

SOLID STATE TESLA COIL

by

Dr. Gary L. Johnson
Manhattan, Kansas

Some years ago I developed an interest in Tesla coils. I was teaching a senior elective course at Kansas State University where we talked about power MOSFETs and topics related to high voltages and currents. I decided to use a Tesla coil as a class project. We would talk about design aspects, then design, build, and test a coil. The best description of the results of that plan was fiasco, or maybe disaster. We had some sparks, but none where they belonged. That was one of the most humiliating experiences of my career.

I learned several things from that experience. One is that the Tesla coil is more complex than I had thought. Another was that there seemed to be a mismatch between theory and experiment. At that time, at least, people would go through pages of high powered mathematics and quit without giving an example of how to use all the formulas. Experimentalists would sometimes make fun of the theorists, and give rules-of-thumb on how to make long sparks. It was like I was hearing a debate on whether the best cooks use recipes or not. My mother never used a recipe and I always enjoyed her cooking. However, my own talents are such that if I am to cook anything fit to eat, I need a recipe.

This book is written for people like me, challenged when faced with doing something without a recipe or complete set of instructions. I will throw in things learned from other Tesla coil builders, but will quickly admit that when it comes to making long sparks, there are many who are far better than I.

I started asking questions about Tesla coils that any electrical engineer would ask. These include:

1. What is the input impedance?
2. What are the fractions of input power that are dissipated in the spark itself, in electromagnetic radiation, the coil wire, the coil form, the toroid, the spark gap, and other

circuit components?

3. Are there circuit models that allow these questions to be answered on the computer before building and testing devices in open air?
4. What are the differences between Tesla coils driven by or through spark gaps, vacuum tubes, or solid state devices?
5. What are the important factors in producing long sparks (energy per bang, power input at spark inception, rate of change of power, the coil, the toroid, etc.)?

One would expect the answers to these questions to come from a mix of theory and experiment. One would develop a theory or model and then go to the laboratory to measure parameters and check performance. The theory would then be adjusted to reflect experimental observations.

We now review a little Tesla coil history and look at the ‘simplest’ model, the lumped circuit element model.

1 History

Nikola Tesla (1856 - 1943) was one of the most important inventors in human history. He had 112 U.S. patents and a similar number of patents outside the United States, including 30 in Germany, 14 in Australia, 13 in France, and 11 in Italy. He held patents in 23 countries, including Cuba, India, Japan, Mexico, Rhodesia, and Transvaal. He invented the induction motor and our present system of three-phase power in 1888 [20]. He invented the Tesla coil, a resonant air-core transformer, in 1891. Then in 1893, he invented a system of wireless transmission of intelligence. Although Marconi is commonly credited with the invention of radio, the U.S. Supreme Court decided in 1943 that the Tesla Oscillator patented in 1900 had priority over Marconi’s patent which had been issued in 1904 [15]. Therefore Tesla did the fundamental work in both power and communications, the major areas of electrical engineering. These inventions have truly changed the course of human history.

After Tesla had invented three-phase power systems and wireless radio, he turned his attention to further development of the Tesla coil. He built a large laboratory in Colorado Springs in 1899 for this purpose. The Tesla secondary was about 51 feet in diameter. It was in a wooden building in which no ferrous metals were used in construction [15]. There was a massive 80-foot wooden tower, topped by a 200-foot mast on which perched a large copper ball which he used as a transmitting antenna. The coil worked well. There are claims of bolts of artificial lightning over a hundred feet long, although Richard Hull asserts that from Tesla’s notes, he never claimed a distance greater than 43 feet. From photographic evidence, the maximum may have been closer to 22 feet [12].

Tesla then abandoned the Colorado Springs Laboratory early in 1900, having learned what he needed from that facility, and also having become somewhat unpopular as a result of frequently knocking the local sub-station off line.

Since that time, it appears that no one has built a Tesla coil of both the size and performance of the Colorado Springs coil. Apparently the only coil of that size was built by Robert Golka at Wendover Air Force Base in Utah [8] and later moved to a facility near Leadville, Colorado [9, 19]. The original purpose of this coil was to produce artificial lightning for testing the effects of lightning striking aircraft in flight. Golka determined that the average voltage produced in Utah was about 10 MV, with the highest voltage observed being 25 MV. Operation was spectacular, even if not quite at the level of the Colorado Springs coil.

When Golka's coil was moved to Leadville, however, it performed very poorly. Golka and his associates were basically unable to properly tune the coil. There has been considerable speculation over the reasons for the difference in performance, but one problem seems to be that we did not have adequate theoretical models for the design and operation of Tesla coils. What appeared to be minor differences in location and construction caused a major decrease in performance. The number of variables was simply too large to allow for a purely experimental optimization of performance before the coil was dismantled and moved early in 1990.

Some work on theoretical models has been performed by high energy physicists [6, 10, 1, 17, 18]. They are interested in high voltage capacitor discharges for research in plasma physics and in the production of pulsed particle or radiation beams. The most common way of producing such high voltage discharges is the Marx circuit, in which capacitors are charged in parallel to a lower voltage and then discharged in series through a number of airgaps. The Marx circuit requires the capacitor bank to be divided into sub-banks well-insulated from each other and from ground. A Tesla coil offers an alternative method of charging the high voltage capacitors. Discharges are reported in the range of 100 kA at 1 MV, with one report of 2.5 MV [10]. These models are all lumped parameter models.

There are a number of experimenters who build Tesla coils as a hobby. The Tesla Coil Builders Association has several hundred members and a quarterly newsletter published by Harry Goldman [7]. Harry has announced plans to stop publishing the newsletter at the end of 2001. The Tesla Coil Builders of Richmond has been a very active local group [11], although their leader Richard Hull has recently become interested in other activities. A number of manuals are available on how to build coils [16, 4, 5]. The one by Lee [16] is especially well illustrated with pictures of capacitors and other components that might be needed for a moderate sized Tesla coil. There is an Internet listserv (www.pupman.com) that has about 700 subscribers, which has been very helpful to me.

The brothers James and Kenneth Corum have done considerable work on distributed models of Tesla coils in the past few years [2, 3]. They argue that lumped parameter models are not adequate for all situations. Sometimes a distributed circuit analysis must be made. In this case, the Tesla coil secondary and another component called the extra coil are considered

as sections of transmission lines. This explains some of the effects in an elegant manner. They have written a sophisticated computer program, TCTUTOR, to analyze Tesla coils. They have also performed considerable historical research into Tesla's notes made on his facility in Colorado Springs [21].

The Tesla coil community is divided over the issue of lumped versus distributed models. A majority favors the lumped model approach. Some are outspoken in their belief that distributed models are useless at best and just plain wrong on important issues. I confess to being somewhere in the middle on this controversy. James Corum and I both have our Ph.D.s in electromagnetic theory, so I can mostly understand what he says, and I therefore have a natural orientation to the distributed approach. In my eyes, I am like a Baptist pastor of a 50 person congregation and James is like Billy Graham. That is, I hold him in awe. I have heard the Corums speak several times, and have gotten caught up in their knowledge and excitement.

On the other hand, I cannot honestly say that TCTUTOR has been helpful to me in building and understanding Tesla coils. I can see significant problems with distributed models, which will be discussed later. And James, like many bright people, has a tendency to talk down to us slow ones. This puts some people off, of course.

In this book we will look at both lumped and distributed models. We will point out difficulties with both. We will look at some data, and ask which approach does best in describing reality.

2 Classical Tesla Coil

A classical Tesla coil contains two stages of voltage increase. The first is a conventional iron core transformer that steps up the available line voltage to a voltage in the range of 12 to 50 kV, 60 Hz. The second is a resonant air core transformer (the Tesla coil itself) which steps up the voltage to the range of 200 kV to 1 MV. The high voltage output is at a frequency much higher than 60 Hz, perhaps 500 kHz for the small units and 80 kHz (or less) for the very large units.

The lumped circuit model for the classical Tesla coil is shown in Fig. 1. The primary capacitor C_1 is a low loss ac capacitor, rated at perhaps 20 kV, and often made from mica or polyethylene. The primary coil L_1 is usually made of 4 to 15 turns for the small coils and 1 to 5 turns for the large coils. The secondary coil L_2 consists of perhaps 50 to 400 turns for the large coils and as many as 400 to 1000 turns for the small coils. The secondary capacitance C_2 is not a discrete commercial capacitor but rather is the distributed capacitance between the windings of L_2 and the voltage grading structure at the top of the coil (a toroid or sphere) and ground. This capacitance changes with the volume charge density around the secondary, increasing somewhat when the sparks start. It also changes with the surroundings of the coil, increasing as the coil is moved closer to a metal wall. This may have been one of the reasons

that Golka's coil worked better in Utah than in Colorado, because the metal walls were closer to the coil in Colorado.





Figure 1: The Classical Tesla Coil

The symbol G represents a spark gap, a device which will arc over at a sufficiently high voltage. The simplest version is just two metal spheres in air, separated by a small air gap. It acts as a voltage controlled switch in this circuit. The open circuit impedance of the gap is very high. The impedance during conduction depends on the geometry of the gap and the type of gas (usually air), and is a nonlinear function of the current density. This impedance is not negligible. A considerable fraction of the total input power goes into the production of light, heat, and chemical products at the spark gap. In any complete analysis for efficiency, an equivalent gap resistance R_{gap} could be defined such that $i^2 R_{gap}$ would represent the power loss in the gap. This would have rather limited usefulness because of the mathematical difficulty of describing the arc.

The arc in the spark gap is similar to that of an electric arc welder in visual intensity. That is, one should not stare at the arc because of possible damage to the eyes. At most displays of classical Tesla coils, the spark gap makes more noise and produces more light than the electrical display at the top of the coil.

When the gap is not conducting, the capacitor C_1 is being charged in the circuit shown in Fig. 2, where just the central part of Fig. 1 is shown. The inductive reactance is much smaller than the capacitive reactance at 60 Hz, so L_1 appears as a short at 60 Hz and the capacitor is being charged by the iron core transformer secondary.

A common type of iron core transformer used for small Tesla coils is the neon sign transformer (NST). Secondary ratings are typically 9, 12, or 15 kV and 30 or 60 mA. An NST has a large number of turns on the secondary and a very high inductance. This inductance will limit the current into a short circuit at about the rated value. An operating neon sign has a low impedance, so current limiting is important to long transformer life. However, in Tesla coil use, the NST inductance will resonate with C_1 . The NST may supply two or three times the NST rated current in this application. Overloading the NST produces longer sparks, but

Figure 2: C_1 Being Charged With The Gap Open

may also cause premature failure.

When the voltage across the capacitor and gap reaches a given value, the gap arcs over, resulting in the circuit in Fig. 3. We are not interested in efficiency in this introduction so we will model the arc as a short circuit. The shorted gap splits the circuit into two halves, with the iron core transformer operating at 60 Hz and the circuit to the right of the gap operating at a frequency (or frequencies) determined by C_1 , L_1 , L_2 , and C_2 . It should be noted that the output voltage of the iron core transformer drops to (approximately) zero while the input voltage remains the same, as long as the arc exists. The current through the transformer is limited by the transformer equivalent series impedance shown as $R_s + jX_s$ in Fig. 3. As mentioned, this operating mode is not a problem for the NST. However, the large Tesla coils use conventional transformers with per unit impedances in the range of 0.05 to 0.1. A transformer with a per unit impedance of 0.1 will experience a current of ten times rated while the output is shorted. Most transformers do not survive very long under such conditions. Golka was not alone in burning out some of his transformers. The solution is to include additional reactance in the input circuit.




Figure 3: Tesla Circuit With Gap Shorted.

The equivalent lumped circuit model of the Tesla coil while the gap is shorted is shown in Fig. 4. R_1 and R_2 are the effective resistances of the air cored transformer primary and secondary, respectively. The mutual inductance between the primary and secondary is shown

by the symbol M . The coefficient of coupling is well under unity for an air cored transformer, so the ideal transformer model used for an iron cored transformer that electrical engineering students study in the first course on energy conversion does not apply here.




Figure 4: Lumped Circuit Model Of A Tesla Coil, arc on.

At the time the gap arcs over, all the energy is stored in C_1 . As time increases, energy is shared among C_1 , L_1 , C_2 , L_2 , and M . The total energy in the circuit decreases with time because of losses in the resistances R_1 and R_2 . There are four energy storage devices so a fourth order differential equation must be solved. The initial conditions are some initial voltage v_1 , and $i_1 = i_2 = v_2 = 0$. If the arc starts again before all the energy from the previous arc has been dissipated, then the initial conditions must be changed appropriately.

The Corums present the necessary solution technique in their manual [3] and also the computer code. The voltages and currents are not single frequency sinusoids. Rather there is a frequency spectrum with one hump for M small and two humps for M large. This is fascinating material for lovers of circuit theory, but is of somewhat limited usefulness in suggesting design changes for better performance.

It appears to this author that the time domain solution is more useful than the frequency domain. We simply examine v_1 , v_2 , i_1 , and i_2 as time increases, either graphically or in some sort of tabular printout. We then change one or more of the energy storage device values and do it again. It is also helpful to calculate the energy stored in each device. If the total energy stored in the circuit is decreasing monotonically with time, at the rate power is being absorbed by R_1 and R_2 , then one can be reasonably confident that the computer code is working correctly.

The time domain solution resembles a drunken walk in that it is difficult to predict what a given value will do next. Energy is moving among storage devices like cannon balls rolling around on the deck of an old sailing ship. Patterns can be changed readily by changing component values. We need a strategy for evaluating each solution for movement toward or away from some optimum. This strategy is developed by recognizing the following facts. After a small number of half cycles of i_1 , the arc will dissipate and the spark gap will again become an open circuit. At this point we want as much energy as possible stored in the secondary, either as $i_2^2 L_2 / 2$ or $v_2^2 C_2 / 2$. Any energy stored in C_1 when the gap opens is not available to

produce the desired high voltages on C_2 .

With proper design (proper values of C_1 , L_1 , C_2 , L_2 , and M) it is possible to have all the energy in C_1 transferred to the secondary at some time t_1 . That is, at t_1 there is no voltage across C_1 and no current through L_1 . If the gap can be opened at t_1 , then there is no way for energy to get back into the primary. No current can flow, so no energy can be stored in L_1 , and without current the capacitor cannot be charged. The secondary then becomes a separate RLC circuit with nonzero initial conditions for both C_2 and L_2 , as shown in Fig. 5. This circuit will then oscillate or “ring” at a resonant frequency determined by C_2 and L_2 . With the gap open, the Tesla coil secondary is simply an RLC circuit, described in any text on circuit theory. The output voltage is a damped sinusoid.




Figure 5: Lumped Circuit Model Of A Tesla Coil, arc off.

Finding a peak value for v_2 given some initial value for v_1 thus requires a two step solution process. We first solve a fourth order differential equation to find i_2 and v_2 as a function of time. At some time t_1 the circuit changes to the one shown in Fig. 5, which is described by a second order differential equation. The initial conditions are the values of i_2 and v_2 determined from the previous solution at time t_1 . The resulting solution then gives the desired peak values for voltage and current. The process is tedious, but can readily be done on a computer. It yields some good insights as to the effects of parameter variation. It helps establish a benchmark for optimum performance and also helps identify parameter values that are at least of the correct order of magnitude. However, there are several limitations to the process which must be kept in mind.

First, as we have mentioned, the arc is very difficult to characterize accurately in this model. The equivalent R_1 will change, perhaps by an order of magnitude, with factors like i_1 , ambient humidity, and the condition, geometry, and temperature of the electrode materials. This introduces a very significant error into the results.

Second, the arc is not readily turned off at a precise instant of time. The space between electrodes must be cleared of the hot conducting plasma (the current carrying ions and electrons) before the spark gap can return to its open circuit mode. Otherwise, when energy starts to bounce back from the secondary, a voltage will appear across the spark gap, and

current will start to flow again, after the optimum time t_1 has passed. With fixed electrodes, the plasma is dissipated by thermal and chemical processes that require tens of microseconds to function. When we consider that the optimum t_1 may be $2\ \mu\text{s}$, a problem is obvious. This dissipation time can be decreased significantly by putting a fan on the electrodes to blow the plasma away. This also has the benefit of cooling the electrodes. For more powerful systems, however, the most common method is a rotating spark gap. A circular disc with several electrodes mounted on it is driven by a motor. An arc is established when a moving electrode passes by a stationary electrode, but the arc is immediately stretched out by the movement of the disc. During the time around a current zero, the resistance of the arc can increase to where the arc cannot be reestablished by the following increase in voltage.

The rotary spark gap still has limitations on the minimum arc time. Suppose we consider a disc with a radius of 0.2 m and a rotational speed of 400 rad/sec (slightly above 3600 rpm). The edge of the disc is moving at a linear velocity of $r\omega = 80\ \text{m/s}$. Suppose also that an arc cannot be sustained with arc lengths above 2 cm. It requires $0.02/80 = 25\ \mu\text{s}$ for the disc to turn this distance. This time can be shortened by making the disc larger or by turning it at a higher rate of speed, but in both cases we worry about the stress limits of the disc. Nobody wants fragments of a failed disc flying around the room. The practical lower limit of arc length seems to be about $10\ \mu\text{s}$. With larger coils this may be reasonably close to the optimum value.

The third reason for concern about the above calculations is that the Tesla coil secondary has features that cannot be precisely modeled by a lumped circuit. One such feature is ringing at ‘harmonic’ frequencies. Neither the distributed or lumped models do a particularly good job of predicting these frequencies. Data will be presented later for a medium sized secondary (operated as an extra coil, explained in the next section), with a high Q resonance at about 160 kHz. When applied power is switched off, the coil usually rings down at 160 kHz. Sometimes, however, it will ring down at $3.5(160) = 560\ \text{kHz}$. A third harmonic appears in many electrical circuits and has plausible explanations. A 3.5 ‘harmonic’ is another story entirely.

These three reasons explain why we never see a paper giving a complete Tesla coil design with experimental data verifying the theoretical design. We get started with theory, but at some point have to move to an experimental optimization. The saying is, “Tune for most smoke”, which Harry Goldman attributes to Bill Wysock and Gary Legel. It is a tribute to the experimentalists that we have coils in existence with names like “Nemesis” that can produce sparks fifteen feet long [11].

3 Magnifier

As mentioned above, the classical Tesla coil uses two stages of voltage increase. Some coilers get a third stage of voltage increase by adding a magnifier coil, also called an extra coil, to their classical Tesla coil. This is illustrated in Fig. 6.




Figure 6: The Classical Tesla Coil With Extra Coil

The extra coil and the air core transformer are not magnetically coupled. The output (top) of the classical coil is electrically connected to the input (bottom) of the extra coil with a section of copper water pipe of large enough diameter that corona is not a major problem. A separation of 2 or 3 meters is typical.

Voltage increase on the extra coil is by transmission line action, rather than the transformer action of the iron core transformer. Voltage increase on the air core transformer is partly by transformer action and partly by transmission line action. When optimized for extra coil operation, the air core transformer looks more like a transformer (greater coupling, shorter secondary) than when optimized for classical Tesla coil operation.

The lumped circuit enthusiast would say that voltage rise is by RLC resonance. Both camps agree that voltage rise in the secondary and especially in the extra coil are not by transformer action.

Although not shown in Fig. 6 the extra coil depends on ground for the return path of current flow. The capacitance from each turn of the extra coil and from the top terminal to ground is necessary for operation. Impedance matching from the Tesla coil secondary to the extra coil is necessary for proper operation. If the extra coil were fabricated with the same size coil form and wire size as the secondary, the secondary and extra coil tend to operate as a long secondary, probably with inferior performance to that of the secondary alone. There are guidelines for making the coil diameters and wire sizes different for the two coils, but optimization seems to require a significant amount of trial and error.

In my quest for a better description of Tesla coil operation, I decided that the extra coil was the appropriate place to start. It looks like a vertical antenna above a ground plane, so there is some prior art to draw from. While the classical Tesla coil makes an excellent driver to produce long sparks, it is not very good for instrumentation and measurement purposes. There are just too many variables. The spark gap may be the best high voltage switch available today, but inability to start and stop on command, plus heating effects, make it difficult to use when collecting data.

I therefore decided to build a solid state driver. Vacuum tube drivers have been used for many years and several researchers have developed drivers using power MOSFETs, so this was not entirely new territory. It turned out to be a long term project. At the beginning, I had little idea about the input impedance of a coil above a ground plane, or how much power would be required to get significant sparks (say, half a meter in length or more). There have been many iterations, but I finally produced a design that would make sparks. Two major disadvantages are that it requires a digital oscilloscope with deep memory for tuning purposes, and one can make longer sparks using a standard spark gap. These disadvantages make it unlikely to sweep the Tesla coil community. There might be situations, however, where this approach would be useful. One is a museum installation, for example, where sparks of 0.5 to 1 meter are acceptable, and long life and low maintenance are critical factors.

The remainder of this document is a collection of my notes on this project, including some deadends. There are discussions on

1. Capacitance
2. Inductance and Transformers
3. Gate Driver and Inverter
4. Lumped Model
5. Experimental Results

Capacitance appears in many different places in the Tesla coil system, in the power supply, the controller, the driver, the coil body itself, and the top toroid or sphere. It therefore gets a lengthy treatment. Other items get a somewhat lesser treatment.

References

- [1] Boscolo, I., G. Brautti, R. Coisson, M. Leo, and A. Luches, “Tesla Transformer Accelerator for the Production of Intense Relativistic Electron Beams”, *The Review of Scientific Instruments*, Vol. 46, No. 11, November 1975, pp. 1535–1538.
- [2] Corum, J. F. and K. L. Corum, “A Technical Analysis of the Extra Coil as a Slow Wave Helical Resonator”, *Proceedings of the 1986 International Tesla Symposium*, Colorado Springs, Colorado, July 1986, published by the International Tesla Society, pp. 2-1 to 2-24.
- [3] Corum, James, F., Daniel J. Edwards, and Kenneth L. Corum, *TCTUTOR - A Personal Computer Analysis of Spark Gap Tesla Coils*, Published by Corum and Associates, Inc., 8551 State Route 534, Windsor, Ohio, 44099, 1988.

- [4] Couture, J. H., *JHC Tesla Handbook*, JHC Engineering Co., 19823 New Salem Point, San Diego, CA, 92126, (1988).
- [5] Cox, D. C., *Modern Resonance Transformer Design Theory*, Tesla Book Company, P. O. Box 1649, Greenville, TX 75401, (1984).
- [6] Finkelstein, David, Phillip Goldberg, and Joshua Shuchatowitz, “High Voltage Impulse System”, *The Review of Scientific Instruments*, Volume 37, Number 2, February 1966, pp. 159-162.
- [7] Goldman, Harry, *Tesla Coil Builders Association News*, 3 Amy Lane, Queensbury, NY, 12804, (518) 792-1003.
- [8] Golka, Robert K., “Long Arc Simulated Lightning Attachment Testing Using a 150 kW Tesla Coil”, IEEE International Symposium on Electromagnetic Compatibility, October 9-11, 1979, San Diego, CA, pp. 150 - 155.
- [9] Grotz, Toby, “Project Tesla - An Update”, *Tesla Coil Builders Association News*, Volume 9, No. 1, January, February, March, 1990, pp. 16-18.
- [10] Hoffmann, C. R. J., “A Tesla Transformer High-Voltage Generator”, *The Review of Scientific Instruments*, Vol. 46, No. 1, January 1975, pp. 1-4.
- [11] Hull, Richard L., Tesla Coil Builders of Richmond, 7103 Hermitage Rd., Richmond, Virginia, 23228.
- [12] Hull, Richard L., “The Tesla Coil Builder’s Guide to The Colorado Springs Notes of Nikola Tesla”, Tesla Coil Builders of Richmond, 1993.
- [13] Johnson, Gary L., “Using Power MOSFETs To Drive Resonant Transformers”, *Tesla 88*, International Tesla Society, Inc., 330-A West Uintah, Suite 215, Colorado Springs, CO 80905, Vol. 4, No. 6, November/December 1988, pp. 7-13.
- [14] Johnson, Gary L., *The Search For A New Energy Source*, Johnson Energy Corporation, P.O. Box 1032, Manhattan, KS 66505, 1997.
- [15] Jones, H. W., “Project Insight - A Study of Tesla’s Advanced Concepts”, Proceedings of the Tesla Centennial Symposium, Colorado Springs, Colorado, August 9-12, 1984.
- [16] Lee, Thomas W., *High Voltage Generation with Air-Core Solenoids*, 8329 E. San Salvador Dr. Scottsdale, Arizona, 85258, (1989).
- [17] Luches, A. and A. Perrone, “Coupled Marx-Tesla Circuit for Production of Intense Relativistic Electron Beams”, *The Review of Scientific Instruments*, Vol. 49, No. 12, December 1978, pp. 1629-1630.

- [18] Matera, Manlio, Roberto Buffa, Giuliano Conforti, Lorenzo Fini, and Renzo Salimbeni, “Resonant Transformer Command Charging System for High Repetition Rate Rare-Gas Halide Lasers”, *The Review of Scientific Instruments*, Vol. 54, No. 6, June 1983, pp. 716-718.
- [19] Peterson, Gary L., “Project Tesla Evaluated”, *Power and Resonance, The International Tesla Society’s Journal*, Volume 6, No. 1, January/February/ March 1990, pp. 25-34.
- [20] Terbo, William H., “Opening Address”, Proceedings of the Tesla Centennial Symposium, Colorado Springs, Colorado, August 9-12, 1984.
- [21] Tesla, Nikola, *Colorado Springs Notes*, A. Marincic, Editor, Nolit, Beograd, Yugoslavia, 1978, 478 pages.

IDEAL CAPACITORS

Capacitors are used in almost every activity of electrical engineering, yet information on capacitor characteristics is scattered through a variety of textbooks, databooks, and manufacturers literature. The following is an attempt to organize some of this information.

A typical capacitor is a two-terminal device consisting of two conductors separated by a dielectric. When a voltage difference V_o is applied to the conductors, a charge of $+Q$ will appear on one conductor and an equal and opposite charge $-Q$ on the other conductor. The capacitance C is defined as the ratio of the charge on one conductor to the potential difference.

$$C = \frac{Q}{V_o} \quad (1)$$

where C is in farads, Q is in coulombs, and V_o is in volts. Actually, one farad is a rather large capacitance, so capacitance values are usually expressed in terms of μF (10^{-6}F) or pF (10^{-12}F).

The total energy stored in a capacitor is

$$W_E = \frac{1}{2} \int_{vol} \epsilon E^2 dv = \frac{1}{2} C V_o^2 = \frac{1}{2} Q V_o = \frac{1}{2} \frac{Q^2}{C} \quad (2)$$

where W_E is in joules, E is the electric field in V/m , and ϵ is the permittivity. The integral expression shows that the energy stored in a capacitor with a fixed voltage difference across it increases as the permittivity of the material increases.

The permittivity is usually expressed as the product of a relative permittivity ϵ_r and the permittivity of free space ϵ_o .

$$\epsilon = \epsilon_r \epsilon_o \quad (3)$$

where

$$\epsilon_o = 8.854 \times 10^{-12} \text{F/m} \quad (4)$$

The relative permittivity is unity for a vacuum and typically in the range of 2 to 6 for most dielectrics, as we shall discuss in more detail later.

Capacitors are frequently used in series in a circuit, as shown in Fig. 1. There will be no actual charge transfer through the dielectric material. However, the electric fields will cause a movement of charge within the series string. The battery supplies a positive charge to the left plate of C_1 . This positive charge attracts an equivalent negative charge on the right plate

of C_1 . The movement of this negative charge leaves behind a positive charge of the same amount, which the electric field will force onto the left plate of C_2 . The process continues until each capacitor has the same charge $Q = Q_s$ on its left plate. That is,

$$Q_s = Q_1 = Q_2 = Q_3 \quad (5)$$

The total voltage across the series combination is

$$V = V_1 + V_2 + V_3 \quad (6)$$

and since

$$V = \frac{Q}{C} \quad (7)$$

$$\frac{Q}{C_s} = \frac{Q}{C_1} + \frac{Q}{C_2} + \frac{Q}{C_3} \quad (8)$$

which can be solved for the series capacitance C_s .

$$C_s = \frac{1}{\frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3}} \quad (9)$$




Figure 1: Capacitors in Series

A circuit of parallel capacitors is shown in Fig. 2. The voltage on each capacitor is the same and the amount of stored charge on each capacitor will be proportional to the individual capacitance values. It is not hard to show that the total parallel capacitance C_p is given by

$$C_p = C_1 + C_2 + C_3 \quad (10)$$




Figure 2: Capacitors in Parallel

1 Capacitance of Common Geometries

The ratio of Q to V_o depends on the geometrical arrangement of the conductors and on the electrical characteristics of the dielectric. The capacitance of a parallel plate capacitor, as illustrated in Fig. 3, is

$$C = \frac{\epsilon A}{d} \quad (11)$$

where A is the area in m^2 and d is the separation between plates. This formula is accurate only when fringing can be neglected, that is, when d is small in comparison with A .




Figure 3: A Parallel Plate Capacitor

Later on, we will be interested in the capacitance of geometries where there are two different dielectrics. The simplest case is shown in Fig. 4. There is a layer of dielectric with relative permittivity $\epsilon_r > 1$ of thickness x , and a second layer of air, with thickness y . The boundary between the two dielectrics can be considered a floating electrode. In fact, we can place a conducting plate on the boundary without changing the results at all. We basically have two capacitors in series. When we solve for the series capacitance, we get

$$C = \frac{\epsilon_o A}{y} \frac{\epsilon_r}{(\epsilon_r + x/y)} \quad (12)$$




Figure 4: Capacitor with Different Dielectrics

Another important geometry is that of a coaxial cable of inner radius a , outer radius b , and length ℓ , which has capacitance

$$C = \frac{2\pi\epsilon\ell}{\ln(b/a)} \quad (13)$$

The geometry for the coaxial cable is shown in Fig. 5.




Figure 5: Coaxial Cable

The common 50 Ω coaxial cable 213/U (RG-8A/U) has a nominal capacitance of 29.5 pF/ft (96.8 pF/m). The small 75 Ω video cable 59B/U has a nominal capacitance of 21.0 pF/ft (68.9 pF/m). Most other 50 and 75 Ω cables will have capacitance values very close to these. Physically larger cables capable of carrying more power will have b and a increased in the same proportion so the ratio b/a and the capacitance will remain the same as that of a smaller cable.

Another geometry of great practical interest is the twin conductor transmission line, shown in Fig. 6a. This is composed of two conductors of radius r , with separation $2h$ between conductor centers. The conductor-to-conductor capacitance C_{cc} is given by

$$C_{cc} = \frac{\pi\epsilon\ell}{\ln[(h + \sqrt{h^2 - r^2})/r]} = \frac{\pi\epsilon\ell}{\cosh^{-1}(h/r)} \quad (14)$$

If the two conductors have a small radius and are located far apart, the expression for capacitance becomes

$$C_{cc} = \frac{\pi\epsilon\ell}{\ln(2h/r)} \quad (15)$$

The error in the approximate expression is only 5.26% when $h = 2r$ and 1.16% when $h = 3r$ so the latter equation really has a wide range of usefulness.

The conductor-to-plane capacitance C_{cp} between a cylindrical conductor of radius r and a conducting plane a distance h from the cylinder, as shown in Fig. 6b, is twice the value given by the previous two equations.

$$C_{cp} = \frac{2\pi\epsilon\ell}{\cosh^{-1}(h/r)} \approx \frac{2\pi\epsilon\ell}{\ln(2h/r)} \quad (16)$$

This equation can be used to find the capacitance between two unequal conductors. We find the capacitance of each conductor to an imaginary ground plane, and then combine the two values for C_{cp} using the formula for capacitors in series.




Figure 6: Twin Conductor Transmission Line

Another geometry of interest is that of a spherical capacitor of two concentric spheres with radii a and b ($b > a$) as shown in Fig. 7. It is not practical to actually build capacitors this way, but the symmetry allows an exact formula for capacitance to be calculated easily. This is done in most introductory courses of electromagnetic theory. The capacitance is given by [4, Page 165]

$$C = \frac{4\pi\epsilon}{1/a - 1/b} \quad (17)$$

If the outer sphere is made larger, the capacitance decreases, but does not go to zero. In the limit as $b \rightarrow \infty$, the *isolated* or *isotropic* capacitance of a sphere of radius a becomes

$$C_{\infty} = 4\pi\epsilon a \quad (18)$$




Figure 7: Spherical Capacitor

C_∞ gives us a lower bound for the capacitance of a spherical top loading element of a Tesla coil with respect to ground. One way of arriving at a reasonable estimate of the actual capacitance is to start with the isotropic capacitance C_∞ and add a correction term, as we shall see in the next section.

We will also be interested in isotropic capacitances of shapes other than spheres. Some will be difficult to impossible to calculate analytically, so we will use the isotropic capacitance of a sphere as a starting point for making an estimate. Rectangular boxes and short cylinders, for example, will have similar isotropic capacitances to a sphere. We just need to find an equivalent radius (or diameter) for these nonspherical shapes. Several different equivalents could be used, such as a geometrical mean equivalent, an arithmetic mean equivalent, or the radius of a sphere with the same surface area as the other shape. It turns out that the simplest method, the arithmetic mean, does quite well [2]. Consider a rectangular box with orthogonal edge dimensions a , b , and c . Define an equivalent sphere diameter ℓ_e where

$$\ell_e = \frac{a + b + c}{3} \quad (19)$$

For other shapes we use the dimensions of the box in which the other shape can be placed. Making the change from radius to diameter, the isotropic capacitance is now

$$C_\infty = 2\pi\epsilon\ell_e \quad (20)$$

This approach will give acceptable results in many cases. But, of course, there is no way of knowing the amount of error, or when some other approach would yield better results. It will get us in the right ballpark, however, and sometimes allow us to determine lower bounds of acceptable values obtained from other techniques.

Suppose, for example, that we wanted the capacitance between two spheres separated by several diameters. We suspect that the parallel plate capacitor formula will not be very accurate and are unable to locate a better formula. What can we do? The lower bound for the capacitance between two spheres is just half of C_∞ , the isotropic capacitance for one sphere, as can be argued from Fig. 8.

The dark point at the center of the figure is obviously ‘the’ point at infinity that mathematicians love to talk about. If the capacitance of each sphere with respect to this ‘point’




Figure 8: Capacitance Between Two Spheres

is C_∞ , then the capacitance between spheres has a lower bound of $C_\infty/2$, from the formula for two capacitors in series. Bringing the spheres closer together will increase the capacitance but they cannot be separated far enough to reduce the capacitance below $C_\infty/2$.

2 Toroid Capacitance

An important emphasis of this book is the analysis of the Tesla coil. Among other things we will be interested in the capacitance of the top element (usually a ‘fat’ toroid) with respect to ground. We will also need the capacitance between adjacent turns (which look like ‘thin’ toroids) and finally the capacitance of a turn with respect to ground. As usual, we will use the published results as much as possible and leave the derivations to others.

The dimensions of a toroid are shown in Fig. 9.




