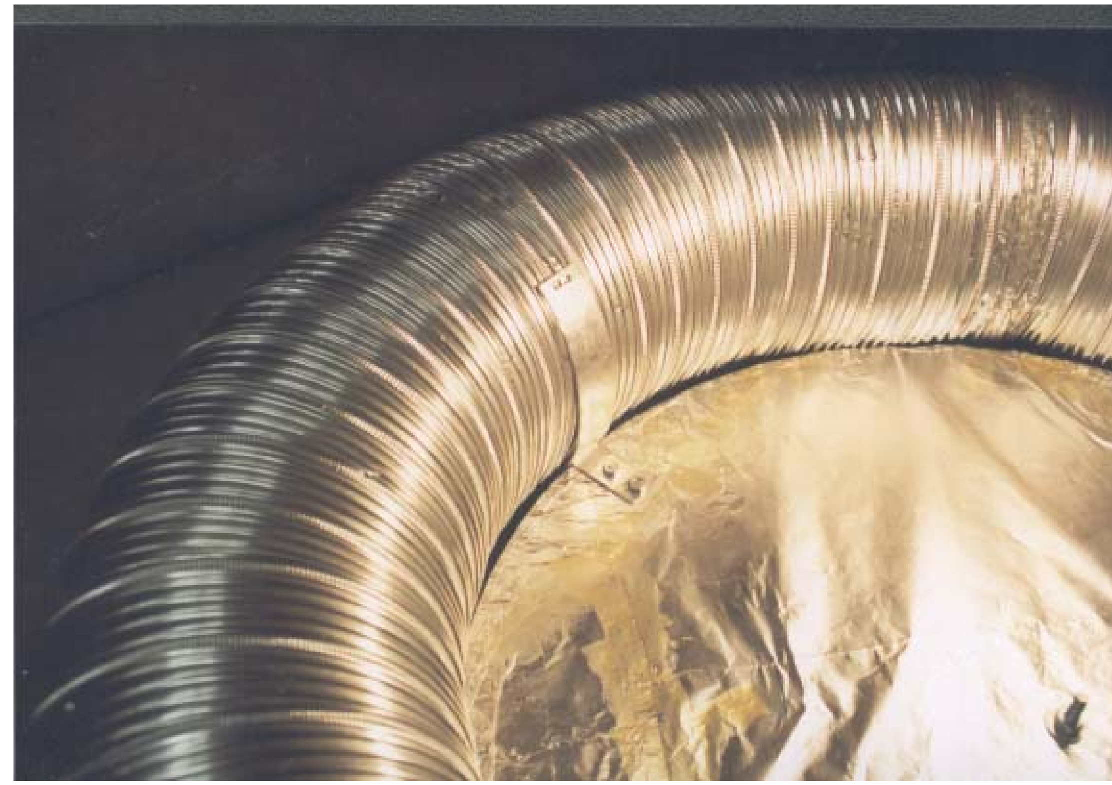

Coil Winder and Voltage Divider

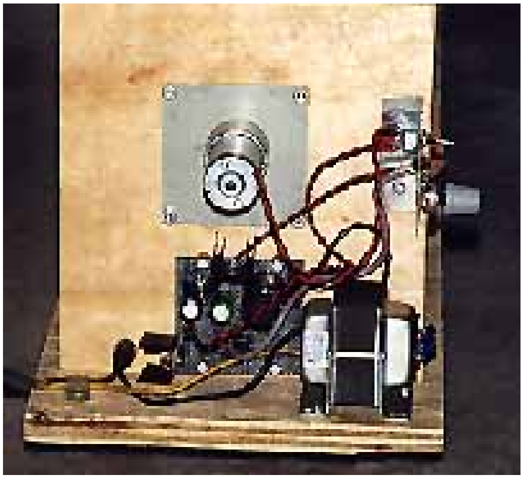

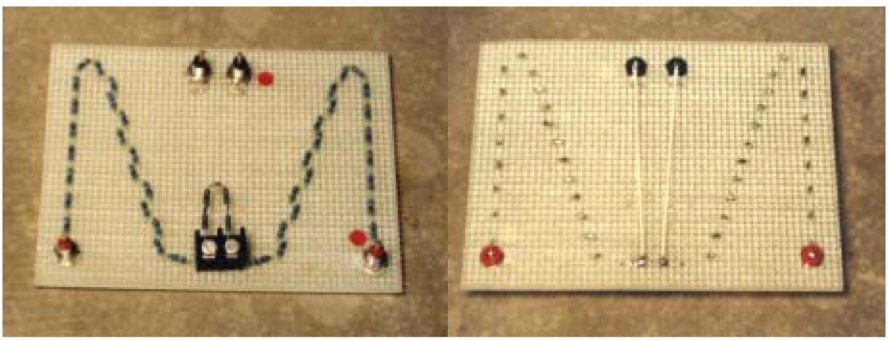

Finished Coil and Builder - Matt Behrend





Ouput Arcs