LUMPED RLC MODEL

We are now at the point of looking at models for the Extra coil. I have already discussed the conflict between the Corum brothers and most of the Tesla coil community regarding distributed versus lumped models. My background is similar to that of James Corum, so I started this research with the assumption that the distributed approach would be superior. I spent considerable time learning some computer codes the Corums wrote, and even supplied them with codes to calculate inductance. But I was unable to predict all the interesting features with the distributed approach, and it appeared that a massive effort would be required if one wanted to do so. So I turned to the lumped RLC model. The inductance and capacitance values seem to be reasonably well understood, as detailed in the last two chapters, but resistance is a big challenge. I do not consider the following to be a definitive work on Tesla coil resistance, but rather a starting point for discussion.

The first circuit we want to model is shown in Fig. 1. The IGBT driver (discussed in the next chapter) is driving a vertically mounted coil with a toroid on top. The driver and coil are physically separated for reasons of health and safety. The wires between the two have some inductance L_1 and capacitance C_1. The obvious choice for these conductors would be a coaxial cable of appropriate voltage and current rating. However, the capacitance C_1 must be charged and discharged by the leading edge of the applied voltage pulse each half cycle. The associated current shows up as switching noise in the current sense resistors. One can reduce C_1 by about half by using an open wire transmission line rather than coax, which is what I did. Open wire lines are often operated as balanced lines (neither wire grounded). In this case, the open line is unbalanced (one wire 'hot' and one wire grounded).

Damping of the high frequency components is improved by greater inductance and greater resistance in the line between driver and coil. Both inductance and resistance increase with a wire of smaller diameter. I tried both 14 ga and 8 ga wire, and chose the 14 ga PVC coated solid copper wire.

The hot conductor must be protected from flashover from the top of the coil. Total protection is probably like *total* protection from lightning or tornadoes, probably not possible and certainly expensive. What I did will provide at least some protection. I took scrap sections of two inch thick styrofoam about one foot wide and cut a small groove down the center. The groove was oriented up and the 14 ga wire was laid in the groove. This would help prevent the 14 ga wire from arcing to the floor. Another section of styrofoam was laid on top the first. Then a length of coax was laid on top of the sandwich, directly over the 14 ga wire. The coax was connected to earth ground. A flashover to ground would normally hit the coax. A flashover to the hot conductor would have to miss the coax, go through two inches of styrofoam and then a layer of PVC to get to the 14 ga wire. This gets the probability of a damaging strike down to an acceptable value.

The driver applies a square wave of voltage between the left end of L_1 and the ground

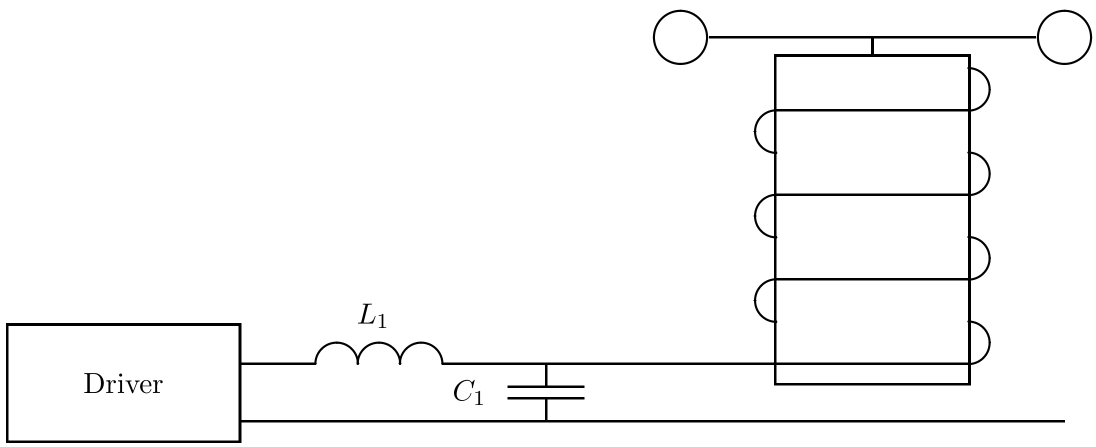

Figure 1: Drive for Tesla Coil

plane of the Tesla coil. This square wave is composed of an infinite series of cosinusoids, the fundamental and all odd harmonics. The peak value of the fundamental is $4/\pi$ times the dc supply voltage. That is, if $V_{dc} = 500$ V, then the instantaneous voltage of the fundamental is $v_i = 637\cos\omega t$. The first three terms of the Fourier series are

$$f(t) = \frac{4}{\pi}[\cos\omega t - \frac{1}{3}\cos 3\omega t + \frac{1}{5}\cos 5\omega t + \cdots] \tag{1}$$