Figure 9: Toroid Dimensions

There are other coordinate systems besides rectangular, cylindrical, and spherical in which variables can be separated and Laplace’s equation solved. One of these is toroidal coordinates. Moon and Spencer use this coordinate system to solve for the capacitance of an isolated toroid [3, Page 375] as

$$C_{MS} = 8a\epsilon \left[\frac{Q_{-1/2}(\cosh \eta_0)}{P_{-1/2}(\cosh \eta_0)} + 2 \sum_{n=1}^{\infty} \frac{Q_{n-1/2}(\cosh \eta_0)}{P_{n-1/2}(\cosh \eta_0)} \right] \quad (21)$$

where P and Q are Legendre functions of first and second order and a and η_0 will be discussed later. The subscript ‘MS’ refers to Moon and Spencer, to distinguish the capacitance obtained from the value to be obtained from some empirical formulas later.

Note that Eq. 21 is the corrected version. Godfrey Loudner [1] found that Moon and Spencer had an extra π in their expression (second equation from the top on p. 375), and also in Eq. 13.29 and Eq. 13.29a. There is also a typo on p. 373, second equation from the top, where $V/2$ should be replaced by V/π . Even the great ones can make a mistake!

The Legendre functions can be written in many different forms, as integrals or infinite series, converging for arguments greater than unity or less than unity, and so on. The non-integer subscript $(n-1/2)$ adds another layer of complexity. Many math books do not mention the non-integer case, so one must be diligent in finding the correct expressions. The order of difficulty is much greater than for the sphere. Our mothers warned us that there would be days like this, but let us proceed.

Moon and Spencer give us an expression for $Q_{n-1/2}(\cosh \eta_0)$ [3, Page 373]

$$Q_{n-1/2}(\cosh \eta_0) = \frac{1}{\sqrt{2}} \int_0^\pi \frac{\cos n\theta \, d\theta}{\sqrt{(\cosh \eta_0 - \cos \theta)}} \quad (22)$$

For some reason, they do not give a similar expression for $P_{n-1/2}(\cosh \eta_0)$. However, they do give plots for both functions, which is convenient for checking computational results.

Smythe [6, Page 159] gives an expression for P in the form

$$\pi P_{n'}^m(x) = (n' + 1)(n' + 2) \cdots (n' + m) \int_0^\pi [x + \sqrt{x^2 - 1} \cos \theta]^{n'} \cos m\theta \, d\theta \quad (23)$$

This expression is valid for $x > 0$ and for any n' , including $1/2$, $3/2$, etc. We are only interested in the case $m = 0$. If we interpret the product of terms ahead of the integral sign as a factorial with zero entries (that is with a value of unity), the expression becomes

$$P_{n'}(x) = \frac{1}{\pi} \int_0^\pi [x + \sqrt{x^2 - 1} \cos \theta]^{n'} \, d\theta \quad (24)$$

These expressions for P and Q can be numerically integrated using any scientific programming language (QuickBasic, etc.).

We now return to our discussion of a and η_0 . These are coordinate values in toroidal coordinates, hence need to be translated into a more familiar coordinate system. The major toroid radius w and the minor radius r are given by

$$w = a \coth \eta \quad (25)$$

$$r = \frac{a}{\sinh \eta} \quad (26)$$

Eliminating a between the two equations yields

$$w = r \sinh \eta \coth \eta = r \cosh \eta \quad (27)$$

so

$$\frac{w}{r} = \cosh \eta = x \quad (28)$$

We then solve for η as

$$\eta = \cosh^{-1} \frac{w}{r} = \ln \left(\frac{w + \sqrt{w^2 - r^2}}{r} \right) \quad (29)$$

and for a as

$$a = r \sinh \eta = \frac{r}{2} (\epsilon^\eta - \epsilon^{-\eta}) \quad (30)$$

This ‘exact’ formulation is of most interest to EM theorists and computer programmers. It will seem like overkill to most Tesla coil enthusiasts who just need to get in the right ballpark with a capacitance estimate. For this reason empirical formulas have been developed which yield an approximate value, adequate for most purposes, but obtained with much less effort. Empirical formulas for the capacitance of a toroid are given by [5]

$$C_S = \frac{1.8(D - d)}{\ln(8(D - d)/d)} \quad (d/D < 0.25) \quad (31)$$

$$C_S = 0.37D + 0.23d \quad (d/D > 0.25) \quad (32)$$

where D is the toroid major diameter, outside to outside, in cm, d is the toroid minor diameter in cm, and the capacitance is given in pF. Table 1 gives some isotropic capacitance values, both from the Moon and Spencer numerical integration and the empirical formulas. The deviation or error of C_S with respect to C_{MS} is given in the last column in percent.

We see that the empirical formulation agrees with Moon and Spencer to within 1% for the case of fat toroids, but gets progressively worse as the toroid gets thinner. The error is within 5% for d down to about 0.5 cm (4 gauge wire). Toroids this thin do not have the mechanical strength necessary to serve as top loading elements of a Tesla coil, so we can conclude that the empirical formulas are quite adequate for most purposes.

Table 1: Isotropic Capacitance of Toroids

w	r	D	d	C_{MS}	C_S	error
meters		cm		pF	pF	%
.3	.15	90	30	40.46	40.20	-0.63
.2	.1	60	20	26.97	26.80	-0.63
.1	.05	30	10	13.49	13.40	-0.63
.2	.08	56	16	24.55	24.40	-0.62
.2	.06	52	12	22.02	21.93	-0.42
.2	.04	48	8	19.28	19.52	+1.23
.2	.02	44	4	16.00	16.43	+2.70
.2	.01	42	2	13.72	14.19	+3.42
.2	.005	41	1	12.00	12.48	+4.05
.2	.0025	40.5	.5	10.62	11.14	+4.97
.2	.001	40.2	.2	9.09	9.76	+7.36

3 Solenoid Capacitance (Medhurst)

The isotropic capacitance of a sphere was given above as a simple formula. We looked at the theoretical formulas for capacitance of a toroid, but basically gave up and went to a simpler empirical version. After that learning experience, we will not even try to write exact equations for the isotropic capacitance of a cylinder. We will immediately write the empirical equations as developed many years ago by a man named Medhurst. These will be expressed in several different versions, to meet different needs. The simplest expression for the isotropic capacitance of a cylindrical coil of wire, with diameter D and coil length ℓ , is

$$C_M = HD \text{ pF} \quad (33)$$

where D is in cm, and H is a multiplying factor that equals 0.51 for $\ell/D = 2$, 0.81 for $\ell/D = 5$, and varies linearly between 0.51 and 0.81 for ℓ/D between 2 and 5. Most coilers prefer values for ℓ/D between 3.5 and 4.5, so this linear range is adequate for most purposes.

An expression for H that works for ℓ/D between 2 and 8 is

$$H = 0.100976 \frac{\ell}{D} + 0.30963 \quad (34)$$

Another expression for H that works for ℓ/D between 1 and 8 is

$$H = 0.0005\left(\frac{\ell}{D}\right)^4 - 0.0097\left(\frac{\ell}{D}\right)^3 + 0.0648\left(\frac{\ell}{D}\right)^2 - 0.0757\left(\frac{\ell}{D}\right) + 0.4723 \quad (35)$$

4 Tesla Coil Capacitance

We now have expressions for the isotropic capacitance C_S of a toroid and the isotropic capacitance C_M of a coil. The next step would be to set the toroid on top the coil and add the two capacitances to get an effective capacitance C_{tc} for the Tesla coil. Unfortunately, this only works when the toroid is far away from the solenoid. As the toroid is brought near the coil form, shielding occurs such that the effective capacitance is less than the sum of the two isotropic capacitances. The Tesla coil capacitance might be written as

$$C_{tc} = C_M + KC_S \quad (36)$$

where $K < 1$. A value of $K = 0.75$ should result in a number for C_{tc} within 20% of the correct value for most Tesla coils. The resonant frequency is related to the square root of C_{tc} so a 20% error in capacitance results in only a 10% error in resonant frequency.

Most readers probably feel disappointed here. We have gone to considerable effort and still come up short of an accurate formula for C_{tc} . Our effort is not entirely wasted because we can do ‘what if’ analyses relatively quickly. Questions about the effect of changing coil diameter, coil length, or toroid diameter can be answered with adequate accuracy.

Someone might suggest using a modern digital capacitance meter to measure C_{tc} . This method would probably have greater error than the above formula, because the leads of the capacitance meter have a similar capacitance value as C_{tc} . Also the presence of the meter and a human body will change the capacitance.

It is possible to calculate C_{tc} numerically using Gauss’s Law. If one is careful about measuring and entering all the dimensions and the locations of grounded surfaces, one should get a value for C_{tc} well within 5% of the correct value. There are programs available in the Tesla coil community that do this.

References

- [1] Loudner, Godfrey. Private communication, December 26, 2001.
- [2] Maruvada, P. Sarma and N. Hyttén-Cavallius, “Capacitance Calculations for Some Basic High Voltage Electrode Configurations”, IEEE Transactions on Power Apparatus and Systems, Vol. PAS-94, No. 5, September/October 1975, pp. 1708-1713.
- [3] Moon, Parry and Domina Eberle Spencer, *Field Theory for Engineers*, D. Van Nostrand Company, Princeton, New Jersey, 1961.
- [4] Plonus, Martin A., *Applied Electromagnetics*, McGraw-Hill, New York, 1978.
- [5] Schoessow, Michael, *TCBA News*, Vol. 6, No. 2, April/May/June 1987, pp. 12-15.

- [6] Smythe, William R., *Static and Dynamic Electricity*, Hemisphere Publishing Corporation, A member of the Taylor & Francis Group, New York, Third Edition, Revised Printing, 1989.

LOSSY CAPACITORS

1 Dielectric Loss

Capacitors are used for a wide variety of purposes and are made of many different materials in many different styles. For purposes of discussion we will consider three broad types, that is, capacitors made for ac, dc, and pulse applications. The ac case is the most general since ac capacitors will work (or at least survive) in dc and pulse applications, where the reverse may not be true.

It is important to consider the losses in ac capacitors. All dielectrics (except vacuum) have two types of losses. One is a conduction loss, representing the flow of actual charge through the dielectric. The other is a dielectric loss due to movement or rotation of the atoms or molecules in an alternating electric field. Dielectric losses in water are the reason for food and drink getting hot in a microwave oven.

One way of describing dielectric losses is to consider the permittivity as a complex number, defined as

$$\epsilon = \epsilon' - j\epsilon'' = |\epsilon|e^{-j\delta} \quad (1)$$

where

ϵ' = ac capacitance

ϵ'' = dielectric loss factor

δ = dielectric loss angle

Capacitance is a complex number C^* in this definition, becoming the expected real number C as the losses go to zero. That is, we define

$$C^* = C - jC'' \quad (2)$$

One reason for defining a complex capacitance is that we can use the complex value in any equation derived for a real capacitance in a sinusoidal application, and get the correct phase shifts and power losses by applying the usual rules of circuit theory. This means that most of our analyses are already done, and we do not need to start over just because we now have a lossy capacitor.

Equation 1 expresses the complex permittivity in two ways, as real and imaginary or as magnitude and phase. The magnitude and phase notation is rarely used. Instead, people

usually express the complex permittivity by ϵ' and $\tan \delta$, where

$$\tan \delta = \frac{\epsilon''}{\epsilon'} \quad (3)$$

where $\tan \delta$ is called either the *loss tangent* or the *dissipation factor* DF.

The real part of the permittivity is defined as

$$\epsilon' = \epsilon_r \epsilon_o \quad (4)$$

where ϵ_r is the *dielectric constant* and ϵ_o is the permittivity of free space.

Dielectric properties of several different materials are given in Table 1 [4, 5]. Some of these materials are used for capacitors, while others may be present in oscillators or other devices where dielectric losses may affect circuit performance. The dielectric constant and the dissipation factor are given at two frequencies, 60 Hz and 1 MHz. The righthand column of Table 1 gives the approximate breakdown voltage of the material in V/mil, where 1 mil = 0.001 inch. This would be for thin layers where voids and impurities in the dielectrics are not a factor. Breakdown usually destroys a capacitor, so capacitors must be designed with a substantial safety factor.

It can be seen that most materials have dielectric constants between one and ten. One exception is barium titanate with a dielectric constant greater than 1000. It also has relatively high losses which keep it from being more widely used than it is.

We see that polyethylene, polypropylene, and polystyrene all have small dissipation factors. They also have other desirable properties and are widely used for capacitors. For high power, high voltage, and high frequency applications, such as an antenna capacitor in an AM broadcast station, the ruby mica seems to be the best.

Each of the materials in Table 1 has its own advantages and disadvantages when used in a capacitor. The ideal dielectric would have a high dielectric constant, like barium titanate, a low dissipation factor, like polystyrene, a high breakdown voltage, like mylar, a low cost, like aluminum oxide, and be easily fabricated into capacitors. It would also be perfectly stable, so the capacitance would not vary with temperature or voltage. No such dielectric has been discovered so we must apply engineering judgment in each situation, and select the capacitor type that will meet all the requirements and at least cost.

Capacitors used for ac must be *unpolarized* so they can handle full voltage reversals. They also need to have a lower dissipation factor than capacitors used as dc filter capacitors, for example. One important application of ac capacitors is in tuning electronic equipment. These capacitors must have high stability with time and temperature, so the tuned frequency does not drift beyond some specified amount.

Another category of ac capacitor is the motor run or power factor correcting capacitor. These are used on motors and other devices operating at 60 Hz and at voltages up to 480 V or more. They are usually much larger than capacitors used for tuning electronic circuits,

and are not sold by electronics supply houses. One has to ask for motor run capacitors at an electrical supply house like Graingers. These also work nicely as dc filter capacitors if voltages higher than allowed by conventional dc filter capacitors are required.

The term *power factor* PF may also be defined for ac capacitors. It is given by the expression

$$\text{PF} = \cos \theta \quad (5)$$

where θ is the angle between the current flowing through the capacitor and the voltage across it.

The capacitive reactance for the sinusoidal case can be defined as

$$X_C = \frac{1}{\omega C} \quad (6)$$

where $\omega = 2\pi f$ rad/sec, and f is in Hz.

In a lossless capacitor, $\epsilon'' = 0$, and the current leads the voltage by exactly 90° . If ϵ'' is greater than zero, then the current has a component in phase with the voltage.

$$\cos \theta = \frac{\epsilon''}{\sqrt{(\epsilon'')^2 + (\epsilon')^2}} \quad (7)$$

For a good dielectric, $\epsilon' \gg \epsilon''$, so

$$\cos \theta \approx \frac{\epsilon''}{\epsilon'} = \tan \delta \quad (8)$$

Therefore, the term *power factor* is often used interchangeably with the terms *loss tangent* or *dissipation factor*, even though they are only approximately equal to each other.

We can define the apparent power flow into a parallel plate capacitor as

$$S = VI = \frac{V^2}{-jX_C} = jV^2\omega C^* = jV^2\frac{\omega A}{d}(\epsilon' - j\epsilon'') = V^2\frac{\omega A}{d}\epsilon_r\epsilon_o(j + \text{DF}) \quad (9)$$

By analogy, the apparent power flow into any arbitrary capacitor is

$$S = P + jQ = V^2\omega C(j + \text{DF}) \quad (10)$$

Table 1: Dielectric Constant ϵ_r , Dissipation Factor DF and Breakdown Strength V_b of selected materials.

Material	ϵ_r 60 Hz	ϵ_r 10^6 Hz	DF 60 Hz	DF 10^6 Hz	V_b V/mil
Air	1.000585	1.000585	-	-	75
Aluminum oxide	-	8.80	-	0.00033	300
Barium titanate	1250	1143	0.056	0.0105	50
Carbon tetrachloride	2.17	2.17	0.007	<0.00004	-
Castor oil	3.7	3.7	-	-	300
Glass, soda-borosilicate	-	4.84	-	0.0036	-
Heavy Soderon	3.39	3.39	0.0168	0.0283	-
Lucite	3.3	3.3	-	-	500
Mica, glass bonded	-	7.39	-	0.0013	1600
Mica, glass, titanium dioxide	-	9.0	-	0.0026	-
Mica, ruby	5.4	5.4	0.005	0.0003	-
Mylar	2.5	2.5	-	-	5000
Nylon	3.88	3.33	0.014	0.026	-
Paraffin	2.25	2.25	-	-	250
Plexiglas	3.4	2.76	0.06	0.014	-
Polycarbonate	2.7	2.7	-	-	7000
Polyethylene	2.26	2.26	<0.0002	<0.0002	4500
Polypropylene	2.25	2.25	<0.0005	<0.0005	9600
Polystyrene	2.56	2.56	<0.00005	0.00007	500
Polysulfone	3.1	3.1	-	-	8000
Polytetrafluoroethylene(teflon)	2.1	2.1	<0.0005	<0.0002	1500
Polyvinyl chloride (PVC)	3.2	2.88	0.0115	0.016	
Quartz	3.78	3.78	0.0009	0.0001	500
Tantalum oxide	2.0	-	-	-	100
Transformer oil	2.2	-	-	-	250
Vaseline	2.16	2.16	0.0004	<0.0001	-

The power dissipated in the capacitor is

$$P = V^2 \omega C'' = V^2 \omega C(\text{DF}) \quad (11)$$

Example

Find the real and reactive power into a ruby mica capacitor with area $A = 0.03 \text{ m}^2$, and a dielectric thickness $d = 0.001 \text{ m}$, if the voltage is 2000 V (rms) at a frequency $f = 1 \text{ MHz}$.

$$\begin{aligned} S &= V^2 \frac{\omega A}{d} \epsilon_r \epsilon_o (j + \text{DF}) \\ &= (2000)^2 \frac{2\pi(10^6)(0.03)}{0.001} (5.4)(8.854 \times 10^{-12})(j + 0.0003) \\ &= j36040 + 10.8 \end{aligned}$$

The capacitor is absorbing 36040 capacitive VARs (Volt Amperes Reactive) and 10.8 Watts. The real power of 10.8 W appears as heat and must be removed by appropriate heat sinks.

The real power dissipation in a capacitor varies directly with frequency if the dissipation factor remains constant, and also with the square of the voltage. At low frequencies, the voltage limit is determined by the dielectric strength. At high frequencies, however, the voltage limit may be determined by the ability of the capacitor to dissipate heat. If the ruby mica capacitor in the previous example could safely dissipate only 10 W, but was to be used at 5 MHz, the operating voltage must be reduced from 2000 V to keep the losses in an acceptable range.

The dissipation factor varies significantly with frequency for some materials in Table 1. The actual variation must be determined experimentally. If interpolation or extrapolation seems necessary to find the loss at some frequency not given in the Table, the best assumption would be that DF varies linearly with $\log f$. That is, if $\text{DF} = 0.02$ at $f = 10^2$ and 0.01 at $f = 10^6$, a reasonable assumption at $f = 10^4$ would be that $\text{DF} = 0.015$.

The basic circuit model for a capacitor is shown in Fig. 1. Any conductor, whether straight or wound in a coil, has inductance, and the capacitor inductance is represented by a series inductance L_s . The effects of conductor resistance and dielectric losses are represented by a series resistance R_s . Leakage current through the capacitor at dc flows through a parallel resistance R_p . In some manufacturer's databooks, these are called ESL, ESR, and EPR, where

$$L_s = \text{ESL} = \text{equivalent series inductance}$$

$$R_s = \text{ESR} = \text{equivalent series resistance}$$

$$R_p = \text{EPR} = \text{equivalent parallel resistance}$$




Figure 1: Equivalent Circuit for a Capacitor

This circuit indicates that every capacitor has a *self-resonant* frequency, above which it becomes an inductor. This is certain to puzzle a student making measurements on a capacitor above this frequency if the student is not aware of this fact. R_s is readily measured by applying this frequency to a capacitor, measuring the voltage and current, and calculating the ratio. The capacitive and inductive reactances cancel at the resonant frequency, leaving only R_s to limit the current. The resistance R_p will always be much larger than the capacitive reactance at the resonant frequency, so this resistance can be neglected for this computation.

The self-resonant frequency of a high capacitance unit is lower than that for a low capacitance unit. Hence, in some circuits we will see two capacitors in parallel, say a $10\ \mu\text{F}$ in parallel with a $0.001\ \mu\text{F}$ capacitor as shown in Fig. 2. At first glance, this seems totally unnecessary. However, the larger capacitor is used to filter low frequencies, say in the audio range, while the small capacitor filters the high frequencies which are above the self-resonant frequency of the large capacitor.




Figure 2: Capacitors in Parallel to Filter Two Different Frequencies

An example of the variation of R_s and self-resonant frequency f_{res} with capacitance value and rated voltage is given in Table 2. These are metallized polypropylene capacitors designed for switch-mode power supplies by the company Electronic Concepts, Inc. This application requires a low R_s and a high f_{res} so capacitors designed for other applications will tend to have higher values of R_s and lower values of f_{res} .

The operating frequency range will obviously be less than the self-resonant frequency. This

Table 2: Capacitor Resistance and Self-resonant Frequency for Electronic Concepts Type 5MP Metallized Polypropylene capacitors.

VDC	C	R_s	f_{res}
volts	μF	Ω	kHz
100	1	0.015	1065
100	2	0.012	703
100	5	0.010	385
100	10	0.009	248
200	1	0.020	861
200	2	0.015	609
200	5	0.011	323
200	10	0.009	200
400	1	0.019	784
400	2	0.015	511
400	5	0.010	283
400	10	0.006	200

could easily be as low as a few kHz for large motor run capacitors.

Capacitors will always be rated for a working voltage as well as a specific value of capacitance. This voltage will always be well under the breakdown voltage of the dielectric. It may be specified either as a dc voltage or an ac voltage, depending on the application. Motor run capacitors are always operated on ac, so the voltage is specified as, say, 370 or 480 VAC. They can also be operated on dc, with a dc rating of at least $\sqrt{2}$ times the ac rating. Electrolytic capacitors that can only be operated on dc will have their working voltage expressed as WVDC. This is the maximum dc voltage, plus the peak of the ac ripple voltage, that should be continuously applied to a capacitor to prevent excessive deterioration and aging. Capacitors that are sometimes used on ac, and sometimes on dc, usually have working voltages expressed as WVDC, as was the case for the polypropylene capacitors in Table 2.

We will now present some more detailed background information on several of the dielectrics shown in Table 1.

Mica is a natural material that can be easily split into thin layers. It is very stable and does not deteriorate with age. The maximum capacitance is on the order of $0.03 \mu\text{F}$. Mica capacitors tend to be quite expensive and the larger sizes are used only in critical applications like radio transmitters.

Glass capacitors were first developed during World War II as a replacement for mica capacitors when supplies of mica were threatened. Glass capacitors exhibit excellent long-

term parametric stability, low losses, and can be used at high frequencies. They are used in aerospace applications where capacitance must not vary. The maximum size is limited to about $0.01\ \mu\text{F}$.

Paper capacitors use kraft paper impregnated with non-ionized liquid electrolyte as the dielectric between aluminum foils. These capacitors can provide large capacitances and voltage ratings but tend to be physically large. The thickness of the paper layers and the particular electrolyte used can be varied to produce a wide range of electrical characteristics, which is the reason for not listing paper in the dielectrics of Table 1.

Ceramic capacitors are made with one of a large number of ceramic materials, which include aluminum oxide, barium titanate, and porcelain. These are very widely used as bypass capacitors in electronic circuits. The older style is a single layer of dielectric separating two conducting plates and packaged in a small disc. The newer style is the monolithic, which appears in a rectangular package. It consists of alternating layers of ceramic material and printed electrodes which are sintered together to form the final package. The self resonant frequency is on the order of 15 MHz for the $0.01\ \mu\text{F}$ size, for a total lead length of 0.5 inch, and on the order of 165 MHz for the $0.0001\ \mu\text{F} = 100\ \text{pF}$ size. Ceramic capacitors are classified into Class 1 ($\epsilon_r < 600$) and Class 2 ($\epsilon_r > 600$) by the Electronics Industries Association (EIA). Class 1 ceramic capacitors tend to be larger than their Class 2 counterparts, and have better stability of values with changes in temperature, voltage, or frequency. The best Class 1 capacitor, with nearly constant characteristic, will be labeled 'C0G' using EIA designators, but is often referred to as 'NP0' which stands for 'negative-positive-zero' (temperature coefficient).

Plastic-film capacitors use extremely thin sheets of plastic film as the dielectric between capacitor plates, usually in a coil construction. They have low losses and good resistance to humidity. Four common materials are:

1. Polycarbonate
2. Polypropylene
3. Polystyrene
4. Teflon

Generally speaking, as one moves down in this list from polycarbonate to Teflon, the capacitors get larger, better, and more expensive.

Historically, low-budget Tesla coilers have used either saltwater capacitors (as Tesla himself did) or homemade rolled polyethylene and foil capacitors. The saltwater capacitors are lossy and heavy. They are one of the first components to be replaced as a new coiler grows with his hobby. The homemade foil capacitors are limited by corona onset to about 7000 V. They are oil filled, so if the container leaks or is tipped over, one has a real mess. Serious coilers would look for commercial high voltage, high current, pulse rated capacitors made by specialty

companies like Maxwell. These worked fine, but were expensive if purchased new, and difficult to find on the used market.

This all changed in 1999 when the Tesla coil community moved en masse to multi mini capacitors (MMC). These are small commercial capacitors that one buys by the sackful from Digi-Key and connects in series and parallel strings to get the required ratings.

There are several different capacitor series, some of which are more suitable for Tesla coil work than others. One should look for capacitor types that are rated for “high voltage, high frequency, and high pulses”. A dissipation factor of 0.1% at 1 kHz is good, as is polypropylene for the dielectric. The ECWH(V) and ECQP(U) series of capacitors are good. The Panasonic ECQ-E capacitors utilize metallized polyester (Mylar) which has a higher dissipation factor (1% at 1 kHz) than the polypropylene capacitors.

Capacitor ratings have received a great deal of attention by Tesla coilers. A given capacitor might be rated at 1600 VDC and 500 VAC. The AC rating is always for rms values, so the peak voltage would be about 700 V for a rating of 500 VAC. A capacitor in a Tesla coil primary experiences full voltage reversal, so it would seem wise to find the AC rating, multiply by $\sqrt{2}$, and divide that into the peak voltage available from the iron core transformer to find the number of series capacitors to use.

Experience of coilers, however, has been that there are sufficient safety factors built into the capacitors that using the DC rating is quite acceptable. For example, a 15 kV transformer has a peak voltage of $15\sqrt{2} = 21.2$ kV. Divide 21.2 by 1.6 for 1600 VDC capacitors to get 13.25. Use either 13 or 14 capacitors in series.

The voltage rating is not determined by the dielectric breakdown voltage as much as it is by the *Ionization or Corona Inception Level*. This refers to the level at which ionization or partial discharges can begin to occur inside a bubble of entrapped air or within an air-filled void within the solid dielectric system. If one had *perfect* dielectrics and could always exclude any entrapped air, derating for this phenomenon would not be necessary.

Corona inside a capacitor will chemically degrade the dielectric material over a period of time until the capacitor ultimately fails. Manufacturers of capacitors might select a voltage rating whereby their capacitors will last for one million hours before this occurs. In Tesla coil use, 100 hours of actual operation is a long time. This helps explain why coilers can exceed the manufacturer ratings without immediate problems.

There is evidence that operating a capacitor above its ac voltage rating will shorten its life by a factor of the overvoltage ratio raised to the 15th power. Suppose we have a capacitor rated at one million hours at 500 VAC. For economic reasons we are thinking about operating it at 700 VAC. The life reduction factor is $(700/500)^{15} = 155.57$. Dividing one million by 155.57 gives an expected life of 6428 hours or almost one year of continuous operation. This is more than adequate for most Tesla coil applications.

One brand of capacitor is the WIMA. Terry Fritz comments about them: “The switching power supplies we build have WIMAs in pulsed duty similar to that seen in Tesla coil use.

The real indicator of how long they live is how warm they get. If they heat to about 5 degrees C above the ambient, then they last forever. At about 8 degrees above ambient, we see occasional failures. At 10 degrees, they get to be a problem. In many situations, we go beyond the WIMA chart derating in pulsed applications with no problem at all, just being sure they don't get hot."

WIMA experts say that partial discharges occur only above a certain voltage level and only on ac. Frequency does not seem to be a factor. They agree that the ac rating can be exceeded in Tesla coil applications, but without a serious investigation, were mentioning factors like 1.25 to 1.5.

Some capacitors that are candidates for a MMC use actual metal foil for electrodes while others use metallized plastic. The metal foil devices are physically larger for a given rating and are much more robust in Tesla coil service.

Dielectric material, electrode thickness, and lead size all limit the rate at which charge can be added to or removed from a capacitor. This limit is expressed as the dV/dt limit. The peak current I_{peak} can be determined from the dV/dt value by the equation

$$I_{peak} = \frac{dV}{dt} C \quad (12)$$

where appropriate multipliers should be used to get everything in volts, seconds, and farads. By way of reference, a 1000 pF capacitor rated for 10,000 V/ μ s will withstand 10 A surges. A capacitor with a dV/dt rating less than 1000 V per 5 μ s probably does not have metal foil endplates and should not be used.

The peak current that a Tesla coil primary capacitor must provide is determined by the capacitor voltage and the surge impedance. If the secondary were removed, the surge impedance would be

$$Z_s = \sqrt{\frac{L_1}{C_1}} \quad (13)$$

where C_1 is the capacitance in the Tesla coil primary and L_1 is the inductance of the primary. When the secondary is in place, the surge impedance will increase from the above value, so the peak current required from the capacitor will be less. Using the above expression for Z_s should be a worst case calculation. If it tells you that there might be 100 A flowing after the gap shorts, and your individual capacitors are rated at 10 A each, then you would need ten parallel strings. Since this figure includes some factor of safety, you might be able to get by with somewhat fewer strings. One WIMA expert said that the maximum amperage rating is not as critical as the Corona Inception Level mentioned above.

The polypropylene capacitors useful for Tesla coil work have a self-healing mode. A partial discharge will eventually cause a local failure. A burst of energy through this short will vaporize everything around the short, effectively removing the short. The capacitor continues

to operate, but with a slightly lower capacitance due to some of the electrodes being removed. It is a good idea to carefully measure the capacitance of the MMC periodically. A decrease of even 1% could indicate that the capacitors are being stressed, and that total failure is possible.

Terry Fritz reports testing some WIMA MKS 4 (K4) capacitors rated at 1 μF and 400 VDC on a large DC supply. These are metallized polypropylene dielectric capacitors that are encapsulated in epoxy filled 3×1 cm rectangles. Starting at about 900 V there are a few little snaps, indicating that the dielectric had punched through and the arc blasted the thin metal layer on the other side. Going to 1200 V gave a few more snaps. Around 1500 V there were many snaps and the capacitor shorted out.

When a capacitor cleanly blows apart, that is a sign that they were destroyed by over voltage. When a capacitor puffs up and looks melted and burned, that is from too much current. It is always a good idea to think about possible failure modes, and put the capacitors in appropriate enclosures to protect bystanders.

I had a student build a small design project once, that illustrates this point. He used an electrolytic capacitor rated at about 25 VDC in some circuit for 120 VAC usage, and I did not notice it before testing. He was proudly demonstrating performance, using a variac to bring up the voltage. At about 90 V, the electrolytic capacitor exploded with a noise somewhere between a pistol and a small shotgun. The contents of the capacitor, about a cubic centimeter of plastic fibers, hit the student in the middle of the chest. He was not injured, but for a brief moment he thought he had been killed by this exploding capacitor. I suspect he still pays more attention to capacitor voltage ratings than most people.

Many coilers recommend placing resistors across the MMC to bleed off the charge after power is removed. This can eliminate some very unpleasant surprises when making adjustments between runs. One can get true high voltage resistors but they are expensive. Most use a long chain of ordinary resistors therefore. Digi-Key sells 0.5 W carbon film resistors rated at 350 volts each for about two cents each in quantity. At this price, there is little point in not using adequate safety factors.

Suppose we feel comfortable in operating a 0.5 W resistor at 0.1 W and at the rated peak voltage of 350 V each. We solve for the resistance as

$$R = \frac{V^2}{P} = \frac{(350/\sqrt{2})^2}{0.1} = 612 \text{ k}\Omega$$

where we round off to the nearest standard value of 620 k Ω .

The number of resistors needed for a string across say a 15 kV transformer secondary would be

$$N = \frac{15000\sqrt{2}}{350} = 60 \text{ resistors}$$

We then check the time constant to get a measure of how long it takes to discharge a

Table 3: Dielectric Constant ϵ'_r and Dissipation Factor DF of Water.

		Frequency in Hz			
Temperature		10^5	10^6	10^7	10^{10}
1.5°C	ϵ'_r	87.0	87.0	87	38
	DF	0.1897	0.01897	0.00195	1.026
25°C	ϵ'_r	78.2	78.2	78.2	55
	DF	0.3964	0.03964	0.00460	0.545
85°C	ϵ'_r	58	58	58	54
	DF	1.2413	0.12413	0.01259	0.259

capacitor chain. If the capacitance happened to be 27 nF, the time constant would be

$$\tau = RC = (60)(620000)(27 \times 10^{-9}) = 1 \text{ second}$$

A capacitor will be mostly discharged after five time constants or 5 seconds in this example.

There are obviously many other design possibilities. If we assumed the resistors would accept a higher voltage before arcing over, then we would want to use higher ohm values to keep heating within bounds. Terry Fritz tested some Yageo 10 M Ω 0.5 W carbon film resistors to see where they would actually fail. At 4000 V the resistor started to turn brown and smoke a little. This should be expected since it is now drawing 2 W, four times its power rating. At 5300 V, it gave a very satisfying crack and arced along its outside surface. After that, the 2 M Ω resistor measured 11.1 M Ω . So a resistor specified at 350 V actually failed at 5300 V, a safety factor of 15. Terry sees no problem with using these 0.5 W resistors up to 1000 V if the resulting wattage is not excessive.

2 Dielectric Loss in Water

Water did not appear in Table 1 because it is so much different from other substances that it needs special treatment. The relative permittivity ϵ'_r and the dissipation factor DF are shown in Table 3 for different values of temperature and frequency.

The columns for 10^5 and 10^6 Hz (100 kHz and 1 MHz) would bracket the range of operation for most Tesla coils. Only the very large coils operate below 100 kHz and the very small coils operate above 1 MHz. The column for 10^7 Hz was included to show that relative permittivity

does not change between 10^6 and 10^7 Hz and that the dissipation factor continues the same decline as between 10^5 and 10^6 Hz. As frequency increases to 10^{10} Hz (the general range of microwave ovens) the relative permittivity decreases somewhat while the dissipation factor becomes quite large. It is obvious that heating food with radio waves is much more effective at 10^{10} Hz than at 10^7 Hz.

