

Development of a portable Tesla coil apparatus

**Kenneth D Skeldon, Alastair I Grant, Gillian MacLellan
and Christine McArthur**

Department of Physics and Astronomy, University of Glasgow, University Avenue, Glasgow G12 8QQ, UK

E-mail: kskeldon@physics.gla.ac.uk

Received 10 October 1999

Abstract. The Tesla coil uses high-frequency transformer action together with resonant voltage amplification to generate potentials in the range of tens to hundreds, or even thousands of kilovolts. We describe a range of experiments designed to investigate the Tesla coil action, ending up with the design and development of a touring Tesla coil with a carefully considered trade-off between portability and performance. The apparatus plugs into a standard 230 V wall socket, drawing an average of 7 A, and features a telescopic, collapsible secondary coil. The assembled system stands almost 2 m high and yields sparks up to 1.5 m long, giving it physical stage presence which makes for memorable demonstrations, even in larger venues.

1. Introduction

The Tesla coil has long been regarded as one of the most effective ways of producing continuous high-potential electrical discharges in air, whether for the purposes of research, teaching or visual display. In this article we present the design and development of a portable Tesla apparatus, concentrating along the way on the various physical aspects associated with the Tesla action: these include electromagnetic induction, coupled resonant systems, power transfer and energy conservation and, ultimately, the discharge phenomena pertaining to electrical currents at high frequency and high voltage. The development of a Tesla coil offers considerable value through both the theoretical and experimental exploration of these concepts [1], while the completed apparatus serves as a spectacular lecture demonstration and exhibit for use in an undergraduate teaching capacity as well as in science outreach activities for the public and in schools [2]†.

The simplest Tesla coil can be described in terms of an air-cored resonant transformer, in which an oscillatory current of very large amplitude flows in a vertically mounted primary coil of comparatively large dimensions, but having only a small number of turns (often of the order of ten or so). The primary current induces a substantial current in the bottom few turns of a free-standing secondary coil which sits inside the large primary coil, but is physically much higher, and also has very many more turns (usually several hundred). The secondary coil is an electrically resonant component by design, because of its self-inductance and capacitance, and is configured to have its top end free and its bottom end grounded. During operation it behaves

† A considerable amount of information is available on the World Wide Web by searching for ‘Tesla coil’; this includes a wealth of educational and technical material.

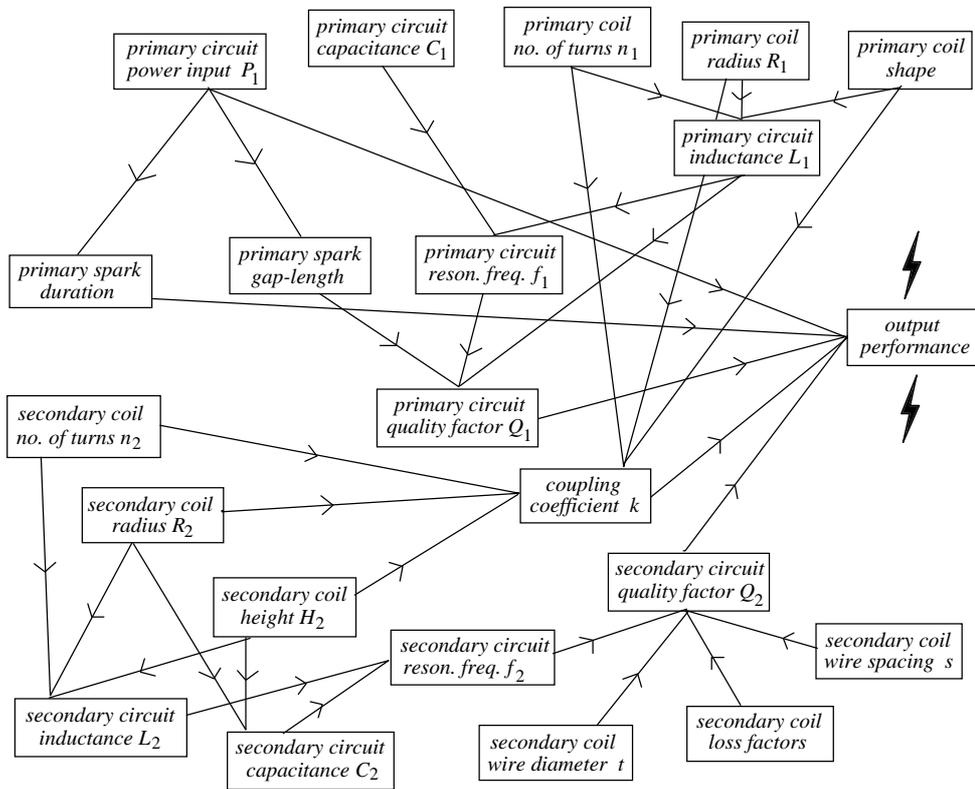


Figure 1. Flow diagram depicting the complex interrelationships between the various Tesla coil design parameters. Every physical design choice eventually has an impact on output performance, some to a far greater extent than others.

like a series LC circuit, being driven by the induced oscillatory current at its base. The result is a very high potential developing between the top end of the secondary coil and ground by a combination of transformer action and voltage magnification. The resulting ionization of the air produces dramatic corona discharge from an appropriately shaped top electrode.

The amount of available source material on Tesla coil theory and construction is now considerable, to the extent that when one searches for literature on the subject, it is invariably the case that no clear consensus for any particular design philosophy actually emerges. Indeed, the researcher will most probably be left seeking any discernable convergence of opinion from a vast set of seemingly contrary viewpoints on exactly what the factors that make for a well-designed coil are. This diversity of opinion highlights what is probably the major design challenge associated with Tesla coil construction, namely that it is an apparatus with a vast set of freely selectable design parameters, as indicated in figure 1. Moreover, alternative sets of design parameters could yield significantly different coil constructions, but with comparable output performance levels. In contrast, parameter sets with only one or two ill chosen values can seriously degrade the achievable output available from a Tesla coil system. It is this issue of the choice of design parameter sets, and how attention to some basic physics can help with this matter, that we wish to address later in our paper. In particular, we shall keep in mind that our aim here is to build a Tesla coil which is portable and able to be used in a wide variety of locations. Also, and perhaps more fundamentally, the design parameters of any proposed system will be influenced by the method we adopt to power the Tesla coil, and so we shall briefly discuss this issue first.

2. Tesla coil excitation methods

The basic Tesla coil configuration is shown in figure 2, alongside example circuits for two possible methods of excitation, these being continuous excitation (by powering the system from a suitably impedance matched rf signal source) and transient excitation (by switching a charged capacitor into the primary circuit). The portable Tesla coil we have built relies on the latter excitation technique; however, in order to explain why this is so, it is instructive first to briefly consider the practicalities of a continuously driven system.

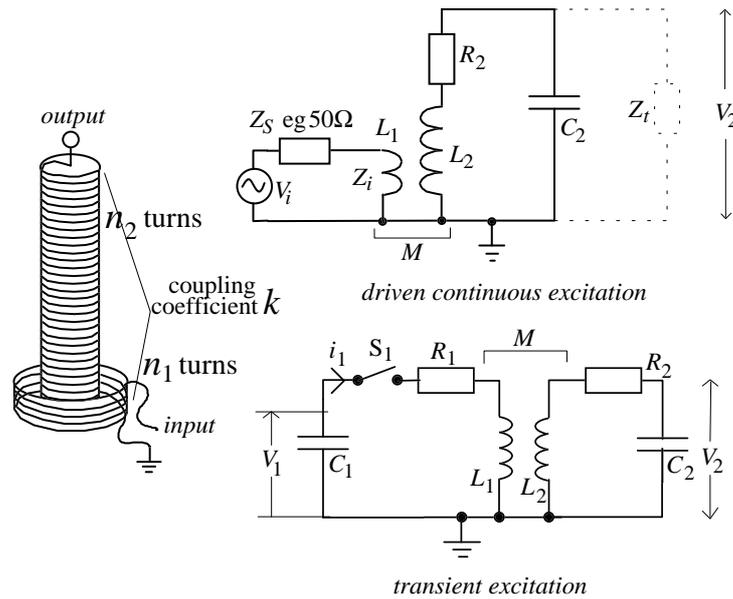


Figure 2. Basic Tesla coil circuit shown with two excitation methods. In the driven arrangement power is impedance matched to the primary circuit from a signal generator. In the transient excitation method a charged capacitor delivers an impulse to the primary coil and the system is energized by the subsequent transient response.