The harmonics of the exciting wave will drive higher order resonances of the Tesla coil, if these resonances are harmonically related. It appears, at least for the coils I have built, that any higher order resonances are not exact multiples of the fundamental frequency. That is, I apply a square wave of voltage to the feed point, and observe a current that looks sinusoidal at the fundamental frequency. The lumped RLC model automatically excludes higher order resonances, so if they are of significance, we must use a distributed model to describe them. My feeling at the time of this writing is that higher order resonances are not a problem, at least not enough of a problem to exclude the use of the lumped model.

1 The RLC Model

The simplest model that can be proposed for the Extra coil is the series RLC circuit shown in Fig. 2. The inductance L_2 is given by Wheeler's formula. It can also be calculated from first principles [2] using elliptic integrals and a computer. The two methods typically agree within one percent. It can be measured to the same accuracy with a hand-held inductance meter. The only thing that changes its value is significant amounts of ferromagnetic materials inside or near the coil. These are avoided because of high losses. The inductance does not change with frequency or when sparks occur. If everything about a Tesla coil were as well behaved as the inductance, modeling would be easy.

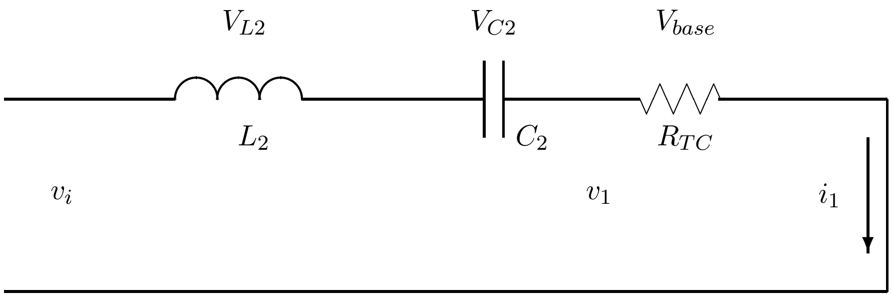

Figure 2: Tesla Coil with Series Resonant LC Circuit

The capacitance C_2 is the modified sum of the capacitance of the coil itself, as given by Medhurst, and the isotropic capacitance of the toroid. The toroid and the coil shield each other, so the effective capacitance is always less than the sum of the Medhurst capacitance and the toroid capacitance. There should be a simple empirical formula to calculate C_2, given the Medhurst and toroid capacitances, but I do not know what it is. C_2 can be calculated numerically using Gauss's Law. Terry Fritz has written a program to do this. If one is careful about measuring and entering all the dimensions and the locations of grounded surfaces, one should get a value for C_2 well within 5% of the correct value. Of course, C_2 increases when the spark occurs due to the size of the toroid becoming effectively larger. This lowers the frequency of operation while the spark is present.

There is really no way to directly measure C_2. C_2 is in the range of tens of pF, the same range as the leads of a capacitance meter. The presence of a meter and operator will change the capacitance of the coil. One can measure L_2 and the resonant frequency and then calculate C_2 from these measured values.

If only a crude approximation is desired, one can use the Medhurst capacitance plus 75% of the isotropic toroid capacitance for C_2. Since resonant frequency is related to the square root of capacitance, a 20% error in capacitance results in only a 10% error in frequency, which is sometimes close enough.

Now we come to the third element of the model, R_{TC}. This is the input impedance of the coil at resonance. We have seen that we can readily get a value for L_2 to within 1%, and a value for C_2 to within 5% with a little more work. Or we can measure frequency and calculate C_2 to within one or two percent. But R_{TC} is another story, as I have mentioned. For now, let us assume we have an appropriate value for R_{TC} and explore the features of the model.