At Tesla coil frequencies, the dissipation factor $(DF)_2$ at some frequency f_2 can be found from the listed value $(DF)_1$ at frequency f_1 by

$$(DF)_2 = (DF)_1 \frac{f_1}{f_2} \quad (14)$$

3 Dielectric Losses of a Tesla Coil

There are dielectric losses primarily in two places in a Tesla coil, the coil form and the insulation around the conductor. There may also be some loss in the air surrounding the coil, in the soil, and in insulating supports near the coil, but the ones of most interest seem to be the coil form and the wire insulation. We saw in Eq. 11 that the loss in a capacitor is equal to the square of the applied voltage times the radian frequency times the capacitance times the dissipation factor. We can make reasonable estimates for voltage, frequency, and DF, but what about the capacitance?

A cross section of the coil is shown in Fig. 3.




Figure 3: Cross Section of Tesla Coil

We show three turns of the several hundred that are typically used. This is a close wound coil where the conductors are immediately adjacent to one another. Space wound coils usually use a spacer of the same size as the conductor, which would result in a cross section showing half as many conductors in a given winding length. The coil diameter is D , the winding length

is ℓ_w and the coil form length is ℓ_{cf} . The wall thickness of the coil form is t . At the top of the coil is a toroid with voltage V_{tor} , while at the bottom the voltage is zero. To make life a little simpler, we assume the toroid can be represented by a flat sheet on top the coil form, and that the zero potential ground plane is directly under the coil.

Most of the dielectric between toroid and ground plane is air, but a small portion is the coil form. We can think of the total capacitance between toroid and ground plane as a coil form capacitance C_{cf} in parallel with an air dielectric capacitor and a winding insulation capacitance C_{wi} . If the parallel plate capacitor formula is valid, we can quickly write an expression for the coil form capacitance.

$$C_{cf} = \frac{\epsilon \pi t D_{cf}}{\ell_{cf}} \quad (15)$$

As stated earlier, the expression is valid when fringing can be ignored. This condition is true when the separation is small compared with the area but here we have a small area and a large separation. So we go back and examine the situation for which fringing is negligible. In any capacitor, there are electric field lines directed from the positive to the negative plate, and equipotential lines (lines or surfaces of constant potential) which intersect each other at right angles. The areas enclosed by these lines are called curvilinear squares. The sides may be curved and of different lengths, but the four corners are all right angles. In a parallel plate capacitor, the equipotential lines are equally spaced straight lines, as are the electric field lines, so the curvilinear squares become true squares. It turns out that inside the coil form of Fig. 3 that electric field lines are mostly straight down and the equipotential lines are nearly equally spaced.

For the equipotential lines to be exactly equally spaced, the voltage must increase linearly from bottom to top. The actual mathematical form of the increase is probably not quite linear, as shown when the turn-to-turn voltage of the top few turns, for example, exceeds the dielectric strength of the insulation and turn-to-turn sparking occurs. But it will not be vastly different from linear, either, so the parallel plate formula should give the proper order of magnitude. I would be surprised if it were in error by more than 20% or 30%.

A circuit model for the equivalent resistances that represent losses in the coil form and in the winding insulation is shown in Fig. 4. The capacitances C_{cf} and C_{wi} are in parallel with other capacitances (not shown) such that the total capacitance is C_{tc} . The power dissipated in the coil form is given by

$$P_{cf} = V_{tor}^2 \omega C_{cf} (\text{DF})_{cf} \quad (16)$$

and the equivalent resistance is given by

$$R_{cf} = \frac{V_{tor}^2}{P_{cf}} = \frac{1}{\omega C_{cf} (\text{DF})_{cf}} \quad (17)$$

Example: What is the power loss and equivalent coil form resistance in a PVC coil form that is 1 m in length, 3 mm wall thickness, and diameter of 200 mm, at a frequency of 200 kHz if the top voltage is 500,000 V? Assume $\epsilon_r = 3.0$ and $DF = 0.015$.

$$C_{cf} = \frac{\epsilon_o \epsilon_r D_{cf} t}{\ell_{cf}} = \frac{(8.854 \times 10^{-12})(3.0)\pi(0.200)(0.003)}{1} = 5.01 \times 10^{-14} \text{ F}$$

$$P_{cf} = V_{tor}^2 \omega C_{cf} (DF) = (500,000)^2 2\pi(200,000)(5.01 \times 10^{-14})(0.015) = 236 \text{ W}$$

$$R_{cf} = \frac{1}{\omega C_{cf} (DF)_{cf}} = \frac{1}{2\pi(200,000)(5.01 \times 10^{-14})(0.015)} = 1.06 \times 10^9 \text{ } \Omega$$

To get the toroid to this voltage probably requires an input power of 10 kW or more, so 236 W is not a large loss, percentage wise. On the other hand, the coil form would soon melt if this power were applied continuously. If maximum efficiency is desired, one can always go to a polyethylene coil form, where DF is less than 0.0002




Figure 4: Equivalent Resistances for Coil Form and Wire Insulation Losses.

We turn now to the consideration of the dielectric loss in the winding insulation. Most Tesla coils are wound with magnet wire, a wire in wide use for winding motors and transformers. It has a dielectric coating that is tough both mechanically and thermally. It will withstand temperatures well above 100°C. Unfortunately, it is not particularly low loss at Tesla coil frequencies. Dissipation factors are similar to PVC or nylon. The coating called Heavy Soderon actually has a thin layer of nylon on the surface, to help withstand mechanical abuse.

I have observed that the losses of tight wound magnet wire coils sometimes increase dramatically with humidity, even at input voltages far below those necessary for breakout. It seems that some magnet wire coatings can absorb moisture, and also some coil forms. I bought two plastic barrels at a recycling place that I thought were polyethylene. Whatever they were, the input resistance of the coil could change by a factor of two or more as humidity changed from 25 to 100%. On the other hand, a coil of some magnet wire on a PVC form would see little change in input resistance, probably less than 10%. So moisture in or on the wire insulation and the coil form certainly has the capability of increasing losses.




Figure 5: Two Adjacent Turns of Magnet Wire

The basic geometry that we will try to solve for power loss is shown in Fig. 5.

This figure shows two adjacent turns of a space wound coil. The conductor has radius r_1 and is surrounded by a dielectric of permittivity ϵ_1 . The turns are then contained in a dielectric of permittivity ϵ_2 , where usually $\epsilon_2 = \epsilon_o$. However, if the winding has been coated with some substance like polyurethane, then ϵ_2 will be different from that of air.

Like many problems in electromagnetics, there is no exact analytic formula for capacitance of the exact structure. We have a choice of an exact solution for an approximate structure, or an approximate solution for an exact structure, or combinations thereof. Fig. 5 shows the parameters for a structure that approximates our coil winding, and has an exact analytic solution. The advantage of an analytic solution is that we can easily see which factors are the most important. Questions about the effect of wire size, insulation thickness, and conductor spacing can probably be answered correctly even if the answers differ from the correct values by 20% or more.

The geometry in Fig. 5 is for a parallel conductor transmission line. The conductor on the right is at potential $V = V_o$ and the one at the left at potential $V = -V_o$. The vertical plane between the conductors is then at potential $V = 0$, by symmetry. The potential is $V = V_x$ (a constant) on the outside of the dielectric layer. If the dielectric is eccentric, then the surfaces lie on the surfaces of constant η of the bycylindrical coordinate system [2, pages 361-366]. There are three factors that make this an approximate geometry: First, the actual dielectric is centered rather than eccentric. Second, the magnet wire will lie against a coil form, which will affect the electric field lines and the capacitance. And third, there are additional conductors on either side of the two shown, which also affect field lines and capacitance. However, I think this geometry is the best that we can do, if we want an analytic solution.

The capacitance between the conductor on the right and the outer surface of its dielectric is

$$C_1 = \frac{2\pi\epsilon_1\ell}{\eta_o - \eta_x} \quad (18)$$

and the capacitance between the outer surface and the $V = 0$ plane is

$$C_2 = \frac{2\pi\epsilon_2\ell}{\eta_x} \quad (19)$$

where ℓ is the circumference of one turn of the winding. To get the power dissipation, we need the voltage $V_o - V_x$ across ϵ_1 and the voltage V_x across ϵ_2 . The solution process is fairly standard for Laplace's Equation. We write general solutions for Laplace's Equation in both dielectrics.

$$V_1 = A_1 + B_1\eta \quad \text{in } \epsilon_1 \quad (20)$$

$$V_2 = A_2 + B_2\eta \quad \text{in } \epsilon_2 \quad (21)$$

The boundary equations are

$$V = V_o \quad \text{at } \eta = \eta_o \quad (22)$$

$$V = 0 \quad \text{at } \eta = 0 \quad (23)$$

$$V_1 = V_2 \quad \text{at } \eta = \eta_x \quad (24)$$

The fourth boundary equation is that the normal electric flux density is continuous on the $\eta = \eta_x$ surface.

$$\epsilon_1 \frac{\partial V_1}{\partial \eta} = \epsilon_2 \frac{\partial V_2}{\partial \eta} \quad (25)$$

After some algebra, we find

$$V_1 = \frac{V_o(\eta_x(\epsilon_1 - \epsilon_2) + \epsilon_2\eta)}{\eta_o\epsilon_2 + \eta_x(\epsilon_1 - \epsilon_2)} \quad (26)$$

$$V_2 = \frac{V_o\epsilon_1\eta}{\eta_o\epsilon_2 + \eta_x(\epsilon_1 - \epsilon_2)} \quad (27)$$

We can find V_x by substituting $\eta = \eta_x$ in either of the last two equations. The power dissipations in the two regions are

$$P_1 = (V_o - V_x)^2 \omega C_1 (DF)_1 = \frac{V_o^2 \epsilon_2^2 (\eta_o - \eta_x) \omega (2\pi \epsilon_1 \ell) (DF)_1}{[\eta_o \epsilon_2 + \eta_x (\epsilon_1 - \epsilon_2)]^2} \quad (28)$$

$$P_2 = V_x^2 \omega C_2 (DF)_2 = \frac{V_o^2 \epsilon_1^2 \eta_x \omega (2\pi \epsilon_2 \ell) (DF)_2}{[\eta_o \epsilon_2 + \eta_x (\epsilon_1 - \epsilon_2)]^2} \quad (29)$$

The voltage V_o is half the voltage between turns. If there are N turns and the voltage at the top of the coil is V_{tor} then

$$V_o = \frac{V_{tor}}{2N} \quad (30)$$

The total power P_{wi} dissipated in the wire insulation and in the gap between the turns is

$$P_{wi} = 2N(P_1 + P_2) \quad (31)$$

The relationships between the variables in Fig. 5 are

$$w_i = \sqrt{a^2 + r_i^2} \quad (32)$$

where $i = 1$ or 2 .

$$\eta_o = \sinh^{-1} \left(\frac{a}{r_1} \right) \quad (33)$$

$$\eta_x = \sinh^{-1} \left(\frac{a}{r_2} \right) \quad (34)$$

$$s = w_2 - w_1 = \sqrt{a^2 + r_2^2} - \sqrt{a^2 + r_1^2} \quad (35)$$

We start with a table lookup to find the radius r_1 of the bare magnet wire (without insulation). Then we use $w_1 = r_1 + \text{build} + \text{space}$ (if any), where build refers to the thickness of the dielectric on the magnet wire, and solve for a and η_o . Then it gets a little tricky. Do we use an r_2 such that the minimum distance between the eccentric cylinders is the build, or the average distance, or something in between? My first impulse is to use the minimum distance, since most of the ‘action’ is between the two conductors, so we need the most accurate description of dielectric in that region. That is, I would write

$$r_2 - s = r_1 + \text{build} \quad (36)$$

This is a transcendental equation in one unknown, r_2 . There is no simple closed form solution, but if one plugs in values for a and r_1 , it is not difficult to get a numeric value for r_2 . Note that if the dielectrics are touching, then r_2 , s , and w_2 all go to ∞ . The model requires that $P_2 = 0$ for this case and P_1 is the dissipation in ϵ_1 if ϵ_1 fills the entire space.

Example: Find the dielectric losses in three different coils built on the same coil form used in the previous example. The winding length is one meter, the coil diameter is 200 mm, the top voltage is 500,000 V, and the toroid size is adjusted as necessary to get a frequency of 200 kHz. We want to compare the cases: (a) a 16 gauge coil tight wound, (b) a 16 gauge coil with a small air gap, 2 mils, between windings, and (c) a 22 gauge coil space wound with the same number of turns as the 16 gauge tight wound coil. These cases all have dry air as the surrounding dielectric. Then we want to examine a fourth case of the same geometry as (b) but with water condensed on the magnet wire surface. We are a long way from having a good model for this situation, but we might get some interesting results if we assume that the 2 mil gap is filled with 10% water and 90% air and use interpolated values for permittivity and dissipation factor. The magnet wire dielectric is Heavy Soderon with thickness 1.65 mils for either wire size. Results are shown in Table 4.

Table 4: Dielectric Loss Example

wire size	16	16	22	16
r_1 , mils	25.4	25.4	12.65	25.4
space, mils	0	1	12.85	1
w_1 , mils	27.05	28.05	27.15	28.05
N , turns	725	700	725	700
L , mH	19.0	17.7	19.0	17.7
a , mils	9.303	11.901	24.023	11.901
ϵ_{r1}	3.39	3.39	3.39	3.39
ϵ_{r2}	1	1	1	7.82
$(DF)_1$	0.03	0.03	0.03	0.03
$(DF)_2$	0	0	0	0.198
r_2	∞	70.32	16.03	70.32
η_o	0.3585	0.4529	1.400	0.4529
η_x	0	0.1684	1.194	0.1684
V_o	345	357	345	357
P_1	2.36	0.352	0.010	2.017
P_2	0	0	0	3.416
$2NP_1$	3423	493	14	2824
$2NP_2$	0	0	0	4782

We see a loss of 3423 W for the first case. This seems high, but the input power to the coil at the moment the top voltage is 500 kV is probably at least 10 kW and perhaps well over 100 kW, so this is not impossible. We also note that small winding irregularities lower this number significantly.

If the average gap between windings is only 2 mils, the loss drops from 3423 to 493 W. It is difficult for amateurs to wind coils with no gap at all, so the effective loss is most likely under 1 kW. If we go to a true space wound coil, as for in the 22 gauge shown in the third column, the losses drop to only 14 watts. If we planned to operate the coil in CW mode, or for very extended periods in disruptive mode, we may want to consider space winding just to keep the coil temperature rise low.

When 10% water is added to the dielectrics for the case of a 2 mil gap, the losses increase dramatically. They jump from 493 W to $2824 + 4782 = 7606$ W. These are not precise numbers, but indicate the trend that moisture increases coil losses substantially.

4 DC Capacitors

In this section we will examine the dc capacitors used to convert time varying voltages into steady dc, called filter capacitors. They usually only experience one polarity of voltage, hence can be of the *polarized* type. Aluminum oxide electrolytic and tantalum capacitors are of this type. The exact value of a capacitor used for filtering is not very critical. One may see tolerances as high as -20% to $+80\%$. That is, a $100\ \mu\text{F}$ capacitor would have an actual value between 80 and $180\ \mu\text{F}$ with this tolerance. Similarly, capacitance variation with temperature is not critical.

The most important goal is a high capacitance in a small volume. We accomplish this by using a material with a high dielectric constant in a structure with a large area and thin dielectrics. One ingenious solution to this goal is the aluminum electrolytic capacitor. This consists of a ribbon of aluminum foil on which a thin film of aluminum oxide has been formed electrochemically, and a water based electrolyte fluid which acts as the opposing plate. A second ribbon of aluminum foil is used to make electrical connection to the electrolyte fluid. The two ribbons of aluminum foil are mechanically separated by a ribbon of porous paper, as shown in Fig. 6. The three ribbons are rolled up and placed in a can filled with the liquid electrolyte.




Figure 6: Cross Section of Aluminum Electrolytic Capacitor

The oxide coated foil is the positive plate, called the anode. The aluminum oxide film is the dielectric, and the fluid electrolyte is the negative plate, called the cathode. The second aluminum ribbon serves only to make good electrical connection to the cathode. It is usually bonded to the aluminum can that houses the capacitor.

The oxide dielectric has a thickness on the order of $0.01\ \mu\text{m}$. The plate separation d has therefore been reduced to a very small value. Also, aluminum oxide has a dielectric constant in the range of 8 to 11, which is relatively high for a dielectric. This combination allows a very high capacitance for a small volume.

The oxide is grown by placing the aluminum foil in an electrolyte solution and applying a dc voltage between the aluminum and the electrolyte. The resulting oxide layer has a thickness proportional to the applied voltage, about $0.0014\ \mu\text{m}$ per volt at room temperature. The voltage at which the oxide is grown must be considerably higher than the proposed operating voltage to provide adequate dielectric strength. The leakage current increases rapidly as the operating voltage approaches the forming voltage value.

The film of aluminum oxide has the property of diode action. That is, current can flow one way through the anode but not the other. The aluminum electrolytic is therefore limited to dc applications. The polarity marked on the can must be carefully observed since a negative voltage more than a volt or two will cause breakdown of the film and destruction of the capacitor. This limitation is overcome in non-polarized electrolytic capacitors intended for ac applications by simply using aluminum oxide layers on *both* aluminum ribbons and facing the ribbons in opposite directions. This produces the same effect as putting two diodes in series, but pointed in opposite directions.

While the capacitance per unit volume is quite high for this capacitor type, there are several disadvantages. One is that the maximum operating voltage is limited to about 450 V. This is obviously not a problem for most computer applications, but can be a limitation in situations requiring high power. Another limitation is that the dissipation factor is higher than for other capacitor types. A third limitation is that there may be significant leakage currents. That is, both R_s and R_p need to be considered in doing loss calculations. A fourth limitation is that the shelf life is somewhat limited. A capacitor with an acceptable leakage current when first made may have an unacceptable leakage current after not being used for a year or two. There is a process called *reforming* where the capacitors are heated to 85°C and a voltage is applied, which can restore acceptable performance when this problem occurs.

Tantalum capacitors usually have solid electrolytes, are smaller, and have better characteristics than aluminum capacitors. The small size of tantalum capacitors is due to the high dielectric constant of the tantalum-oxide dielectric. In addition to being small, tantalum capacitors have relatively low leakage currents, good resistance to vibrations, minimal capacitance and ESR variations with changes in temperature and no major parametric changes over the life of the device.

It is interesting to compare a high quality ac capacitor with a Tantalum. According to Table 2, a $1\ \mu\text{F}$ 100 V polypropylene capacitor has a value of $R_s = 0.015\ \Omega$. A MEPCO

Series 40YW Tantalum capacitor rated at 1 μF and 50 VDC has a typical R_s of 16 Ω , and is still within specification at 80 Ω . This does not mean the Tantalum capacitor is inferior to the metallized polypropylene capacitor, but rather that the Tantalum capacitor should not be used in applications requiring large ripple currents.

Example

Calculate the power dissipation in a capacitor described by $R_s = 0.1 \Omega$, $R_p = 10^6 \Omega$, $L_s = 10 \mu\text{H}$, and $C = 400 \mu\text{F}$ if the applied voltage is $v = 200 + 6 \sin \omega t$, with $\omega = 2\pi(120) = 744 \text{ rad/sec}$. This would describe a 200 VDC supply with a 6 V ripple voltage.

We need to apply superposition to solve for the dc power dissipation and then the ac power dissipation. At dc, C represents an infinite impedance, so we have R_s , L_s , and R_p in series. This series impedance is essentially just R_p . The average power dissipated in R_p is

$$P_{dc} = \frac{V_{dc}^2}{R_p} = \frac{(200)^2}{10^6} = 0.04\text{W}$$

At 120 Hz, the capacitive reactance is a few ohms in parallel with $R_p = 10^6 \Omega$, so the capacitor model essentially becomes R_s , L_s , and C in series. The series impedance of this circuit at 120 Hz is

$$Z_{ac} = R_s + j\omega L_s - \frac{j}{\omega C} = 0.1 + j0.00744 - j3.316\Omega$$

with magnitude $|Z_{ac}| = 3.310$. Obviously, the capacitive reactance would have been an excellent approximation to the total series impedance. The rms current is then

$$I_{ac} = \frac{6}{\sqrt{2}(3.31)} = 1.28\text{A}$$

The average power due to this current flowing through R_s is

$$P_{ac} = I_{ac}^2 R_s = 0.16\text{W}$$

The total power being dissipated in the capacitor is then the sum of P_{dc} and P_{ac} .

$$P_{tot} = P_{dc} + P_{ac} = 0.04 + 0.16 = 0.2\text{W}$$

We see that the ac ripple of only 6 volts causes four times the heating of the 200 volts of dc. Also the capacitor leads must be built to carry a current of 1.28 A. This would not be a problem, but if the capacitance were increased by a factor of ten, wire size would start to become important.

5 Pulse Capacitors

Applications for pulse capacitors include particle accelerators, metal forming, laser drivers, and X-ray generators (and Tesla coils, of course). In these situations, discharges are often highly

oscillatory, and voltage reversals up to 90% can be seen. Flash tubes and electrical impulse welders have relatively longer discharge times and smaller reversals. Pulse-discharge or energy-storage capacitors are usually charged over a relatively long time and discharged in a short time (microseconds or less). Voltages and currents may be very high. The inexpensive aluminum and tantalum capacitors do not meet these requirements, so special techniques are used to fabricate these capacitors. The traditional materials used include castor oil/Kraft paper, plastic film, ceramic, and mica, although the polypropylene capacitors considered earlier in the chapter would also work. This section is included to help coilers determine the suitability of surplus pulse capacitors for Tesla coil use.

An important parameter is the pulse repetition rate. If this is faster than, say, 1 pulse/minute, heating effects become important. Overheating can shorten the lifetime of a pulse capacitor considerably.

The lifetime is typically expressed in number of pulses or *shots*. Capacitors are readily available with a lifetime of greater than 100,000 shots, with voltages up to 100,000 V and peak currents up to 50,000 A, if the peak current of the first negative peak of the decaying sinusoid is no more than 30% of the first positive peak.

There are some rather interesting limits to the speed of discharge. One is that the capacitor plates act like a parallel plate transmission line. A charged transmission can only discharge in a time equal to the two-way transit time, t , along the line where

$$t = \frac{2\ell}{c} \sqrt{\epsilon_r} \quad (37)$$

where ℓ is the greatest distance from the capacitor lead to the end of the plate, and c is the speed of light in vacuum. Light propagates at a speed of 0.3 m/ns in vacuum so it is quite possible to exceed 100 ns discharge time with a rolled parallel plate construction. A capacitor built with 0.1 μF of capacitance and 10 nH of inductance will not ring down in 100 ns because of this effect, even if that is the value predicted from the (low-frequency) circuit theory model. The only way to make a fast capacitor is to manufacture it with short sections such that this transmission line transient is much shorter than the desired lumped element discharge time.

This transmission line discharge is a good way to get a single pulse of carefully controlled voltage and period. A coaxial cable is charged to a desired voltage and then discharged through the load, say with a mercury switch. This is an inexpensive method of getting a pulse of say 500 V and 40 ns.

References

- [1] Maruvada, P. Sarma and N. Hyllén-Cavallius, "Capacitance Calculations for Some Basic High Voltage Electrode Configurations", IEEE Transactions on Power Apparatus and Systems, Vol. PAS-94, No. 5, September/October 1975, pp. 1708-1713.

- [2] Moon, Parry and Domina Eberle Spencer, *Field Theory for Engineers*, D. Van Nostrand Company, Princeton, New Jersey, 1961.
- [3] Plonus, Martin A., *Applied Electromagnetics*, McGraw-Hill, New York, 1978.
- [4] Reference Data for Radio Engineers, Fifth Edition, Howard W. Sams & Co., 1968.
- [5] Sarjeant, W. James and R. E. Dollinger, *High-Power Electronics*, TAB Books, 1989.
- [6] Schoessow, Michael, *TCBA News*, Vol. 6, No. 2, April/May/June 1987, pp. 12-15.
- [7] Smythe, William R., *Static and Dynamic Electricity*, Hemisphere Publishing Corporation, A member of the Taylor & Francis Group, New York, Third Edition, Revised Printing, 1989.
- [8] Terman, Frederick Emmons, *Radio Engineers Handbook*, McGraw-Hill, 1943.

INDUCTORS AND TRANSFORMERS

We now consider the inductor. Like the capacitor, it appears in all sizes. There are significant differences, however. One is that the commercial availability of prefabricated inductors is much less than for capacitors. Engineers dealing with inductors will routinely buy coil forms and wire, and assemble their own device. With capacitors we can get by with only a vague notion of how they were designed and built. With inductors, we must do the design, which forces us to have a deeper understanding of the steps involved.

A second difference is the amount of electromagnetic theory required. With capacitors, we could escape with a mention of electric field and permittivity. Inductors require that we jump right into some challenging concepts of magnetic fields and energy, and even set up some line or volume integrals. We will refer back to a first course in electromagnetic theory as needed, and even pull out some results from more advanced courses.

1 Definitions

Consider a coil of wire as shown in Fig. 1. The resistance of the wire can be modeled as a separate lumped device so we can think of the coil as being perfectly conducting. If i is a finite dc current, the voltage v will be zero, as would be expected across a perfect conductor. If a time varying current is applied, however, there will be a related voltage observed across the coil. From the circuit theory viewpoint, the relation is given by

$$v = L \frac{di}{dt} \quad (1)$$

where L is the *inductance* of the coil in henrys (H).




Figure 1: Simple Inductor

We will also observe a time varying voltage when we pass a magnet by the coil, even if the current is zero. (Faraday was the first to observe this, in 1831). Moving a magnet is philo-

sophically quite different from applying a current, but it turns out that we can mathematically describe both situations with the electromagnetic equivalent of Eqn. 1.

$$v = N \frac{d\phi}{dt} \quad (2)$$

where N is the number of turns on the coil and ϕ is the magnetic flux passing through the coil. The direction of a flux ϕ that is produced by a current i is determined by the right hand rule. That is, if you curl the fingers of your right hand in the direction of the current flow, the thumb will point in the direction of the flux.

Setting Eq. 1 and Eq. 2 equal to each other and integrating to remove the differential operator yields

$$Li = N\phi \quad (3)$$

This equation has circuit quantities on the left and field quantities on the right, allowing us to move back and forth between the two ways of thinking. Solving for the inductance L gives

$$L = \frac{N\phi}{i} = \frac{\lambda}{i} \quad (4)$$

where we have introduced the *flux linkage* λ . This term better describes the case where ϕ is not constant between adjacent turns. The flux linkage can be considered as the equivalent flux which gives all the correct results when passing through a single turn coil.

The flux ϕ and the flux linkage λ are proportional to the current i . The relationship is linear if there are no ferromagnetic materials in the vicinity, which then gives us a constant value for L independent of the actual value of i . Thus for the air-cored coil, L is just a function of the geometry of the coil, much like our expression for capacitance that was calculated from area and separation of plates in the previous chapter. Unfortunately, inductance formulas tend to be much more complicated than the formula for a parallel plate capacitor.

The electric power input to the inductor is

$$p = vi = Ni \frac{d\phi}{dt} \quad (5)$$

There are no losses in our perfectly conducting coil, so whatever power flows in at one time must flow out at another time. In the sinusoidal case, power flows in for half a cycle and back out the next half cycle. The power flow results in stored magnetic field energy in the coil. The differential energy input during the differential time dt is

$$dW = p dt = Ni d\phi \quad (6)$$

where W is the stored energy in the field.

We want to relate this stored energy to the field quantities B and H , where

$$B = \mu_r \mu_o H \quad \text{T} \quad (7)$$

B is the magnetic flux density in tesla (T) or webers/m² (Wb/m²), H is the magnetic intensity in A/m, μ_r is the relative permeability (= 1 for vacuum), and μ_o is the permeability of free space in henrys per meter.

$$\mu_o \equiv 4\pi \times 10^{-7} \quad \text{H/m} \quad (8)$$

The relative permeability is very close to unity for all materials except for the ferromagnetic materials iron, cobalt, nickel, and a number of special alloys. For these materials, μ_r may range from 10 to 10⁵. The relative permeability is also a function of magnetic intensity in ferromagnetic materials, making what would be a linear problem into a nonlinear one.

The magnetic flux density B may also be expressed in gauss, where 10⁴ gauss = 1 tesla. The earth's magnetic flux density varies from 0.2 to 0.6 gauss, depending on location. The 60 Hz magnetic flux density in a home or office is usually less than a few milligauss except near a source (electric heater, computer monitor, electric razor, blow dryer, etc.) where it may be a few tens of milligauss or even more. A modern well-designed 60 Hz power transformer will probably have a magnetic flux density between 1 and 2 T inside the core. It requires considerable effort and special designs to get much above 2 T. The necessary current density causes heating in the conductors, unless, of course, the conductors are cooled into the superconductor region. Fluxes as high as 8 to 16 T have been used in accelerators and energy storage systems.

Other conversion factors which might be needed are:

- 1 Oersted = 250/ π = 79.6 ampere-turns/meter
- 1 Tesla = 10,000 gauss

The flux ϕ passing through an area A is the integral of the magnetic flux density B over that area.

$$\phi = \int B dA \quad \text{Wb} \quad (9)$$

which becomes simply $\phi = BA$ if B is constant over the area.

We also need Ampere's circuital law

$$Ni = \oint \mathbf{H} \cdot d\boldsymbol{\ell} \quad (10)$$

This states that the current enclosed by any arbitrary path is given by the integral of the dot product of the vector $\mathbf{H}$ and a differential length $d\ell$ along that path. If we extend our coil around into a toroid shape, H (the magnitude of $\mathbf{H}$) will be essentially constant inside the toroid and Ampere's circuital law becomes

$$Ni = H\ell \quad (11)$$

where ℓ is the length of a circle in the toroid. The energy stored in the magnetic field is now

$$dW = Ni d\phi = \frac{H\ell}{i} i A dB = (\text{Vol}) H dB \quad (12)$$

where $\text{Vol} = A\ell$ is the volume where the magnetic energy is stored. The total energy can be found by integration.

$$W = (\text{Vol}) \int_0^B H dB \quad \text{Joules} \quad (13)$$

If the permeability is constant (the magnetic circuit is linear) the integral can be quickly evaluated.

$$W = \frac{(\text{Vol}) B^2}{\mu_r \mu_o} \frac{1}{2} = (\text{Vol}) \mu_r \mu_o \frac{H^2}{2} \quad \text{J} \quad (14)$$

This equation has much important information in it. Suppose that we have a magnetic circuit that is entirely ferromagnetic. H is determined by the current and is independent of the permeability. $B = \mu H$ is large and the total energy stored is large. Suppose now that we cut a small air gap across the magnetic circuit. The flux ϕ drops substantially because of the increased reluctance of the magnetic circuit. B will have about the same (smaller) value in both the iron and the air gap, so $H = B/\mu$ will be much larger in the air gap because of the lower permeability. The total integral of $\mathbf{H} \cdot d\ell$ stays the same but a large fraction of the integral comes from the air gap portion. So the total energy stored decreases as the air gap length increases, and the fraction of the total energy stored in the air gap increases dramatically.

Although not as instructive, the total energy can also be expressed in circuit quantities as

$$W = \frac{1}{2} Li^2 \quad \text{J} \quad (15)$$

We can also use the result from Ampere's circuital law to determine the flux ϕ . In the simple case of uniform flux density B and no air gap, this becomes

$$\phi = BA = \mu_r \mu_o HA = \frac{\mu_r \mu_o NiA}{\ell} \quad (16)$$

2 Ferromagnetic Losses

Two things happen when a time varying magnetic field exists inside a ferromagnetic material. Magnetic domains rotate in the material to align with the magnetic field and the Faraday voltage induced inside the material produces what are called eddy currents. The rotation of domains each cycle produces a frictional type loss called the hysteresis loss P_h . Experimentally, we find that

$$P_h = K_h f B_{max}^z \quad (17)$$

where K_h is an empirical property of the material, f is the frequency, B_{max} is the maximum flux density, and z is an empirically determined value, usually between 1.6 and 2.0 for power frequency transformer steels.

The hysteresis loop of the material is obtained by plotting the magnetic flux density B against the magnetic intensity H as shown in Fig. 2. The area inside the hysteresis loop is the energy dissipated each cycle. Some ferromagnetic materials have very thin hysteresis loops, resulting in low losses, while others have relatively fat loops and correspondingly high losses.




Figure 2: Hysteresis Curve

As the driving magnetic field increases, the resulting flux density increases at a slower rate as more domains become aligned with the magnetic field. This phenomenon is called *saturation*. If linearity is desired, then the transformer should be operated at low flux levels where the hysteresis loop is nearly linear. For power transfer, however, it is most cost effective to operate the device into its saturated region. The exact amount is a matter of engineering judgment.

We might define at least two definitions for permeability, from which we can get some

guidelines for saturation. These are the dc and ac permeabilities

$$\mu_{dc} = \frac{B}{H} \quad (18)$$

$$\mu_{ac} = \frac{\Delta B}{\Delta H} \quad (19)$$

The two permeabilities are identical for very low drive levels that are symmetric about zero. If a dc bias exists, the ac permeability will always be smaller than μ_{dc} . The ac permeability is the main parameter of interest to filter choke designers. At least some engineers consider a material to be saturated when the dc permeability has dropped to half its initial value, or when the ac permeability has dropped to one-eighth of its initial value [12, page 26,28].

A magnetic circuit with eddy currents is shown in Fig. 3. The current is inversely proportional to the resistance seen by the induced Faraday voltage. In a large piece of steel the resistance can be very low, even though the resistivity of steel is not very low compared with a good conductor like copper. For this reason, magnetic circuits at power frequencies are usually made of thin sheets of steel, called laminations, which are on the order of 0.5 mm thick.

The equation for eddy current loss has the form

$$P_e = K_e f^2 B_{max}^2 \quad (20)$$

The experimentally determined constant K_e depends on the resistivity and the dimensions of the material. A detailed analysis shows that K_e is proportional to the square of the lamination thickness, so it is important to keep the laminations as thin as possible. Eddy current losses can be kept acceptably low at 60 Hz with little difficulty, but become excessive at a few kHz, even with very thin laminations. Therefore, inductors or transformers built for operation above 1 kHz are rarely made of laminated material. Instead, they are made of even smaller pieces of ferromagnetic material, typically powdered iron or ferrites. Powdered iron suffers from low permeability and low resistivity compared with ferrites, so we shall concentrate on the latter.

3 Ferrites

Ferrites were developed during and after World War II. The chemical formula for ferrites is ZFe_2O_4 , where Z stands for any of the divalent ions: zinc, copper, nickel, iron, cobalt, manganese, or magnesium, or a mixture of these ions. The bulk resistivities are in the range of 10^2 to 10^9 ohm-cm, compared with 10^{-5} ohm-cm for powdered iron. This very high resistivity reduces the eddy current losses so that ferrites can be used for frequencies up to 20 MHz or even more.




Figure 3: Eddy Currents in Magnetic Circuit

Ferrites are basically a form of ceramic, made by mixing fine powders of appropriate oxides, compressing the mixture, and firing it in carefully controlled atmospheres at temperatures of about 1100°C to 1200°C . The most common ferrites are mixtures of two ferrite powders, either manganese-zinc or nickel-zinc. Magnetic properties can be varied over a significant range by changing the ratio of the two divalent ions and by changing the processing conditions. Each material is given a unique code number. Ferroxcube, for example, assigns a “3” as the first digit of its MnZn materials (3C85, 3B7, 3D3, etc.) and a “4” as the first digit of NiZn materials (4C4, 4A, 4A6, etc.)