Continuously driven coil arrangement

The process of impedance matching is important in many areas of physics where we wish to transfer as much power as possible from one place to another. In order to introduce the impedance matching concept here, consider the photograph in figure 3 which shows a small Tesla coil system we wound for test purposes. In order to obtain the corona discharge shown in the inset picture, we matched about 36 W of power into the primary coil of the device from a 50 Ω source impedance signal generator. The problem with this system is that the input impedance (Z_i in figure 2) is dependent on the termination impedance (Z_t in figure 2), which is that existing between the discharge points and ground. However, Z_t is a very variable quantity, and hence so too is Z_i : if the corona discharge is arrested, for example by covering the top of the coil with an insulating sleeve, then Z_t , together with the impedance matching, will change considerably. The same is true if an arc is allowed to strike. In order to achieve optimum impedance matching, the source impedance (Z_s in figure 2) must be the same as the load impedance Z_i , which clearly cannot be the case for all possible discharge conditions simultaneously. In a large coil, with high electrical powers, the drive electronics are likely to be readily damaged by the reflected rf energy caused by such volatile impedance matching



Figure 3. A small continuously driven Tesla coil system operating at a frequency of approximately 4 MHz. The inset photograph shows the corona produced when the coil draws approximately 36 W of power.

conditions. Nevertheless, given a steady corona discharge it is possible to achieve a nominal matching for the incident power, which we shall discuss briefly.

There are four distinct circuit equations which can be solved to yield an expression for Z_i , the input impedance of the system. (An exact expression for the input impedance can be written

$$Z_i = R_1 + \frac{\omega^2 M^2 R_2 (1 + (Z_t/R_2) + \omega^2 Z_t^2 C_2^2)}{R_2^2 + 2R_2 Z_t + Z_t^2 (1 - (\omega^2/\omega_0^2))^2 + \omega^2 L_2^2 (1 + (R_2^2 Z_t^2/Z_0^4))} + j \left[\omega L_1 - \frac{\omega^3 M^2 L_2 (1 - (Z_t^2/Z_0^2) + \omega^2 Z_t^2 C_2^2)}{R_2^2 + 2R_2 Z_t + Z_t^2 (1 - (\omega^2/\omega_0^2))^2 + \omega^2 L_2^2 (1 + (R_2^2 Z_t^2/Z_0^4))} \right]$$

whereupon the conditions of real (resistive) input termination conditions can be deduced by setting the imaginary part of Z_i to zero). This expression is quadratic in ω^2 and there are two distinct frequencies where the impedance measured across the primary coil is purely real, the upper of which is the high-admittance state. At this resonance frequency, and if the secondary coil had little or no loss, then we would have $Z_i \propto Z_0^2/Z_t^\dagger$. However, the secondary coil loss R_2 cannot be neglected and will even dominate near the condition of infinite termination impedance, giving $Z_i \propto Z_0^2/Z_t + R_2$. The primary-to-secondary coupling is essentially an air-core rf transformer; indeed the freedom to change the ratio of turns here is advantageous as an impedance matching aid. The impedance transfer will occur in proportion to the inductances of the two coils, or equivalently, to the square of the turns ratio. The coupling factor will enter as a flux linkage deficit and will have an effect in proportion to the number of turns, giving the

[†] In many respects, the results for the Tesla system are similar to those for a quarter-wave transmission line transformer, where a shorted termination impedance is transformed to a high input impedance and vice-versa. See for example [3].

much simpler formula

$$Z_i \sim \frac{1}{k^2} \left(\frac{n_1}{n_2} \right)^2 \left[\frac{Z_0^2}{Z_t} + R_2 \right] \quad (1)$$

for the input impedance of the system at the voltage resonance of the secondary circuit. In order to test the impedance matching techniques we can perform some simple calculations on our small test coil. First, we can simplify matters by stating that all the 36 W of power is matched into the secondary coil, yielding the output corona. Meanwhile, the spark striking distance suggests a potential of approximately $V_2 \sim 10$ kV at the output[†]. By equating the power to V_2^2/Z_t we deduce that $Z_t \sim 3$ M Ω . Thus, for the test coil we have $Z_0^2/Z_t \sim 200$ Ω , R_2 is small compared to this and is neglected, $k \sim 0.1$, $n_1 = 8$ and $n_2 = 400$ giving $Z_i \sim 80$ Ω . Given the source impedance of 50 Ω this would imply a voltage standing-wave ratio (VSWR) of about 1.6 compared with a measured VSWR of 2 which is in reasonable agreement considering the approximations made.

The extrapolation to larger systems is evident. In order to match sufficient power into a larger coil one would require a combination of a low-loss, high- Q secondary component, high termination impedance and a primary rf supply capable of delivering many kilowatts of power. Of course, the impedance matching can always be arranged to favour lower voltage drives operating at higher currents, which is probably more readily realizable in practice, but even so, the electronics would require careful design and considerable expertise to implement. We have run circuit simulations[‡] which suggest that in order to produce approximately 100 kV across a termination impedance of 30 M Ω using a secondary coil with $Q \sim 500$, $Z_0 \sim 50$ k Ω and $f_0 \sim 100$ kHz (optimal power transfer conditions) one requires approximately 800 W matched to the system. In terms of power rating, this is within the capability of relatively inexpensive, commercially available amplifiers designed for power audio. However, some modification would be necessary to obtain the broader frequency response, and there is always the problem of the changing termination conditions.

Transient excited coil arrangement

In a transient excited system, the design criterion compared to the driven case changes for various reasons. For example, the requirement for a high- Q secondary coil in the driven case is essential, given its role as a low-loss voltage magnifier like a series LCR circuit. However, the resonance effect does not play such an important role in the secondary voltage generation of a transient excited system. For instance, the peak voltage attained does not come anywhere close to the maximum available via true steady-state resonant excitation, given a high- Q coil with favourable termination conditions. Nevertheless, it is true that once a certain peak energy has been attained by the secondary coil, the subsequent decay of this energy will take longer with a higher- Q component. However, there are other factors such as energy exchange back to the primary circuit, and energy loss through corona discharge itself, which can outweigh the high- Q benefit [4].

An example of a standard spark-gap switched primary circuit is shown in figure 2. The behaviour of the generator in the absence of any termination impedance can be deduced by solving two circuit equations [4]. The basic operation can be described as follows. The high-voltage capacitor C_1 is charged to a voltage V_1 on a relatively long timescale compared to the rf oscillation period. When the spark-gap S_1 fires, C_1 discharges and energy rings between it and the primary coil L_1 , inductively coupling energy into the secondary circuit. The secondary coil attains its peak voltage after a short time, dependent largely on the mutual inductance M but typically over a few rf cycles; we should again point out that the timescale here is much

[†] We have taken the estimate of 10 kV cm⁻¹ in view of the curvature of the discharge points. The electric-field breakdown value for dry air is nearer 30 kV cm⁻¹.

[‡] The simulations have been done both in MAPLE using idealized coupled- LCR circuit representations, and also in a circuit-analysis package called *LISO* that takes into account the sort of electrical circuit loadings found in practice.

shorter than the natural time constant of the secondary coil given its Q , and, therefore, Q_2 and voltage specification should be considered with caution in a transient excited system (indeed during the secondary coil excitation phase, it is more pertinent to be concerned about Q_1 , i.e. the losses in the primary circuit). To the observer of such a system, the corona that is produced seems continuous and rf in nature, but in reality there is significant output present for only a very small percentage of the time, and the major part of the operation cycle is spent waiting for C_1 to be recharged. If the charging device is a mains-powered eht transformer, there will most probably be a burst of rf activity every 0.01 s or so (determined by the line frequency). However over this time, an input power transformer rated at say 1 kW will have had the opportunity to deposit an extremely large amount of energy in C_1 . Herein lies the simplicity and also the cleverness of the transient excited circuit: that it can still deliver an extremely impressive peak output which will not fail to impress any spectator, albeit at the expense of a much poorer time-integrated performance, which will probably be of scant concern to the spectator.

