build a Tesla lighting plant

Son of TESLA COIL

by George Trinkaus

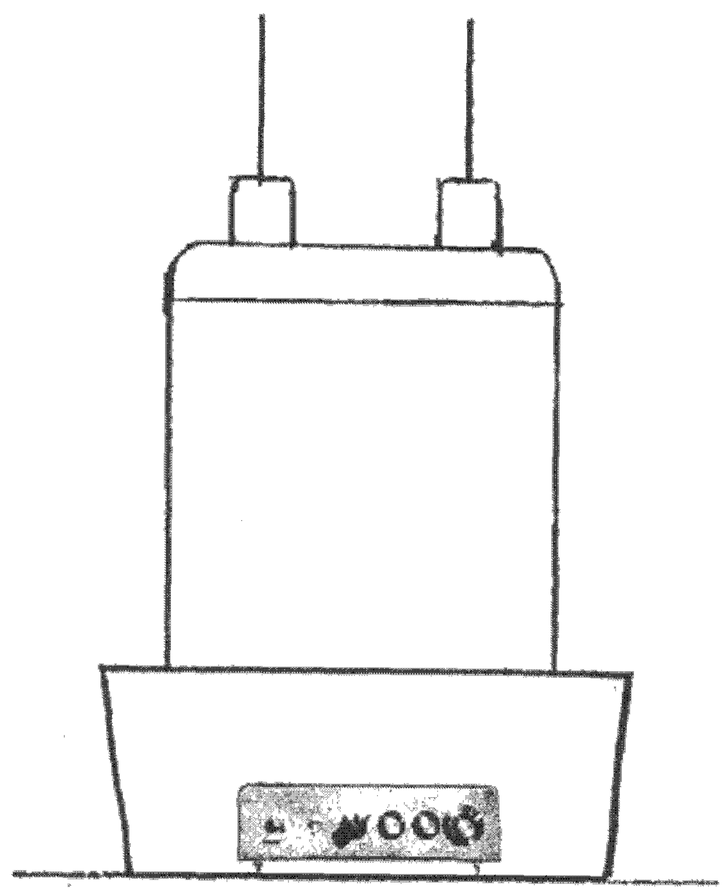

We've thrown lots of sparks, but is that all Tesla had in mind? In this sequel to the classic *Tesla Coil,* we venture into the utilitarian Tesla. In clear English and 40 illustrations, we ponder the questions: Can Tesla coils do real work? Can they power an alternative fluorescent lighting? Build a "third generation" Tesla coil. Discover the magic of oil immersion, of close coupling, and of precise pulsing. Learn how to drive a Tesla coil from low-voltage battery power. See how the new neon technology can be adapted for Tesla off-the-shelf. Build a solid-state pulse generator and a Tesla lighting plant. Then hook up your ammeter and see what it says about Tesla efficiencies. Can battery-powered Tesla technology come to our aid in the next big blackout? Will Tesla technology be the next big step forward in home power?

1. A Third Generation?

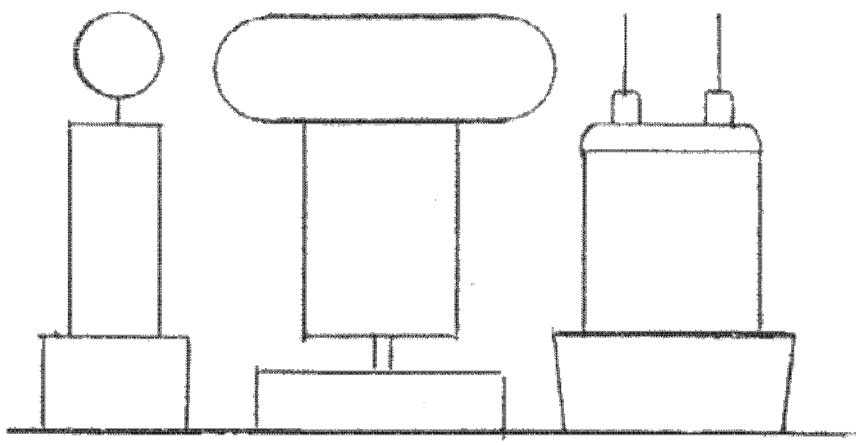

You can find all three generations of Tesla coil in Tesla's own work, but one particular configuration, the secondary wound with a single layer of fine wire on an elongated cylinder, has dominated Tesla coil building by experimenters for nearly a century.