3.1 Ferrite Temperature Limits

The Curie temperature (the temperature at which the ferrite becomes nonmagnetic) can be relatively low. The Ferroxcube 3E5 ferrite may have a Curie temperature as low as 120°C . Other Ferroxcube materials have Curie temperatures up to 300°C . If there is any possibility of the ferrite device operating at a temperature above 120°C due to internal losses, a ferrite material with an adequate Curie temperature must be selected.

The copper wire used in winding inductors or transformers remains mechanically stable at temperatures far above the ferrite Curie temperature, so the wire itself is not of concern. However, the insulation on the wire may fail at relatively low temperatures. Polyvinyl chloride (PVC) insulated wire typically has a maximum temperature rating between 80°C and 105°C , for example. Magnet wire has a somewhat higher rated temperature, such as the Belden polythermaleze coating rated at 180°C . Belden also makes Teflon coated wires rated at 200°C and 260°C .

Heat is dissipated from the surface of the inductor or transformer by a combination of radiation and convection. Heat radiated depends on the device surroundings, while convection depends on air flow over the device, so it is very difficult to accurately predict temperature rise in most installations. In any case, both radiation and convection will be directly proportional to the total exposed surface area of the core and windings. We can therefore describe the independent variable as total watts dissipated in core and copper per unit area of the device. The dependent variable, temperature rise, is directly proportional to the dissipation in W/cm^2 . One manufacturer (Magnetics) has calculated a temperature rise of 10°C for a surface dissipation of $0.01 \text{ W}/\text{cm}^2$ and a rise of 100°C for a surface dissipation of $0.1 \text{ W}/\text{cm}^2$, given some reasonable assumptions. Each increment of $0.01 \text{ W}/\text{cm}^2$ results in a temperature rise of 10°C .

Example.

An inductor has a total heat dissipation of $0.06 \text{ W}/\text{cm}^2$ and is in an ambient temperature of 50°C . What is a reasonable estimate of inductor temperature?

Based on the Magnetics guideline, $0.06 \text{ W}/\text{cm}^2$ should yield a 60°C temperature rise above the ambient, so the inductor temperature will be $60 + 50 = 110^\circ\text{C}$. PVC insulated wire should not be used in this situation.

4 Mutual Inductance

Consider two inductively coupled coils as shown in Fig. 4. The current i_1 produces a flux ϕ_{11} that links with i_1 . Part of ϕ_{11} is lost as *leakage flux* $\phi_{1\ell}$ and part of it, flux ϕ_{21} links both currents i_1 and i_2 . Current i_2 likewise produces a flux ϕ_{22} with part of it, flux ϕ_{12} , that links both currents. The relationship among these fluxes is

$$\phi_{11} = \phi_{21} + \phi_{1\ell} \quad (21)$$

$$\phi_{22} = \phi_{12} + \phi_{2\ell} \quad (22)$$

The self-inductance of circuit 1 is

$$L_{11} = \frac{N_1 \phi_{11}}{i_1} \quad (23)$$

and similarly for L_{22} . The mutual inductance of circuit 1 with respect to circuit 2 is based on the flux in circuit 1 that is produced by the current in circuit 2.

$$L_{12} = \frac{N_1 \phi_{12}}{i_2} \quad (24)$$

and similarly for L_{21} .




Figure 4: Mutual Inductance

It can be shown that $L_{12} = L_{21}$ in a homogeneous medium of constant permeability. To emphasize this fact we define a new symbol M for the mutual inductance.

$$M = L_{12} = L_{21} \quad (25)$$

The maximum value for M is $\sqrt{L_{11}L_{22}}$. We define the *coefficient of coupling* k as

$$k = \frac{M}{\sqrt{L_{11}L_{22}}} \quad (26)$$

The coefficient of coupling can reach values as high as 0.998 in iron-core transformers. It is difficult to make k much above 0.5 in air-core transformers.

The voltage v_2 produced by the primary current i_1 is given by

$$v_2 = M \frac{di_1}{dt} \quad (27)$$

If the two inductors of Fig. 4 are connected in series, the total inductance is

$$L = L_{11} + L_{22} \pm 2M \quad (28)$$

where the plus or minus is determined by whether the mutual flux tends to reinforce or cancel the fluxes of the individual coils. This is a convenient method to measure the mutual inductance. Just measure the series inductance twice, once with one coil reversed, subtract one result from the other, and solve the resulting expression for M .

5 Inductance Formulas

Let us now examine the inductance formulas for some simple geometries. First we will look at the inductance of a nonmagnetic coaxial transmission line. The radius of the inner conductor is a , and the inside radius of the outer conductor is b . From Ampere's circuital law, it is easy to show that, for $a < r < b$,

$$H = \frac{I}{2\pi r} \quad \text{A/m} \quad (29)$$

and therefore

$$B = \mu_o H = \frac{\mu_o I}{2\pi r} \quad \text{T} \quad (30)$$

We cannot use $\phi = BA$ since B varies from inner to outer conductor. Instead, we integrate to find the flux crossing any radial plane extending from $r = a$ to $r = b$ and from, say, $z = 0$ to $z = \ell$.

$$\phi = \int B dS = \int_0^\ell \int_a^b \frac{\mu_o I}{2\pi r} dr dz = \frac{\mu_o I \ell}{2\pi} \ln \frac{b}{a} \quad \text{Wb} \quad (31)$$

The flux links the current once, so $N = 1$. From Eqn. 4 the inductance for this length ℓ is

$$L = \frac{\phi}{I} = \frac{\mu_o \ell}{2\pi} \ln \frac{b}{a} \quad \text{H} \quad (32)$$

Suppose now that we try to use this expression to find the inductance of a segment of isolated straight conductor. As the radius of the outer conductor $b \rightarrow \infty$, the corresponding inductance also becomes infinite. What this result tells us is that we never actually have a isolated straight conductor carrying a current without some return path. We must always consider the return path for current if we expect to get inductance values that have any relationship to reality.

In the situation of a toroidal coil of N turns and a current I , as shown in Fig. 5, the magnetic flux density is

$$B = \frac{\mu_r \mu_o N I}{2\pi r} \quad \text{T} \quad (33)$$

For a toroid cross section that is rectangular, as shown in Fig. 6, the integration is straightforward.

$$\phi = \int_{r=a}^b \int_{z=0}^T \frac{\mu_r \mu_o N I}{2\pi r} dr dz = \frac{\mu_r \mu_o N I T}{2\pi} \ln \frac{b}{a} \quad (34)$$

We then multiply the flux by N to get the total flux linkages, and divide by I to get the inductance. For the rectangular cross section case, this is

$$L = \frac{\mu_r \mu_o N^2 T}{2\pi} \ln \frac{b}{a} \quad \text{H} \quad (35)$$




Figure 5: Toroidal Coil




Figure 6: Toroidal Coil Cross Section

5.1 Tesla Coil Inductance (Wheeler)

Many times we just want to make a quick estimate of the inductance of some simple structure without making extensive calculations. Many approximate formulas have been developed in the days before hand calculators and computers, of which one will be given here. Watch out for the fact that the above formulas are given in the standard SI units, but the following formula is in the English system.

The low-frequency inductance of a single-layer solenoid is approximately [17, p. 55].

$$L_w = \frac{r^2 N^2}{9r + 10\ell} \mu\text{H} \quad (36)$$

where r is the radius of the coil and ℓ is its length in inches. This formula is accurate to within one percent for $\ell > 0.8r$, that is, if the coil is not too short. It is known in the Tesla

coil community as the Wheeler formula. The structure of a single-layer solenoid is almost universally used for Tesla coils, so this formula is very important. In normal conditions (no other coils and no significant amounts of ferromagnetic materials nearby) it is quite adequate for calculating resonant frequency.

Example

What is the approximate inductance of an air-cored solenoid with $r = 8$ inches, $\ell = 30$ inches, and $N = 175$ turns?

$$L = \frac{(8)^2(175)^2}{9(8) + 10(30)} = \frac{1960000}{372} = 5270 \text{ } \mu\text{H}$$

If one needs the inductance for other geometries, Terman [17] has a number of expressions. A somewhat more recent paper by Fawzi and Burke [3] gives formulas for calculating the self and mutual inductances of circular coils in a form suitable for computer calculation.

6 High Frequency Transformers

Conventional low frequency transformers consist of coils of wire wound around steel laminations. As mentioned earlier, the losses, especially the eddy current losses, become excessive at frequencies above a few kHz with this technology. Transformers built for operation at frequencies above a kHz or so are built around ferrite cores or air cores. Ferrite cores yield very compact, efficient transformers. Saturation limits operation to moderate power levels, however. If extremely large currents or powers are involved, then the air core transformer may be the logical choice. We shall discuss both types.

The symbolic construction of a transformer is shown in Fig. 7. A voltage v_1 produces a current i_1 which in turn produces a flux ϕ_1 . Part of ϕ_1 links the secondary winding and produces a voltage v_2 by Faraday's Law. If some load is connected, a current i_2 will flow. This current produces a flux which opposes the original ϕ_1 according to Lenz's Law. This reduces the voltage induced in the primary so that more primary current will flow for a given source voltage. A complete description of transformer action in terms of field quantities has been developed only rather recently [4, 6, 9, 10, 16].

7 The Ideal Transformer

It will be convenient to describe the actual transformer in terms of an *ideal* transformer. This is a transformer with no copper losses, no hysteresis or eddy current losses, and perfect magnetic coupling between primary and secondary. For such a device, the relationships between input and output voltages and currents are




Figure 7: A Two Winding Transformer

$$\frac{v_1}{v_2} = \frac{N_1}{N_2} \quad (37)$$

$$\frac{i_1}{i_2} = \frac{N_2}{N_1} \quad (38)$$

The relationship between the input and output apparent power is

$$v_1 i_1 = v_2 \frac{N_1}{N_2} i_2 \frac{N_2}{N_1} = v_2 i_2 \quad (39)$$

The input impedance is

$$Z_1 = \frac{v_1}{i_1} = \frac{v_2 N_1 / N_2}{i_2 N_2 / N_1} = Z_2 \frac{N_1^2}{N_2^2} \quad (40)$$

The ideal transformer thus changes the level of voltage, current, and impedance between primary and secondary.

8 The Actual Transformer

A complete circuit model of the actual transformer is shown in Fig. 8.

In Fig. 8, R_1 and R_2 are the resistances of the primary and secondary windings, L_1 and L_2 are the leakage inductances, R_m is an equivalent resistance representing the hysteresis and eddy current losses, and L_m is the magnetizing inductance. The circuit is usually simplified by eliminating the ideal transformer and replacing all the impedances, voltages, and currents




Figure 8: Transformer Model

on the secondary side with their equivalent values as seen by the primary. If we define the turns ratio a as

$$a = \frac{N_1}{N_2} \quad (41)$$

the simplified circuit is as shown in Fig. 9.




Figure 9: Actual Transformer Referred to Primary

We have also shifted to the phasor notation in Fig. 9, replacing inductances by their equivalent reactances ($X_1 = \omega L_1$, etc.), and the instantaneous voltages and currents by the phasor quantities. The notation can be shortened even more by defining an equivalent resistance and reactance

$$R_{eq} = R_1 + a^2 R_2 \quad (42)$$

$$X_{eq} = X_1 + a^2 X_2 \quad (43)$$

The load current is given by

$$\frac{I_2}{a} = \frac{V_1}{R_{eq} + jX_{eq} + a^2 Z_L} \quad (44)$$

The copper losses are given by

$$P_{copper} = (I_2/a)^2 R_{eq} \quad (45)$$

and the core losses are given by

$$P_{core} = \frac{V_1^2}{R_m} \quad (46)$$

Other conducting materials can be used, such as aluminum, but it is tradition to refer to these series losses as copper losses regardless of the conducting material. It is good practice to keep the copper losses and core losses within a factor of two of each other, at least on large power transformers. The two losses tend to work against each other in a design. The core losses are reduced by reducing the maximum magnetic flux density in the core, which requires either a larger core cross-sectional area or more turns on each winding. Either approach requires more wire, which increases the copper loss. There may be instances where the core losses in a ferrite core used at high frequencies are much higher than the copper losses in reasonably sized wire. In such cases, one should go ahead with proper sizes rather than try to reduce the wire size and increase the copper losses to attempt to maintain some arbitrary parity.

The air core transformer has no core losses, of course. R_m can be removed from Fig. 9 in such cases.

9 Transformer Design

Now we are ready to design a simple transformer. We want to select wire sizes, core size, core material, and number of turns on primary and secondary so that the transformer will meet the requirements without overheating, but without being so large that it is more expensive than necessary.

The most important rating is the required voltage of operation. This determines core size and material, and the number of turns. The wire size is then selected to handle the transformer current rating without excessive copper losses. As we have seen before, the voltage is related to the flux by Faraday's Law.

$$v_2 = N_2 \frac{d\phi}{dt} \quad (47)$$

To a good approximation, the input and output voltages are sinusoidal. That is,

$$v_2 = \sqrt{2}V_2 \cos \omega t \quad (48)$$

where V_2 is the rms output voltage.

The flux is then found by integration.

$$\phi = \frac{1}{N_2} \int v_2 dt = \frac{\sqrt{2}V_2}{\omega N_2} \sin \omega t \quad (49)$$

The flux density B is given by $B = \phi/A$, where A is the cross-sectional area of the core. The maximum flux density, B_{max} , is found when $\sin \omega t = 1$.

$$B_{max} = \frac{\sqrt{2}V_2}{\omega N_2 A} \quad (50)$$

B_{max} is determined either by published data for a particular type of magnetic material, or by measurement. It is typically in the range of 1 T for low frequency laminated transformer steel, and in the range of 0.1 to 0.3 T for ferrite cores. The most efficient use of the magnetic material occurs when B_{max} is slightly above the knee of the magnetization curve. As the material saturates, the permeability $\mu = B/H$ starts to decrease. The inductance is directly proportional to permeability, so the inductance starts to decrease also. But the inductance is defined as $L = N\phi/i$. The flux ϕ is proportional to the sinusoidal voltage so it does not saturate. Therefore, as the permeability and the inductance decrease, the current i must increase. The peak of the magnetizing current increases rapidly above the knee of the magnetization curve, and can exceed the peak of the rated current if the transformer voltage is increased too far. The magnetizing current becomes very nonsinusoidal, with a high harmonic content, at higher voltages. As a rough guideline, the peak of the magnetizing current should not exceed perhaps 10 % of the peak of the rated current. That is, if the rated current were 5 A rms, with a peak of $\sqrt{2}(5) = 7.07$ A, then the peak of the magnetizing current should not exceed about 0.7 A.

Once we know B_{max} , the minimum number of turns can be determined from the above equation.

$$N_{2,min} = \frac{\sqrt{2}V_2}{\omega B_{max} A} \quad (51)$$

Example.

A Ferroxcube 204XT250-3F3 ferrite core is to be used for a transformer at 50 kHz. The rms input and output voltages are to be 10 V. Determine the proper number of turns on each winding.

We first examine a published hysteresis curve for this material, which indicates that saturation is acceptable up to about 0.3 T, at least for low frequency operation. We then check a chart of core loss versus flux density which shows a recommended operating range of 100 to 300 mW/cm² for this material. If the frequency is above about 25 kHz, then B_{max} must be reduced to maintain the total heating in this range. At 50 kHz, $B_{max} = 0.2$ T causes a heating of slightly under 200 mW/cm², which is deemed acceptable. From the published mechanical data, the area A is 0.148 cm². The minimum number of turns is then

$$N_{2,min} = \frac{\sqrt{2}(10)}{2\pi(50 \times 10^3)(0.2)(0.148 \times 10^{-4})} = 15.2 \text{ turns}$$

We would normally round up to the next higher integer, 16 turns.

Suppose we were interested in using the same ferrite core for the same 10 V transformer at 60 Hz. Assuming $B_{max} = 0.3$ T, the minimum number of turns is

$$N_{2,min} = \frac{\sqrt{2}(10)}{2\pi(60)(0.3)(0.148 \times 10^{-4})} = 8450 \text{ turns}$$

Any attempt to wind this many turns through a toroid opening of only 0.312 inches inside diameter would be frustrating at best. This points out the fact that low frequency transformers must be relatively large.

The power rating of a transformer can be determined from Ampere's circuital law and Faraday's law, which state, for the sinusoidal case,

$$H\ell = Ni \tag{52}$$

$$v = N \frac{d\phi}{dt} = N\phi_{max}\omega \cos \omega t \tag{53}$$

Converting these equations to rms values yields the apparent power

$$S = VI = \omega NBA \left(\frac{H\ell}{N} \right) = \omega BHA\ell = \omega BH(Vol) \tag{54}$$

where B and H are rms values and Vol is the volume of the magnetic material. We see that the apparent power is directly proportional to the frequency and to the volume of the transformer. This helps explain why aircraft use 400 Hz rather than 60 Hz. The transformer volume and mass are reduced by the same ratio, thus increasing the aircraft payload.

Example.

The Ferroxcube core of the previous example has a volume of 0.462 cm³ and a relative permeability of 1800. What is the power rating at 50 kHz and a $B_{max} = 0.2$ T?

$$VI = \omega \frac{B_{max}}{\sqrt{2}} \left(\frac{B_{max}}{\sqrt{2}\mu_r\mu_o} \right) (Vol) = \frac{2\pi(50,000)(0.2)^2(0.462 \times 10^{-6})}{(2)(1800)(4\pi \times 10^{-7})} = 1.283 \text{ W}$$

References

- [1] Dexter Magnetic Materials Division Catalog, printed about 1992. 6730 Jones Mill Court, Norcross, GA 30092, (404) 448-4998. Includes catalogs from Fair-Rite, Ferronics, Magnetics, and Siemens.
- [2] Fair-Rite Soft Ferrites Catalog, 12th Edition, January, 1993, P.O. Box J, One Commercial Row, Wallkill, NY, 12589, (914) 895-2055.
- [3] 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.
- [4] Herrmann, F. and G. Bruno Schmid, "The Poynting Vector Field and the Energy Flow Within a Transformer", *American Journal of Physics*, Vol. 54, No. 6, June, 1986, pp. 528-531.
- [5] Hoffmann, C. R. J., "A Tesla Transformer High-Voltage Generator", *The Review of Scientific Instruments*, Vol. 46, No. 1, January 1975, pp. 1-4.
- [6] Lorrain, Paul, "The Poynting Vector in a Transformer", *American Journal of Physics*, Vol 52, No. 11, November, 1984, pp. 987-988.
- [7] Magnetics Catalog of Ferrite Cores for Power and Filter Applications, printed about 1989. 900 E. Butler Road, P.O. Box 391, Butler, PA 16003, (412) 282-8282.
- [8] Magnetics Catalog of Powder Cores—MPP and High Flux Cores for Filter and Inductor Applications, printed about 1995. Same address and phone as above.
- [9] Morton, N., "The Poynting Vector Distribution in a Simple Transformer", *American Journal of Physics*, Vol. 55, No. 5, May, 1987, pp. 472-474.
- [10] Newcomb, William A., "Where is the Poynting Vector in an Ideal Transformer?", *American Journal of Physics*, Vol. 52, No. 8, August, 1984, pp. 723-724.
- [11] Pauly, Donald E., "Power Supply Magnetics—Part I: Selecting Transformer/Inductor Core Material", *PCIM*, January, 1996, pp. 23, 24, 26, 28, 30, 32, 34, 38.
- [12] Pauly, Donald E., "Power Supply Magnetics—Part II: Selecting Core Material", *PCIM*, February, 1996, pp. 36, 38, 40, 42-49.

- [13] Pauly, Donald E., "Power Supply Magnetics—Part III: Selecting High Frequency Core Material", *PCIM* March, 1996, pp. 20, 22, 24, 26, 28, 30-33, 37, 38, 40, 41.
- [14] Philips Ferrite Materials and Components Catalog, Elna Ferrite Laboratories, 234 Tinker Street, Woodstock, NY 12498, (914) 679-2497.
- [15] Reference Data for Radio Engineers, Fifth Edition, Howard W. Sams & Co., 1968.
- [16] Siegman, A. E., "Letter to the Editor", *American Journal of Physics*, Vol. 51, No. 6, June, 1983, p. 492.
- [17] Terman, Frederick Emmons, *Radio Engineers Handbook*, McGraw-Hill, 1943.

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

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

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,

$$\mathbf{V}_1 = \mathbf{V}_L + \mathbf{V}_C + \mathbf{V}_R \quad (2)$$




Figure 3: Series Resonant LC Circuit

Using Ohm's Law in phasor form gives

$$\mathbf{V}_1 = \mathbf{I}_1 \mathbf{Z} = \mathbf{I}_1 (j\omega L + \frac{1}{j\omega C} + R) \quad (3)$$

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

$$\omega_0 L = \frac{1}{\omega_0 C} \quad (4)$$

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

$$\omega_0 = \frac{1}{\sqrt{LC}} \quad (5)$$

The resonant frequency in Hertz is

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \quad (6)$$

The series impedance, expressed as a phasor, is

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

where

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

and

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

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° . 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

$$\beta = \omega_2 - \omega_1 = \frac{R}{L} \quad (10)$$

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

$$Q = \frac{\omega_0}{\beta} = \frac{\omega_0 L}{R} = \frac{1}{\omega_0 C R} = \frac{1}{R} \sqrt{\frac{L}{C}} \quad (11)$$

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 $25\ \Omega$. 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

$$L \frac{di}{dt} + v_c = V_d \quad (12)$$

The component equation for the capacitor is

$$C \frac{dv_c}{dt} = i \quad (13)$$

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

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

and

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

where the angular resonance frequency ω_0 is

$$\omega_0 = 2\pi f_0 = \frac{1}{\sqrt{LC}} \quad (16)$$

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

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

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° , 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° . On one occasion I calibrated the various phase shifts in my instruments (described in the following) and arrived at a phase

shift of 94° 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° . 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° 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° , 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° to zero.

But how does one actually measure the phase shift? A well done measurement can be more convincing than 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.

RESISTANCE OF COIL

The resistance R_{TC} in the RLC model is an effective or equivalent resistance which represents all the losses in the Tesla coil. It includes

1. Ohmic or copper losses
2. Dielectric losses, coil form and conductor insulation
3. Eddy current losses in toroid, strike ring, and soil
4. Radiation losses
5. Losses in the spark

It is surprising how difficult it is to calculate these various losses. Making meaningful measurements can also be challenging. If we operate at low input voltage so we are below spark breakout, then we can ignore the last term for the moment. I will ramble through some considerations for the other losses.

1 Temperature Effects

Almost all coils are wound with copper wire. It is moderately priced and widely available. There might be an occasional aluminum coil, usually from some ‘bargain’ at a surplus auction. Aluminum has higher resistivity than copper, so to get a given resistance the wire must be physically larger. We saw earlier that to get a high toroid voltage we needed a coil with large L and/or small R . If we use larger wire to keep the resistance the same, the coil must be physically larger and the inductance will decrease. We would expect therefore that aluminum coils would always be inferior to copper coils.

Example

You are given a choice between two spools of magnet wire, each 1000 feet in length. The copper is 24 gauge, with nominal resistance 25.67Ω , while the aluminum is 22 gauge, with nominal resistance 26.46Ω . You have a piece of 5 inch diameter PVC pipe and are considering using the entire 1000 ft to wind a coil. You wish to compare the inductances of the two prospective coils.

You assume a build (thickness of dielectric) of 1.65 mils. The diameter of the 22 gauge wire is then $23.35 + 2(1.65) = 28.65$ mils and the diameter of the 24 gauge wire is $20.1 + 2(1.65) = 23.4$ mils. The nominal number of turns would be

$$N = \frac{1000}{\pi(5/12)} = 764$$

The winding length for the 22 gauge wire would be

$$\ell_w = 764(28.65) = 21,890 \text{ mils} = 21.89 \text{ inches}$$

and 17.19 inches for the 24 gauge wire.

By Wheeler's formula, the inductance for the 22 gauge coil would be

$$L = \frac{r^2 N^2}{9r + 10\ell_w} = \frac{(2.5)^2 (764)^2}{9(2.5) + 10(21.89)} = 15,110 \text{ } \mu\text{H}$$

The 24 gauge coil has an inductance of 18,120 μH . For this particular example where r and N are held fixed, ℓ_w is proportional to wire diameter and since it appears in the denominator, a smaller wire diameter results in a larger inductance.

If we hold the winding length fixed, say at 17.88 inches, then the 1000 ft of aluminum wire will fit on a coil of 624 turns and 6.121 inches diameter. The inductance is now 17,670 μH , larger but still less than that of the copper coil. The aluminum coil has more resistance and less inductance than the copper coil, both undesirable. As a side issue, if we hold the winding length fixed, the coil gets to be relatively short and fat. Standard wisdom is that the ratio of length to diameter needs to be on the order of four or five, so this is another negative factor. You should buy the copper wire.

Aluminum forms a coating of aluminum oxide when exposed to air. This oxide is not a conductor (copper oxide is) so there is always a problem in making electrically solid connections. This is not an unsolvable problem since electric utilities use aluminum wire almost exclusively, for cost and weight reasons. The trained individual with the right tools might be able to get acceptable performance from an aluminum coil. The rest of us should stick with copper.

Silver is slightly more conductive than copper, but the price is much higher. Copper losses are not considered that big a factor in coil performance, so there is little incentive to go to silver wire.

The dc resistance of copper wire is determined by table look-up or by measurement with an ohmmeter. The tables typically give the resistance of 1000 ft of wire at a specified temperature, say 20°C or 68°F. The tables give the resistance to four significant places, but using more than three is somewhat of a joke because of the resistance variation with temperature. If the table resistance is R_1 at a temperature T_1 , then the resistance R_2 at some other temperature T_2 is

$$R_2 = R_1 \frac{T_2 - T_i}{T_1 - T_i} \quad (1)$$

where T_i is the *inferred absolute zero temperature*, -234.5°C for copper. If the table resistance is given for 20°C, then the resistance of a copper wire at temperature T_2 is

$$R_2 = R_1 \frac{T_2 + 234.5}{20 + 234.5} \quad (2)$$

The resistance has dropped to 90% of the tabulated value at -5.45°C (not impossible for an unheated shop in some parts of the world) and is at 110% of the tabulated value at 45.45°C . The wire could certainly rise to this temperature during operation even on a pleasant day.

The inferred absolute zero temperature does not have much to do with absolute zero (-273°C). Pure iron has an inferred absolute zero of -162°C and Manganin has an inferred absolute zero of $-167,000^{\circ}\text{C}$. The fact that the inferred absolute zero of copper is not drastically different from absolute zero makes it useful as a temperature sensor. If one is curious about how warm a Tesla coil got during operation, the best way of measuring the temperature is to measure the dc resistance and work backwards through the above equations.

2 Skin Effect

At dc, current flows uniformly across the entire cross section of the conductor. As frequency increases, however, a phenomenon called skin effect causes less of the total current to flow in the center of the wire. Having less conductor available causes the resistance to increase.

An expression for skin depth can be derived as

$$\delta = \frac{1}{\sqrt{\pi f \mu \sigma}} \quad (3)$$

where f is the frequency in Hz, μ is the permeability of the conductor ($4\pi \times 10^{-7}$ for nonferromagnetic materials), and σ is the conductivity.

The skin depth for copper at 20°C is

$$\delta = \frac{0.066}{\sqrt{f}} \text{ m} \quad (4)$$

Most introductory electromagnetic theory books derive the expression for ac resistance as

$$R_{ac} = R_{dc} \frac{b}{2\delta} \quad (5)$$

where b is the radius of the wire and δ is the skin depth, in consistent units. This equation is only valid for $\delta \ll b$. As might be expected, this excludes most Tesla coils, so we must find other expressions.

A number of more advanced electromagnetic theory books derive an expression for the ac resistance. Typical is the treatment by Ramo, Whinnery, and Van Duzer [6]. They start with Maxwell's Equations, and write a differential equation for the current density $J_z(r)$ inside an isolated straight cylindrical conductor centered on the z axis.

$$\frac{d^2 J_z}{dr^2} + \frac{1}{r} \frac{dJ_z}{dr} + T^2 J_z = 0 \quad (6)$$

where

$$T^2 = -j\omega\mu\sigma = -j(2\pi f\mu\sigma) = -j\frac{2}{\delta^2} \quad (7)$$

The solution for current density is

$$J_z = AJ_0(Tr) \quad (8)$$

where $J_0(Tr)$ is the Bessel function of the first kind and zero order.

The infinite series for this Bessel function is

$$J_0(Tr) = \sum_{m=0}^{\infty} \frac{(-1)^m \left(\frac{Tr}{2}\right)^{2m}}{(m!)^2} = \sum_{m=0}^{\infty} \frac{(-1)^m (T^2)^m r^{2m}}{2^{2m} (m!)^2} \quad (9)$$

When we insert the above expression for T^2 , the series becomes

$$J_0(Tr) = \sum_{m=0}^{\infty} \frac{(-1)^m (-j)^m \left(\frac{2}{\delta^2}\right)^m r^{2m}}{2^{2m} (m!)^2} \quad (10)$$

Note that

$$(-1)^m (-j)^m = (j)^m \quad (11)$$

$$\frac{2^m}{2^{2m}} = \frac{1}{2^m} \quad (12)$$

$$\frac{r^{2m}}{\delta^{2m}} = \left(\frac{r}{\delta}\right)^{2m} \quad (13)$$

With these substitutions, we can write

$$J_0(Tr) = \sum_{m=0}^{\infty} \frac{(j)^m \left(\frac{r}{\delta}\right)^{2m}}{2^m (m!)^2} \quad (14)$$

The series can now be split into its real and imaginary parts.

$$\text{Real}[J_0(Tr)] = 1 - \frac{\left(\frac{r}{\delta}\right)^4}{2^2(2!)^2} + \frac{\left(\frac{r}{\delta}\right)^8}{2^4(4!)^2} - + \cdots \quad (15)$$

$$\text{Imag}[J_0(Tr)] = \frac{\left(\frac{r}{\delta}\right)^2}{2(1!)^2} - \frac{\left(\frac{r}{\delta}\right)^6}{2^3(3!)^2} + \frac{\left(\frac{r}{\delta}\right)^{10}}{2^5(5!)^2} - + \cdots \quad (16)$$

These are related to the ber and bei functions as follows.

$$\text{Real}[J_0(Tr)] = \text{ber}(\sqrt{2}r/\delta) \quad (17)$$

$$\text{Imag}[J_0(Tr)] = \text{bei}(\sqrt{2}r/\delta) \quad (18)$$

There are several ways to proceed to find the ac resistance. One is to use one of Maxwell's curl equations to find the magnetic field inside the wire, a constant times the derivative of the electric field. This magnetic field is used to find the total current in the wire. The electric field times the wire length ℓ is the total voltage. The ac resistance is just the real part of the ratio of voltage to current.

To my knowledge, every method of finding the ac resistance is tedious. I will proceed with a basic Circuit Theory I type approach. Consider the straight wire as formed of N concentric cylinders, as shown in Fig. 1.




Figure 1: Wire Formed of Concentric Cylinders

The wire radius is b . The radius to each boundary between cylinders is designated by r_k , where $r_N = b$, $r_1 = b/N$, $r_0 = 0$, etc. The cross sectional area of cylindrical shell k is

$$\text{Area} = \pi(r_k^2 - r_{k-1}^2) \quad (19)$$

The dc resistance of cylindrical shell k is

$$R_{dck} = \frac{\rho \ell}{\pi(r_k^2 - r_{k-1}^2)} \quad (20)$$

where ρ is the resistivity and ℓ is the length of the wire. The dc resistance of the entire wire is

$$R_{dc} = \frac{\rho \ell}{\pi b^2} \quad (21)$$

The average current density in a shell is approximately the current density at the midpoint of a shell. That is,

$$J_{zk} = AJ_0(T_r) \quad \text{for } r = \frac{r_k + r_{k-1}}{2} \quad (22)$$

The current in shell k is

$$I_k = J_{zk} \pi(r_k^2 - r_{k-1}^2) \quad (23)$$

Note that I_k is a phasor. Both the magnitude and the phase of the current and current density vary from the center to the surface of the wire.

The total current is given by

$$I = \sum_{k=1}^N I_k \quad (24)$$

where all the real parts of the various I_k are added together, and all the imaginary parts, to form the phasor I . The power dissipated in shell k is

$$P_k = I_k I_k^* R_{dck} \quad (25)$$

where $I_k I_k^*$ indicates that the current has been multiplied by its complex conjugate to form $|I_k|^2$. The total power dissipated in the wire is

$$P = \sum_{k=1}^N P_k \quad (26)$$

But by definition

$$P = |I|^2 R_{ac} \quad (27)$$

Therefore,

$$R_{ac} = \frac{P}{|I|^2} = \frac{\sum_{k=1}^N P_k}{|I|^2} \quad (28)$$

and

$$\frac{R_{ac}}{R_{dc}} = \frac{\sum_{k=1}^N P_k}{R_{dc}|I|^2} \quad (29)$$

We need to do one more step to get this formula into computational form. The ratio of resistances is independent of a specific current flow so we can set $A = 1$ in Eq. 22. For shorthand we will write the real part of J_{zk} as $\text{RE}J_k$ and the imaginary part as $\text{IM}J_k$. After pages of algebra, all the terms like $\rho\ell$ cancel out and we have the final result.

$$\frac{R_{ac}}{R_{dc}} = \frac{\sum_{k=1}^N [\text{RE}J_k^2 + \text{IM}J_k^2] \left(\left(\frac{r_k}{b} \right)^2 - \left(\frac{r_{k-1}}{b} \right)^2 \right)}{\left[\sum_{k=1}^N \text{RE}J_k \left(\left(\frac{r_k}{b} \right)^2 - \left(\frac{r_{k-1}}{b} \right)^2 \right) \right]^2 + \left[\sum_{k=1}^N \text{IM}J_k \left(\left(\frac{r_k}{b} \right)^2 - \left(\frac{r_{k-1}}{b} \right)^2 \right) \right]^2} \quad (30)$$

The algebra basically normalizes itself to use a cylinder of unit radius. The variable is skin depth, which appears in the argument of the Bessel function as r/δ . If the wire radius is one skin depth, the ratio $R_{ac}/R_{dc} = 1.020$, so the dc resistance can be used for any wire for which the radius is less than a skin depth. Other values are given in Table 1.

Table 1: R_{ac}/R_{dc} for various values of b/δ .

b/δ	R_{ac}/R_{dc}
1	1.020
2	1.263
3	1.763
4	2.261
5	2.743
6	3.221
7	3.693
8	4.154

This analytic technique of finding the ac resistance has been around for a long time and is described in many books. Before the advent of computers, it was difficult to use. One would use tables of Bessel functions which listed real and imaginary components separately, such as Jahnke and Emde [4]. The notation among different authors differed, making it difficult to compare results. Of course, computers have now made the task of calculating these infinite series much more manageable. Actually, if the wire radius is not more than about 8 skin depths, the first 12 terms of Eq. 14 are more than adequate.