RLC Circuit With Sinusoidal Source

We have seen that a square wave can be decomposed into a sine wave (or cosine wave) of fundamental frequency plus an infinite series of harmonics. If the harmonics do not have much effect upon operation, we can safely model the square wave as a sine wave of the same frequency, at least for some purposes. For brevity, we will call the three circuit elements just R, L, and C, and redraw the circuit slightly as shown in Fig. 3. If only a single frequency sine wave is present, we can use phasor analysis. The input phasor voltage $\mathbf{V_1}$ drives a current $\mathbf{I_1}$ through the circuit. By Kirchhoff's Law,

$$\mathbf{V_1} = \mathbf{V_L} + \mathbf{V_C} + \mathbf{V_R} \tag{2}$$

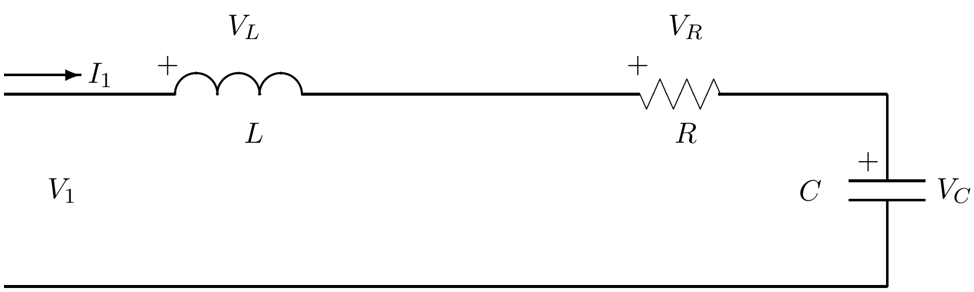

Figure 3: Series Resonant LC Circuit

Using Ohm's Law in phasor form gives

$$\mathbf{V_1} = \mathbf{I_1 Z} = \mathbf{I_1}(j\omega L + \frac{1}{j\omega C} + R) \tag{3}$$

The voltage across the inductor leads the current through the inductor by 90^o, while the voltage across the capacitor lags the current by 90^o. The two voltages tend to cancel each other, and actually do cancel at the resonant frequency ω_0, as defined by

$$\omega_0 L = \frac{1}{\omega_0 C} \tag{4}$$

We then solve for the resonant frequency in radians per second as

$$\omega_0 = \frac{1}{\sqrt{LC}} \tag{5}$$

The resonant frequency in Hertz is

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \tag{6}$$

The series impedance, expressed as a phasor, is

$$\mathbf{Z} = Z\angle\phi = R + j\left(\omega L - \frac{1}{\omega C}\right) \tag{7}$$

where

$$Z = \sqrt{R^2 + \left(\omega L - \frac{1}{\omega C}\right)^2} \tag{8}$$

and

$$\phi = \arctan\frac{\omega L - (1/\omega C)}{R} \tag{9}$$

At resonance, $\mathbf{Z} = R$, a real number, so the input current $\mathbf{I_1}$ is in phase with the input voltage $\mathbf{V_1}$. This means that the voltage $\mathbf{V_C}$ will *lag* the input voltage $\mathbf{V_1}$ by 90^o. We will say more about this in a moment.

The bandwidth of the series circuit is defined as the range of frequencies in which the amplitude of the current is equal to or greater than $1/\sqrt{2}$ times the maximum amplitude. This current produces a heating effect of half the maximum value, so the frequencies at the bandwidth limits are called the half-power frequencies ω_1 and ω_2. With a little algebra, the bandwidth is shown to be

$$\beta = \omega_2 - \omega_1 = \frac{R}{L} \tag{10}$$

The quality factor Q is defined as the ratio of the resonant frequency to the bandwidth, so for the series circuit we have

$$Q = \frac{\omega_0}{\beta} = \frac{\omega_0 L}{R} = \frac{1}{\omega_0 CR} = \frac{1}{R}\sqrt{\frac{L}{C}} \tag{11}$$

The amplitude of the voltage across either the inductor or the capacitor at the resonant frequency ω_0 is Q times the amplitude of the source voltage.

The above is standard Circuit Theory I. We see that to get a large voltage in our Tesla coil that it needs to be high Q, having large L and/or small R. With this model, a space wound coil of 14 gauge wire (called 14S) that I use for testing would have a Q of about 500. The inductance is about 17 mH, the resonant frequency is 160 kHz, and the input impedance R_{TC}

is about 25 Ω. If I apply a sinusoidal source voltage with an rms value of 1000 V, the voltage across the inductance should rise to 500,000 V. The voltage does not actually reach this value because a spark occurs first, but if the toroid were large enough and smooth enough, it should reach this voltage.