Why? Perhaps the inspiration came from the publication of the sensational Century Magazine photos of 1900 showing Tesla's Colorado Springs helical coils. This image, in the absence of others that might have been published had Tesla's work not been otherwise suppressed, may have dominated the public consciousness to the exclusion of other designs that might have caught on.

Generation I

A long, narrow helical secondary coil; a helical primary of several turns of heavy wire set far apart from the secondary (loosely coupled); the terminal, a sphere or whatever, relatively small in diameter: with some notable exceptions, this basic format dominated Tesla coil building through much of the 20th century.

Generation II

The Gen II coil evolved in the 1980's and is now the dominant fashion among the cognoscenti of Tesla-coiling who keep up with the Tesla Coil Builders Association News. Championed in that publication, especially by the amazing Tesla Coil Builders of Richmond, Virginia, this evolution retains the traditional single-layer helical secondary, but now there is a greater appreciation of the larger diameter secondary in proportion to length. The primary is likely to be a flat spiral but may be a Gen I helical or cone-like, expanding outward as it rises.

A loose coupling between primary and secondary is appreciated more than ever, and, if it can be mechanically devised, the separation between primary and secondary is changeable, a variable coupling. The argument for loose coupling is not just that it provides insurance against arcing between the coils but that in closer coupling the action of the primary has a damping effect on secondary oscillations. The coils should be isolated so that they can swing freely without interference. This is the argument. (An exception is the magnifying transmitter whose primary and secondary are closely coupled while the third "extra" coil is isolated.)

The most distinctive feature of the Gen II coil, though, is its huge terminal, which is donut-shaped, a toroid. This is fitted snugly down over the top of the secondary instead of ascending upward like the old Gen I ball.

Gen II builders appreciate that the larger the terminal load the higher the voltage output. Huge amplifications are possible with this technique. And, as the terminal is enlarged and the voltage magnifies, the increased capacity afforded by the larger terminal brings the resonant frequency down. This is desirable because experimenter spark systems are rarely capable of vibrations above one hundred kilocycles, and the same also argues, of course, for a larger secondary coil as well. My