Because of the difficulties of using this analytic technique before computers, people developed lookup tables and empirical models to find R_{ac}/R_{dc} . For historical interest and for the benefit of any reader adverse to writing computer code, we will present the approximations developed by Terman [8], who has a detailed discussion of this topic. He defines R_{ac} in terms of a parameter x , where

$$x = \pi d \sqrt{\frac{2f}{\rho(10^7)}} \quad (31)$$

for nonmagnetic materials. Here, d is the conductor diameter in meters, f is the frequency in Hz, and ρ is the resistivity in ohm meters. As x gets very small, due to either low frequency or small wire, the ac resistance approaches the dc resistance. Above about $x = 3$, R_{ac}/R_{dc} varies essentially linearly with x according to the expression

$$\frac{R_{ac}}{R_{dc}} = 0.3535x + 0.264 \quad (x > 3) \quad (32)$$

Terman gives the following tabular values of R_{ac}/R_{dc} for x between 0 and 3.

Table 2: R_{ac}/R_{dc} for various values of x .

x	R_{ac}/R_{dc}	x	R_{ac}/R_{dc}
0	1.0000	1.5	1.026
0.5	1.0003	1.6	1.033
0.6	1.0007	1.7	1.042
0.7	1.0012	1.8	1.052
0.8	1.0021	1.9	1.064
0.9	1.0034	2.0	1.078
1.0	1.005	2.2	1.111
1.1	1.008	2.4	1.152
1.2	1.011	2.6	1.201
1.3	1.015	2.8	1.256
1.5	1.020	3.0	1.318

For copper, $\rho = 1.724 \times 10^{-8}$ ohm meters. For those of us still using wire tables in English

units, where wire diameters are given in mils (1 mil = 0.001 inch), Terman [8] has reduced the expression for x to

$$x = 0.271d_m\sqrt{f_{MHz}} \quad (33)$$

where d_m is the wire diameter in mils and f_{MHz} is the frequency in MHz.

For example, the 14 ga coil being examined here has a dc resistance of $3.99\ \Omega$ at 20°C . For a large toroid, the resonant frequency is 160 kHz. The nominal diameter of 14 ga wire is 64.08 mils. We calculate x as

$$x = 0.271(64.08)\sqrt{0.16} = 6.946$$

The ac resistance of the wire in the coil (assuming the wire is uncoiled and is supported in one straight line) is then

$$R_{ac} = 3.99(0.3535(6.964) + 0.264) = 10.85\ \Omega$$

We see that skin effect makes a significant difference in resistance, especially where larger wire sizes or higher frequencies are used.

The two methods presented here (Terman and the one using Bessel functions) should yield very nearly the same results if we use the relationship

$$x = 1.412\frac{b}{\delta} \quad (34)$$

For example, if $b/\delta = 3$, then $x = 4.236$ and

$$\frac{R_{ac}}{R_{dc}} = 0.3535(4.236) + 0.264 = 1.761 \quad (35)$$

which is reasonably close to the value of 1.763 given in Table 1.

3 Proximity Effect

The effect of adjacent turns in the coil causes the current density to be even more nonuniform than for the straight wire, which raises the resistance even more. This effect is called the *proximity effect*. Terman [8] has a curve for two straight parallel cylinders carrying current the same direction that shows an increase of about 33% for the wires touching physically (but not electrically), and an increase of about 10% for the case when the two wires are separated by a gap equal to the wire diameter.

Medhurst

But our problem is not that of two parallel wires going to infinity, but that of hundreds of turns of wire in a finite space. Perhaps the first person to examine this problem experimentally was Medhurst [5]. We saw his results for the self-capacitance of a coil back in Chapter 2. The same paper gives tables for ϕ_M , the increase in resistance over R_{ac} as a function of coil length over coil diameter and wire diameter over wire spacing. Since such tests are critical to our Tesla coil model, it seems appropriate to discuss his paper in some detail.

Medhurst used between 30 and 50 turns of bare copper wire wound in grooves of a low loss dielectric former. The former material was Distrene, which he claimed to have a power factor of about 0.0003. The coil and former were dried carefully before testing. This means that dielectric losses in the coil form, winding insulation, and humidity should be negligible. He did not mention the possibility of eddy current losses, but I suspect these were negligible as well. Testing was done at low power, so corona or spark losses would have been zero. His test method did not separate ohmic (heating) losses from radiation losses. Standard wisdom is that radiation from small coils at relatively low frequencies is negligible, and I have seen no contrary evidence, so it will be assumed that his results apply to ohmic heating.

He used two different wire sizes, 18 and 20 s.w.g. I assume these refer to the British Standard Wire Gauge, where 18 s.w.g. has a diameter of 48 mils (1.219 mm) and 20 s.w.g. has a diameter of 36 mils (0.9144 mm). For the tight wound case, he used double-silk-covered (dsc) or single-silk-covered conductors. Such coatings are not readily available today, so it would be hard to replicate his findings.

A confusing aspect of his paper is that the symbol H is used in four different ways. It is used for magnetic field and for the unit of inductance, both of which are immediately obvious and not a problem. It is used for a function of coil length over coil diameter as we saw in Chapter 2, where $C_M = HD$ for the self-capacitance (or Medhurst capacitance) of a coil. Then the symbol is also used for a function of wire diameter and skin depth in determining the ac resistance. Some other symbol would have been more appropriate.

The coils Medhurst tested are not ‘typical’ Tesla coils. Tesla coils usually have far more than his 30 - 50 turns. He tested at frequencies of around 1 MHz, somewhat above most Tesla coil operating frequencies. He restricted himself to cases where the skin depth is a small fraction of the wire diameter.

We might define a Medhurst resistance R_M , similar to the Medhurst capacitance C_M , where

$$R_M = R_{ac}(1 + k_f(\phi_M - 1)) \quad (36)$$

where k_f is a monotonic function of frequency that is zero for very low frequencies and unity for very high frequencies. That is,

$$R_M = \phi_M R_{ac} \quad f \text{ large} \quad (37)$$

and

$$R_M = R_{ac} \approx R_{dc} \quad f \text{ small} \quad (38)$$

Medhurst makes no attempt to find k_f . Therefore we will have to restrict ourselves to the case of high frequencies, where the proximity effect is fully ‘saturated’, or where the wire radius is several skin depths thick. ϕ_M is given in his Table VIII. A portion of that table is shown in Table 3.

Table 3: Experimental values of ϕ_M , the ratio of high-frequency coil resistance to the resistance at the same frequency of the same length of straight wire.

	ℓ_w/D						
d/z_1	1	2	4	6	8	10	∞
1	5.55	4.10	3.54	3.31	3.20	3.23	3.41
0.9	4.10	3.36	3.05	2.92	2.90	2.93	3.11
0.8	3.17	2.74	2.60	2.60	2.62	2.65	2.81
0.7	2.47	2.32	2.27	2.29	2.34	2.37	2.51
0.6	1.94	1.98	2.01	2.03	2.08	2.10	2.22
0.5	1.67	1.74	1.78	1.80	1.81	1.83	1.93
0.4	1.45	1.50	1.54	1.56	1.57	1.58	1.65
0.3	1.24	1.28	1.32	1.34	1.34	1.35	1.40

In this table, ℓ_w is the coil winding length, D is the coil diameter, d is the diameter of the copper wire, and z_1 is the center-to-center spacing between adjacent turns, all in consistent units.

This table indicates that the proximity effect can easily double or triple the measured input resistance over that predicted by R_{ac} for a straight wire of the same length. In the previous subsection, a 14 gauge coil was mentioned which had a dc resistance of 3.99 Ω , and an ac resistance of 10.85 Ω at 160 kHz. By interpolation in Table 3, ϕ_M is found to be 1.85. Therefore, the predicted ohmic resistance of the coil would be $(1.85)(10.85) = 20 \Omega$ (at sufficiently high frequencies). The measured input resistance is about 23 Ω at resonance. This measured resistance includes dielectric losses and eddy current losses in addition to ohmic losses. The difference of 3 Ω would represent eddy current, dielectric, and transmission line losses, which is probably not far from reality. If the proximity effect is not fully ‘saturated’, the predicted resistance would be less, and the difference between predicted and measured would be greater.

Poynting

I then started thinking about methods to explain and perhaps even calculate ϕ_M . Note that we are still operating in the circuit theory mode here, such that the current is the same in every turn of the coil. The effect of distributed capacitance, which causes the current to be different in different turns will be discussed later. I worked on two different methods that I had not seen in the literature. One was to use Poynting's vector and calculate power flow into the copper of the coil from the total magnetic and electric fields. The other was to split the current into filaments and require a distribution of filament currents that would minimize the magnetic field in the center of the wire. Both methods worked to some degree if ℓ_w/D was not too short. After discussing these two methods, I will present some results from a recent paper [2]. This paper is computationally superior to my two methods, but still is not all that great for short coils. I have come to realize that finding the resistance of a short coil is a tough problem. I am tempted to take up golf or something less frustrating!

The Poynting vector is defined as

$$\mathbf{S} = \mathbf{E} \times \mathbf{H} \quad \text{W/m}^2 \quad (39)$$

where $\times$ refers to the cross product of two vectors.

A student in Circuit Theory I will hear the instructor say that power flows down a wire *inside* the wire (electrons bumping along). The student will walk next door and hear the EM Theory instructor say that *obviously* power flows down the *outside* of a wire in the form of electric and magnetic fields. The speed of propagation is determined by the dielectric constant of the insulating material, not by any property of the conductor, hence the action must be happening outside the copper. The wire gets hot due to some leakage of the fields from outside to inside through the mechanism of Poynting's vector.

Most students compartmentalize this information, so think power flows one place in one class and the opposite place in another class, and rarely ask where the power really flows. The problem is that no one really knows. Correct answers are obtained with either approach.

There is a problem, in that Poynting's vector does not appear to work in some cases. Consider the following example.

Example

On a clear day the average electric field is 130 V/m at the surface of the earth, directed down (the earth is negatively charged with respect to space). At the earth's magnetic equator, the earth's magnetic flux density $B = \mu H$ is horizontal, with magnitude about half a gauss, or 0.5×10^{-4} T. If $\mathbf{E}$ is directed 'down' and $\mathbf{H}$ is directed 'south', then the Poynting vector $\mathbf{S}$ is directed 'west'. The magnitude is

$$S = EH = \frac{EB}{\mu_o} = \frac{130(0.5 \times 10^{-4})}{4\pi \times 10^{-7}} = 5200 \quad \text{W/m}^2$$

A power density of 5200 W/m² would be very noticeable, but our senses tell us that nothing is actually flowing. No one has built a machine to extract and sell this power. A student who asks about this problem is likely to get one of two responses. The first is a raised voice and some subtle ridicule for asking such a dumb question. The second is a variant of “Trust me. Sometimes the electron acts like a wave, sometimes like a particle. We who have already been initiated into the scientific priesthood will train you to know when it acts like what.” But I digress.

Standard wisdom among the priesthood is that Poynting’s vector works when $\mathbf{E}$ and $\mathbf{H}$ are the cause and effect of each other, or are both produced by the same source. If current is flowing in a wire, it produces a magnetic field around the wire and a voltage drop (electric field) along the wire, so this is a proper application.

For a straight conductor centered on the z axis, the magnetic field $\mathbf{H}$ is given by

$$\mathbf{H} = \frac{I}{2\pi r} \mathbf{a}_\phi \quad (40)$$

where I is the current flowing in the conductor, r is the distance from the z axis, and $\mathbf{a}_\phi$ is the unit vector in the ϕ direction around the conductor. The current I produces a voltage drop along the conductor and the electric field is $\mathbf{E} = IR' \mathbf{a}_z$ where R' is the resistance per unit length. The cross product $\mathbf{a}_z \times \mathbf{a}_\phi = -\mathbf{a}_r$, so the Poynting vector is directed into the conductor.

The magnitude of the Poynting vector entering a straight wire of radius b is

$$S_b = EH = IR' \frac{I}{2\pi b} \quad (41)$$

The total power entering the conductor is determined by integrating the Poynting vector over its surface. If the wire length is ℓ and the radius is b , the power is

$$P_b = \int_0^{2\pi} \int_0^\ell (IR') \left(\frac{I}{2\pi b} \right) b d\phi dz = I^2 R' \ell \quad (42)$$

which is exactly what is predicted by standard circuit theory.

The magnetic field of a filamentary loop, carrying a current I , centered on the z axis and located in the xy plane is [7]

$$H_\rho = \frac{I}{2\pi} \frac{z}{\rho \sqrt{(a+\rho)^2 + z^2}} \left[-K + \frac{a^2 + \rho^2 + z^2}{(a-\rho)^2 + z^2} E \right] \quad (43)$$

$$H_z = \frac{I}{2\pi} \frac{1}{\sqrt{(a+\rho)^2 + z^2}} \left[K + \frac{a^2 - \rho^2 - z^2}{(a-\rho)^2 + z^2} E \right] \quad (44)$$

where K and E are the complete elliptic integrals of the first and second kind. The general structure is shown in Fig. 2, which shows three turns out of an N -turn Tesla coil. Coil diameter

is D and the radius of the individual wires is b . The distance between adjacent turns is z_1 and the overall winding length is ℓ_w . Current is flowing into the paper on the right, and out of the paper on the left. The electric field in each conductor is in the same direction as the current.




Figure 2: Three Turns of Tesla Coil

The elliptic integrals K and E are functions of a parameter k , where

$$k^2 = \frac{4a\rho}{(a + \rho)^2 + z^2} \quad (45)$$

One can either look up the values for K and E in a math table, or can evaluate their infinite series representation. The series converges relatively slowly, so computation of the first one or two thousand terms is not unreasonable, using double precision numbers.

As might be expected, people have developed techniques for calculating these elliptic integrals with less computational effort. Gauss's method using the arithmetic-geometric mean method converges very fast, usually in less than 20 terms for 10-digit accuracy [1]. The incremental H field on the inside surface of turn i that is produced by current I in turn j is given by

$$\Delta H_{z1i} = \frac{I}{2\pi} \frac{1}{\sqrt{(2a - b)^2 + (i - j)^2 z_1^2}} \left[K + \frac{a^2 - (a - b)^2 - (i - j)^2 z_1^2}{b^2 + (i - j)^2 z_1^2} E \right] \quad (46)$$

and similarly for the field on the outside of the coil, ΔH_{z2i} . The total field H_{z1i} is found by summation.

$$H_{z1i} = \sum_{j=1}^N \Delta H_{z1i} \quad (47)$$

Similar expressions hold for H_{z2i} , $H_{\rho 1i}$, and $H_{\rho 2i}$. These four fields are tabulated in Table 4 for a 21 turn coil with $\ell_w/D = 1$ and $2b = z_1$. This corresponds to the upper left corner of Table 3.

Table 4: Predicted values of H_{z1i} , H_{z2i} , $H_{\rho 1i}$, and $H_{\rho 2i}$ in A/m for a 21 turn coil with $\ell_w/D = 1$ and $d/z_1 = 1$. $D = 8$ inches and $d = 2b = 0.4$ inches.

i	H_{z1i}	H_{z2i}	$H_{\rho 1i}$	$H_{\rho 2i}$
1	65.16	-19.79	-66.47	-35.09
2	73.92	-23.69	-35.09	-24.72
3	77.62	-23.76	-24.72	-18.60
4	79.92	-23.19	-18.60	-14.32
5	81.58	-22.52	-14.32	-11.06
6	82.82	-21.90	-11.06	-8.42
7	83.76	-21.36	-8.42	-6.20
8	84.46	-20.94	-6.20	-4.26
9	84.93	-20.64	-4.26	-2.49
10	85.21	-20.46	-2.49	-0.82
11	85.31	-20.40	-0.82	0.82
12	85.21	-20.46	0.82	2.49
13	84.93	-20.64	2.49	4.26
14	84.46	-20.94	4.26	6.20
15	83.76	-21.36	6.20	8.42
16	82.82	-21.90	8.42	11.06
17	81.58	-22.52	11.06	14.32
18	79.92	-23.19	14.32	18.60
19	77.62	-23.76	18.60	24.72
20	73.92	-23.69	24.72	35.09
21	65.16	-19.79	35.09	66.47

We see that the magnetic field is up on the inside of the coil and down on the outside. Likewise, it is directed radially inward over the bottom half of the coil (turns 1-10) and outward over the top half of the coil. The z component is largest in the middle of the coil, while the ρ component is largest at the ends.

The total power flowing into the hollow cylinder bounding the coil is then

$$\begin{aligned}
P &= \int S d(\text{area}) \approx (\text{Area})(E)(\sum H) \\
&= \ell_t 2b(IR')(H_{\rho t} - H_{\rho b}) + \ell_t z_1 (IR') \left(\sum_{i=1}^N (H_{z1i} - H_{z2i}) \right)
\end{aligned} \tag{48}$$

$$(49)$$

where ℓ_t is the length of one turn and I have added the magnetic fields $H_{\rho t}$ and $H_{\rho b}$ to include the end effects at the top and bottom of the coil. The minus sign in front of the H_{z2i} term is necessary to indicate that the Poynting vector is into the copper on the outside of the coil as well as on the inside, and likewise for the $H_{\rho b}$ term.

The power flowing into the same length of straight wire would be

$$P_{st} = (\text{Area})(E)(H) = (N)(\ell_t)(2\pi b)(IR')\left(\frac{I}{2\pi b}\right) \tag{50}$$

Power is proportional to resistance so the ratio of P to P_{st} is another ϕ , call it ϕ_p to distinguish it from the Medhurst ϕ . The ϕ_p is

$$\phi_p = \frac{\ell_t 2b(IR')(H_{\rho t} - H_{\rho b}) + \ell_t z_1 (IR') \left(\sum_{i=1}^N (H_{z1i} - H_{z2i}) \right)}{N \ell_t 2\pi b IR' (I/2\pi b)} \tag{51}$$

Canceling the common terms gives

$$\phi_p = \frac{H_{\rho t} - H_{\rho b} + \sum_{i=1}^N (H_{z1i} - H_{z2i})}{1/2\pi b} \tag{52}$$

For the case shown in Table 4,

$$\phi_p = \frac{2273.92}{(21)(31.33)} = 3.456 \tag{53}$$

We see that this approach yields an increase in resistance of the coil as compared with the same amount of straight wire of 3.456, as compared with the Medhurst prediction of 5.55. This is within a factor of two, as mentioned earlier, but is not really close enough for any definitive computations. After spending considerable time looking at the problem, I decided that one significant flaw in the analysis is that it assumes the actual current flow throughout the conductor cross section can be modeled by a filamentary current (of the same total magnitude) flowing exactly in the center of each conductor. The actual current density is higher near the surface of the conductor due to skin effect, but will not be symmetric inside to outside, or top to bottom, hence the equivalent current filament will not be in the exact

center. The computations are sensitive to the position or location of the current, and if this is not accurately known, it will not be possible to obtain highly accurate results with this Poynting vector approach.

Boundary Conditions

My next analytic attempt was to split the current flow into four filaments, located on the surface at the top, bottom, inside, and outside of each turn. The magnetic field components H_z and H_ρ were calculated at the center of the wire cross section, using Eqns. 43 and 44. The computations are very similar to those of the previous Poynting vector effort. Instead of finding the fields on the wire surface from a current at the wire center, I find the fields at the wire center from current filaments on the surface.

In the high frequency limit, the magnetic field at the center of a conductor will be zero. Currents will distribute themselves to meet this boundary condition. The problem then becomes one of iteration to find a current distribution that will cause the magnetic field to be zero at the center of the conductor. For example, suppose we start with a total current flow of 4 A, or four filaments each carrying 1 A. We calculate the sum of H_z and H_ρ at the center of each turn. We then move a portion of the current in one filament, say 0.1 A, to another filament, and recalculate the sum of the fields. If the sum is smaller, we are headed the right direction, and try moving some more current from the first filament to the second. When we reach the zero field condition, we have found a solution to Maxwell's Equations.

One obvious problem with this simple algorithm is that the limit is found when all the current flows in one filament. If we think of the wire as four quadrants in parallel, but no current is flowing in three of the quadrants, it is like removing three out of four parallel resistors. The resistance of one quadrant is four times the resistance of the entire wire. This approach will yield a maximum ϕ_M of 4.00, not the 5.55 appearing in the upper left corner of Table 3. One can get past this limit by splitting the conductor into more than four filaments, say 8 or 16, or by allowing negative current in some of the filaments.

I tried 8 filaments, but this did not immediately fix the problem. It did not appear to be possible to get to the case of zero magnetic field in the center of the wire while requiring all currents to be non negative. If we allow currents to be negative, such that current is flowing in one direction on one side of the conductor and in the opposite direction on the other side, there must be some lateral or azimuthal flow of current. This violates our original assumption of current flowing only in the direction of the conductors. My mother told me there would be days like this!

While this is certainly an interesting problem in its own right, it has only limited application to Tesla coils which usually have a length/diameter ratio on the order of four. The maximum ϕ_M is now 3.54, from Table 3, which can be analyzed with four filaments and non negative currents. Results of the analysis are shown in Table 5. The column labeled ϕ_M is from Table 3 while the column labeled ϕ_J is for the predicted resistance ratio from this

iterated computer program.

Table 5: Comparison of ϕ_M with ϕ_J for a coil with $\ell_w/D = 4$

d/z_1	ϕ_M	ϕ_J
1.0	3.54	3.36
0.9	3.05	3.11
0.8	2.60	2.63
0.7	2.27	2.30
0.6	2.01	1.99
0.5	1.78	1.80
0.4	1.54	1.60
0.3	1.32	1.40

It can be seen that agreement is quite good between the Medhurst measured values and the calculated values using the four-filament approach. I feel that the Medhurst table has been theoretically validated. I am sure that there are other methods of theoretically determining ϕ_M , perhaps easier and more accurate, but this approach convinced me that proximity effect does indeed have a classical explanation, if one wants to spend the time and effort necessary to find it. In the meantime, Table 3 should be quite adequate for anyone looking for an estimate of the high frequency resistance of a coil of wire.

Fraga

After preparing the above material, a reviewer [1] informed me of a paper by Fraga, Prados, and Chen on this topic [2]. An examination of the paper indicated that the authors seemed to be restricting themselves to the following case.

1. *Long* solenoids
2. Tight wound solenoids
3. Solenoids with negligible distributed capacitance
4. Low frequency ($b \leq \delta$)
5. Multilayered coils

A typical Tesla coil has a length/diameter ratio of about four (not necessarily *long*). Some coils are space wound. Distributed capacitance is always a problem. It will be discussed in the next section. Tesla coils are usually operated at frequencies where the wire radius will be between one and five skin depths, not less than a single skin depth. And Tesla coils are almost never multilayered. So it was not obvious that the paper would be of much use to the Tesla

coil community. On the other hand, Tesla coils are infamous for using things outside their normal operating range (e.g. pulling 60 mA from a neon sign transformer rated at 30 mA), so I looked some more. The authors have developed a relatively simple closed-form expression for the resistance of a coil, not requiring summations or interpolations, so it would be very nice if it worked for the typical Tesla coil.

Their Eq. 35, expressed in my notation for the single-layer coil is

$$R_s = \frac{2\pi N^2 \rho_{\text{eff}} a [\sinh 2\theta + \sin 2\theta]}{\delta_s \ell_w [\cosh 2\theta - \cos 2\theta]} \quad (54)$$

where N is the number of turns, a is the radius of the coil, and ℓ_w is the winding length of the coil. The other terms are described in the following.

One of the difficulties in the analysis of a coil is that the conductor surfaces are not at a fixed distance from the axis of the coil. A fingernail pressed against the coil moves in and out as the hand moves down the coil. Boundary conditions cannot be simply expressed in this geometry. These authors circumvent this problem by converting round wire to square wire with the same cross-sectional area. They deal with the gap between the wires, due to insulation, by elongating the square conductor toward its neighbor until it touches mechanically (but not electrically). The resistivity is increased a proportional amount such that the resistance remains the same.

Let the radius of a cylindrical wire be b , covered by a dielectric coating of thickness s . Define a square conductor with side y with the same area. Then

$$y^2 = \pi b^2 \quad (55)$$

or

$$y = \sqrt{\pi} b \quad (56)$$

The length of one turn is $2b + 2s$. The effective resistivity for this rectangular wire of thickness y is

$$\rho_{\text{eff}} = \rho \frac{2b + 2s}{y} = \frac{2\rho}{\sqrt{\pi}} \left(1 + \frac{s}{b}\right) \quad (57)$$

where ρ is the usual resistivity of the conductor (1.724×10^{-8} ohm meters for copper at 20°C).

They then define an effective skin depth

$$\delta_s = \sqrt{\frac{2\rho_{\text{eff}}}{\mu_0 \omega}} = \sqrt{\frac{4\rho(1 + s/b)}{\sqrt{\pi} \mu_0 (2\pi f)}} = \sqrt{\frac{4(1.724 \times 10^{-8}(1 + s/b))}{\sqrt{\pi}(4\pi \times 10^{-7})(2\pi)f}} = 0.0702 \sqrt{\frac{1 + s/b}{f}} \quad (58)$$

for copper. The last variable, θ , is then defined as

$$\theta = \frac{y}{\delta_s} = \frac{\sqrt{\pi}b}{\delta_s} \quad (59)$$

One big advantage of this approach is that both the skin depth and the proximity effect are included in one equation that would fit on most programmable calculators. Input quantities are wire radius, coil radius and length, number of turns, and frequency, all readily available. Using the Medhurst approach outlined earlier requires one to first find the ac resistance and then multiply this by a factor ϕ_M found by interpolation in Table 3.

A disadvantage of the Fraga method is that the resistance is equivalent to the last column in Table 3. That is, we find the resistance of a section of an infinitely long coil. This method is not capable of properly dealing with the short coil. A comparison of ϕ_M and what I call $\phi_F = R_s/R_{ac}$ is given in Table 6. The Fraga formula was evaluated for 18 gauge wire at 1 MHz to make it as comparable as possible to the Medhurst experiment.

Table 6: Comparison of ϕ_M with ϕ_F for a coil with $\ell_w/D = \infty$

d/z_1	ϕ_M	ϕ_F
1.0	3.41	3.19
0.9	3.11	3.03
0.8	2.81	2.86
0.7	2.51	2.67
0.6	2.22	2.48
0.5	1.93	2.26
0.4	1.65	2.03
0.3	1.40	1.75

We see that the Fraga formula agrees best with Medhurst for a wire diameter to wire spacing ratio of between 0.8 and 0.9, which happens to be the typical ratio for tight wound magnet wire. Therefore, there is hope for the Fraga formula for typical Tesla coils (tight wound magnet wire with a length/diameter ratio of at least four).

I then proceeded to compare the coil resistance R_s predicted by Fraga with the coil resistance R_M predicted by Medhurst and with the measured resistance R_{TC} for my coils. Results are given in Table 7.

This table gives the wire diameter $d = 2b$ in both mils and mm, the coil radius a , the winding length ℓ_w , the total length of wire used in the coil, the number of turns N , the Wheeler inductance L , the Medhurst capacitance C_M , the Medhurst factor ϕ_M interpolated from Table 3, and the dc resistance R_{dc} at 20°C. Measured inductance is always close to the calculated value from Wheeler's formula. Measured resistance agreed with the calculated value to within 1% or so, when corrected for temperature.

Table 7: Predicted and Measured Coil Resistance for Several Coils

	12T	14S	14T	16B	18B	18T	20T	22T	22B
$2b$, mils	80.81	64.08	64.08	50.82	40.30	40.30	31.96	25.35	25.35
$2b$, mm	2.053	1.628	1.628	1.291	1.024	1.024	.8118	.6439	.6439
a , meters	.231	.198	.107	.24	.235	.107	.107	.107	.24
ℓ_w , meters	1.613	1.166	1.392	.7	.47	.881	.952	.945	.7
wire, meters	509	482	623	675	617	534	702	424	675
turns N	351	387	797	445	418	794	1052	631	445
L , mH	14.2	17.2	19.34	49.4	55.02	29.1	47.3	17.2	49.4
C_M , pF	30.58	23.95	20.70	22.66	21.22	15.45	16.14	16.08	22.66
ϕ_M	1.68	1.85	3.03	4.20	4.25	3.15	3.02	1.62	1.54
R_{dc}	2.66	3.99	4.45	8.89	12.9	11.2	23.4	22.4	35.7
f_0 , kHz	246.4	247.6	251	153.5	147.3	236.9	181.1	301.5	153.5
R_s	22.83	29.99	44.11	51.13	62.49	67.55	97.33	65.94	69.00
R_M	17.96	24.7	46.35	80.93	95.61	76.50	111.39	58.38	66.09
R_{TC}	18.7	24.5	43.5	93.1	118.0	70.5	94.2	47.0	73.0
f		217.3	211.3	128.4	123.6	176.1	135.9	227.9	129.5
R_s		28.10	40.48	46.76	54.04	58.23	84.47	56.76	62.23
R_M		23.13	41.67	74.97	88.92	66.52	98.37	51.57	55.00
R_{TC}		25.6	42.3	87.1	103.0	65.9	88.0	46.7	69.5
f		158.8	151.8	91.9	86.2	122.5	94.2	158.1	93.1
R_s		24.02	34.31	39.56	46.90	48.57	70.40	45.76	51.74
R_M		20.06	35.86	64.76	76.65	57.03	85.04	44.56	55.00
R_{TC}		24.5	39.6	75.0	78.9	58.1	78.7	42.4	66.8

The suffix S refers to space wound, while T refers to a tight wound coil. The suffix B refers to a barrel whose sides are not perfectly straight, for coils 16B and 22B. The winding is tight wound on the flat portions of the barrel and space wound on the transition portions. The barrels were assumed to be made of polyethylene when I purchased them at the local recycling plant. The barrels used for 12T and 18B are straight sided and have thinner walls than the 16B and 22B barrels. They were once used as tanks for water softeners. Coil 14S is on a coil form built from a 0.125 inch sheet of polyethylene. Coils 14T, 18T, 20T, and 22T are on PVC forms. Coil 20T has 3 layers of polyurethane on it, the others have no coating. Coil 12T is made in two sections for ease of handling.

The 12 ga wire is Essex type USE-2 (or type RHH or type RHW-2), which are types specified in the National Electrical Code. This type has a relatively thick insulation, which spaces the conductors farther apart than the thin insulation of magnet wire. It has a nominal insulation thickness of 45 mils as compared to a 15 mil thickness for the more common Type THHN. The 14, 16, 18, and 20 ga wires are magnet wires coated with Heavy Soderon. This

Essex coating has a top layer of nylon. Magnet wires are available with one, two, three, and four layers of insulation. Heavy Soderon is equivalent to a two layer coating. The 22 gauge wire is not magnet wire but has a yellow insulating jacket, that I assume is PVC. The wire was acquired surplus and the numbers are difficult to trace.

The Medhurst capacitance and the Wheeler inductance are used to calculate the resonant frequency f_0 . The actual resonant frequency was measured with a HP54645 digital scope. Agreement was within 5%. The operating resonant frequency will always be lower than the unloaded value due to the toroid on top, so this number is of mostly academic interest anyhow.

The observed resistances include displacement current effects and any dielectric losses, which might increase the actual resistance well above that due to copper losses alone. The highest frequency in the table refers to the case of no top load, while lower frequencies were obtained with different sizes of toroids mounted on top the coil.

Coil 14T has the greatest length/diameter ratio (6.47) of the group, so we would expect Fraga's formula to work the best for this case. Indeed, it predicts a resistance only about 4% less than that predicted by Medhurst, both close to the experimental values.

As we move to coil 18B, which has the smallest length/diameter ratio (1.0), Fraga's formula significantly under predicts the resistance, being about 55% of observed and about 63% of the Medhurst values.

On the other hand, coil 22T has a reasonable length/diameter ratio (4.42), but greater dielectric thickness s , and Fraga's formula over predicts the resistance, being about 124% of observed and about 109% of Medhurst.

For the seven coils 14T–22B, Fraga's formula predicts resistance values about 90% of what Medhurst predicts. All things considered, this is pretty good. If we stick with coils close wound with magnet wire, and a length/diameter ratio of 4 or more, Fraga's formula should be quite acceptable.

The next step was to determine the character of the transition of resistance values as frequency is raised from dc to those frequencies where the proximity effect is fully implemented. That is, there is no proximity effect at dc. The measured resistance of a coil is the same as the same length of straight wire. As frequency increases, however, the skin effect causes resistance to increase, and the proximity effect causes resistance to increase even more. Presumably, at a high enough frequency where the skin depth is a small fraction of the wire radius, the proximity effect saturates. The measured resistance continues to increase, but only due to the skin effect.

I measured the input resistance as a function of frequency for seven of my coils, 14T, 16B, 18B, 18T, 20T, 22T, and 22B. This includes all five coils with a Medhurst factor $\phi_M = R_{TC}/R_{ac}$ of 3.0 or more. I used a standard bench function generator with sine wave output, followed by a linear amplifier. The amplifier is a simple single-stage inverting op-amp, using the Apex PA-19, rated at 4 A, ± 36 V, slew rate 900 V/ μ s. Voltage was measured with a standard scope probe, current with a Philips PM9355 current probe. The output of the

current probe was fed to the second channel of the scope, a HP54645D digital oscilloscope capable of calculating rms values of measured waveforms. I applied a voltage between about 2.5 and 10 V (rms) to the base of the coil, tuned the function generator for resonance, observed the current value, and calculated the resistance as the ratio of voltage to current.

The test location was inside a 54 by 90 ft metal building (manufactured by Morton) that is typically used for livestock or storage of ag equipment. One 15 by 30 ft corner was framed in, insulated, and equipped with furnace and air conditioner. Inside this instrument room was a double copper wall Faraday cage with footprint 8 by 12 ft. Electricity was provided to this screen room through an isolation transformer and power line filters. The computer, oscilloscope, and other sensitive equipment were located in this screen room. It is interesting to note that reception on a transistor AM radio was not affected by moving it from outside to inside the Morton building, but reception was impossible inside the screen room, even with the massive copper door open. The AM band starts at 550 kHz, slightly above most of my testing in the 100–300 kHz range, so it would appear that the Faraday cage was effective in the necessary frequency range. I never observed any failures in electronic equipment inside the cage that I could link to transient high fields outside the cage.

This Morton building has a wood frame and wood trusses supporting the roof. There is no ceiling so the wood trusses and roof are visible inside the building. The bottoms of the trusses are about 16 ft above the floor. Depending on the location, the roof will be about 17 to 25 ft above the floor.

Except for the screen room, all electrical outlets in the building were wired in the conventional manner for North America, with the third wire connected to utility ground. This ground wire is connected to earth at every power pole on the utility system. I installed my own ground system under the dirt floor of the Morton building, consisting of three lengths of copper tubing buried a foot or so below floor level. The soil is very hydrophilic, so when I water the grass on the outside of the building, the copper tubing is located in wet earth. The measured resistance between this local ground and utility system ground was on the order of 1 Ω , a quite acceptable value. Open circuit voltage between the two grounds may be as high as several volts, and current flow between grounds may be as high as several amps. The utility ground is a major source of noise in sensitive measurements, so all measurements were made using only the local ground. The Faraday cage was connected to this local ground, as well as the metal skin of the Morton building.