The RLC model predicts such a voltage, but how do we check the model for validity? Does the Tesla coil toroid actually reach 500 kV? If we try to measure the toroid voltage directly, we have to attach some sort of probe, which must have some capacitance and resistance. The probe capacitance lowers the resonant frequency. At least part of R_{TC} is frequency dependent. The probe may have frequency dependence as well. We can measure the toroid voltage with the new Q and new R_{TC} but we cannot be positive about the toroid voltage when the probe is removed.

I have spent considerable effort to measure the toroid voltage with a battery operated, fiber optic coupled, electric field probe. Results are very repeatable and show some interesting features, but there is always the question about absolute calibration. To calibrate the electric field probe, I have to connect a voltmeter to the toroid at some point. In addition to changing the resonant frequency and R_{TC}, the voltmeter probe will also change the electric field pattern.

So we agree that we cannot measure the toroid voltage to within 1%. But what is a reasonable estimate of the accuracy of measurement? I would estimate that the uncertainty of the absolute value of toroid voltage is about 10%, almost certainly no worse than 20%. I used to say that such an accuracy was close enough for government work, but such a statement would be politically incorrect today. Actually, I quit using the phrase after my uncle, a long time employee of the Corps of Engineers, became angry when I said it.

Sparks are a statistical phenomena, anyhow. Even when the toroid voltage is accurately known, say in a Marx generator, there is uncertainty about whether the spark will occur and about how long it will be. One voltage may produce a spark one time out of ten, and a higher voltage may see a spark nine times out of ten. These two voltages may easily differ by more than 10%. So if I know the toroid voltage to within 10%, it should be quite adequate.

Another way of estimating the toroid voltage is to observe the spark length and multiply by some constant. Two polished brass spheres separated by one cm will break down in dry air at 30 kV. So, the assumption goes, if we have a 10 cm spark from toroid to air, the toroid voltage must have been 300 kV. This is one of those great ideas that just do not work in practice. One of the main problems is that spark length is a nonlinear function of voltage. Once a spark is started, it does not take as much voltage to keep it going.

An example of the error involved is given by the following: I was applying a square wave of voltage with an rms value of about 1200 V to the Tesla coil. With a Q of 500, the maximum possible toroid voltage would be 600 kV. Sparks were occurring at lower voltages, probably not over 500 kV. I observed a spark 143 cm long to a ground rod. Using 30 kv/cm would predict a toroid voltage of 4350 kV, about a factor of nine too high. This technique certainly gives one the satisfaction of claiming big voltages, but doesn't have much relation to reality. One could scale the constant down to say 3 kV/cm and be closer to the true value, but the nonlinearity of the spark keeps this from being very good either. Measuring the electric field with a fiber optic coupled probe should get much better results for the toroid voltage than the technique of multiplying by a constant.

Anyhow, I believe the RLC model does a reasonable job of predicting the toroid voltage before breakout, to within perhaps 10 to 20% of the 'correct' value. The streamer adds resistance to the system, which lowers Q and toroid voltage. The model and experimental results seem to be in reasonable agreement on these features.

RLC Circuit With Square Wave Source

The sinusoidal steady state analysis of the previous section does not tell us how the input current builds up when voltage is applied. For that we need to go back to the time domain. We will use the actual square wave input voltage and ignore the resistance for the first few cycles. Until the current builds up, the voltage across the resistor will be much smaller than the voltage across either the inductor or capacitor, so this is not a bad approximation.

The input voltage v_i is equal to V_d the first half cycle and $-V_d$ the second half cycle. The current everywhere in the circuit is i and the capacitor voltage is v_c. The initial current is I_0 and the initial capacitor voltage is V_{c0}. Kirchhoff's voltage law is

$$L\frac{di}{dt} + v_c = V_d \tag{12}$$

The component equation for the capacitor is

$$C\frac{dv_c}{dt} = i \tag{13}$$

The solution of this set of equations for $t \geq t_0$ is as follows:

$$i(t) = I_0 \cos\omega_0(t - t_0) + \frac{V_d - V_{c0}}{Z_0}\sin\omega_0(t - t_0) \tag{14}$$

and

$$v_c(t) = V_d - (V_d - V_{c0})\cos\omega_0(t - t_0) + Z_0 I_0 \sin\omega_0(t - t_0) \tag{15}$$