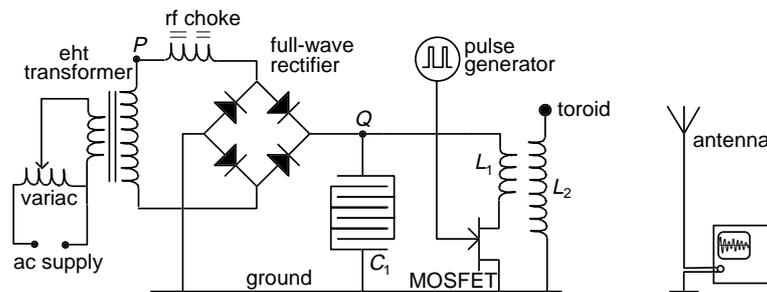


Figure 4. Circuit schematic for a transient excited Tesla coil system, tested first at low voltage with a MOSFET transistor to trigger the primary impulse, instead of the spark-gap used in the high-voltage design.

In order to obtain some insight into the behaviour of the transient excited system, consider the schematic in figure 4. This circuit represents the final arrangement chosen for our coil system which we describe in the next section. The only real difference from the final mains-powered system is that here the output of the eht transformer is full-wave rectified, enabling the role of the spark-gap to be conveniently replaced by a MOSFET transistor switch. By supplying the circuit from a low-voltage supply we can safely investigate the system performance for a range of MOSFET gating rates/duty-cycles (equivalent to spark-gap duration and frequency), mutual inductances through the coupling coefficient of primary and secondary coils, and coil designs. A similar experiment to demonstrate Tesla coil theory was conducted by Bruns [5]. The important point here is that the components involved in our experiment, such as the eht transformer, control variac, and even the capacitor value and coil configurations, will be exactly those that will feature in the high-voltage state, whereupon the linear nature of the system should allow the eventual performance to be predicted. The measurement of the primary voltage was made with a high-impedance probe at point Q in figure 4. Measurement of the secondary coil voltage is a two-stage procedure. A true measurement is made non-invasively using an antenna placed approximately 4 m from the coil as illustrated in figure 4, and a calibration measurement is also made with the antenna and a high-impedance probe connected directly to the secondary coil. The MOSFET was controlled by a pulse generator which allowed control over the gate pulse duration and frequency. In this way, spark-gap switching properties could be investigated.

Some results from this low-voltage experiment are shown in figure 5(a)–(d). In (a) we show a MAPLE simulation of the voltages across the primary and secondary coils after a single initial impulse. Graph (b) shows the actual measurements made using the set-up of figure 4. Given the antenna calibration shown in graph (b) we deduce reasonable primary-to-secondary voltage gain agreement with the MAPLE model of graph (a). Graphs (a) and (b) clearly show

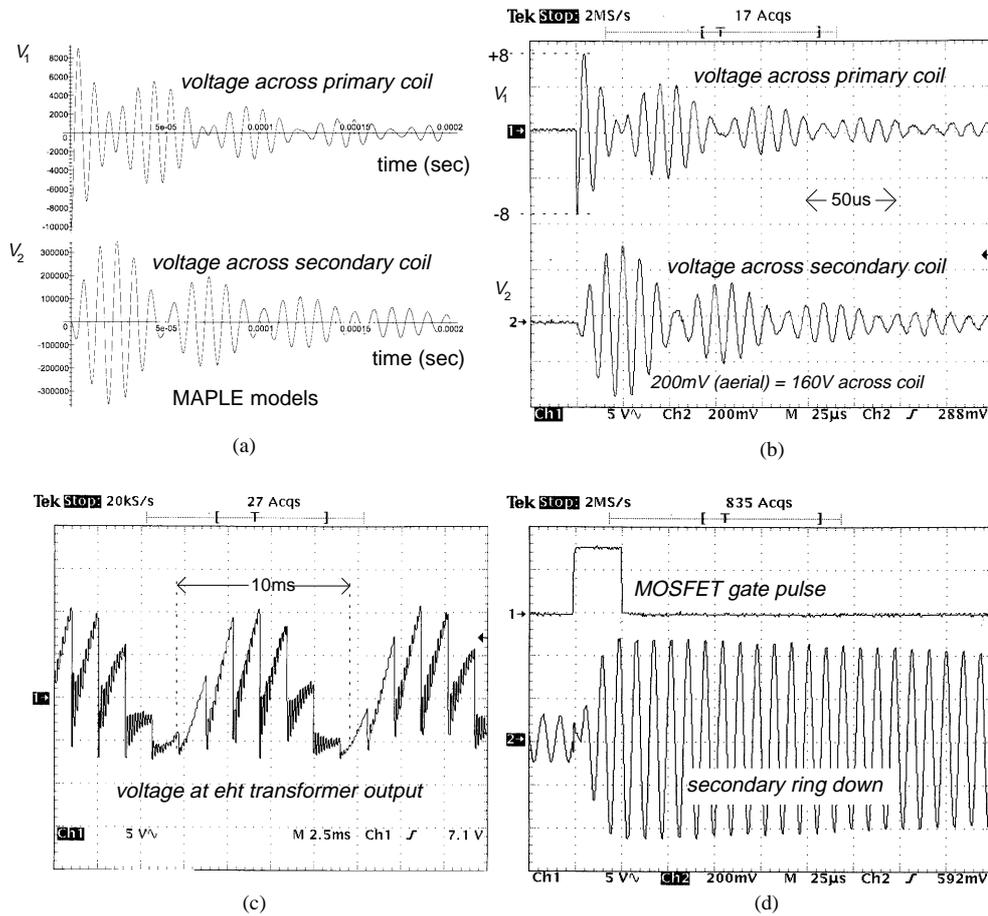


Figure 5. Some graphs from the low-voltage investigation of the transient excited Tesla coil system: (a) MAPLE simulation of the system behaviour; (b) measured response with same circuit parameters as in (a); (c) the voltage measured at point P in figure 4; (d) the secondary coil behaviour subsequent to the action of switching out the primary circuit just as all the energy transfers out of it.

how the energy in the system is exchanged between the primary and secondary circuits. The beat frequency at which this energy exchange occurs is determined by the mutual inductance M , or coupling coefficient $k = M/\sqrt{L_1 L_2}$, which, very usefully, can be accurately calculated using graph (b) ahead of final construction[†]. If the MOSFET switch is gated several times per mains cycle we obtain trace (c) for the voltage measured at point P in figure 4. This would represent the switching pattern of a rotary spark-gap where the spark frequency and duration can be controlled. The charging pattern of C_1 can be seen in (c) and this depends on the impedance of the eht transformer. Also evident from graph (c) is some rf contamination from the primary oscillation. However, this is at the considerably suppressed level afforded by the

[†] The beat frequency f_b is related to the coupling coefficient k via

$$f_b = \left[\frac{1}{(1-k)^{1/2}} - \frac{1}{(1+k)^{1/2}} \right] f_0$$

where f_0 is the nominal resonant frequency of the system, given for example by $f_0 = 1/(2\pi\sqrt{L_2 C_2})$. See for example [4].

rf choke, which is a particularly important component in the high-voltage circuit, where the rf signals could otherwise damage the transformer insulation. Graph (d) applies to the situation where the MOSFET gate pulse is quenched at the end of the first secondary voltage rise, so that the secondary coil rings down at a rate determined by its own Q .

3. Parameter set for a collapsible Tesla coil

Although the transient mode of operation distances the Tesla coil designer from the problems of impedance matching and drive electronics, there does remain the problem of choosing design parameter sets which best utilize a given input power for a lengthy output spark discharge. Looking back at figure 1 we see how dependent all the various design parameters are on one another, and so it is important that we do not make a poor initial choice which may have an adverse effect on performance. We decided to design the secondary coil first, given its key role in determining the output. This will also set the resonance frequency which can be met in the subsequent development of the primary circuit.

Secondary coil parameters

The only preliminary factors influencing the secondary coil design are the input power of the system and the anticipated potential across the secondary windings. Since it is important that the apparatus can plug into a standard UK 13 A ‘ring-mains’ socket, it must require no more than approximately 3 kW and preferably draw much lower power than this on average. It has been said that a well-designed coil should be able to generate an output spark length of 30 cm for each kilowatt of input power[†].