A damp earth floor made the interior of the building too humid, so I covered the floor with polyethylene sheets, and covered those with about four inches of milled asphalt. This is a product obtained when asphalt roads are recycled. If carefully packed, it forms almost as nice a surface as concrete, and is much less expensive. It is also much easier to penetrate if one wanted to install something in the earth.

Even though the local ground is very satisfactory for making sparks with Tesla coils, I was concerned that variations in earth moisture would affect my proximity effect measurements. I therefore installed a metal ground plane on top of the asphalt millings, about 8 by 16 ft,

consisting of sheets of aluminum siding bonded together with copper flashing and sheet metal screws. This would obviously not be suitable for spark tests when hundreds of kV and tens of amps are present, but for low level measurements of a few volts and less than one amp should be ok.

The coil under test was placed on a piece of 2 inch thick blue styrofoam on this metal ground plane. Location was about 10 ft from the wall of the Morton building and about 10 ft from the wall of the instrument room. A length of 50 Ω coaxial cable was run under the ground plane, through the wall of the instrument room, through the wall of the Faraday cage, and to the function generator. The shield of the coax was connected to the ground plane close to the base of the coil. The center conductor of the coax was connected to the base of the coil, so the coil was driven like a vertical helix above a ground plane. With no toroid or other connection to the top of the coil, the inductance of the coil would resonate with the Medhurst capacitance. This would be the maximum frequency at which the impedance could be measured.

Capacitance could be increased and frequency lowered by placing larger and larger toroids on top the coil. However, it is not feasible to have toroids large enough to lower the frequency to a few kHz, as required for this test. Different size toroids also change the local field distribution, which might affect the results. So I decided to connect a capacitor to the top of the coil, and the other lead of the capacitor to a heavy ground wire that went up from the top of the coil to about 10 ft above the floor, over to the wall of the instrument room, and down to the floor where it was connected to both the metal ground plane and to the copper tubing of the local ground. Electrically, this forms a simple series RLC circuit. If large value capacitors are used, the resonant frequency can be under 1 kHz.

Like many other things about Tesla coils, using the wrong type of capacitor can lead to surprises. The voltage rating is important, as a little thought will reveal. The Q of these coils is high, perhaps as high as 500. Driving the coil with 10 V would then put as much as 5000 V across the capacitor. My first effort was to use variable capacitors here. I had a ham type variable capacitor that would handle the voltage, but receiving type variable capacitors would arc over between the plates. Even when the voltage was acceptable, the resistance of the sliding contacts in the capacitors was too high (and too erratic) at the required current levels.

The final solution was to place an aluminum cake pan, flat bottom up, on top of the coil and electrically connected to it. A half spun aluminum toroid was connected to the hanging ground wire. Polyethylene sheets were placed on the cake pan and the half spun toroid would be placed on the sheets, flat side down, to form a simple parallel-plate capacitor. Different thicknesses would allow for the resonant frequency to be lowered to about one fourth of the maximum value. For even lower frequencies, lumped capacitors were used.

Figure 3 shows the variation of R_{TC}/R_{ac} versus frequency for the coil 14T, plotted solid with diamond symbols. It also shows the Fraga resistance R_s divided by R_{ac} , plotted dashed. Theory and experiment match extremely well. Again, the Fraga formula works well for long

coils that are tight wound with magnet wire.





Figure 3: R_{TC}/R_{ac} for coil 14T

I then performed the same measurements of R_{TC} for coil 18B, the shortest coil in my collection. Results are shown in Fig. 4. The Medhurst ϕ_M is 4.25 for this coil where ϕ_F is about 3. The predictions by Fraga are below the measured values over the entire range of frequencies. Even worse, the experimental and theoretical start to diverge at about 50 kHz.

I have performed other tests which show that this particular magnet wire absorbs moisture, and this moisture will cause increased losses. Dielectric losses are entirely separate from the proximity effect. It appears that the dielectric losses become significant above 50 kHz for coil 18B, causing the ratio R_{TC}/R_{ac} to go well above the value ϕ_M predicted by Medhurst for the case with negligible dielectric losses.

Figs. 3 and 4 also show another effect, a very interesting concept that is otherwise difficult to explain. This concept is that there is little penalty in performance if one uses a smaller wire in a coil. That is, the effect on spark length is not as strongly related to the wire resistance as one would expect.

If one looks closely at these figures, it is evident that the transition region moves higher in frequency as the wire gets smaller. At a given frequency, the larger wire will always be closer to saturation. Consider an extreme example where two coils are each wound with 1000 ft of magnet wire, one with 14 ga and the other with 28 ga. Assume the coils are resonant at 70 kHz and that $\phi_M = 3$. The dc resistance of the 14 ga coil is 2.525 Ω while the dc resistance of the 28 ga coil is 64.9 Ω , a factor of 25.7 greater. The ac resistances at 70 kHz are, from Eqs. 30 or 32, $(1.888)(2.525) = 4.768 \Omega$ for the 14 ga coil and $(1.0034)(64.9) = 65.1 \Omega$ for the 28 ga coil. The 14 ga coil has reached full saturation from the proximity effect at 70 kHz, so

Figure 4: R_{TC}/R_{ac} for coil 18B

the effective R_{TC} is $(4.768)(3) = 14.3 \Omega$. The 28 ga coil has not started into the proximity effect yet at 70 kHz, so its resistance is still just 65.1 Ω . The ratio of resistances at 70 kHz is $65.1/14.3 = 4.55$, a considerable reduction from 25.7.

It would have been nice to finish this treatment of copper losses with a formula that was accurate to within 5% for any frequency, any length/diameter ratio, and any wire diameter to wire spacing ratio. But that remains for someone else. I hopefully have described the problem so that the reader can get within perhaps 20% of the correct value, and sometimes even better.

4 Displacement Current Effect

We have examined two methods for empirically or theoretically determining the copper loss in a coil, the methods of Medhurst [5] and Fraga [2]. These both assume that the conduction current is the same at all points in the coil, which is the usual case for circuit theory type RLC models. But this is not necessarily the case for a Tesla coil. This is one place where the lumped circuit models just cannot go. We finally have to use a distributed approach. If the conduction current is less than the input current in part of the coil, we would expect the effective resistance to also be less. Likewise, if the conduction current is greater than the input current, then we would expect the effective resistance to increase. Conduction current can be greater than the input current in one part of a coil, and less in another part, due to displacement currents. We will try to illustrate this concept with Fig. 5. A Tesla coil is connected to a toroid with a switch S_1 . The input current at the base is the conduction

current i_{in} .

Every part of the coil has a capacitance to every other part. We show four capacitors in this four turn coil, C_{21} from turn 2 to turn 1, C_{42} from turn 4 to turn 2, C_{t2} from the toroid to turn 2, and C_{2g} from turn 2 to ground. Each capacitor has a current flow in it, with the same subscripts.




Figure 5: Conduction Current

The incremental current entering turn 2 is

$$\Delta i = i_{42} + i_{t2} - i_{21} - i_{2g} \quad (60)$$

If the two currents i_{42} and i_{t2} are greater than the other two currents, then Δi is positive. If the current in turn 1 is i_{in} and Δi is added in turn 2, then the current in turn 3 is greater than i_{in} . This situation is quite possible for the lower part of the coil where many turns above the turn in question are adding currents and only a few turns below are subtracting currents. It is more likely to occur if C_{2g} is small, that is, if the coil is mounted well above a ground plane.

On the other hand, as we get toward the top of the coil, there are many turns below the turn in question that are subtracting currents from Δi and only a few turns above that are adding currents. The minimum current occurs in the top turn of the coil. If switch S_1 is open (no toroid), this minimum current is zero. So we have the situation where the current increases in the lower part of the coil, will hit a maximum probably somewhere in the middle third of the coil, and then start to decrease toward the minimum current value at the top of the coil.

Possible variations of conduction current in the coil winding as a function of position y is shown in Fig. 6. The current marked “ S_1 closed” is a possible current with toroid while the

current marked “ S_1 open” is a possible current without a toroid.




Figure 6: Conduction Current in a Coil

There are several qualitative factors that can be deduced from these curves, without going to the effort to do the full distributed analysis. We started on this quest by asking “Is the lumped RLC model for the Tesla coil a useful concept, or must we proceed immediately to the distributed model?” If this distributed effect is not too big, such that any fudge factor is no more than a few percent, then we can still use the lumped model. If the effect is too big, then we will be inclined to use the distributed model in order to get decent results.

We note that the current in the coil is greater than the input current in the lower part of the coil and less over the upper part of the coil. These tend to offset each other so that the effective resistance may not be much different from the resistance calculated for the uniform current case.

As larger toroids are added, the current in the coil increases, and likewise the effective resistance. However, the resonant frequency decreases with larger toroids, which lowers the effective resistance for the uniform current case. This predicts a smaller change in resistance with frequency than predicted by either Medhurst or Fraga.

The final qualitative factor is that if a coil has a geometry that causes the peak current in the coil to be well above the input current (say 30% or more), the effective resistance will also be well above that for the uniform current case. We can now go back to Table 7 and check out these predictions.

We note that for coil 14T, the measured resistance is *below* both R_s and R_M for a coil without a toroid and *above* both R_s and R_M for a coil with the largest toroid. The measured resistance holds more nearly constant with variation in frequency than is predicted by Medhurst or Fraga, a very consistent observation throughout all my testing.

It appears that as coils get shorter and fatter, the interior current in the coil gets larger and the effective resistance increases as compared with the predictions of Medhurst and Fraga. Consider coils 22T and 22B. Coil 22T (relatively long and thin) has a measured resistance below the predictions, while coil 22B (short and fat) has a measured resistance above the predictions. The other two short coils (16B and 18B) also have measured resistances above

the predictions.

The average ratio of R_{TC} over R_s for the coils 14T, 18T, 20T, and 22T was 1.004, while the average ratio of R_{TC} over R_M for these four coils was 0.927. It appears to me that for normal Tesla coil geometries (length/diameter = 4 or more) that the uniform current assumption is not too bad. It appears to yield accuracies within $\pm 10\%$, which should be acceptable in most applications. We conclude that the displacement current effect is very real and easily observed in data sets like Table 7, but the errors involved in ignoring it are not so severe that we cannot use the lumped model.

5 Dielectric Losses

We turn now to losses in the coil form and in the wire insulation. Both the coil form and the wire insulation form a part of the coil capacitance. By Gauss's Law the coil capacitance is the sum of electric flux lines leaving the coil, divided by the coil voltage. Some flux lines go from toroid to earth through the coil form, some from turn to turn through the wire insulation, and some through air. These can be considered as three capacitors in parallel. Since the volumes of the coil form and the wire insulation are much smaller than the volume of air, and the relative permittivity is only two or three times that of air, their capacitances will be a small part of the total. The losses may still be significant, of course.

If we assume a linear increase of voltage along the coil, the flux lines in the coil form will be uniform from top to bottom (no fringing) just like the parallel plate capacitor. We should be able to use the formula for the parallel plate capacitor without great error.

Dielectric losses are usually modeled by a resistor in parallel with the capacitance, rather than in series. They are related to the capacitor voltage rather than to the capacitor current. We have two dielectrics, the coil form and the wire insulation, so we have two resistors in parallel with the Tesla coil capacitance. We saw equations for the power dissipation in the coil form, P_{cf} , and in the wire insulation, P_{wi} , back in Chapter 3. If the toroid voltage is V_{tor} , then the parallel resistances can be defined as

$$R_{cf} = \frac{V_{tor}^2}{P_{cf}} \quad (61)$$

$$R_{wi} = \frac{V_{tor}^2}{P_{wi}} \quad (62)$$

A somewhat more detailed RLC model is shown in Fig. 7. The resistances directly related to current are shown as R_M , R_{eddy} , R_{spark} , and R_{rad} , where the last three items are the equivalent resistances representing losses to eddy currents in toroid and ground plane, the spark itself (when present), and any losses due to radiation. We considered the Medhurst

resistance R_M earlier in the chapter. You can replace R_M by R_s from Fraga's formula if you prefer. The resistances related to voltage are R_{cf} and R_{wi} .




Figure 7: Detailed Lumped Model of Tesla Coil

For single-frequency, steady-state operation, the parallel combination of a capacitor and two resistors can be modeled as a series capacitor and resistor, call it R_{die} . This is straightforward Circuit Theory I, but a bit tedious. We write an expression for the parallel impedance of R_{cf} , R_{wi} , and the capacitance, rationalize it, and simplify the real term. We assume that the parallel resistances are much larger than the capacitive reactance, as they will be for any coil with acceptable losses. The algebra goes as follows:

$$Z = \frac{1}{1/R_{wi} + 1/R_{cf} + j\omega C_{tc}} = \frac{1}{G + j\omega C_{tc}} = \frac{G - j\omega C_{tc}}{(G + j\omega C_{tc})(G - j\omega C_{tc})} \quad (63)$$

We define R_{die} as the real part of Z .

$$R_{die} = \text{Real } Z = \text{Real} \left(\frac{G - j\omega C_{tc}}{G^2 + \omega^2 C_{tc}^2} \right) \approx \frac{G}{\omega^2 C_{tc}^2} = \frac{P_{cf} + P_{wi}}{V_{tor}^2 \omega^2 C_{tc}^2} \quad (64)$$


To make things even more complicated, we will define R_{die} as the series combination of three resistors, the series coil form resistance R_{cfs} , the series wire insulation resistance R_{wis1} , and the series coating resistance R_{wis2} . If there is no coating on the coil (polyurethane or equivalent) and if the wire and atmosphere are dry then $R_{wis2} = 0$.

$$R_{die} = R_{cfs} + R_{wis1} + R_{wis2} \quad (65)$$

The revised circuit model for the Tesla coil is shown in Fig. 8.

We now proceed to get specific equations for these three series resistors, referring back to Chapter 3 for the power dissipated in terms of V_{tor} .

$$R_{cfs} = \frac{P_{cf}}{V_{tor}^2 \omega^2 C_{tc}^2} = \frac{V_{tor}^2 \omega C_{cf} (DF)_{cf}}{V_{tor}^2 \omega^2 C_{tc}^2} = \frac{C_{cf} (DF)_{cf}}{\omega C_{tc}^2} \quad (66)$$

Figure 8: Lumped Model of Tesla Coil with Series R_{die}

$$R_{wis1} = \frac{\epsilon_2^2(\eta_o - \eta_x)(\pi\epsilon_1\ell_t)(DF)_1}{N\omega C_{tc}^2[\eta_o\epsilon_2 + \eta_x(\epsilon_1 - \epsilon_2)]^2} \quad (67)$$

$$R_{wis2} = \frac{\epsilon_1^2\eta_x(\pi\epsilon_2\ell_t)(DF)_2}{N\omega C_{tc}^2[\eta_o\epsilon_2 + \eta_x(\epsilon_1 - \epsilon_2)]^2} \quad (68)$$

It would appear that R_{die} decreases as $1/\omega$ or as $1/f$. This is certainly contrary to our intuition. What is even more surprising is that the dissipation factor of water is also proportional to $1/\omega$ over the typical Tesla coil frequency range. $(DF)_{\text{water}} = 0.396$ at $f = 10^5$ Hz and 0.0396 at $f = 10^6$ Hz [3]. This would make R_{die} vary as $1/\omega^2$. I have not convinced myself that this is correct, but can see no flaw in the above analysis. I will terminate the theoretical discussion of dielectric losses on that note.

There are many things going on that make it difficult to be precise about frequency variation of the other losses, but generally speaking, R_{ac} increases as $\sqrt{f}$, and R_{eddy} increases as f^2 . R_{rad} will increase at a rate somewhere between f and f^2 . R_{spark} can be ignored below the spark inception voltage. Depending on which terms are dominant loss terms, we may not see a pronounced change in input impedance with frequency. That has been my experience. Input impedance will drift from day to day, (mostly with humidity), but there is no obvious frequency dependence. Of course, other things are happening. We know that R_{ac} increases with temperature, while R_{die} increases with humidity. If these were the only factors, we would expect a cold, dry winter day to have the lowest impedance, and a hot, muggy day to have the highest impedance and the worst performance. In cases where moisture is a factor, performance might improve after a period of operation which caused the coil form to heat up and dry out.

Moisture has been a very frustrating factor in my testing. Coils are located inside a metal skin building with no climate control, in eastern Kansas. Temperatures can vary from below freezing to 40°C (104°F) or more. Relative humidities vary from 25% to 100%. Typical of my measurement problems are two sets of input impedance data in Table 8 for 3/17/01 and 4/6/01. The 3/17/01 data were collected when the bay temperature was about 9°C and the

relative humidity was about 28%. On 4/6/01 the temperature was about 17°C and the relative humidity was 100%. It had been damp all week with heavy fog the day before.

Table 8: R_{TC} Measured on Two Different Days

Coil	14S	14T	16B	18T	18B	20T	22T	22B
frequency, kHz	249.1	266.1	145.5	242.6	145.6	183.6	307.0	148.3
R_{TC} 3/17/01	24.5	43.5	93.1	70.5		94.2	47.0	73.0
R_{TC} 4/6/01	27.6	45.5	175.8	75.7	127.7	100.0	53.4	126.7

We see that two of the coils experienced large changes in R_{TC} , coils 16B and 22B. Both coils used a plastic barrel as a coil form that I thought was polyethylene. I got the barrels at the local recycling plant. Coil 22B used the same type of wire as coil 22T which was wound on a piece of PVC, so the difference in R_{TC} between these two coils had to be the coil form. These results indicate that some coil forms are worse than others. These barrels evidently soak up water in amounts sufficient to raise the input impedance by a factor of two.

Coils 14T, 18T, 20T, and 22T were wound on PVC while coils 14S and 18B were wound on polyethylene. Only coil 20T had any type of coating put on top the winding (polyurethane). Both PVC and polyethylene appear to have about the same increase in R_{TC} with humidity, so it is hard to argue that one should spend more money on the more expensive polyethylene.

As long as one stays with good quality coil forms, it appears that high humidity will raise R_{TC} by 5–10% from the low humidity case. One could reverse engineer Eq. 67 and find an effective dissipation factor $(DF)_1$ that would be an appropriate function of humidity, but I am not sure it would be worth the trouble.

The eddy current loss will be a strong function of how near the conducting material is located to the coil. A coil sitting on a ground plane would have a much larger R_{eddy} than one sitting on a one meter high stack of Styrofoam blocks. If soil moisture affects the eddy current loss in the earth beneath the coil, then this term could vary widely from day to day. If tests are being done inside a metal building, then the walls and roof of the building would contribute to the eddy current loss.

Many coils have a strike ring located around their base, to intercept sparks before hitting the feed line or other components. There is general agreement in the Tesla coil community that this ring should be open rather than shorted. There have been observations where a shorted copper ring has significantly degraded spark length. A spun aluminum toroid also presents a shorted path for eddy currents, but there is less agreement that this represents a significant loss. It is argued that conduction currents are smaller at the top of the coil, therefore induced currents must be less.

I built a toroid of 0.25 inch copper tubing pieces on insulating disks, connected together at one point by a conducting disk. The ends were placed into heat shrink tubing, which was then shrunk to hold the ends a fixed small distance apart. This toroid was then compared with a

spun aluminum toroid of similar capacitance, and also with a smaller toroid made of one inch copper tubing with diameter slightly greater than that of the coil form. The smaller toroid was an attempt to get a shorted turn as near to the coil as possible. It lacked the capacitance to be an effective toroid for long sparks, of course. I could not find any significant difference in input impedance between the insulated toroid and the spun aluminum toroid. The shorted copper ring, however, had about 10% higher input impedance than the toroid that was not a shorted turn. This suggests that you would not notice any improvement if you cut your beautiful spun aluminum toroid into pieces to eliminate eddy currents. The effect is there, and can be measured if one really works at it, but is not that significant in most situations.

Overall, my tests indicated that R_{eddy} is no more than a few percent of R_{TC} . If a little thought is given to separation of conducting materials from the immediate vicinity of the coil, eddy current losses can be ignored. Likewise in all my tests, R_{rad} is very close to zero. I was unable to detect a signal from the coil more than perhaps 100 m away. At worst, it would be a number like $0.01\ \Omega$, which is a negligible portion of a typical measured resistance of 25 to 50 Ω .

6 Conclusion

I believe that Fig. 8 is a reasonable model for a Tesla coil. There is scientific basis for calculating (or estimating) R_M , R_{cfs} , R_{wis1} , and R_{wis2} , and for ignoring R_{eddy} and R_{rad} . It has the proper indication for changes in the model when a spark occurs. The model indicates that when a spark occurs, the equivalent resistance R_{spark} increases from zero to some finite value, so the input resistance *increases* during a spark. This is exactly what happens experimentally.

Unfortunately, great precision is difficult to impossible to obtain. R_M or R_s can be calculated to within a few ohms given the techniques in this chapter. I consider this a vast improvement over my state of knowledge when I started this project. Trying to get more accuracy is probably not warranted because of the strong influence of moisture on coil resistance. If R_M is $50 \pm 5\ \Omega$, and R_{die} might vary from 0 to 5 Ω or more as humidity goes from 0 to 100%, there is little point in reducing the uncertainty on R_M .

In my opinion, a complete distributed model will not be any better in dealing with skin effect, proximity effect, and dielectric losses, and would certainly be more of a programming problem. The one thing that this lumped model cannot deal with directly is the displacement current effect. A distributed model can determine the actual current distribution, which can then be used to find a predicted effective resistance of the coil.

Both approaches (lumped and distributed) have advantages. I believe the lumped approach is better at determining resistance. However, the lumped approach will not show anything about resonances at harmonic frequencies, and cannot deal with things like the current distribution. Hopefully there will be peaceful coexistence, where each method will be

used to its full advantage.

References

- [1] Benson, Barry, bensonbd@erols.com, Private communication, November, 2001.
- [2] Fraga, E., C. Prados, and D.-X. Chen, “Practical Model and Calculation of AC Resistance of Long Solenoids”, *IEEE Transactions on Magnetics*, Vol. 34, No. 1, January, 1998, pp. 205–212.
- [3] Harrington, R. F., *Time-Harmonic Electromagnetic Fields*, McGraw-Hill, New York, 1961. p.455
- [4] Jahnke, Eugene and Fritz Emde, *Tables of Functions*, Dover, 1945.
- [5] Medhurst, R. G., “H.F. Resistance and Self-Capacitance of Single-Layer Solenoids”, *Wireless Engineer*, February, 1947, pp. 35-43, and March, 1947, pp. 80-92.
- [6] Ramo, Simon, John R. Whinnery, and Theodore Van Duzer, *Fields and Waves in Communication Electronics*, Second Edition, John Wiley, 1984.
- [7] Smythe, William R., *Static and Dynamic Electricity*, Hemisphere Publishing Corporation, A member of the Taylor & Francis Group, New York, Third Edition, Revised Printing, 1989.
- [8] Terman, Frederick Emmons, *Radio Engineers Handbook*, McGraw-Hill, 1943.

TESLA COIL DRIVER

In this chapter, we deal with the design of the Tesla coil driver. We will discuss several components or subsystems, including the power supply, the controller, the gate driver and inverter, the fast fuse, and the current sense resistors. There are a limited number of interconnections between the subsystems. We will try to be clear as to what the inputs and outputs are for each subsystem.

1 Power Supply

Initially, I used a circuit without a step-up transformer. I could then apply up to ± 170 V to the inverter. This allowed for testing of many concepts, but the voltage was not high enough to make significant sparks. I then decided to use a step-up transformer ahead of the inverter. It appeared that a factor of ten ($10(\pm 170) = \pm 1700$ V) would give me enough range. So I bought a 240/2400 transformer from T&R Electric. This is an oil filled transformer, sometimes called a pole pig. Some older distribution systems used a 2400 V delta three phase configuration. These are typically being upgraded to 7200 V wye, so the 240/2400 V transformers should be available on the used market for some time. I got one rated at 5 kVA, so the rated input current is $5000/240 = 20.8$ A. This is a very conservative rating, for continuous operation in a 40°C ambient. I would expect it to run four hours at 30 A without a problem, and 10 minutes at 40 A. The high side rating is 2400 V and 2.08 A.

The power circuit is shown in the figure Fig. 1. The part ahead of the 240/2400 V transformer T1 is very common with classical Tesla coils. There is a switch SW1, which is actually a 50 A circuit breaker back at the main electrical service entrance. Then SW2 is a 40 A breaker in a metal box near the transformer. Inside the metal box is a power line filter which hopefully keeps some of the trash from getting back to the utility (or affecting other circuits in my lab). There are two inductors consisting of four turns of heavy wire on iron powder toroids, about $80\ \mu\text{H}$ as measured on an Escort ELC-120 LCR meter. Then I used $0.039\ \mu\text{F}$, 1600 V polyethylene capacitors on both sides of the inductors. These capacitors have a self resonant frequency of about 5 MHz. Most of my testing was done at frequencies between 120 kHz and 240 kHz, so the capacitors should help for the fundamental and several harmonics. But really high frequency RFI will go right on by. Electrical wiring is very lossy for frequencies above a few MHz, so I would not expect a problem, and indeed have not experienced anything.

I added a $6\ \mu\text{F}$, 370 VAC motor run capacitor to provide more filtering for the low frequency noise. There is also a MOV (Metal Oxide Varistor), the 20K43IU, to clip any really nasty pulses. The circle marked PL is a red panel light that comes on when the 40 A breaker is turned on. This is one of many safety features that could be added.




Figure 1: Tesla Coil Power Supply

The transformer output is rectified by a full-wave bridge rectifier made from 6 kV diodes I found in my junk box. The current rating is unknown, but almost certainly inadequate to charge the capacitor bank if full transformer voltage is abruptly supplied to discharged capacitors. The variac is always turned to zero before SW2 is turned off, and is slowly brought to the desired setting after SW2 is turned on.

Current through a rectifier into a capacitor bank tends to be very peaked unless a smoothing inductor is used. I used powdered iron toroids, the Amidon T520-26, with 200 turns of 20 ga enameled wire on each one, for the inductor. One toroid had an inductance of 14 mH. Four in series were used to get a total of 56 mH. A reviewer asked how I arrived at this inductance. I was thumbing through a power electronics book [1] and found an example where average voltage was plotted against the product of inductance and load current for a 230 VAC input. The peak voltage would be $230\sqrt{2} = 325.27$ V. They selected an operating point of 288.4 V, which required a LI product of 15 mH-A. In my case I had 2400 VAC rather than 230 VAC, which made my $LI = (2400/230)(15) = 156.5$ mH-A. The rated current for my transformer is 2.8 A at 2400 V, so

$$L = \frac{156.5}{2.8} = 56 \text{ mH}$$

The circuit will work without any inductor (if peak current for my junkbox diodes is not exceeded) and will work with two or three times this much inductance. It really is not very critical. One has to watch out for saturation and voltage withstand capability, of course. This particular inductor had a measured inductance of 45 mH for 1 A of current, 55 mH for 2 A, and 52 mH for 4 A, which I considered close enough to the design value of 56 mH.

The capacitor bank was formed of 16 electrolytic capacitors rated at 1400 μ F and 450 V. Four were placed in series to get a string rated at 350 μ F and 1800 V. Two strings were then paralleled to get a rated capacitor C_3 of 700 μ F and 1800 V. Two more strings were added to get another capacitor C_3 for the negative supply. The figure shows the measured value of the capacitors, 740 μ F rather than 700 μ F.

Finally, resistors R_1 were added to discharge the capacitor bank. These are 35 k Ω , 100 W resistors, two in series on each side of the supply. At full voltage of ± 1700 V, the power dissipation on each side is

$$P = \frac{V^2}{R} = \frac{(1700)^2}{70000} = 41 \text{ W} \quad (1)$$

or about 20 W per resistor. Normally, I do not exceed ± 1200 V, or a power of about 10 W per resistor. The resistors are mounted at an angle in free air, so air can flow by convection through the interior of the resistors. No evidence of overheating (discoloration, odor, smoke) has been noticed.

The nominal rated voltage difference between $V-$ and $V+$ is $2400\sqrt{2} = 3394$ V, or about ± 1700 V. It appears the maximum voltage available under load at the inverter output will be somewhat less than this amount, in the range of ± 1500 to ± 1600 V. It turns out that ± 1200 V is adequate to get interesting sparks, so most of my testing is done at that level or below.

Power was run into and out of the capacitor bank with wire made by the Alpha Wire Corporation, No. 1638, 18 ga, with 0.090 inch rubber insulation. There are three wires, one each for $V+$, $V-$, and neutral. The transformer provides electrical isolation, so this neutral is not the same as the utility neutral. The neutral is connected to a buried ground plane through a current sense resistor (shown on a later figure) so it is not precisely at ground potential. The neutral must be kept electrically isolated from earth and from any connection that might yield a ground loop.

2 Gate Driver And Inverter

There have been several generations of driver and inverter. One early generation used two power MOSFETs, the IRF450, rated at 500 V and 13 A, in a half bridge configuration. The 240 VAC line and a variac were used to drive the rectifier without a step-up transformer. The maximum rms voltage at the output of the MOSFETs was about 170 V in this configuration. (Note that a square wave of amplitude ± 170 V has an rms value of 170 V). I could certainly measure input impedance of the Tesla coil under nonspark conditions with this voltage, but could not do much in the way of producing sparks.

A following generation had a 3:1 high frequency autotransformer at the MOSFET output, which raised the voltage to $3(170) = 510$ V. This allowed me to produce sparks up to 10 inches long. Most of the time, I had to put a small bump on the toroid to get breakout. I could now measure input impedance under sparking conditions, but I did not have enough voltage range to see how impedance varied with voltage. I could not look at spark length as a function of voltage either. The leakage inductance of the autotransformer, only a few μH , also caused some ringing problems with stray capacitances in the circuit.

The last generation is shown in Fig. 2. I needed more voltage (and power) than was readily

available with MOSFETs so I decided to try IGBTs (Insulated Gate Bipolar Transistors) in series. There are four Harris HGTG18N120BND IGBTs in series in each leg. These are nominally rated at 1200 V and 18 A. The 1200 V is a rating that we *never* want to exceed. However, the device will switch up to 100 A until losses cause it to overheat. In a Tesla coil application where we operate in single-shot mode or in a low duty cycle, say, 10 ms on and 1 sec off, the device should function with rms currents up to 50 or 60 A. At this writing, the inverter has withstood 30 A rms for short periods without difficulty.

When one string of four IGBTs is on, the other string sees the entire voltage, plus to minus, of the capacitor bank. If we were operating at the capacitor rating of ± 1700 V, the four IGBTs must withstand $(2)(1700) = 3400$ V, or 950 V per device. Voltages will never divide evenly, and there will always be some overshoot or ringing due to stray inductance, so this is about the limit of a safe design.

Signals to turn on a string of IGBTs arrive from the controller by way of the fast fuse (which includes the CD4073, a three-input AND gate). R_1 , the $43\ \Omega$ resistor at the left was selected experimentally to get a fast rise time with minimal overshoot.

There are differences among components which will cause the IGBTs in a string to turn on and off at slightly different times. This causes the voltages to not be shared equally. Even a few ns makes a significant difference. At one time I tried to deal with the turn-on and turn-off times separately. I put 1N4151 diodes across the $1\ \text{k}\Omega$ pots to allow all IGBTs to be turned on as fast as possible. The $1\ \text{k}\Omega$ pots then allowed the turn-off times to be adjusted between IGBTs. I later removed them when I could not see much benefit. Without the diodes, the $1\ \text{k}\Omega$ pots slow down both the turn-on and turn-off when the resistance is increased.

The 4420 is a non-inverting driver for gates of MOSFETs and IGBTs. It can supply up to 6 A while charging a capacitive load. Each gate is charged through a pulse transformer. I used a ferrite toroid with 5 turns of 20 ga. magnet wire for the primary and 10 turns for the secondary. I ordered the Philips TX22/14/6.4-3F3 (former part number 846XT250-3F3) but I think I received the TX22/14/6.4-3E2A. Most ferrite toroids I have seen have no markings on them whatsoever. You have to keep the toroids in a container with the description. However, Philips coats the different ferrite materials with different colored epoxy paints. The 3F3 material is painted orange while the 3E2A material is painted white, and I got white toroids. The materials are not vastly different, and I suspect either one would work in this application.

The numbers 22, 14, and 6.4 in the part number refer to the outside diameter in mm, the inside diameter, and the height or thickness of the toroid, respectively. The windings need to be physically separated enough to withstand up to 1700 V, so I put the primary on one side of the toroid and the secondary on the other side. The leakage inductance would be smaller if the windings are on top on one another, but then I would need to place additional insulation between windings. This can be a challenge in relatively small toroids, so I opted for air insulation between totally separate windings.




Figure 2: Gate Drive and IGBT Circuit

The gate needs at least 10 V to turn on, and 12 to 15 V is better, to assure some noise immunity. The gate voltage should not exceed ± 20 V on a repetitive basis, and the absolute maximum is ± 30 V. When the first pulse is applied, the capacitor C_7 is discharged, and its voltage does not change instantaneously. If 15 V is applied to the left capacitor terminal, the right side of the capacitor will also try to have 15 V on it. However, the two back-to-back 1N4737 zeners will clamp the voltage at about 8 V. There will then be 7 V across the impedance of the 4420, the resistance of circuit traces, and C_7 . The pulse transformer will put out about double the input voltage or 16 V. If anything else is happening, the bidirection zener, P6KE18, clamps the gate voltage at 18 V.

While current is flowing out the pulse transformer secondary to the IGBT gate and the bidirectional zener, it is also flowing into C_7 . With a pulse train of approximately 50% duty cycle, C_7 will charge to 7.5 V after the first few cycles. Once C_7 is charged, voltage excursions of +15 to 0 at the 4420 output will be seen as a pulse train of ± 7.5 V at the pulse transformer input. The gate would then see ± 15 V. The bidirectional zener and the 1N4737 zeners will not conduct after the first few cycles.

There are many tradeoffs in the gate driver circuit. C_7 and the leakage inductance form a resonant circuit which is excited by the first pulse. As energy trades back and forth between C_7 and the inductance, the gate voltage can drop below the threshold value, causing the IGBT to try to turn off when it should be on. This effect seems to be helped by keeping the leakage inductance small, that is, by having as few turns as possible on the pulse transformer. On the other hand, the time that a pulse transformer can hold a voltage is proportional to the number of turns, so for low frequency operation, one needs many turns.

One should not use ceramic capacitors for C_7 . Some of them change capacitance value with temperature, and the initial pulse will heat the capacitor enough to cause a thermal resonance in the gate voltage in the 3 to 10 kHz range. A polyethylene capacitor rated at 50 V was much better than a ceramic capacitor.