where the angular resonance frequency ω_0 is

$$\omega_0 = 2\pi f_0 = \frac{1}{\sqrt{LC}} \tag{16}$$

and the characteristic impedance (sometimes called the surge impedance) Z_0 is

$$Z_0 = \sqrt{\frac{L}{C}} \quad \Omega \tag{17}$$

We assume we are operating at resonance, ω_0 of the applied square wave = ω_0 of the RLC circuit. The time axis is selected so $t_0 = 0$ and $t = 0$ at the start of the square wave. The initial current through L and the initial voltage on C are both zero. The equations for i and V_c become

$$i(t) = \frac{V_d}{Z_0} \sin \omega_0 t$$

and

$$v_c(t) = V_d - V_d \cos \omega_0 t$$

The current starts and ends at zero, with a peak value of V_d/Z_0. The initial current for the next half cycle is zero. Capacitor voltage starts at zero but builds to a value of $2V_d$, which is the initial value for the next half cycle. Waveforms are shown in Fig 4.

For the second half cycle, the applied voltage is $-V_d$. The starting time t_0 is half the period of the square wave. The current is

$$i(t) = \frac{-V_d - 2V_d}{Z_0} \sin \omega_0 (t - t_0)$$

and the capacitor voltage is

$$v_c(t) = -V_d - (-V_d - 2V_d) \cos \omega_0 (t - t_0)$$

The current starts and ends at zero, but this half cycle it reaches a peak value of $-3V_d/Z_0$. The factor $t - t_0$ starts at zero at the beginning of the second half cycle so the cos function starts at $+1$ and goes to -1. The capacitor voltage starts at $2V_d$, as it must, and ends at $-4V_d$, which is the initial capacitor voltage for the third half cycle.

One can continue on with this solution process as long as desired. The peaks for current are V_d/Z_0, $-3V_d/Z_0$, $+5V_d/Z_0$, $-7V_d/Z_0$, etc., occurring at the midpoints of each half cycle. The peaks for capacitor voltage are $2V_d$, $-4V_d$, $6V_d$, $-8V_d$, etc., occurring at the ends of each half cycle. The capacitor voltage lags the current waveform by 90^o, the standard relationship between voltage across and current through a capacitor.

This suggests one way of measuring Z_0 without measuring the resonant frequency. We apply a half cycle of voltage (actually somewhere between one third and two thirds of a cycle), and measure the peak current. The problem is that the peak current during the first half cycle