Therefore, knowing nothing of the potential efficiency of the system we might end up with, we can cite this number to estimate that with an upper bound of 3 kW we should aim for 90 cm spark length. At 10 kV per centimetre, this would suggest a voltage of approximately $V_2 \sim 900$ kV across the secondary coil. The distribution of voltage along the secondary coil follows the first quadrant of a sine function, a potential node at the grounded bottom end and a potential anti-node at the top discharge terminal. The greatest voltage gradient ΔV_{\max} therefore appears at the lower turns and depends on the number of windings n_2 , being given by $\Delta V_{\max} = \pi V_2 / 2n_2$ [4]. This factor should be kept in mind when designing the secondary coil.

Since a major feature of the Tesla coil project was the construction of a portable apparatus, we began to look at the possibility of constructing a sectioned, collapsible secondary coil. We could not find any references to this having been tried before so we set about designing our own component, by first investigating the Q of various types of secondary coil design. While the electrical loss of the secondary coil is not as crucial in the spark-gap excited system, it still seemed like good practice to build-in high Q if at all possible. The Q for a single-layer coil has been the subject of much deliberation in the past and is difficult to calculate, since it involves not only the inductance and capacitance of the coil, but the wire type and spacing, the operation frequency, and losses such as resistive, dielectric and inductive damping [6–9].

Using the set-up of figure 4 we made measurements of Q for several coils. Each coil was excited by a loosely coupled primary coil, in turn triggered by a single MOSFET gate pulse of precisely the correct duration to open the primary circuit when all the energy is exchanged to the secondary coil, as was the case for the trace of figure 5(d). The decay time τ (1/e time) for the voltage oscillation of each secondary coil was recorded on a storage oscilloscope and gives the Q via the simple relation $Q = \omega\tau$ where $\omega = 2\pi f$ is the angular resonant frequency of the coil, also deducible from the oscilloscope trace. We observed a trend that very wide coils with a height at least four or five times their base diameter had the highest

[†] This number has been cited in several WEB-based resources, presumably having being established over a number of years by various independent experimenters.

Q values. Tall thin coils do not make for good secondary designs. This may be understood, given that the Q should be proportional to $\sqrt{L/C}$ where C scales roughly linearly with coil radius[†] while L scales as the square of the coil radius[‡]. But there are other factors, and Q is inversely proportional to the coil resistance; however, while the DC resistive loss certainly scales linearly with coil radius, the influence of the skin effect is reduced for coils with lower resonance frequency[§]. Since $\omega \approx 1/\sqrt{LC}$, wider coils would still have the advantage in terms of achieving higher Q . Our measurements also suggest that a higher turn density is preferable, presumably because the inductance rises as the square of this while resistive effects again rise only linearly. Meanwhile, damping losses associated with being close to the ground plane reduce the working Q of shallow coils, leaving close-wound, wide but reasonably high designs as the overall optimal choice. Conveniently, this type of secondary coil also maximizes the flux linkage with a wide primary coil, and a reasonably coupled system can be achieved with a fairly low profile primary coil. We have noted no clear advantage in any particular wire type or insulation, although the coils which exhibited very high Q were wound on largely air-frame or tuffnol formers. Air-frame formers have the additional advantage of being low mass and more compatible with portability.

With this in mind we designed a five-section secondary coil, each section constructed from two wooden rings connected by eight plastic tubes approximately 30 cm high. The top ring of each section has a diameter 5 cm smaller than that of the bottom ring giving a tapered side profile with an angle to the vertical of approximately 5° . The top ring of each section has a diameter equal to that of the bottom ring of the next section, yielding an easy-to-assemble conical coil. The bottom diameter of the lowest section is 60 cm, chosen to be the largest dimension still compatible with a reasonably portable apparatus. Each section has 176 turns of close-wound, insulated, stranded wire of 16/0.2 gauge. The sections have 4 mm plugs fitted on to their top ends with mating sockets on their undersides. These electrically define the ends of the coil for each section and also act as mechanical keys so that the coil can be easily assembled while automatically yielding its electrical continuity. The entire coil stands about 1.8 m high and is mounted on a base unit that is designed to accept the primary components. The voltage gradient at the bottom of the coil is 1.6 kV per turn which is within the dielectric strength of the insulation on the 16/0.2 wire. In addition, the total height of the top terminal above the ground should prevent flashover to any of the primary components and eliminate the need for a guard ring.

Finally, and actually rather importantly for the overall usefulness of the Tesla coil, the sections of the secondary coil, when inverted, fit inside one another making for a very portable apparatus as shown in figure 6. We also constructed a metal toroid and sphere which mount to the top section to serve as discharge terminals. The sphere is particularly interesting and will be discussed again later. The secondary coil has a resonance frequency of 110 kHz with the toroid and sphere mounted and a Q of 420 as measured using the ring-down time.

[†] Medhurst's empirical formula for the capacitance C of a single-layer coil is given by

$$C = \left[11.26l + 16r + 76.4 \frac{r^{3/2}}{l^{1/2}} \right] \text{ pF}$$

where l and r are the length and radius of the coil respectively. An even simpler, but less accurate, rule of thumb is to consider the coil diameter in centimetres as the capacitance in picofarads. We have tested many single-layer coils of various dimensions and winding densities and have yet to discover one whose capacitance is not given by the above Medhurst formula to within 10%.

[‡] Wheeler's formula for the inductance of a single-layer coil of length l , radius r and n turns is

$$L = \frac{\mu_0 \pi r^2 n^2}{l + 0.9r}.$$

[§] The skin effect constrains a current at a high frequency f to flow within a certain depth δ from the surface of a conductor given by $\delta = \sqrt{\frac{\rho}{\pi \mu_0 f}}$, where ρ is the resistivity of the material. See for example [10].

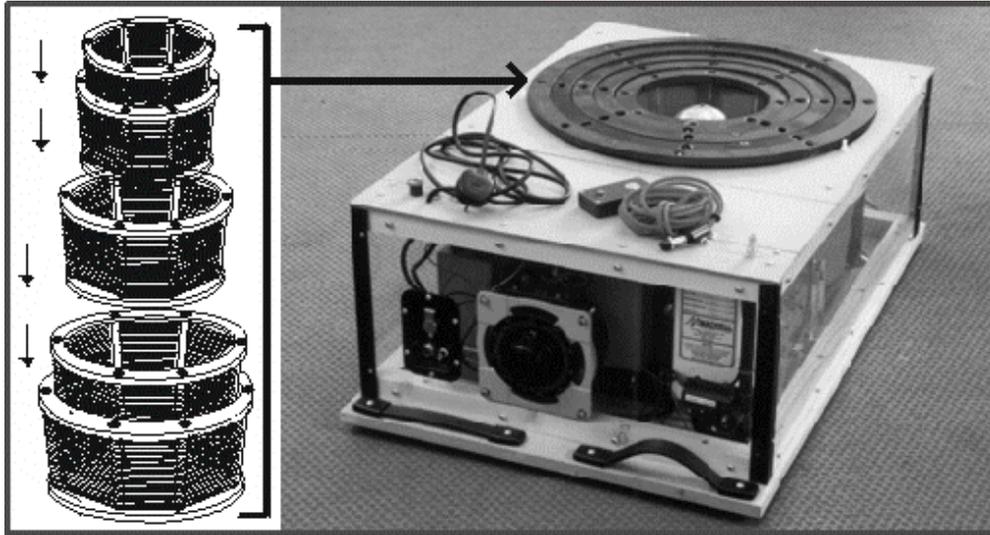


Figure 6. Illustration showing how the collapsible coil sections can be inverted and placed inside one another. The photograph shows the complete apparatus when fully collapsed.