If it were not for the initial transient, design of the driver would be very straightforward. I could not find a book or paper which discussed the design of this particular configuration. Modeling would be difficult because the current and voltage output of the 4420 under overload conditions the first pulse would be strongly dependent on the exact construction of the printed circuit board. The leakage inductance of the pulse transformer would vary with the specific toroid and also with the hand winding technique. I did tests with various values of C_7 and R_9 , and basically quit when I got something to work. There could easily be much better designs than this one.

Circuit board layout of the drivers and IGBTs is also critical to success. The system needs to be as compact as possible, consistent with adequate heat sinking and voltage separations. The three dimensional aspects are difficult to present in two dimensional drawings, so I will attempt a verbal description. All the components shown in Fig. 2 are mounted on a double sided PC board, approximately 6 by 6 inches. Four IGBTs are mounted along one edge, and their 4420 drivers toward the center. The other drivers are located the other side of center

and the other IGBTs on the opposite edge. Drivers and IGBTs are on the ground side of the PC board, with their pins extending through to the trace side. The ground foil is only on the interior of the board, under and around the drivers. The IGBTs are mounted on separate heat sinks and placed on the board where the ground foil has been removed. The heat sinks are finned black aluminum, approximately 1 by 1.6 by 2 inches outside dimensions. There are two pins about an inch apart that go through the board along with the IGBT pins to provide some mechanical stability.

Conventional two-sided board practices are used in the interior of the board. One gets to ground by simply going through a via. The board becomes a single sided board when we get to the outside, where the heat sinks and IGBTs are located. The board itself will easily handle the maximum of ± 1700 V in this application, but holes through the board could compromise this integrity. So the board in the heat sink region is treated as a surface mount board. Components like C_7 , R_9 , D_2 , and the pulse transformer are soldered to pads on the board rather than placed in holes drilled through the board and then soldered. There are obviously many ways to lay out a board, and all designers have their own preferences, but this particular layout seems like a reasonable approach.

The IGBTs are mounted on their heat sinks without an insulating pad. This maximizes the heat transfer, but with the disadvantage that the heat sinks are electrically charged. I think one set of IGBTs was destroyed due to a discharge from a heat sink to the grounded case. I now set the PC board, oriented with heat sinks down, on a 1/16 inch thick layer of polyethylene lying on the bottom of the case.

Next to the MOSFET symbol in Fig. 2 is an upward directed diode. This is built into these particular IGBTs. These diodes are an important part of the circuit. They allow for current to continue flowing after one set of IGBTs turns off and before the other set turns on. They also allow for energy stored in the coil to flow back into the capacitors when the controller is turned off and all gate pulses cease. The waveforms of voltage and current look almost the same after the controller is turned off, except for a 180° phase shift in voltage. While the gate pulses are applied, power flows from IGBTs to Tesla coil. When pulses are removed, power flows from coil to IGBTs.

During inverter operation, voltages divide across the IGBTs according to the integral capacitances from drain to source. Adding some capacitance helps to swamp out individual differences among IGBTs. The energy stored on these capacitors has to be dumped through the IGBTs when they are turned on, so the capacitors need to be kept small. At one time, I put 270 pF across each IGBT, but eventually removed them, except for the case where one IGBT takes more than its share of voltage. This means that its integral capacitance is smaller than the others. In such a case, the voltage across the IGBT that is taking more than its share can be reduced by adding some capacitance in parallel to that IGBT, say 50 or 100 pF.

A capacitor rating of 1000 V should be adequate. A rating of 1000 V is less than the IGBT rating of 1200 V, however, the capacitors have more derating built in than the IGBTs. The IGBTs will be failing quickly at 1300 V, while 1000 V capacitors would probably survive

in this application at 2000 V or more.

Static voltage balance during long off periods is achieved by resistors R_{11} . Without them, the drain-to-source capacitances will charge to random values from the small (and different) leakage currents of the IGBTs. A 430 k Ω resistor has a much greater current than the IGBT leakage current, so static voltage balance is quite good. Each resistor carries 1/8 of the total supply voltage during off periods. During inverter operation, each resistor carries 1/4 of the total supply voltage while its IGBT is off, and 0 while its IGBT is on. The power dissipation for a total voltage of 3400 V is

$$P = \frac{V^2}{R} = \frac{(3400/8)^2}{430000} = 0.42 \text{ W} \quad (2)$$

during the off period, and double this amount, or 0.84 W during inverter operation. I use 2 W resistors and have not noticed any problem with them.

Shown to the right of the 430 k Ω resistors is a long string of 1 M Ω resistors and 180 pF capacitors. This is a voltage divider which outputs one tenth of the input voltage when a X10 scope probe is attached. My scope can measure up to ± 200 V at the tip of its X10 probe, so with this divider I can measure up to ± 2000 V at the input to the Tesla coil.

The divider is a resistive divider at dc and low frequencies, and a capacitive divider at higher frequencies. At 160 kHz, for example, the capacitive reactance of a 180 pF capacitor is about 5500 Ω , much less than the 1 M Ω resistors. If an rms voltage of 1700 V was applied to this capacitor chain, the current would be about 18 mA and the total reactive power would be about 31 VA. In many applications this would be too great a load for the source, but for this situation it should be acceptable.

I bought a X100 probe before I built the divider, hoping the divider would be unnecessary. However, when I received the probe and looked at the specifications, I realized it would not work because it was rated for only 450 VAC.

3 Current Sense Resistors

In the metal case containing the inverter, there are several other components besides those shown in Fig. 2. These include the current sense resistors, a filter for the current waveform to the scope, a small current transformer for supplying current information to the controller, and some high frequency capacitors for voltage and current support. These are shown in Fig. 3. The switches SW1-SW8 represent the 8 IGBTs in the inverter.

When the IGBT string (SW1-SW4 or SW5-SW8) turns on, the circuit needs to be able to supply tens of amperes in a time of 2 or 3 μ s, if the Tesla coil is operating in the range of 200 kHz. There is no way that electrolytic capacitors shown in Fig. 1 can do this, so additional high frequency capacitors are installed close to the IGBTs to support the voltage. I used two strings of capacitors on each side, one string made of 15 polyester capacitors rated at 0.22 μ F




Figure 3: Support Functions for IGBT Circuit

and 1250 V, connected in parallel, and the other string made of 15 polyethylene capacitors rated at $0.039 \mu\text{F}$ and 1600 V, likewise connected in parallel. The result was approximately $3.9 \mu\text{F}$ on each leg of the power supply. The polyester capacitors will be operated above rating at high variac settings but this should be within the safety factor of the capacitors. Polyester capacitors are more lossy than polyethylene capacitors, so continuous operation at 1250 V and a high ripple factor could cause these capacitors to overheat. Like several other components, we are hopeful that a low duty cycle will keep these capacitors from overheating.

The switcher provides a square wave voltage to the Tesla coil input. A square wave can always be composed into a fundamental and a series of odd harmonics of sine waves. The magnifier is not resonant at exactly three (or five, or seven) times the fundamental frequency, so the harmonics always face a very high surge impedance. The current will build up at the resonant frequency but not at the harmonics. This means that a square wave of applied voltage will produce only a sine wave of current. This sine wave will be in phase with the voltage at resonance, will lag above resonance, and lead below resonance.

IGBT switching at resonance means the IGBTs turn on and off when the current through them is small. This is a desirable feature, in that it reduces the losses in the IGBTs. It cannot be too small, however.

It is often not emphasized in power electronics books, but power MOSFETs and IGBTs in this circuit configuration essentially require an inductive load. With a lagging current, when one string of IGBTs is turned off, current will continue to flow, charging and discharging the

drain-to-source capacitances, until the voltage across the opposite string has been reduced to zero. If the voltage reaches zero while the devices are off, current will start to flow through the diodes of the other string. When a gate pulse arrives at those IGBTs, the devices already have zero voltage across them, so the turn-on process works very smoothly. On the other hand, a leading current will result in each string turning on with full voltage across them. Very large transient currents will flow, with steep rise times. Stray inductances will cause substantial ringing in voltage due to these transient currents. A leading current is a disaster on its way to happen in an inverter circuit like this.

We need to be able to tune the Tesla coil across resonance. One way to assure the IGBTs are always operating into a lagging load, even if the Tesla coil happens to be drawing a leading current, is to add an additional inductance in parallel with the Tesla coil. This is shown as L_1 in Fig. 3. This inductor current also provides for relatively soft switching of the IGBTs when the Tesla coil current is very small.

The difference between hard and soft switching is shown in Fig. 4. We have two plots of an (approximately) square wave of voltage applied to a Tesla coil, with the resulting (approximately) sinusoidal current. In the Dec 19 plot, the voltage takes about 300 ns to make the transition. There is little ripple and the current waveform is reasonably smooth. In the Sep 30 plot, however, the voltage transition lasts only about 100 ns. This means that the voltage waveform contains more high frequency components, which can excite various resonances in the Tesla coil itself, the transmission line, and the IGBT driver. There is some ripple on the voltage and more on the current. The Dec 19 plot is obviously the preferred method of operation.

There are two necessary conditions to get soft switching. The current flowing through one bank of IGBTs must be sufficient at the time of turn-off, and the dead band (the time lapse before the other bank of IGBTs turn on) must be adequate. The controller chip (discussed later) can have a dead band as low as 80 ns, but is set for about 400 ns in the Dec 19 plot.

Assume in Fig. 3 that string SW5-SW8 have been on, the Tesla coil current is small, and inductor current is flowing to the left through SW5-SW8. The Tesla coil voltage is V_- . IGBTs SW5-SW8 are now turned off. Current continues to flow to the left, charging the drain-to-source capacitances of SW5-SW8. A similar current flows to the left in IGBTs SW1-SW4, which has the effect of discharging their drain-to-source capacitances. An (approximately) constant current flowing into (approximately) constant capacitances will cause the voltage to slew at an (approximately) constant rate.

If the current is too small, the voltage will not have reached V_+ when SW1-SW4 turn on. The remainder of the transition will be very rapid, with associated ripple. On the other hand, if the current is too large, the entire transition will also be rapid, with similar results. Actually, the Sep 30 plot looks like a case where the inductor current is larger than necessary. When an IGBT string is turned on by the gates during a voltage transition, the remaining transition will be nearly instantaneous.




Figure 4: Soft switching (Dec 19), harder switching (Sep 30).

There is therefore a range of inductance values that will yield acceptable results. For these IGBTs operating at 150 kHz, 122 μH and 210 μH worked. A value of 317 μH would give even nicer waveforms but tuning was very critical. Any small adjustment in phase would yield really ugly waveforms.

Note that the peak inductor current is the critical factor, not the Tesla coil impedance. At higher frequencies, the peak current in a given inductor will decrease, requiring the use of a smaller inductor. If 200 μH works at 150 kHz, then we would expect a 300 kHz coil to need a 100 μH inductor. A coil wound on a ferrite toroid would be an attractive solution to the need for an inductor except for two problems: The number of turns would have to be large in order to avoid saturation, and we would need to build several such inductors, swapping them in and out of the circuit until we found one with approximately the proper inductance. I decided to build a variable inductor instead. I wound 140 turns of 22 gauge yellow insulated hookup wire on a length of 3/4" CPVC pipe. The pipe O.D. is approximately 0.9". The winding length is about 8.4". I then placed three pieces of Amidon R33-050-300 ferrite rods inside a length of 1/2" CPVC pipe. The rods are 0.5" O.D. and 3" long. The O.D. of the rods and the I.D. of the pipe are the same, so the rods fit snugly. Inductance is then varied by sliding the smaller pipe inside the larger one. Minimum inductance is about 54 μH and the maximum is about 1240 μH .

Returning to our discussion of Fig. 3, the current to the Tesla coil is measured as the voltage across R_1 , a 0.02 Ω resistor. R_1 is fabricated from 11 surface mount resistors, 0.22 Ω and 2 W each. They are mounted in parallel across a gap in two sections of 2 inch wide copper foil. Where possible, the ground path is made of wide conductors to lower the stray inductance of the circuit.

When a voltage pulse train of resonant frequency is applied to the Tesla coil, current will build up until the coil contains enough energy to drive the spark. The current amplitude at spark remains relatively constant as input voltage is increased further. A ± 1200 V input will have about the same maximum current as a ± 900 V input. A given value of current will be reached more quickly with a higher voltage of course. A spark that occurs 1.9 ms into the pulse train with ± 900 V applied might occur 1.2 ms after start with ± 1200 V applied. A current of about 30 A (rms) just before spark was observed frequently. The power dissipated in R_1 with this current in it is

$$P = I^2 R = (30)^2 (0.02) = 18 \text{ W} \quad (3)$$

which is within the range of a 22 W resistor. The most critical limit in the circuit is that of the IGBTs, which is 100 A peak in short bursts. If the peak current is 100 A in a sinusoidal waveform, then the rms current limit is 70 A. At 70 A, the power dissipation in R_1 is

$$P = (70)^2 (0.02) = 98 \text{ W} \quad (4)$$

which is a little over four times the steady state rating. For the intermittent operation used here, R_1 should last indefinitely, at least until the IGBTs blow. If the IGBTs fail in a shorted

state, then the stored energy in the capacitor bank is available to turn the $0.22\ \Omega$ resistors, traces on the printed circuit board, and anything else in the way, into fuses. An IGBT that fails shorted will eventually be blown apart and become an open circuit. A metal cover over the inverter is important at high power testing to keep shrapnel from flying around the room. Hearing protection is not a bad idea either.

A current transformer is mounted on the ‘hot’ lead of the Tesla coil, opposite R_1 . I made it from two stacked ferrite toroids, the TX22/14/6.4-3E2A mentioned earlier. The toroids fit over a straight 14 ga wire connecting the inverter to the transmission line. This wire forms a single turn primary. There are 25 turns of 20 ga magnet wire wound on the toroids. If the two ends of this winding were shorted, and if everything were properly designed, a current of 25 A to the Tesla coil should result in a 1 A current in this winding.

We already have a means of measuring Tesla coil current, with R_1 . What is needed is a voltage waveform at the controller that looks just like the current waveform, except cleaner and with sharper transitions. That is, we want to convert the sine wave in Fig. 4 into a square wave. That current wave has a peak value on the order of 100 mV, far too low to clean up with diodes. We need a square wave with amplitude of a volt or two.

A current transformer is always operated with a short across the secondary for two reasons. One is linearity, which we are not particularly interested in here. The other is that if the secondary is open circuited, the secondary voltage can rise to extremely high values, enough to damage the insulation in current transformers built for 60 Hz service. The second reason is used to advantage here. The secondary is loaded with two chains of two diodes each, the chains pointed in opposite directions. As the Tesla coil current passes through zero, the secondary voltage also goes through zero but very rapidly. Two diodes in series have a very high impedance when the voltage across each diode is less than 0.5 V or so. Then while the Tesla coil current is still near zero, the secondary voltage will be clamped at two diode drops, say 1.4 V. The secondary voltage is not a perfect flat top square wave, but will reach 1 V for a Tesla coil current of hundreds of mA, and then rise slowly to 1.4 V or a little more for Tesla coil currents of tens of A. This approach works surprisingly well for the purpose.

There are two additional resistors in Fig. 3, labeled R_2 , that are used for overcurrent protection discussed in the next section. These resistors are made of eight $0.22\ \Omega\ 2\text{ W}$ resistors in parallel. When SW1-4 are on, current flows to the left through R_1 and up through the R_2 attached to x . When SW5-8 are on, current flows to the right through R_1 and up through the lower R_2 . With these current directions, point x will be negative with respect to ground while point y will be positive.

As mentioned, good sparks are obtained with a Tesla coil current on the order of 30 A, while the IGBTs are rated for 100 A in short bursts. We will select 70 A as the overcurrent set point. This will hopefully be high enough to prevent nuisance trips and low enough to protect the IGBTs. At 70 A, the voltage at y will be $(70)(0.0475) = 3.32\text{ V}$, and a -3.32 V at x .

4 IGBT Overcurrent Protection

After one explosion of my inverter, I decided to add overcurrent protection. It is well known in the power electronics business that silicon is a faster acting fuse than any traditional device sold as a fuse. That is, a ten dollar IGBT will always act to protect a ten cent fuse. At least, it seems that way. My only hope seemed to be to remove the gate pulses to the IGBTs as rapidly as possible.

The circuit used is shown in Fig. 5. Gate pulses from the controller (discussed in the next section) are brought to the inverter board through three input AND gates, the 4073. Current through IGBTs 1-4 is measured separately from current through IGBTs 5-8. If either current exceeds a pre-set limit, one input to each three input AND gate will go low, interrupting all gate pulses immediately.




Figure 5: IGBT Overcurrent Protection

Starting at the top left of Fig. 5, we see a voltage divider chain, R_1 , R_2 , and R_3 . Zener D_1 sets the top of R_2 at 12 V. During fault events and even during normal operation, the +15V line is likely to be noisy, so this adds some noise suppression. Pin 5 of the LM319 comparator will be held at about 1/3 of 12 V, or about 4 V. When point y exceeds 4 V, the comparator output goes high, resetting the 4013 and causing its output at pin 13 to go low.

The negative going voltage at x causes more of a problem for a system that uses only +15 V and ground. Instead of setting the inverting input at a fixed value and waiting for the noninverting input to go higher than the inverting input, we set the noninverting input and wait for the inverting input to go lower. D_2 is a nominal 3.9 V zener, so pin 9 is held at this value. Pin 10 is held to a value not over 6.8 V, the rating of D_5 . As x goes below zero, the voltage at pin 10 drops, making D_5 effectively into an open circuit. The voltage on pin 10 is then determined by voltage division across R_7 and R_5 .

Resistors R_{8Y} and R_{8X} are pullup resistors required by this comparator. C_2 provides a small amount of noise suppression. C_1 and R_9 provide an effective short pulse at startup to make sure each 4013 flipflop wakes up with output Q high. D_3 and D_4 help to keep what should be a circuit trace with a digital signal of 0 to +15 V between -0.3 and about +15 V in a noisy environment.


When this circuit acts, the +15 V supply must be turned off and back on to reset the fuse action.

5 Tesla Coil Controller

The controller described in Fig. 6 has taken an amazing amount of time to get to its present state. I have worked on it intermittently for ten years, visiting with many students and faculty along the way, before getting something that basically does the job. It still seems that it should not be that hard. I fully expect some reader to glance at the circuit, get the basic concept, and redesign to achieve something far superior, in a short period of time. But that is the way of science.

The heart of the circuit is the Motorola 34066. It puts out complementary gate drive pulses on pins 12 and 14 at a frequency proportional to the current into pin 3. The gate pulses have a minimum dead band of about 80 ns (so both strings of solid state switches cannot be on at the same time). The dead band can be increased by the resistor R_{DT} . Initially, I used $R_{DT} = 0$ but eventually I put in a 1 k Ω pot for R_{DT} . I would adjust this pot to get minimum transition noise in the current flow to the Tesla coil, which probably occurred at 200-300 ns. A dead band of up to 400 ns will allow for somewhat more variation in other parameters.

The 34066 is turned on by an ENABLE signal at pin 9. The ENABLE signal is supplied in one of two ways: a pulse from a momentary contact switch SW1 or a pulse from an external pulse generator to the port marked BNC. The switch SW1 sets the 4013 flipflop so the 34066 is on until power is removed. The BNC port allows the controller to be turned on and off for specified periods as determined by the pulse generator. With adequate voltage applied, a spark from the Tesla coil would occur in typically 0.5 to 3 ms, so the on period was typically set at 4 to 6 ms. The off period would then be set at 1 second, so we would have 1 pps (pulse per second). On occasion, I would adjust that up to 10 pps, to observe the effect.



Solid State Tesla Coil by Dr. Gary L. Johnson
Figure 6: Tesla Coil Controller

December 27, 2001

It is important that the 4013 flipflop U1A between switch SW1 and the 34066 always wakes up with Q in the low state. Otherwise the controller may start operation immediately after turn on, with possible negative consequences. C_6 keeps the Set low and C_5 and R_{46} keep the Reset of the 4013 high during turn on, so the chip always starts in the right mode.

One of the 34066 outputs (pin 12) is used for a clock pulse into U1B, the other half of the 4013. The data pin is always high, so Q is set high every clock pulse if ENABLE is high. If ENABLE is low, Reset is always high, which holds Q low.

There are two modes of operation, manual or automatic, depending on SW2. In manual, the current into the 34066 pin 3 is determined by the settings of R_{49} and R_{51} . Each mode has two submodes, depending on whether ENABLE is on all the time or just when the pulse generator sends a pulse. Normally in manual, SW1 is also flipped, so we have continuous operation. Current and voltage to the Tesla coil are monitored on a scope, and R_{51} is adjusted to give current in phase with the voltage. This is a resonant point, so current is maximum when it is in phase with the voltage.

The manual mode is not suitable for long term operation. Minor differences in power supply voltage, heating effects in R_{49} and R_{51} , drift in the 34066, the movement of a person watching the sparks, or the change in capacitance due to the space charge of a spark, will cause the frequency to shift by a few hundred hertz. For a high Q Tesla coil, this small shift in frequency can reduce the current by a factor of two, or more. We therefore need a circuit that will monitor the current and automatically adjust the input to pin 3 of the 34066 to keep the Tesla coil current in phase with the voltage. It was this task that turned out to be a challenge.

An important task is to get a voltage waveform into the controller that is in phase with the current. At first I used two Rogowski coils next to the ground lead from the IGBTs. These consisted of many turns of 30 ga wire wrapped around a rectangular block of polyethylene. Each block was 0.5 inch thick by 1 inch wide. The winding length of the coil was about 2 inches. The ground lead is a three inch wide piece of copper foil at this point of the circuit. One coil was pressed against the foil from the bottom, while the other coil was just laid on the foil from above. The two coils were connected in series. Measured inductance of the coils in place was about 400 μH . Since the polyethylene is nonmagnetic, one did not have to worry about saturation and other magnetic effects.

The current signal from the Rogowski coils was brought into the controller enclosure and onto the printed circuit board at the point marked i1, top left corner of Fig. 6. Noise was filtered by R_1 , R_2 , and C_1 , then the signal was limited by the two diodes D_1 and D_2 . Above some minimum current, the voltage across these two diodes would start to look like a square wave, except with sloped sides and rounded shoulders.

The problem with this technique was that the signal looked too much like the current. During buildup, the signal would look sinusoidal, with a peak value well under one diode drop. As the current increased, the slope of the sides of the signal would also increase, changing the apparent phase detected by the following circuit by an excessive amount. The

Rogowski coils also had their own resonant frequencies in the range of a few hundred kHz, which would tend to amplify the noise at those frequencies.

We therefore went to a small current transformer that was described earlier. This current transformer, when terminated by diodes, acts more like a comparator in that a few mV swing on the input would result in maximum swing on the output. It presents a fairly good square wave to port i1 starting with the first pulse to the Tesla coil, many pulses before any current is visible on a scope.

The next step is to convert this small analog signal (plus and minus one diode drop) to a clean digital signal, +15 to zero, a nontrivial task. The solution is clever, even elegant. I wish I could say it was original, but I saw it in an application note. Two pins of the 4046 phase-lock loop are used, one for input and the other for output. There is no feedback. The 4046 converts the analog into a digital signal very nicely. It has an internal biasing network so it will work if a capacitor C3 is used to allow a dc offset. It was my experience that the duty cycle of the resulting square wave was not close enough to 50%, so I added an external bias network R_{17} , R_{18} , and R_{19} . R_{18} needs to be adjusted so the output at pin 2 of the 4046 is near a 50% duty cycle. This is only done once, in initial testing of the circuit.

The current signal is passed through an LM319 comparator and a 4049 inverter to provide additional buffering. R_{14} and C_{12} provide an adjustable amount of phase delay. This delay allows the operator to tune through resonance a significant amount. Without the adjustable delay, one might just be able to tune to resonance, or even not quite to resonance. There is also the option of using another 4049 inverter to get another 180° of delay. One could accomplish the same function by reversing the leads from the current transformer.

U1B, the 4013 flipflop, now has the voltage waveform (actually the gate signal to IGBTs 1-4) applied to the Clock and the current waveform applied to the Reset. Q is set high by the voltage and low by the current. While Q is high, the switch U2C, a 4066, is turned on, causing the capacitor C_{11} to discharge a little through resistor R_{20} . When Q is low, C_{11} charges through the resistors R_9 and R_{35} . One looks at Tesla coil voltage and current on the scope and adjusts R_{20} until the current is approximately in phase with the voltage and has acceptable smoothness. If the current looks like the Sep 30 plot in Fig. 4, I would look for ways to smooth it out.

At resonance, the duty cycle on Q should be on the order of 0.3 to 0.5. As the frequency drifts below resonance, the duty cycle decreases, R_{20} is not connected as much, the voltage on C_{11} increases, increasing the current to pin 3 of the 34066 and thus increasing the frequency. The opposite happens when the frequency drifts above resonance. If the duty cycle at resonance is 0.4, things need to work over a duty cycle range of say 0.2 to 0.6. If we start off with a resonance duty cycle of 0.1 or 0.9, we just do not have the tuning range that we need.

The voltage on C_{11} is buffered through U6, a LM324 op amp in a voltage follower configuration. This is an old, slow op amp, but great speed is not needed in this application. A precision op amp would probably be better, but a quick look failed to turn up a single supply op amp that met all the other requirements. I did not want to add a -15 V supply to the

printed circuit board just for the op amp if this single supply version would do the job, and it does reasonably well.

There are some patches around the op amp that I threw in to get it to work in a minimum fashion. One challenge was that the firing of the IGBTs would shift under load. For example, the controller might be carefully tuned so current was exactly in phase with voltage at an applied voltage of say ± 100 V. The variac setting would be increased so the applied voltage became ± 200 V. The current would now be lagging the voltage by perhaps 10° . The controller was sending out a gate drive signal at the same phase point with respect to the current as before, but the lag between the gate signal and the zero crossing of the voltage waveform to the Tesla coil was changing with current. What was needed was for the current into C_{11} to decrease slightly as Tesla coil current increased, similar to a small adjustment on R_{20} . But it needed to be dynamic, during the growth of Tesla coil current under spark conditions.

The patch was the two op amps at the top of the figure. The voltage across diodes D_1 and D_2 will increase slightly with current. Diode D_{13} pulls off the negative peak, so the junction between D_{13} and R_3 becomes slightly more negative as Tesla coil current increases. The pin 7 output of the inverting op amp will become slightly more positive.

D_4 is a 7.5 V zener that limits the allowable voltage to the op amp and to the 34066.

In one-shot mode, ENABLE is off until the spark is desired. The ENABLE signal is run through D_{10} , R_{34} , an LM 319, and a 4049 inverter to switch U2D, which forces the voltage on C_{11} to be that set on R_{11} when the switch is on. When ENABLE is turned on, the switch turns off in a time determined by the time constant of R_{34} and C_4 . This causes the circuit to always start hunting for resonance from the right starting point. This has to be retuned for each different coil or other adjustment which changes the resonance frequency.

Actually life is even more complicated. I would occasionally see a problem with hunting or oscillation about the resonant frequency. R_{20} would be set correctly, but current would drift from a 20° lag with respect to the applied voltage to a 20° lead and back. In this high Q circuit, this is too much phase shift. The IGBTs would go into a hard switching mode, stray inductances would ring, and the current waveform would get really ugly. To fix this problem, I added the compensating capacitor C_{15} . Now the current would stay in phase with the voltage to within a few degrees. But a new problem popped up. With ENABLE off, pin 3 of the 34066 shows two pn junctions in series to the outside world. With R_{41} in place, the voltage of pin 3 is about one volt. When ENABLE turns on, pin 3 wants to accept current at 1.4 volts. So the pin 3 voltage needs to change by 0.4 volts quickly. However, the voltage across the compensating capacitor C_{15} cannot change instantly. There will therefore be a few cycles where the pin 3 voltage is too low, such that the output frequency will be too low as well.

The patch for this problem is the RC time constant formed by R_{31} and C_4 . When ENABLE goes high, pin 12 of the 4049 stays high for 10 to 20 cycles. R_{11} controls frequency while the 34066 is settling down. Hopefully, the current to the Tesla coil will build up enough during this time that when pin 12 of the 4049 goes low, frequency control is immediately assumed

by a pulse train of about the right duty cycle to the 4066 switch controlling R_{20} .

Another patch is the combination of R_{34} , R_{21} , C_2 , C_9 , and the three switches U9A, U9B, and U9C connected to pin 16 of the 34066. This is the one shot control input of the 34066. Frequency and pulse width output of the 34066 are controlled by external components attached to pins 1 and 2 (oscillator) or pin 16 (one shot). In our case, we select values so the oscillator input does the controlling and the one shot input just floats along, doing basically nothing. The only exception to this is the first two pulses out of the 34066. The IGBT gate drive circuits have capacitors that need to be charged by the first few pulses. One effect of starting operation with a discharged capacitor is that with a standard length gate pulse, the first bank of IGBTs will still be on when the second bank turns on. This forms a short across the main power supply. I lost at least one inverter to the high currents of this short. My fix was to use the one shot input to make the first pulse to each bank of IGBTs shorter than the following pulses.

The patch works in the following manner: When ENABLE goes high, C_2 charges through R_{34} . While the voltage on C_2 is still below the transition voltage of the 4049 inverter U4E, pin 2 of the 4049 will be high and the 4066 switch U9C will be on. This switch connects pin 16 of the 34066 to two variable resistor strings through two other switches, U9A and U9B. One switch is controlled by pin 12 of the 34066 and the other is controlled by pin 14. For the first pulse at pin 12, switches U9C and U9A are on, and the resistor string R_{36} and R_{37} are connected to pin 16 of the 34066. The parallel combination of the resistor string and C_9 determines the discharge time for the one shot function of the 34066. R_{37} is adjusted so the one shot turns off the pin 12 pulse at the desired time, instead of the longer time of the 34066 oscillator function, controlled by R_{45} and C_7 . R_{34} is adjusted so that after the first pulse at pins 12 and 14, switch U9C is turned off. The one shot time constant is now determined by C_9 and R_{21} , which is longer than the oscillator time constant, so this part of the circuit is no longer doing anything.

This patch actually has two functions: As mentioned earlier, slightly shorter pulses allow the gate drive capacitors to charge up without the ugly feature of having both sides of the IGBT inverter on at the same time. And the shorter pulses also improve symmetry of the inductor current. In steady state operation, the IGBTs switch when the inductor current is at a peak. The voltage across the variable inductor is approximately a square wave, so the inductor current is basically a saw tooth wave. This saw tooth should be symmetric about zero. The current for the first pulse starts at zero rather than a negative peak, so if the first pulse is the same length as the following ones, the inductor current will rise to about double the steady-state peak. The current waveform will then drift down with time as the losses in the circuit cause a change toward symmetry. The IGBT switching current is too high for half the IGBTs and too low for the other half during the first few cycles. This can easily cause ringing and IGBT destruction. The current waveform is helped immensely by shortening the first pulse until the peak inductor current is about equal to the long term peak (rather than double). R_{39} is then adjusted so the peak inductor currents during the first few cycles are not drastically different from the steady-state peaks. It is not essential that the peaks be exactly

those of steady state during the first few cycles. Peak values between 70% and 130% of the long term peaks seem to work in an acceptable fashion.

It should be mentioned that driving a digital chip like the 4049 with an analog signal from an RC network is considered poor design. The RC voltage goes through the transition voltage of the inverter slowly, so there will be an interval when the inverter output is neither high nor low. This confuses any following chips. While the control voltage to the 4066 switch is at mid range, the resistance through the switch will be in the $k\Omega$ range, which adds to the resistance of the attached string (i.e. R_{36} and R_{37}), and changes the one shot timing. A better way of doing it is with a comparator, similar to what is shown in the lower right corner of the TCC601 schematic. But poor designs sometime work, as is the case here, and saved me from building yet another PCB with another chip on it.

Yet another patch to deal with a peculiarity of the 34066 is R_{33} , D_9 , and the pulse transformer connected to pin 10, the FAULT input. The 34066 starts immediately when ENABLE is applied, but may put out another hundred pulses after ENABLE is removed. Residual energy stored in the Tesla coil will ring down over a similar time period. If there are no gate pulses on the IGBTs, and if energy is adequate, the IGBT diodes will conduct on the half cycle as ringing occurs. This is not a problem. Power is flowing from the Tesla coil to the power supply rather than the other way around. An increase in voltage on the power supply capacitors can be noted.

If gate pulses continue during the ENABLE off period, when the frequency control has been disabled, we have the possibility of gate pulses getting out of phase such that one bank of IGBTs will turn on while the diodes of the other bank are conducting. This is not as bad as the gates of both sides being on at the same time, but still results in hard switching and massive ringing within the inverter and power supply. This ringing can destroy IGBTs being operated close to their rating. We therefore want the 34066 to turn off when ENABLE goes low.

The 34066 has a FAULT input where a 1 V pulse will shut down the output pulse train within 70 ns. So a positive pulse is needed from the negative going edge of the ENABLE pulse. I first considered using an unused inverter in the 4049. This meant that the FAULT input would be high before ENABLE turned on, and for a few ns afterwards. This would very probably mean that the 34066 would never function.

My next choice was to use a pulse transformer. I had a good supply of the 204XT250-3F3 ferrite toroids on hand, so I wound two windings of 30 ga. wire on a toroid, one about 60 turns and the other about 20 turns. The winding with 60 turns is connected to R_{33} . When ENABLE goes high, a negative pulse appears at the FAULT input, limited by the forward voltage drop across D_9 . When ENABLE goes low, a positive pulse appears at the FAULT input, limited by the zener voltage of D_9 . The 1N4730 is a 3.9 V zener. Voltage limits on the FAULT input are -1 to $+6$ V, so we should be within these limits. No effort was made to do a detailed design of the pulse transformer. Current has to be low, 1 mA in this case, to avoid overloading the 4071. Plenty of turns are needed for this low current. A step down ratio of

say three to one would be appropriate, so 1 mA in the primary would result in approximately 3 mA in the secondary. My first attempt worked, with a peak voltage of about 2 V at the FAULT input, so I did not proceed further.

One other patch, R_{43} and R_{44} between pins 1 and 12 of the 34066, attempts to fix a very subtle problem. C_{11} needs to be small in order to have a high slew rate when a spark occurs. A spark increases the effective size of the Tesla coil top load, which increases the effective secondary capacitance, which lowers the operating frequency. The controller needs to hold frequency at resonance within perhaps 100 Hz (better than 100 parts per million) before spark, and then slew at rates up to 100 Hz per μ s when the spark occurs. This controller will do this, but at the expense of significant ripple on C_{11} . The current into pin 3 of the 34066, that controls frequency, will also have this ripple present.

There will be a double frequency ripple at pin 1, one hump controlling gate pulse length at pin 12 and the following hump the length at pin 14. The ripple on pin 3 will cause one of the gate pulses to be longer than the other by a few percent. The Tesla coil sees a slightly non symmetrical waveform with the time spent at $V+$ somewhat different from the time at $V-$. This is not a particular problem to the Tesla coil, but the inductor across the Tesla coil will integrate this non symmetry over the duration of the ENABLE pulse. Instead of an inductor current varying symmetrically from +1 to -1 A, it might vary from +0.6 to -1.4 A. We might have not enough current to properly switch the IGBTs at one extreme, and too much at the other extreme. The current waveform might show more noise on the positive going edge than the negative going edge, or vice versa.