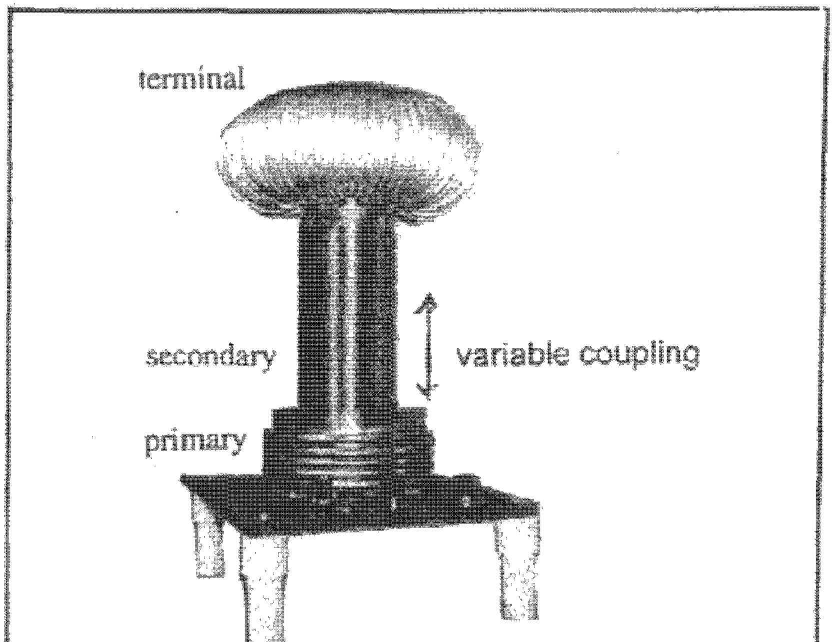

author's Gen II coil

Gen II coil shown here, however, is a miniature (secondary only 8 inches long by 3-inch diameter), and with the toroid is resonant up at 540 kc, so it has never seen one-to-one pulsing on a spark system.

As Tesla coil building evolves, there is greater appreciation of matching of the fundamental frequency of the primary pulsing system with the resonant frequency of the secondary. This may be more manageable with solid-state pulsing systems than with spark-gap systems, as we shall see.

Generation III

The Tesla coil I'm calling Gen III is inspired by the coil demonstrated by Nikola Tesla at his London Lecture of 1892, and, who knows, if this lecture and the drawings published with it (in 1894 by James Commerford Martin) had been more widely circulated through

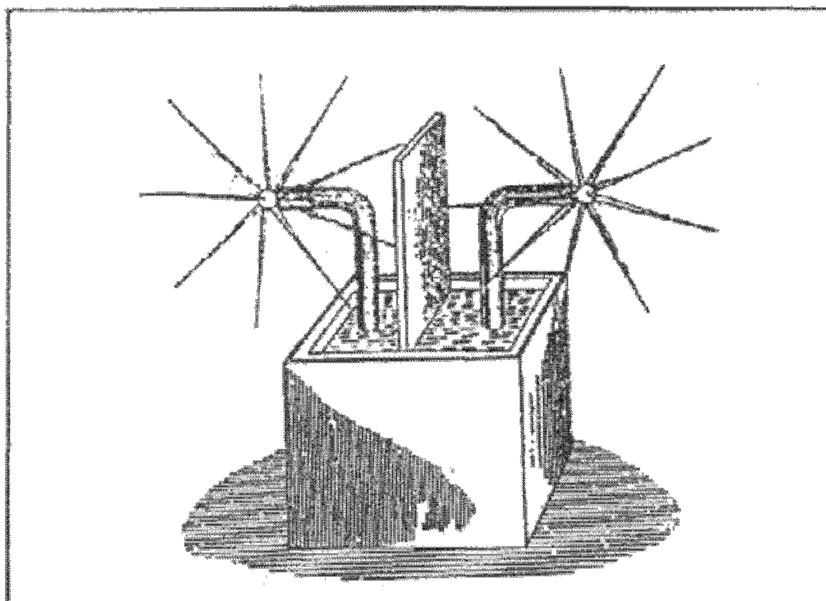

Tesla's London coil

subsequent decades, then this model, through more difficult to construct, may have become the experimenter's standard.

The proposed Gen III Tesla coil is a radical departure. Its format has more in common with the old induction coil than with any of the traditional single-layer helicals. It is a multi-layer coil. The secondary is wound right on top of the primary, closely coupled, not loosely.

How can this be?

My solid-state mentor, Jim Campos (who has worked with me on spark systems as well as transistor, and from whom we will hear much more later) argues that loose coupling is only necessary to compensate for the vagaries of spark. If the primary circuit goes completely off, then it is "isolated." There is no damping of the secondary oscillation.

From the standpoint of efficient transfer of energy, of course, a close coupling is desirable. A close coupling is especially desirable when pulsing the primary coil with relatively low voltages, as will be the case when pulsed from solid state circuitry.

oil immersion

With a multilayer secondary and close coupling, insulation becomes critical. The solution is immersion in oil. Says Tesla, "As sparks would soon destroy the insulation, it is necessary to prevent them. This is best done by immersing the coil in a good liquid insulator, such as boiled-out oil. Immersion in a liquid may be considered almost an absolute necessity for the continued and successful working of such a coil."