Primary circuit parameters

The decision to choose an input power level first, rather than start with a preferred output spark-length and work back to the input side is born largely from experience, which has taught us that this is by far the best way to work. In order to provide some headroom we chose 1.5 kW as the target maximum input power. In the UK this translates to our 230 V mains supplying an average current of approximately 6.5 A, although it should be noted that there will conceivably be transient currents higher than this. The next step is to have flexibility over this input power; therefore, we employed an 8 A variac to control the mains input voltage. The output of the variac feeds an eht transformer rated at 1.5 kW for non-continuous use (the specifiable duty cycle here considerably reduces the bulk of the transformer). Perhaps the first non-obvious parameter choice is the output voltage rating of the eht transformer. Having settled for a sensible fixed power level, we must accept more voltage at this point only at the expense of reduced current. The question of what is best here depends most immediately on C_1 . For instance, the larger C_1 is the more readily it will charge from a lower-impedance supply transformer. To go to an extreme, there is no point in choosing C_1 so large that it cannot be completely charged within half the mains supply period. On the other hand, a very small C_1 will charge more rapidly but will have less energy to dispense into the Tesla system upon discharge. The peak output voltage V_2 across the Tesla secondary coil must be in proportion to the square root of C_1 ; in fact basic energy conservation implies that

$$V_2 \propto \left(\frac{C_1}{C_2} \right)^{1/2} V_1 \quad (2)$$

with the constant of proportionality being dictated by the power-coupling efficiency of the system. Meanwhile, the output voltage of the eht transformer should most probably change inversely with C_1 . Thus, for example, if we think about operating on the basis of one complete C_1 charge per half mains cycle, then in spite of equation (2), quadrupling C_1 will actually halve the attainable V_2 because of our fixed-input-power condition; this is like a type of power matching, or balancing, for the spark-gap system. A brief, but sensibly cautious corollary here, would be that a large C_1 is not necessarily a good choice. On the other hand, aiming for a low value for C_1 implies specifying a larger working voltage and this has availability

and cost implications. As an example, a capacitance of 100 nF would present an impedance of about 32 k Ω at 50 Hz. For the transformer secondary impedance to be of this order at 1.5 kW would require the voltage to be $\sqrt{1500 \times 32\,000}$, a voltage of 6.9 kV. If we reduce the capacitance by a factor α then the voltage increases by α . Given that there may be transient voltages higher than the working voltage of the capacitor, we decided on a voltage rating of 10 kV which suggests $C_1 \sim 56$ nF. In the end, we settled for a 60 nF, 40 kV, rapid discharge capacitor from Maxwell[†] and a 10 kV specification for the 1.5 kW intermittent duty-cycle eht transformer. From graph (a) in figure 5 we can predict that the peak output voltage should be approximately 400 kV.

The value of C_1 defines the value of L_1 since the resonance given by $1/2\pi\sqrt{L_1C_1}$ should be 110 kHz in order to couple efficiently to the secondary coil. The degrees of freedom are lessening by this stage and the only real question is how to physically construct the primary coil and what coupling coefficient to strive for. Primary windings can grow in a simple helix, or a flat spiral, or may even have a component of both yielding an inverted cone shape. We chose the flexible approach of winding a shallow helical primary coil that was concentrically slightly wider than the secondary coil in order that it can be lowered and raised to adjust coupling. We used thick, insulated, stranded copper wire of 100/0.25 gauge to handle the large peak currents that flow in this coil when C_1 discharges. The last turn of the coil is a solid winding made of aluminium which is connected by means of a moveable clip to aid with fine tuning of the primary circuit resonance frequency. The coil had approximately six turns in total and measured 5 cm high by 65 cm in diameter with an inductance of 3.5 mH $\pm 10\%$. The large area cross section together with the high winding density of the secondary coil make it possible to achieve coupling coefficients of 0.15 to 0.2 with the primary coil completely concealed in the base unit underneath the collapsible secondary.

The spark-gap, while being a very convenient high-voltage switch, is also a resistive loss in the primary circuit. The length of the overall arc should be as small as possible, while the breakdown voltage must be kept at a level just lower than the peak eht transformer voltage. For this reason, curved or flat electrodes are commonly chosen because they minimize electric field strength in the gap. In addition, the arc should quench soon after forming in order to minimize the wasteful energy exchange between the primary and secondary as demonstrated in figure 5(a) and (b). Often, a series of metal plate gaps is used to maximize the arc cooling. Alternatively, a rotary wheel can be used to offer a configuration of spinning electrodes which forces arc quenching at a rate determined by the linear speed of the wheel edge. However, as we might expect, the particular choice of spark-gap is not totally independent from the design parameter set and, in particular, it depends greatly on the initial few decisions regarding the eht transformer and the value of C_1 . Out of interest, we built a 4000 rpm rotary gap for our coil, with the expectation that, while it might improve the arc-quenching time, it would upset the power-transfer conditions that led to the determination of C_1 and the eht rating. Interestingly, the rotary gap did not improve output performance, and was even detrimental to the longest arc length we could achieve. The action of the rotary gap in quenching the primary spark can be seen from figure 7 which shows the beating of energy between the primary and secondary circuits ceasing at the third secondary voltage rise. In order for the rotary gap to be useful with the present eht transformer, C_1 would have to be reduced in value, yielding a more continuous output from the secondary coil, albeit at the expense of less intense arc discharges. Alternatively, the rotary gap could be used with the 60 nF capacitor but the power requirement would increase and a new transformer would be mandatory.

[†] We chose Maxwell Technologies (energy products) model 31393 (AMS Electronic Ltd, Devon, UK supplier) which can handle peak currents of 25 kA. Care must be taken not to choose a capacitor intended primarily for DC use where the power dissipation within the capacitor is typically specified too low for repetitive discharge use. On the other hand, pulse capacitors are not designed for prolonged DC use, but are ideal for the Tesla coil application. It is advisable to check also that sudden polarity reversal is acceptable, since some pulse capacitors have a non-100% tolerance for this, which is clearly unsuitable for Tesla system use. Finally, it should be noted that the life of a pulse capacitor in terms of charge/discharge cycles is very strongly increased the larger is the headroom between the applied voltage and the rated voltage.

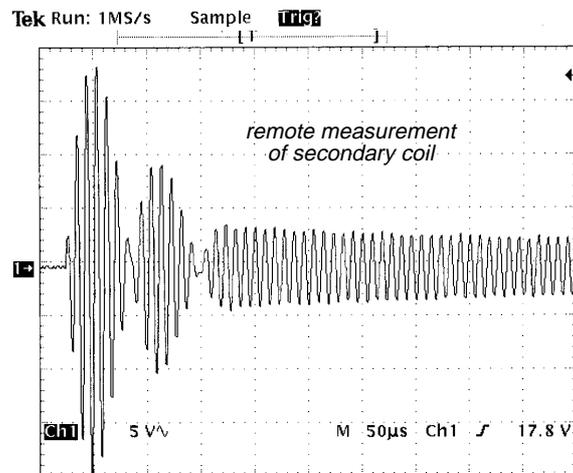


Figure 7. The output of the Tesla coil while operating and producing corona discharge. The measurement is made with a remote antenna placed approximately 10 m from the coil. The action of the rotary gap can be seen via the sudden cessation of energy exchange around the third secondary rise suggesting a primary spark duration of about $150 \mu\text{s}$. We normally operated the Tesla coil with the rotary gap stationary, since on the whole it was detrimental to output performance with our chosen design parameters.

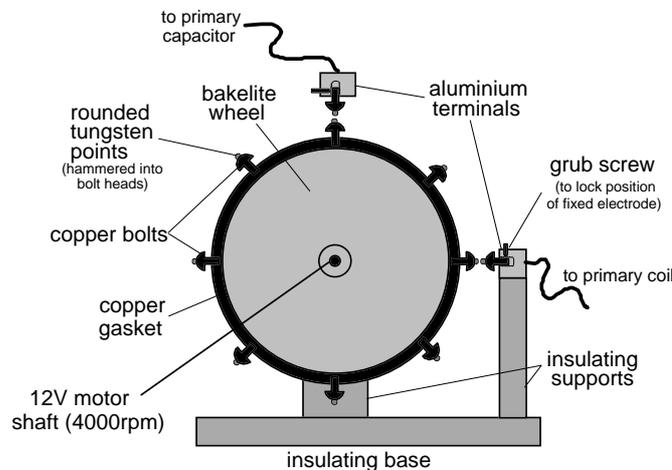


Figure 8. The eight-point rotary spark-gap, which can also be clamped to form a series two-stage stationary gap.

The single most significant improvement we made on the primary circuit side was in taking the effort to change the spark-gap electrodes from conveniently available rounded copper bolt heads to tungsten points. The tungsten points were formed from a length of welding rod and embedded into copper bolts to fix them. The final spark-gap, shown in figure 8, retains the rotary design but can be clamped to form a dual stationary gap. In changing the spark-gap material to tungsten the long-term degradation of the points is much reduced, but even in short-term operation the use of a material that does not burn so readily in the arc makes a significant improvement to the output corona stream. We should point out that readers interested in specifics regarding the control electronics for the Tesla coil, including our remote control

design and measures taken to avoid switching surges, are welcome to contact the authors. A detailed circuit schematic of a similar basic control arrangement can be found in [4].