A value of 1.5 M Ω for $R_{43} + R_{44}$ was found empirically to balance the current into the current mirror of the 34066 and give equal length output pulses. R_{44} , a 500 k Ω pot, allows a small adjustment if it seems necessary. One would look at the inductor with a current probe (Hall effect so it shows the total current, not just the ac component), and adjust the pot until the inductor current is symmetrical about zero.

6 Initial Tuning

People not building a controller board can skip this section. I include it as a personal check list. I lay out the circuit to fit on a double sided printed circuit board, 4.5 by 6 inches, using a software package (HiWire in my case). I etch and drill my own boards, but am thinking about using the experts to do this. I would certainly get a better board to work with.

After soldering on the components, but before mounting the chips in the sockets, I apply +15 and ground to the board. This would detect some solder bridges and some gross errors in component values. For the board TCC601, I measured 11 mA in the +15 V line, which seemed reasonable. I then inserted the chips and applied +15 V again. I measured 39 mA, which still seemed reasonable.

I then put the board into Manual Mode, and push the Enable button, so the board should

be operating at fixed frequency. I then check that there are complementary signals at pins 12 and 14 of the 34066, that R_{DT} changes dead time, and that R_{49} and R_{51} will adjust frequency.

The variable inductor and the Tesla coil are connected to the IGBT inverter, and the variac is adjusted to some low voltage like 6% of rated. In Manual Mode the Enable button is pushed and R_{49}/R_{51} are tuned for resonance. Heating effects on the board will probably prevent one from holding the exact resonant frequency, but we should be able to tune through resonance. With the frequency near resonance, R_{18} is adjusted so pin 2 of the 4046 has a 50% duty cycle.

With the circuit still operating, the duty cycle at pin 13 of the 4013 is checked. At resonance, the duty cycle appeared to be about 30%, and would vary from perhaps 10% to 50% as R_{49}/R_{51} are varied across resonance. If this general condition was not observed, I would try adjusting R_{14} , and failing that, would reverse the input current leads at port i1 or add (or subtract) an inverter stage between the 4046 and the 4013.

With the circuit still operating, measure the voltage at the wiper of R_{51} (2.09 V for a coil operating at 177 kHz). Set R_{11} so the wiper has the same voltage. This should start the circuit at the resonant frequency in Automatic Mode, so the current will build up during the first few cycles.

Turn the board off, switch to Automatic Mode, and apply +15 V pulses to the port marked BNC from a pulse generator in single-shot mode. Adjust R_{37} and R_{39} so the first pulse of both pin 12 and pin 14 is shorter than the following pulses. More precise adjustments will be made later.

Check that pin 12 of the 4049 goes low about 60 μ s after Enable is turned on. This allows current to build up for 10 to 15 cycles under control of R_{11} before control is changed to the Tesla coil current as monitored at port i1.

Check that pin 12 of the 34066 goes low (output pulses stop) when Enable goes low. This confirms that the fault transformer is connected correctly.

In Automatic Mode and with the variac set at about 6% we want to adjust the circuit so current will build up to its steady state value while a short pulse (1 to 5 ms) is applied to the BNC input. Set R_{20} to midrange (say 25 k Ω). Current should build up for a few cycles under control of R_{11} , and then may decline after that due to improper adjustments. Adjust R_8 until it appears current is continuing to build after R_{11} loses control. Look at the Tesla coil voltage and current and continue to adjust R_8 until current and voltage are approximately in phase. R_8 is set only once (for each coil). Phase adjustments during regular operation are made with R_{20} .

R_8 is a part of two LM324 opamp stages which help the controller maintain the proper phase relationship as the TC current builds (or declines during the discharge), and as applied voltage is changed. When the current waveform is a good sinusoid and is in phase with the voltage at steady state, the current waveform should also be a good sinusoid when the current is at 20% of the steady state value, even though the phase between voltage and current will

be somewhat different. The resistors R_{20} , R_8 , and R_{10} all interact, making a good strategy for tuning important. I set R_{20} at midrange, about 25 k Ω , and $R_8 = 0$. I then rig up a test fixture of a 1.5 V battery and a pot and temporarily attach it to the junction of R_3 and R_4 so I can apply a small negative voltage (about 0.2 V) to this junction. I then adjust R_{10} until the voltage at pin 14 of U6 is about 12 V. I remove the test fixture, apply voltage to the IGBTs, and adjust R_{20} for the best current waveform at full current. If the current waveform at 20% current is not acceptable, I increase R_8 to perhaps 10 k Ω or 20 k Ω , put the test fixture back on, and readjust R_{10} for the same voltage as before at pin 14 of U6. Remove the test fixture and try again. This approach should change the phase shift between small currents and large currents without affecting anything else.

For the next step, we remove the Tesla coil entirely, and drive only the variable inductor. First we look at pin 16 of the 34066 and adjust R_{34} until U9C is only on for the first two triangular pulses. This may require some adjustment of R_{37} and R_{39} . The waveform on pin 16 should show a rapid rise to 5.1 V and a linear decline to 3.6 V, followed by a constant voltage of 3.6 V, for the first two pulses. Later pulses should show a rapid rise to 5.1 V and a linear decline to perhaps 4.5 V. If pin 16 does not drop to 3.6 V, then this pin does not control the turnoff of gate pulses.

Now we hook a current probe on a lead to the inductor. Adjust R_{37} and R_{39} until the inductor current is somewhat symmetric. The first pin 14 pulse should terminate when the inductor current is about half the steady state peak-to-peak value. The first pin 12 pulse should terminate when the inductor current is about the same negative value as the previous positive value. The inductor current will not be perfectly symmetric at this point, so adjustments should continue until the inductor current is as close to being symmetric as possible. The first few peaks should be within ± 20 or $\pm 30\%$ of the steady state peaks.

7 Waveform Plots

Perhaps a few plots of voltage and current will help understanding. In the following figures, A1 and A2 refer to the two analog channels of the HP 54645D oscilloscope. A1 is the voltage waveform applied to the Tesla coil, measured at the output of the inverter. There is a X10 voltage divider in addition to the probe, so a scale of 20 V/div is actually 200 V/div. A2 is the current, as represented by the voltage across a 0.02 Ω resistor. A scale of 100 mV/div would be 5 A/div.

The plots at the bottom of the figures are five of the sixteen digital channels of the scope. Channel 0 is the clock input of the 4013, an inverted output of the 34066. Channel 1 is the output of the 4046. Channel 2 is the Reset input to the 4013. Channel 3 is the output of the 4013 and the input to the 4066 that controls frequency by controlling voltage on the LM324 op amp. Channel 4 is pin 12 of the 4049. While it is high, frequency is controlled by R_{11} .

Fig. 7 shows the various waveforms for the first 180 μ s. The first voltage pulse is shortened,

as mentioned earlier. The current waveform is mostly switching noise for the first few cycles, although we start to see some current signal by the third or fourth cycle. The 4046 in Channel 1 appears to synchronize with the voltage immediately. So the controller has a decent digital signal for current even when the analog current waveform is still small and noisy during the first few cycles.




Figure 7: Controller Waveforms at Start

When Channel 4 goes low the controller is allowed to dynamically control frequency. The circuit is properly adjusted so this transition occurs smoothly, with the current waveform becoming more sinusoidal and continuing to grow.

Fig. 8 shows the same waveforms after 1 ms of operation (except for Channel 4, which is no longer of interest to us). The voltage is a square wave with relatively slow transitions from one rail to the other. The current is an almost perfect sine wave. If the voltage makes more rapid transitions from one rail to the other, switching noise will appear on the current. For this particular case, the current is lagging the voltage by about 7° and the duty cycle of signal on Channel 3 is about 28%, both values in a quite acceptable range.

The gate pulses from the 34066 have a duty cycle less than 50% because of the dead band. The inverted gate pulses in Channel 0 therefore has a duty cycle greater than 50%. This is not a problem in the tuning operation. The square wave representing current, Channel 1, does




Figure 8: Tesla Coil and Controller Waveforms after 1 ms

have a duty cycle close to 50%, however, since R18 has been adjusted to this effect. Channel 2 shows the Channel 1 signal after 2 inversions and a phase shift. The time delay between Channel 0 and Channel 2 is essential to the functioning of the circuit. If more lag is desired, R_{14} can be increased, although a better design might be to eliminate the analog delay and use the delay in two more inverter stages.

The frequency control signal, Channel 3, goes high when Channel 0 goes high, and goes low when Channel 2 goes high. If the current starts to lead the voltage, the Channel 2 waveform shifts to the left. The control signal (Channel 3) is not high as long. The voltage of the LM324 drifts up a little to bring the current back into phase.

Fig. 9 shows the voltage applied to the Tesla coil and the resulting current for about 4.5 ms. The voltage is A1 at the top of the figure and the current is A2 at the bottom. Operation is well below breakout. The ENABLE pulse is actually about 2.1 ms long. Current builds up as predicted by either transmission line theory or lumped RLC analysis. Current builds nearly linearly at first, then the rate of increase slows as losses increase. Current will hit its limit when the input power is equal to the losses, in this case between 2 and 3 ms after start.

As current increases, the voltage drop across the IGBTs and the droop in the electrolytic capacitors becomes greater. The decrease in voltage applied to the Tesla coil is not great,




Figure 9: Tesla Coil Voltage and Current Waveforms

but is most easily noticed when ENABLE is turned off. Energy stored in the Tesla coil has to be dissipated somehow. Instead of power flowing from inverter to coil, it now flows in the opposite direction. Voltage is constrained by the built-in diodes of the IGBTs. Power supply capacitors are now being charged instead of discharged. All the voltage drops in wiring and the IGBTs reverse in sign. We therefore see a small step increase in voltage when ENABLE is removed.

After about 0.6 ms, the stored energy is no longer able to force the IGBT diodes into forward conduction. Without a power supply affecting the circuit, we now have a classic RLC ring down. Both voltage and current are decreasing exponentially. During this portion of the cycle, the IGBTs are acting as capacitors.

Fig. 10 shows a closeup of the waveforms when ENABLE goes low. Both outputs of the 34066 go low, so Channel 0 goes high and stays there. This is the clock input to the 4013, which does not respond to constant voltage levels. The next reset pulse (Channel 2) sets the 4013 output low (Channel 3), where it remains. The controller is still measuring Tesla coil current, as seen by Channels 1 and 2, but is no longer controlling frequency.

Frequency is now being controlled by the Tesla coil. Instead of the controller and inverter trying to apply power to the Tesla coil at the right frequency and phase, the Tesla coil is




Figure 10: Voltage Phase Shift at ENABLE Turnoff

applying power to the inverter. Frequency control is no longer a problem. Voltage and current will always be in the proper phase for this high Q resonant circuit during ring down.

As power flow reverses in the line between Tesla coil and inverter, either voltage or current (or both) need to shift phase accordingly. It would be an interesting quiz question to ask students in a power electronics course to predict the top two waveforms in Fig. 10. It turns out that the current never misses a beat. The voltage has a cycle that is half the normal length, such that at the end of the short cycle, voltage is out of phase with current. One is reminded of the old teaching aid for electricity where current flow is equivalent to water flow. Moving water has inertia, and any pressures or heads will be the result of the inertia in the water. I would not want to carry the analogy very far, but in this case it gives the right answer. Current appears to have inertia and voltage follows along as best it can.

Fig. 11 shows the voltage and current waveforms later in the cycle. Stored energy is no longer able to drive the IGBT diodes into conduction. The relative phase between voltage and current has shifted about 90°. One possible explanation is that Fig. 10 is showing voltage and current of a stiff power supply while Fig. 11 is showing voltage and current of a capacitor. The stiff power supply being recharged acts like a resistor as far as phase angles are concerned. A power supply will have real power going out (normally) or in (in this case), hence must

have voltage and current at least somewhat in phase. When the Tesla coil stored energy is inadequate to connect to the power supply, we just measure voltage across and current through the intrinsic capacitance of the IGBTs. A pure capacitance always has current leading the voltage by 90° . Therefore, a phase shift between Fig. 10 and Fig. 11 is not too surprising.




Figure 11: Voltage and Current During Ring Down

References

- [1] Mohan, Ned, Tore M. Undeland, and William P. Robbins. *Power Electronics: Converters, Applications, and Design*, John Wiley, New York, 1989.

SPARKS




Figure 1:

Coil

The coil form for 14S was fabricated of 1/8th inch polyethylene sheet, spliced with hot glue. There were 387 turns of 14 gauge copper magnet wire space wound on the form. Coil diameter was 15.6 inches (39.64 cm) and the winding length ℓ_w was 45.9 inches (116.6 cm). Calculated inductance according to Wheeler's formula was 17.22 mH. The measured inductance was 16.94 mH with a ELC-120 meter. The Medhurst capacitance was calculated to be 23.95 pF.

The resonant frequency is then

$$f = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{(17.22 \times 10^{-3})(23.95 \times 10^{-12})}} = 247.8 \text{ kHz} \quad (1)$$

Measured resonant frequency of the coil with no top load of any type was 248 kHz. All results reported here for this coil are with a half-spun aluminum toroid, nominally 24 inches outside diameter by 6 inches, for which the resonant frequency was 207 kHz.




Figure 2:

We want the input impedance of the Tesla coil, the ratio of input voltage over input current. The combination of a square wave voltage and a sinusoidal current is different from what we are used to, so we need to go back to circuit theory to make sure we have all the correct multiplying factors.

A voltage square wave of value $\pm V_p$ has the same heating capability when applied to a non-inductive resistor as a dc voltage V_p , hence is said to have an rms value of $V_{ac} = V_p$. The corresponding dc current in a resistor would have an rms value of $I_{ac} = I_p$. The average power delivered to the resistor is the product of the rms voltage and the rms current,

$$P_{ave} = V_p I_p = V_{ac} I_{ac} \quad (2)$$

Suppose now we apply a sinusoidal voltage $V_p \sin \omega t$ to a non-inductive resistor. The resulting current is $I_p \sin \omega t$. The average power is

$$P_{ave} = \frac{1}{\pi} \int_0^\pi V_p I_p \sin^2 \theta d\theta = \frac{V_p}{\sqrt{2}} \frac{I_p}{\sqrt{2}} = \frac{V_p I_p}{2} = V_{ac} I_{ac} \quad (3)$$

When the square wave voltage produces a sinusoidal current, the integral for apparent power becomes




Figure 3:

$$S = \frac{1}{\pi} \int_0^{\pi} V_p I_p \sin \theta d\theta = \frac{2}{\pi} V_p I_p = \frac{2}{\pi} V_p (\sqrt{2} I_{ac}) = 0.9 V_{ac} I_{ac} \quad (4)$$

For this case, the apparent power is no longer the simple product of rms voltage and rms current (as for dc and single frequency sinusoids), but has a 0.9 multiplying factor. The difference is due to the fact that the square wave voltage is composed of an infinite series of harmonics (fundamental, third, fifth, etc.). Each harmonic contributes to the rms value of the square wave. The current has no harmonics, so the higher voltage harmonics do not produce any contribution to the apparent power.

If the current is in phase with the voltage then the apparent power S is equal to the average power P_{ave} . If there is a phase shift, S will be greater than P_{ave} . S is readily determined from the product of rms voltage and rms current (and appropriate multiplying constants). P_{ave} requires a wattmeter for measurement. Building an accurate wattmeter in the hundreds of kHz range is a challenge, and was not done for this research.

At resonance, when the input impedance is real, it can be defined as

$$Z_{in} = \frac{P_{ave}}{I_{ac}^2} = \frac{2\sqrt{2}V_p I_{ac}}{\pi I_{ac}^2} = \frac{2\sqrt{2}V_p}{\pi I_{ac}} = 0.9 \frac{V_{ac}}{I_{ac}} \quad (5)$$

The results are given in the following table. Units are volts, amps, watts, and volt-amps,




Figure 4:

except for the voltage at the Tesla coil toroid, which is in kV.

	V_{dc}	I_{dc}	V_{ac}	I_{ac}	V_{top}	P_{dc}	S_{ac}	Z_{in}
	40.7	0.94	37.0	1.10	21	38.3	36.6	30.3
	82.5	1.94	76.4	2.25	42	160	155	30.6
	121	2.85	111	3.35	58	345	335	29.9
	160	3.73	145	4.43	75	599	580	29.5
Table 1. Tesla coil measurements.	218	2.28	208	2.85	47	497	533	65.7
	283	2.20	270	2.75	46	623	668	88.4
	346	2.15	332	2.75	46	744	822	108.6
	199	2.50	190	2.78		498	475	61.6
	218	2.50	214	2.78		545	534	69.4
	283	2.43	274	2.75		688	677	89.8

The first four rows are for non-breakout conditions, while the last six were taken with about a 6 inch long violet plume (brush discharge) projecting from a bump on the toroid. Below breakout, the coil acts just like a monopole above a ground plane, with input impedance about $30\ \Omega$. If the antenna were a vertical wire a quarter wavelength (90 electrical degrees) long, and had a sinusoidal current distribution, the input impedance (radiation resistance) would be $36.5\ \Omega$. An input impedance of $30\ \Omega$ corresponds to an electrical length of about 85 degrees.




Figure 5:

The coil appears to operate as a constant-current device in breakout. Both the input dc current and the ac current into the coil stay about constant, perhaps even decreasing a little, as the ac voltage applied to the coil increases from 208 to 332 V. The top voltage appears to be proportional to the current, hence also stays about constant. The input impedance increases proportional to the voltage, in order for the current to remain constant.

It may be noticed that the apparent power S_{ac} is less than the input dc power P_{dc} below breakout (rows 1-3 in Table 1), by amounts that would be reasonable for losses in the MOSFETs, if the apparent power were numerically equal to the real power (current exactly in phase with voltage). Above breakout, apparent power is greater than the input dc power (for rows 4-7 in Table 1) and increases with applied voltage. This implies that there is now a phase shift between voltage and current. This appears to be due to the leakage reactance of the autotransformer. When the input impedance of the Tesla coil is around $30\ \Omega$, the autotransformer is heavily loaded, and the leakage reactance is swamped. Above breakout, however, the input impedance of the Tesla coil rises rapidly, so leakage effects in the autotransformer become more noticeable. The leakage reactance is in series with a high Q RLC equivalent circuit for the Tesla coil. Increasing R changes the resonant conditions slightly, requiring some retuning for maximum current. Because of resonance, the voltage at the Tesla coil terminals will be *higher* than the voltage at the input to the leakage reactance (which cannot be measured). This resonant effect for this high Q circuit is easily enough to produce the observed increase in apparent power.




Figure 6:

To test this hypothesis, a second autotransformer was fabricated from a Phillips 4229 C81 ferrite pot core using 30 turns of 18 ga magnet wire, center-tapped. The initial permeability of this material is about 2700, hence the leakage inductance should be far less than for the first autotransformer. Data for rows 8-10 of Table 1 were taken with this autotransformer. As expected, the apparent power into the Tesla coil is now less than the average dc power into the MOSFETs. No significant changes occur in the other numbers.

OLD DATA.TEX, NOW TCCHAP9.TEX

1 Calorimeter

Ohmic losses, dielectric losses, and eddy current losses in the toroid all go into heat in the Tesla coil. Eddy current losses in the building and soil, plus any radiation losses, appear as heat elsewhere. Before discovering the major effect of proximity, I was able to explain only between 10% and 30% of the input impedance with ohmic losses. Eddy current losses seemed unlikely to contribute more than a few percent to the losses. The choice seemed to reduce to dielectric losses versus radiation. In an effort to distinguish between these effects, I built a ‘poor man’s’ calorimeter. This is basically an insulating box made of two inch thick blue styrofoam, with inside dimensions about 6 ft by 6 ft by 6.5 ft high. Seams were closed with duct tape. One panel was removable to change Tesla coil configurations inside. A small fan was mounted inside a section of 5 inch aluminum irrigation pipe and mounted vertically in a corner to provide air circulation. A Radio Shack indoor/outdoor thermometer was used to measure temperature. The outdoor temperature probe was threaded through a 10 ft length of copper water pipe and mounted inside the irrigation pipe. All pipes were electrically bonded and connected to the local ground.

Calibration was obtained by applying 60 Hz directly to the coil while measuring power with a wattmeter. Total input energy in kWh was calculated by multiplying power by time. A calibration constant could be obtained by dividing temperature rise by energy input. Each coil has a somewhat different thermal mass, so each has its own calibration constant.

One interesting observation was the difference in temperature increase between non sparking and sparking conditions. In the non sparking case, the coil gets hot and then heats the air. There is a noticeable lag between when power is applied and when temperature starts to rise. Likewise there will be a temperature overshoot for a few minutes after power is removed. In the sparking case, however, the spark heats the air directly. The coil is heated by the air, rather than the air being heated by the coil. Air temperature increases rapidly after power is applied, and decreases almost as rapidly when power is removed. The rate of decrease will change when thermal equilibrium is reached, when the coil temperature is equal to the air temperature.

For the same amount of energy dissipated inside the calorimeter, the air temperature will be higher for a longer period of time under spark conditions. This means more energy will be lost through the walls of the calorimeter. The styrofoam is rated as R10, so this wall loss becomes significant at temperature differences of more than 6 or 8°C. One should not get too excited about small differences between spark and non spark results, therefore.

From these experimental results it was apparent that most of the losses were occurring

inside the calorimeter. It was also noticed that input impedance was changing significantly from day to day. For example, on 12/29/00, I measured the input impedance of the 16 gauge coil as $87\ \Omega$, and on 1/15/01 as $122\ \Omega$, using identical procedures. The temperature where the coils are located was -4°C on 12/29/00 and -1°C on 1/15/01. This small difference would not cause such a difference in input impedance.

I then decided to check the effects of relative humidity. I placed some desiccant on the floor of the calorimeter and watched the input impedance as the space dried out. The impedance dropped from $122\ \Omega$ to $100\ \Omega$ within a few minutes, until the desiccant saturated. I then put a humidifier inside the calorimeter to raise the humidity to 100%. The impedance rose from $100\ \Omega$ to $200\ \Omega$ in 10 minutes and from $200\ \Omega$ to $300\ \Omega$ in the next 10 minutes, and finally reached $350\ \Omega$. Moisture has a major effect on input impedance for this particular coil!

I then did the same humidifier test with the 22 gauge coil that is dimensionally the same as the 16 gauge coil. The impedance rose from $90\ \Omega$ to $150\ \Omega$ as the humidity rose to 100% for that coil.

For coil 14T, the impedance rose only from $43.7\ \Omega$ to $45.5\ \Omega$ (corrected for temperature) as the humidity rose to 100%. What can we conclude from these strange observations? As was mentioned in Chapter 6, coils 16B and 22B were wound on plastic barrels found at the recycling place. These were the only two coils that showed such a strong variation in impedance with moisture, so it seems reasonable to argue that the barrels must somehow soak up water and become much more lossy as relative humidity increases.

On the other hand, the impedance of coil 16B increases by $250\ \Omega$ while the impedance of coil 22B increases by only $60\ \Omega$ as humidity goes to 100%. The barrels appear to be identical. The two coils have identical numbers of turns and resonant frequencies. The Fraga resistance is $R_s = 51\ \Omega$ for 16B and $69\ \Omega$ for 22B, but the measured resistance is always higher for coil 16B. Either the barrels are not actually identical or there is something about the wire on 16B that is lossier than the wire on either 14T or 22B. Both the 14 gauge and 16 gauge wires are Essex Heavy Soderon, the 14 gauge made in 1994 and the 16 gauge in 1997. Did Essex change their recipe for wire coating between 1994 and 1997? If not, then the barrels are actually different in water absorption properties or something else is happening.

I ran out of ambition before finding a firm answer to this question. I would suggest staying with PVC or polyethylene forms to help keep the water losses to a minimum. Once a coil is built and the impedance is measured, the coil can be misted with distilled water and remeasured. If the impedance changes more than perhaps 5%, expect results to vary with humidity. If this is a problem, use a different coil form and/or different wire.

2 Spark Length

Results

Figure 2 shows the voltage and current waveforms into the Tesla coil when operated at $V_{dc} = 100$ V, which is well below breakout (no spark). The voltage wave is ideally a square wave. The ripple is due to MOSFET switching and resonances formed by stray inductances and capacitances. The main feature is that an applied square wave voltage will produce a sinusoidal current. At resonance, this current will be in phase with the voltage. The coil is high Q so tuning is very sensitive. Variation of a few hundred Hz is sufficient to get operation well away from resonance.

Figure 2. Voltage and current into Tesla coil, $V_{dc} = 100$ V.

We want the input impedance of the Tesla coil, the ratio of input voltage over input current. The combination of a square wave voltage and a sinusoidal current is different from what we are used to, so we need to go back to circuit theory to make sure we have all the correct multiplying factors.

A voltage square wave of value $\pm V_p$ has the same heating capability when applied to a non-inductive resistor as a dc voltage V_p , hence is said to have an rms value of $V_{ac} = V_p$. The corresponding dc current in a resistor would have an rms value of $I_{ac} = I_p$. The average power delivered to the resistor is the product of the rms voltage and the rms current,

$$P_{ave} = V_p I_p = V_{ac} I_{ac} \quad (1)$$

Suppose now we apply a sinusoidal voltage $V_p \sin \omega t$ to a non-inductive resistor. The resulting current is $I_p \sin \omega t$. The average power is

$$P_{ave} = \frac{1}{\pi} \int_0^\pi V_p I_p \sin^2 \theta d\theta = \frac{V_p}{\sqrt{2}} \frac{I_p}{\sqrt{2}} = \frac{V_p I_p}{2} = V_{ac} I_{ac} \quad (2)$$

When the square wave voltage produces a sinusoidal current, the integral for apparent power becomes

$$S = \frac{1}{\pi} \int_0^\pi V_p I_p \sin \theta d\theta = \frac{2}{\pi} V_p I_p = \frac{2}{\pi} V_p (\sqrt{2} I_{ac}) = 0.9 V_{ac} I_{ac} \quad (3)$$

For this case, the apparent power is no longer the simple product of rms voltage and rms current (as for dc and single frequency sinusoids), but has a 0.9 multiplying factor. The difference is due to the fact that the square wave voltage is composed of an infinite series of harmonics (fundamental, third, fifth, etc.). Each harmonic contributes to the rms value of the

Table 1. Tesla coil measurements.

V_{dc}	I_{dc}	V_{ac}	I_{ac}	V_{top}	P_{dc}	S_{ac}	Z_{in}
40.7	0.94	37.0	1.10	21	38.3	36.6	30.3
82.5	1.94	76.4	2.25	42	160	155	30.6
121	2.85	111	3.35	58	345	335	29.9
160	3.73	145	4.43	75	599	580	29.5
218	2.28	208	2.85	47	497	533	65.7
283	2.20	270	2.75	46	623	668	88.4
346	2.15	332	2.75	46	744	822	108.6
199	2.50	190	2.78		498	475	61.6
218	2.50	214	2.78		545	534	69.4
283	2.43	274	2.75		688	677	89.8

square wave. The current has no harmonics, so the higher voltage harmonics do not produce any contribution to the apparent power.

If the current is in phase with the voltage then the apparent power S is equal to the average power P_{ave} . If there is a phase shift, S will be greater than P_{ave} . S is readily determined from the product of rms voltage and rms current (and appropriate multiplying constants). P_{ave} requires a wattmeter for measurement. Building an accurate wattmeter in the hundreds of kHz range is a challenge, and was not done for this research.

At resonance, when the input impedance is real, it can be defined as

$$Z_{in} = \frac{P_{ave}}{I_{ac}^2} = \frac{2\sqrt{2}V_p I_{ac}}{\pi I_{ac}^2} = \frac{2\sqrt{2}V_p}{\pi I_{ac}} = 0.9 \frac{V_{ac}}{I_{ac}} \quad (4)$$

The results are given in the following table. Units are volts, amps, watts, and volt-amps, except for the voltage at the Tesla coil toroid, which is in kV.

The first four rows are for non-breakout conditions, while the last six were taken with about a 6 inch long violet plume (brush discharge) projecting from a bump on the toroid. Below breakout, the coil acts just like a monopole above a ground plane, with input impedance about 30Ω . If the antenna were a vertical wire a quarter wavelength (90 electrical degrees) long, and had a sinusoidal current distribution, the input impedance (radiation resistance) would be 36.5Ω . An input impedance of 30Ω corresponds to an electrical length of about 85 degrees.

As the voltage is increased, a point is reached where the coil goes into breakdown. For this coil and these power levels, the discharge is about six inches long into air, violet in color, and somewhat bushy (called a brush discharge). The current into the coil drops, as does the top voltage.

The coil appears to operate as a constant-current device in breakout. Both the input dc current and the ac current into the coil stay about constant, perhaps even decreasing a little, as the ac voltage applied to the coil increases from 208 to 332 V. The top voltage appears to be proportional to the current, hence also stays about constant. The input impedance increases proportional to the voltage, in order for the current to remain constant.

The resistive losses and radiated power remain about the same, in the range of 200 to 300 W, as input voltage is increased in breakout, so the additional power evidently goes into the discharge. If the total ac input power were 700 W, we would have 50 W (or less) into coil resistance, perhaps 200 W into the radiated field, and 450 W into the discharge. As the input voltage is increased even more, the resistive losses and radiated power would stay about the same while the power into the discharge increases, so the *percentage* of power supplied to the discharge increases with input voltage.

It may be noticed that the apparent power S_{ac} is less than the input dc power P_{dc} below breakout (rows 1-3 in Table 1), by amounts that would be reasonable for losses in the MOSFETs, if the apparent power were numerically equal to the real power (current exactly in phase with voltage). Above breakout, apparent power is greater than the input dc power (for rows 4-7 in Table 1) and increases with applied voltage. This implies that there is now a phase shift between voltage and current. This appears to be due to the leakage reactance of the autotransformer. When the input impedance of the Tesla coil is around $30\ \Omega$, the autotransformer is heavily loaded, and the leakage reactance is swamped. Above breakout, however, the input impedance of the Tesla coil rises rapidly, so leakage effects in the autotransformer become more noticeable. The leakage reactance is in series with a high Q RLC equivalent circuit for the Tesla coil. Increasing R changes the resonant conditions slightly, requiring some retuning for maximum current. Because of resonance, the voltage at the Tesla coil terminals will be *higher* than the voltage at the input to the leakage reactance (which cannot be measured). This resonant effect for this high Q circuit is easily enough to produce the observed increase in apparent power.

To test this hypothesis, a second autotransformer was fabricated from a Phillips 4229 C81 ferrite pot core using 30 turns of 18 ga magnet wire, center-tapped. The initial permeability of this material is about 2700, hence the leakage inductance should be far less than for the first autotransformer. Data for rows 8-10 of Table 1 were taken with this autotransformer. As expected, the apparent power into the Tesla coil is now less than the average dc power into the MOSFETs. No significant changes occur in the other numbers.

Figure 3 shows voltage and current during discharge. V_{dc} was doubled from the data in Figure 2, from 100 to 200 V, and the corresponding input impedance was also doubled, from about $30\ \Omega$ to about $60\ \Omega$. The Tesla coil is acting as a constant current sink.

The phase shift of the top voltage with respect to the input voltage was also checked. If the Tesla coil was acting as a quarter-wave antenna, we would expect the top voltage to lag the input voltage by 90° , due to the time required to propagate along a transmission line. If, on the other hand, the Tesla coil is considered as the secondary of a transformer (where the

primary consists of a few turns of heavy conductor), we would expect the top voltage to be in phase with the input voltage, just like any other transformer. A measured phase shift of 94° was obtained. This does not answer all the interesting questions about electrical length. For example, if the electrical length based on input impedance is 85° , does this imply that the phase shift from bottom to top should also be 85° ? But it supports the concept that a coil driven as a monopole above a ground plane has a phase shift from bottom to top that is close to what would be expected for a quarter-wave antenna and not that of a transformer.

A logical next step for research would be to look at the input impedance of a spiral primary driving the same coil as a secondary in a classic Tesla coil configuration. The input impedance would be on the order of one Ohm, so I_{ac} would have to be above 25 A to get enough power into the this particular coil to get breakout. The primary and secondary would need to be undercoupled so the system would have a single resonant frequency. These features raise some interesting issues, but it should be possible to make such measurements.

Figure 3. Voltage and current into Tesla coil. $V_{dc} = 200$ V.

Discussion of Results

In the first part of this paper (titled simply “Input Impedance of a Tesla Coil”) some data were presented, with very little in the way of explanation. The following is a first attempt at such an explanation.

It is assumed that a direct driven (cw) Tesla coil above a ground plane can be modeled as a lumped series RLC circuit, as shown in Fig. 4. Some aspects of a Tesla coil may require transmission line theory, but many Tesla coil experts would argue that the lumped model is adequate for almost all observations, and certainly for the input impedance as measured at the base.

A classic Tesla coil, with a primary driving the secondary, has another L and another C in its model. However, when the spark quenches, the circuit of Fig. 4 applies, with V_{ac} shorted. The energy stored in L and C at the point of quenching then rings down according to the usual circuit theory equations. It appears that the R measured during cw operation should be the same as the R existing during the disruptive coil ringdown, for similar voltages and currents. If they are different, it would be nice to know why.

The coil in Part I was operated at resonance and the input voltage and current were measured. The ratio was then the input impedance at resonance, shown as R .

Below breakout, for the coil in Part 1, we measured a constant R as V_{ac} increased. The current I_{ac} therefore increased linearly with V_{ac} and the power P increased as a square law curve. At breakout, I_{ac} dropped by about a factor of two, R increased by a similar factor, and therefore $P = I_{ac}^2 R$ decreased by a factor of two. This observation can only be explained with the series model of Fig. 5.

Coil	x	ϕ	f_r	R_{dc}	R_{ac}	R_M	R_{TC}
14S	8.69	1.85	250.6	3.99	13.32	24.70	24.5
14T	8.96	3.03	266.5	4.451	15.28	46.35	43.5
16B	5.27	3.12	146.3	8.891	18.91	59.00	93.1
18B	4.18	4.25	146.8	12.90	22.50	95.61	118.0
18T	5.38	3.15	242.9	11.19	24.26	76.50	70.5
20S	4.39	1.77	256.7	18.36	33.34	59.01	43.0
20T	3.71	3.02	183.9	23.38	36.88	111.39	94.2
22T	3.81	1.62	307.7	22.40	36.10	58.38	47.0
22B	2.66	1.54	149.1	35.70	42.92	66.09	73.0

Below breakout, the current at the top of the coil is zero. This is the current actually flowing into the surrounding space. If we had an ideal quarter-wave antenna, the current distribution would be sinusoidal along the antenna, maximum at the base and zero at the top. With capacitive top loading, the current becomes more linear than sinusoidal. A linear current is plotted in Fig.8, more because it was easy to plot than because it was highly accurate.

Once the discharge starts, the current becomes greater than zero at the top. This would force some sort of redistribution of current along the coil. Since the input current was observed to decrease while the top current rises above zero, a line with less slope was drawn on Fig. 8 to represent the current for the discharge case.