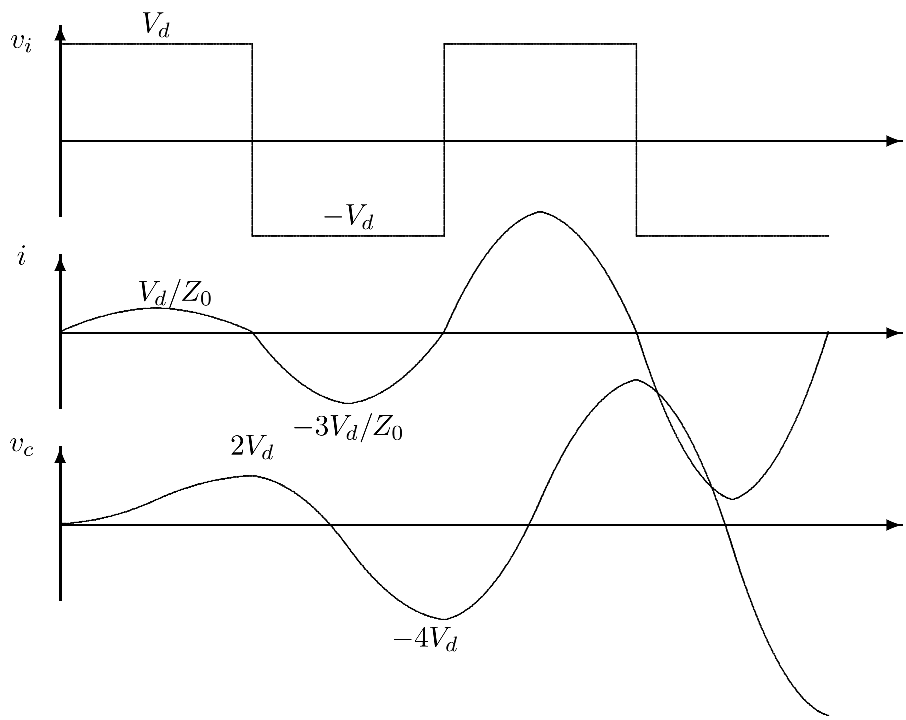

Figure 4: First Two Cycles of Current and Capacitor Voltage in RLC Circuit

may be much smaller than the charging current of the transmission line to the Tesla coil, which would mask the desired effect. Since the toroid voltage is only double the input voltage after the first half cycle, and electric fields are small, we might be able to put a fiber optic coupled pulse generator right at the input terminals of the Tesla coil, and minimize the effects of stray capacitance. On the other hand we could very well discover that the distributed capacitance effects will always mask the desired measurement.

As mentioned earlier, these waves will not increase forever, but will asymptotically approach values controlled by the circuit resistance, (if a spark does not occur first). At steady state we will see the relationships found in the previous section.

Phase Shift, Input to Toroid

One interesting issue is the amount of phase shift from bottom to top of a Tesla coil. There are strongly stated opinions among Tesla coilers as to what this shift is, or should be. Some believe the phase shift is zero. I believe the toroid voltage (with respect to ground) lags the input voltage (also measured with respect to ground) by 90^o. On one occasion I calibrated the various phase shifts in my instruments (described in the following) and arrived at a phase shift of 94^o from base to toroid. This fits the RLC model well within experimental error.

There are several different ways of thinking about the phase shift, some quite detailed and prone to error. Any of us can get sloppy on notation and make such an error. So I will try to be precise.

Earlier in the chapter I showed that current is in phase with the input voltage for the lumped RLC model at resonance. The model predicts that the toroid voltage will lag the input voltage by 90^o. One concern here is that part of the total capacitance is internal to the coil, which confuses the computation of the toroid voltage magnitude. However, if the current entering the toroid is in phase with the current entering the base of the Tesla coil, as the RLC model requires, then this current flows through a pure capacitance from the toroid to ground. So the *phase* of the toroid voltage would be 90^o behind the base voltage, even if the *magnitude* is in error.

We get the same result if we think of a Tesla coil as a quarter-wave antenna above a ground plane. If the antenna impedance matches the transmission line impedance (quite possible) then only a traveling wave exists on the line. The output of a quarter-wave section of transmission line always lags the input by 90^o, and likewise the top of a quarter-wave vertical antenna with respect to the base. The Tesla coil is not the same as a quarter-wave whip antenna because of the mutual inductance and mutual capacitance, but I think this is not adequate to change the phase shift from 90^o to zero.

But how does one actually measure the phase shift? A well done measurement can be more convincing that a basket of theoretical arguments. As discussed earlier, the presence of a probe presents some serious problems with loading the coil and affecting both magnitude and phase. It appears that the next best thing to having a physical probe attached to the toroid is to measure the electric field. With adequate calibration, this should answer most of our questions.

Suppose we have our Tesla coil inside a grounded metal building. Some flux lines go up and hit the roof. Some go sideways to the walls. Some go down to the earth. Each flux line (or tube) can be considered to be between a portion of a capacitor plate at Tesla coil voltage and a similar portion at ground potential. That is, the capacitance of the Tesla coil can be thought of as hundreds of sub-capacitors going in all directions, and connected in parallel. If we take one of these sub-capacitors and put an electrode in it, (not at ground or toroid potential), we have just split the capacitor into two series capacitors. Voltages divide across capacitors inversely as the capacitances. The voltage between the electrode and the adjacent surface can be readily measured by a standard electronics circuit if its capacitance is much larger than the value of the sub-capacitor between toroid and ground. The concept for an electrode near ground is shown in Fig. 5.

The electrode can also be mounted on the toroid, with the transmitter inside the toroid. If the electric field is measured at the toroid, it is essential that the electronics be battery powered, and the signal transmitted away from the coil by an optical fiber. It is not as essential if the voltage is measured near ground potential, but is still a good idea. The alternative,

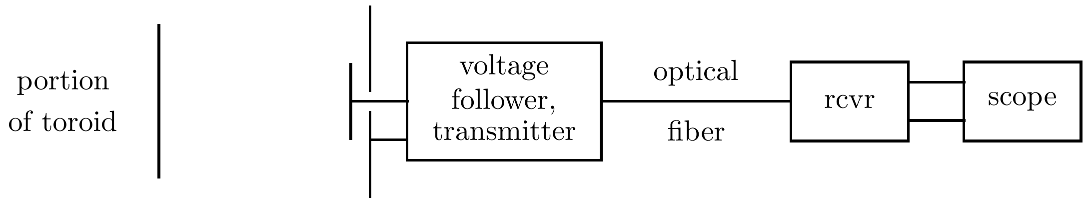

Figure 5: Concept of Electric Field Measurement

a coaxial cable to the electronics, will pick up some of the signal from the Tesla coil on its ground braid. RF on a ground lead can cause erroneous results.