4. Demonstrations with the Tesla coil

Having dealt with the design challenges of Tesla coil construction, we now describe some demonstrations and experiments that can be performed with the apparatus. We also comment on some practical and safety issues, because although Tesla coils seem to carry an air of impunity when it comes to matters concerning risk, they still encompass the dangerous combination of high voltage and high current, and care must always be taken.

Discharge characteristics and electrode types

The discharges produced by the Tesla coil depend to an extent on the electrodes fitted to the top of the secondary coil. We had two electrodes available, a toroid and a sphere, either or

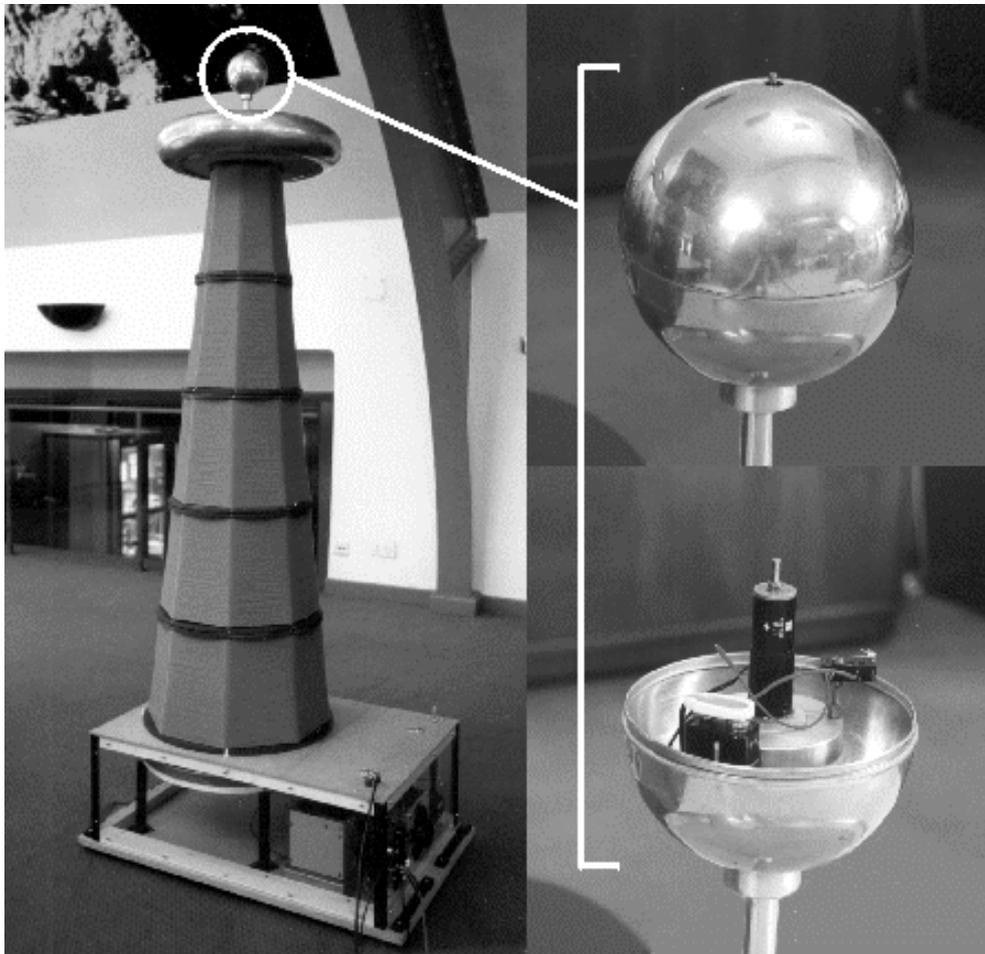


Figure 9. Picture of the assembled Tesla coil and the Faraday sphere showing its interior. The electric motor is switched on by rotating the two halves of the closed sphere with respect to one another.

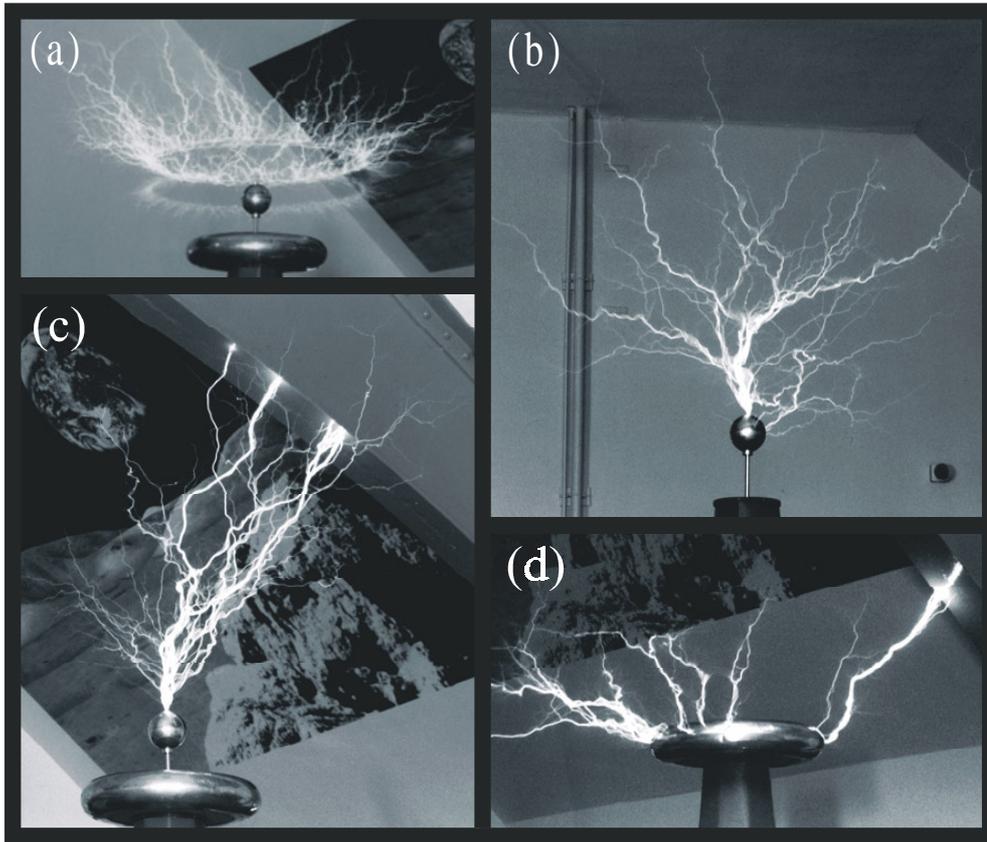


Figure 10. Some illustrations of the Tesla coil output: (a) the ring of corona fire produced by a rotating antenna; (b) corona winding into the space around the Faraday sphere, the roof is approximately 3–4 m above the discharge terminal; (c) intense arcs strike a steel girder at the top of the original Lord Kelvin Lecture Theatre in our department; (d) the presence of the toroid streamlines the corona discharge into a horizontal pattern which is all the more dramatic for spectators.

both of which could be fitted at one time. The assembled coil is shown in figure 9 alongside a more detailed view of the discharge sphere. We call this the Faraday sphere because it contains a small electric motor that rotates a protruding metal shaft. Digital watches or other sensitive pieces of electronics belonging to audience members can be inserted into the sphere and the Tesla coil powered up, whereupon the Faraday cage effect protects the electronics from damage. The purpose of the rotating shaft is to turn specially designed electrodes at the top. For example, figure 10(a) shows our corona antenna fitted to the shaft, giving a spectacular rotating 'ring of fire' effect. The photograph in figure 10(b) shows the corona discharge emanating from the sphere itself. This display was obtained with the variac turned to 80% of maximum. If the discharge electrode is positioned too near a good ground, then an intense discharge occurs where all the energy from the system empties suddenly. An example of this is shown in figures 10(c) and (d) where arcs 1.5 m in length jump to a nearby girder. This type of discharge, although impressive, differs considerably from the more sustained rf corona that it is possible to achieve with the Tesla coil. There are some practical factors associated with such impressive demonstrations. For example, just as the photograph in figure 10(c) was being taken the fire-alarm sirens in our department began to sound, not because of heat or smoke

emission, but rather due to electromagnetic interference with the sensor cabling. Such emi problems with electrical systems are a common consideration when we take our portable Tesla coil on the road to various venues. Indeed, we have found the role of the variac in limiting the input power to be essential to allow the use of the coil in certain locations.

Risk evaluation and the skin effect

The skin effect is often mentioned in connection with a Tesla coil output. It is sometimes even claimed that the output of spark-gap systems is made safe because the skin effect causes the current to flow over the surface of the body, rather than through it. We would take issue with this, primarily because in a spark-gap excited system there is a significant transient component at low frequency caused by the repetitive firing of the spark-gap. If an arc from the secondary coil strikes a person directly this is equivalent to the system suddenly emptying its energy through that person. The skin-effect consideration is no more relevant than if we were considering an individual being struck by a bolt of lightning, for example. It is instructive to consider the consequences of taking a direct spark from the secondary coil. The factors of interest are the voltage and capacitance of the system that discharges the arc, for these will define the energy released and, ultimately, the current that might flow. The secondary coil has a capacitance of approximately $C_2 = 30$ pF. Therefore, with a peak voltage of $V_2 = 500$ kV we might expect a stored energy E_2 given by $E_2 = \frac{1}{2}C_2V_2^2$ which for these values equates to 3.75 J of energy. The current that might flow as a result of this discharge depends on the impedance Z_t encountered to ground. The time taken for the discharge would be of order $Z_t C_1$ and so the power P_t is given by $P_t = E_2/(Z_t C_1)$. The current i_t , given by $\sqrt{P_t/Z_t}$ can then be estimated as $i_t = \sqrt{E_2/C_1}/Z_t$. We may estimate the resistive impedance of the body to be approximately 50 k Ω yielding, with the numbers above, a current of $i_t = 7$ A discharging in a time of approximately 1 μ s. Furthermore, since all of this is averaged over the discharge duration, we will have a peak current even higher than this. There are many assumptions made here; most fundamentally, the arc channel having small impedance compared to the body impedance, the body impedance being measured through the skin, and being purely resistive in nature. Nevertheless, the current is significant and it would be inadvisable to draw a spark from the Tesla secondary coil without the use of a well-grounded discharge wand. Perhaps the most similar situation for a one-off discharge would be the Van de Graaff generator where, in the smaller models designed for schools, one quite routinely takes a spark discharge to the hand. The voltage here is of a similar magnitude to our Tesla coil output, but the metal dome has even smaller capacitance than the Tesla coil and, crucially, the spark frequency is much lower, typically only one intermittent discharge every few seconds, rather than the one hundred per second that could potentially be present in the case of the Tesla coil firing once per mains cycle. By contrast, our continuously driven coil of figure 3, operating as it does at 4 MHz, should exhibit the steady-state behaviour that sustains the skin effect. We have shown this to be the case by taking sparks directly to metal objects held in the hand with no sensation whatsoever. Care must be taken here also, since there is the possibility of rf skin burning. This is particularly undesirable since skin burns obtained in this way can take a long time to heal. We made a key template from card and drew a discharge to the keys held in the hand; the result is shown in figure 11.

Faraday cage demonstrations

The Tesla coil is ideal for giving demonstrations based on Faraday's ice-pail experiment which relies on Gauss's Law to show that the interior of a closed metal container placed in a region of potential gradient is itself electric-field free. As well as the Faraday sphere which serves as one of the top electrodes of the Tesla coil, we have constructed two mesh cages. One is just large enough to fit a watch inside, and is suspended from a tall metal retort stand. In this demonstration, shown in action in figure 12(a), an audience member is asked to donate his or



Figure 11. The heat contained in an rf discharge can be quite intense, although the skin effect defends against risk of electric shock.

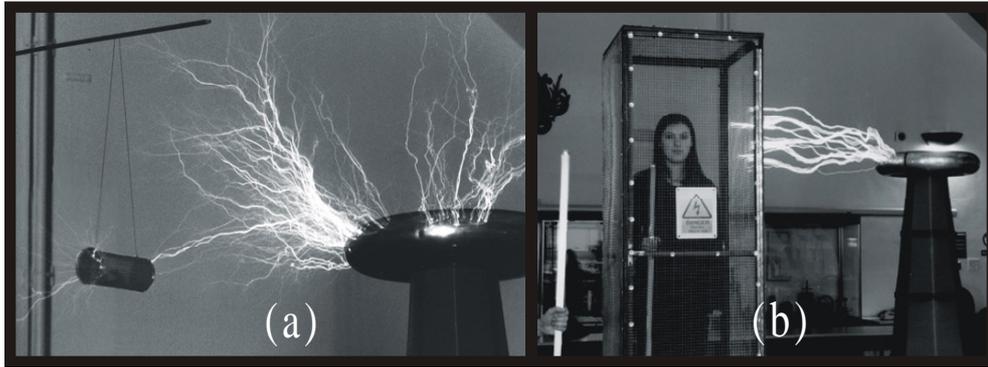


Figure 12. (a) Demonstration of the Faraday cage effect using a watch suspended in a can prepared from steel mesh. (b) A much larger steel cage is used to accommodate a brave volunteer.

her digital watch, a sensitive piece of electronics, and the watch is then placed in the suspended can. The Tesla coil will corona discharge to the can, or else a spark can be drawn via the can to a grounded wand. Either way, the spectacle is in stark contrast to the conclusion whereby the volunteer is given back his or her watch, in full working order! The other cage we have constructed is a collapsible, human-sized Faraday cage. The picture in figure 12(b) shows one of our brave volunteers, Abi Graham, standing inside the cage. The demonstration is further highlighted by the behaviour of the fluorescent lights. The light being held outside the cage illuminates brightly in the presence of the strong electric field around the Tesla coil, while the lamp being held by Abi inside the cage remains unlit.

Further experiments and demonstrations

As well as the usual demonstrations of corona and spark discharge, we set up some other experiments. Although our calculations of the current present in the spark discharge deter us from doing experiments involving discharges contacting individuals, we were still interested in

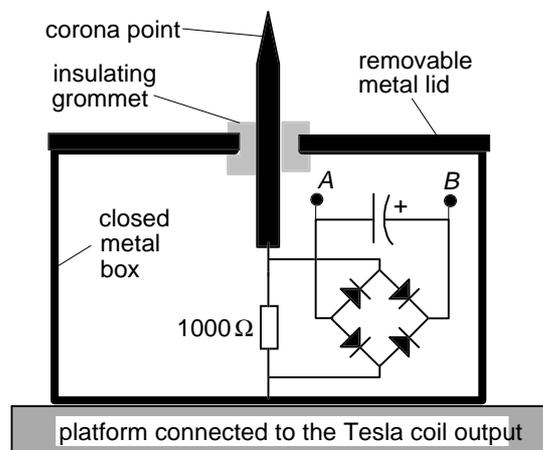


Figure 13. The box designed to measure the integrated current in the rf corona discharge emanating from a point. After running the Tesla coil for a few seconds, the box is opened and the voltage between points *A* and *B* measured.

probing the average current present in a corona discharge. Corona discharge does not produce the type of plasma track that creates a low impedance route for the large flow of current, as in the case of a spark from the secondary coil. However, it is still interesting to investigate the current that is present in the corona, and the experiment we attempted involved placing a measuring device, shown in figure 13, on top of a custom-built insulated platform with a metal floor, shown connected to the Tesla coil in figure 14(a). The device in figure 13 is basically a closed metal box with curved edges to prevent corona leakage. The only opening is in the lid where a sharp metal point protrudes through an insulating grommet. There is a simple internal circuit as shown in figure 13. The current developed by the corona travels through a 1 k Ω resistor across which an alternating voltage develops. Although the Tesla output is typically hundreds of kilovolts, the vast majority of that potential is between the corona point and ground, and the actual voltage across any of the components within the box should be very small. Nevertheless, we used components rated at hundreds of volts to be on the safe side. The voltage developing across the resistor is rectified and goes on to charge a 1000 μ F capacitor which integrates and averages the voltage. The experiment is shown in progress in figure 14(a). After running the Tesla coil for a few times longer than the *RC* time constant, the box is quickly opened and a high-impedance voltmeter takes a reading between points *A* and *B*. The current is then estimated simply as $V_{AB}/1000$. Using this apparatus we measured average currents in the range 10 mA to 100 mA which, although smaller than the peak currents associated with spark discharges, are still significant. In addition, the current is likely to be rather more continuous than with the sparks, possibly lasting for several rf cycles every 10 ms or so.