Actually there are four distinct locations where the electric field might be sensed.

1. Next to toroid
2. At bottom of coil
3. At grounded surface above toroid
4. At toroid height by grounded wall

Each location has its own advantages and disadvantages. I started experimentation with fiber optic systems by putting the transmitter in the toroid. This has the advantage of being in a high electric field. The electrode can be physically small and/or located close to the toroid. One disadvantage is that a spark at that location could destroy your electronics. Another disadvantage is that the smaller toroids are difficult to modify and do not have enough room for the electronics. A half-spun toroid worked quite well, but the classic toroids did not.

My next attempt was to place the electronics at the bottom of the coil. I put a metal disk at the top of the coil form, about two thirds of the diameter of the secondary, and electrically connected to the toroid. I then put a polyethylene bracket across the bottom of the coil form with a smaller disk facing up. A wire from this disk ran through the bracket and into an aluminum box holding the electronics, which was mounted to the bottom of the bracket. This worked quite well for testing at lower powers. The electric field was strong and was not affected by movement of people or things external to the coil. One concern was that the sensor would 'see' the sides of the coil. The voltage would be lower on the sides but the distance would be smaller, so the influence would be similar. This posed the most problem for looking at the toroid voltage during discharge. The sensor would show a voltage decay lasting for five to ten cycles at discharge. The toroid voltage may have dropped much faster than that, but one could not be sure because the coil sides took a longer time to discharge.

The main problem, however, was the occasional spark from top electrode to the grounded aluminum box through the inside of the coil form. This spark was about twice as long as the sparks to air at that time, hence was a bit of a surprise. On at least one occasion, that spark to ground upset something in my driver that caused my IGBTs to give up. I made a number of changes to my system after that, including removing the grounded aluminum box. I also added some polyethylene disks to block the inside of the coil. Now a toroid to ground strike has to go through several layers of polyethylene plus two or three feet further to get to a real ground, if it tries to travel on the inside of the coil.

Philosophically, the best place seems to be above the toroid, at ceiling level. Directly above the coil, the toroid appears as a large disk. The influence of the toroid would be maximum while the influence of the coil sides would be minimum. It is also the most inconvenient location regarding the changing of batteries, switching ranges, and the like. For this to work, the ceiling must be conductive, and immune to corona. That is, a large area of the ceiling must be covered with a good conductor (like copper sheeting), or a large grounded toroid must be used, with the sensor at its center. I will consider doing this if I ever build another lab. For now I will place the sensor at about toroid height, and located next to the grounded metal wall of the building.

It has been suggested [1] that the grounded toroid placed directly above the coil could be on a rope and pulley to make it easier to change batteries. In my case I already have a rope and pulley in place to raise and lower a large and awkward toroid into position on top the coil. If the grounded toroid was not too large and heavy, it could be clipped on the rope and raised into position after the high voltage toroid was lowered into place. I may try that.

It was also suggested [1] that the system could be calibrated by charging the high voltage toroid to some known voltage and discharging it directly to earth through a low impedance path. The coil would be physically in place but disconnected at top or bottom or both. The capacitance of the toroid would resonate with the inductance of the shorting path at some frequency well into the MHz range. If the electronics were still linear at this high frequency, this might work very nicely.

In summary, the simple RLC model appears to do reasonably well in predicting

1. Resonant frequency
2. Toroid voltage magnitude, given the resistance
3. Toroid voltage phase

References

[1] Benson, Barry, benson@erols.com, private communication, November, 2001.

[2] Fawzi, Tharwat H. and P. E. Burke, *The Accurate Computation of Self and Mutual Inductances of Circular Coils*, IEEE Transactions on Power Apparatus and Systems, Vol. PAS-97, No. 2, March/April 1978, pp. 464-468.