Before performing this experiment we had, tentatively, prepared some gloves with metal points on the ends. The demonstration we had in mind was to wear the gloves while standing barefoot connected to the Tesla coil; hence the metal platform. The current propagating through the body of the 'volunteer' will be due in part to the capacitance of the body with respect to the surroundings, and to the current channelled into the corona. The previous experiment was conducted because we thought that the latter contribution would be the dominant current. The results were just at the level to be considered acceptable for a low-power test, and the principal author himself volunteered to try the gloves out. Certainly, one can experience the concept of capacitance at first hand with such an experiment, since before the gloves were even exposed, there is a noticeable feeling of discomfort primarily in the feet! The corona gloves experiment

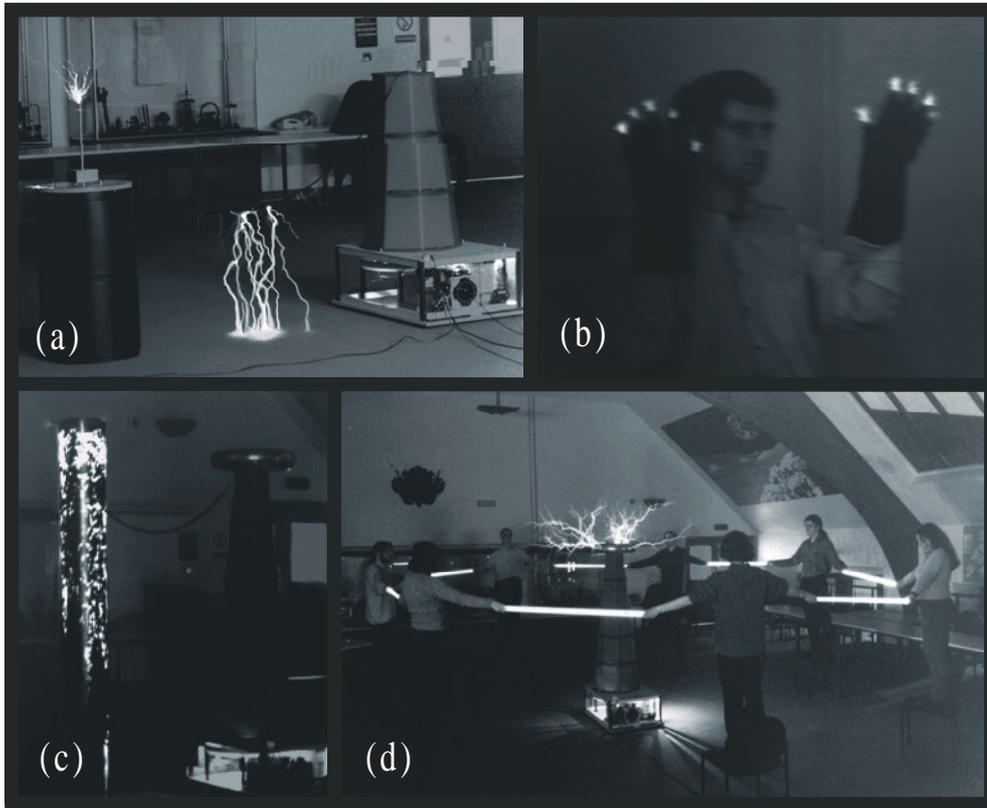


Figure 14. Some pictures of the further experiments we performed. (a) Partially assembled Tesla coil capacitively balanced through being connected to a metal floored insulating platform. Some corona can be seen emanating from the current measuring box mounted on this platform, as well as some unintentional flashover to the floor. (b) Metal-spiked gloves being worn while standing on the insulating platform. The motion blur is hard to avoid, since a rather long exposure is required to capture the corona glow on camera. (c) The Tesla coil used to create a luminous column effect. (d) A human loop interspersed with fluorescent tube lamps.

can be seen in figure 14(b), but it requires a very dark room and is probably much suppressed in view of the capacitive leakage associated with the body. This is not surprising, since if we simplistically think of the human being as an isolated sphere of radius a , then the capacitance of the person with respect to the surroundings is given by $C = 4\pi\epsilon_0 a$, and is seen to offer only a modest impedance to the component of the current at the rf frequency f , given by $Z_C = 1/(8\pi^2\epsilon_0 a f)$. The current estimated from $i = V_2/Z_C$ is then seen to be greater than the corona current measured above, although clearly, there is considerable uncertainty in such simple calculations.

Another experiment is shown in figure 14(c) where a plastic pipe has been very roughly coated with graphite paint. When the Tesla coil secondary coil is connected across the tube, the electricity appears to cause the surface of the tube to glow as the various discharges bridge the gaps in the patchy conductive paint.

The final experiment we report on is an elaboration on the well tried and tested theme of lighting fluorescent lamps by holding them in the electric field surrounding a Tesla coil. The experiment we tried is shown in figure 14(d) and involves a ring of seven people each electrically joined to the next person by means of a fluorescent lamp. The human ring acts as

an inductive loop to an extent and the fluorescent lamps light due to circulating current. If the loop is broken, by somebody letting go of one end of a tube then the tube brightnesses change, although they still glow due to being present in the electric field generated by the secondary coil.

5. Conclusions

We have performed experiments to investigate both continuously driven and transiently excited Tesla coil systems. We have followed up these investigations by designing and constructing a portable Tesla apparatus for use in schools and public venues as well as within the University as a lecture demonstration and teaching aid. Consideration of the physics associated with parameter choices has guided the way in which our apparatus has evolved. The Tesla coil implements the novel feature of a collapsible secondary coil and is excited by a spark-gap employing a rotary wheel which can be rotated or clamped to demonstrate the different effects on output performance. We have conducted low-voltage experiments which are very useful in showing how the energy exchange in a transient excited system occurs and have confirmed the behaviour of the high-voltage system as predicted by MAPLE models. Finally, our Tesla coil has now been used in over 40 public and schools science shows, and we have described some of the experiments and demonstrations that help make seeing a working Tesla coil such an unforgettable experience.

Acknowledgments

We should like to thank the West of Scotland Science and Technology Regional Organisation (SATRO) for its interest in and support for this project. We should also like to thank Colin Craig and Peter Barbour for their valuable technical assistance during the course of the work. The Tesla coil electronics and base unit design were originally developed by Sian Scott, Kenneth Skeldon and Alastair Grant during a vacation project. Financial assistance for the portable Tesla coil was provided by SATRO West Scotland, the Royal Society of Edinburgh (RSE) and the University of Glasgow. Kenneth Skeldon is in receipt of a BP/RSE Research Fellowship while Alastair Grant is an Honorary Research Associate in our department. Gillian MacLellan and Christine McArthur were partially supported by student support grants. We should also like to thank Colin Craig and Peter Barbour for their valuable technical assistance during construction, and Alban Chadenas, Abi Graham, Iain McVicar, Janet Milne, John Nelson and Philipp Steinmann for their help with the demonstrations and photography.

References

- [1] Kelly J B and Dunbar L Sr 1952 The Tesla coil *Am. J. Phys.* **20** 32–35
- [2] The electricity show that we present to schools and the public can be read about in: Skeldon K D *Arcs and Sparks, Catalyst* (Glasgow: Glasgow University Publishing)
- [3] Bleaney B I and Bleaney B 1965 *Electricity and Magnetism* (Oxford: Oxford University Press)
- [4] Skeldon K D, Grant A I and Scott S A 1997 A high potential Tesla coil impulse generator for lecture demonstrations and science exhibitions *Am. J. Phys.* **65** 744–54
- [5] Bruns D G 1992 A solid state low-voltage Tesla coil demonstrator *Am. J. Phys.* **60** 797–803
- [6] Medhurst R G 1942 HF resistance and self-capacitance of single-layer solenoids *Proc. IRE* **30** 412–24
- [7] Callendar M V 1947 Q of solenoidal coils *Wireless Eng.* **24** 185
- [8] Terman F E 1955 *Electronic and Radio Engineering* (New York: McGraw-Hill)
- [9] Anacin B A, Davidovic D M, Karanovic P, Miljevic V M and Radojevic V 1997 Circuit properties of coils *IEE Proc. Sci. Meas. Technol.* **144** 234–9
- [10] Ramo S and Whinnery J R 1945 *Fields and Waves in Modern Radio* (Chichester: Wiley)