Liquid insulation is self-repairing. This self-repairing property can be obtained to a degree even with your conventional open-air helical if you spray on heavy layers of silicone. Spray silicone can be found in the lubricant section of your hardware store.

Oil immersion is feasible for a Gen II helical secondary providing that a second concentric tube can be slipped over it, sealed, and oil poured in between. I have that plan for my little Gen II helical shown, although the variable-coupling feature makes sealing in the oil an engineering challenge.

Before a spell-bound audience of electrical engineers in London in 1892, Tesla drew from his oil coil streamers of fire and strange white luminous brushes. "A most curious form of discharge is observed with such a coil," he noted.

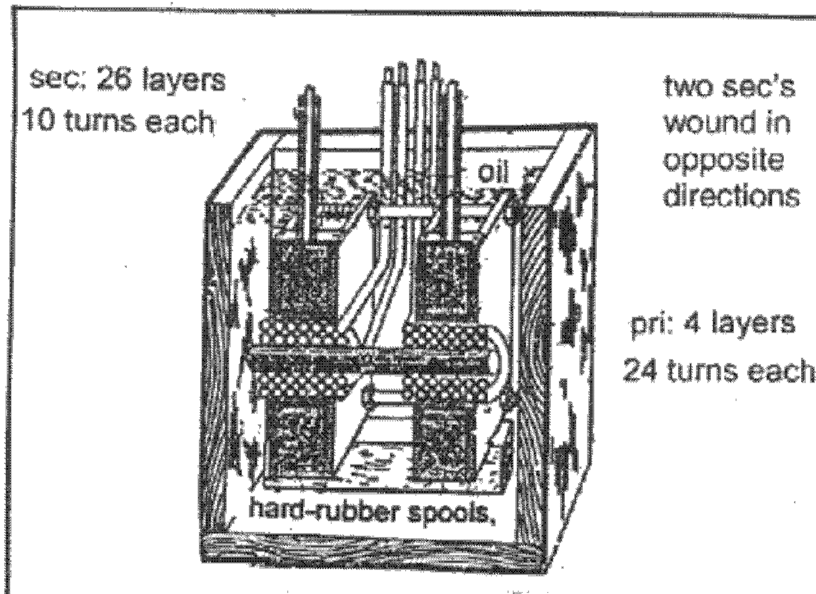

Tesla's London coil, detail

In praise of oil, Tesla said, "I am led to believe that in our future distribution of electrical energy by currents of very high tension, liquid insulation will be used." Tesla envisioned a high-voltage coaxial cable which contained an oil insulation. "The cost is a great drawback," he continues in his lecture, "but if we employ an oil as insulator the distribution of electrical energy with something like 100,000 volts, and even more, becomes at least with higher frequencies, so easy that it could hardly be called an engineering feat (even) at distances of as much as a thousand miles."

Experience with my own oil coil suggests that an empowerment is obtained exceeding that obtained by the mere insulating effect of oil. There is a dividend, a special advantage, secured by oil's dielectric dynamics.

capacitive action

Is the Tesla coil really a capacitor? Conventional electrophysics construes the Tesla coil in electromagnetic terms, the primary and secondary as inductances. And, oh yes, a whole mathematics is worked out accordingly. But Tesla researcher Gerry Vassalatos, who has studied some of Tesla's unpublished notes, says Tesla early on experimented with "coils" that were capacitive, being smooth copper cylinders. These he later spiral-grooved on a lathe, increasing greatly the area and hence capacity. Next, a similar and greater effect from bare wire windings. Did the helical thus evolve?

The dielectric constant of oil is more than twice that of the driest air, the puncture voltage ten times that of air.

In his lecture Tesla also recommends for his oil coil a variable condenser (capacitor), which "should be an oil condenser by all means, as in using an air condenser considerable energy might be wasted." In his patents Tesla shows the entire system, including the spark gap, in flowing oil.

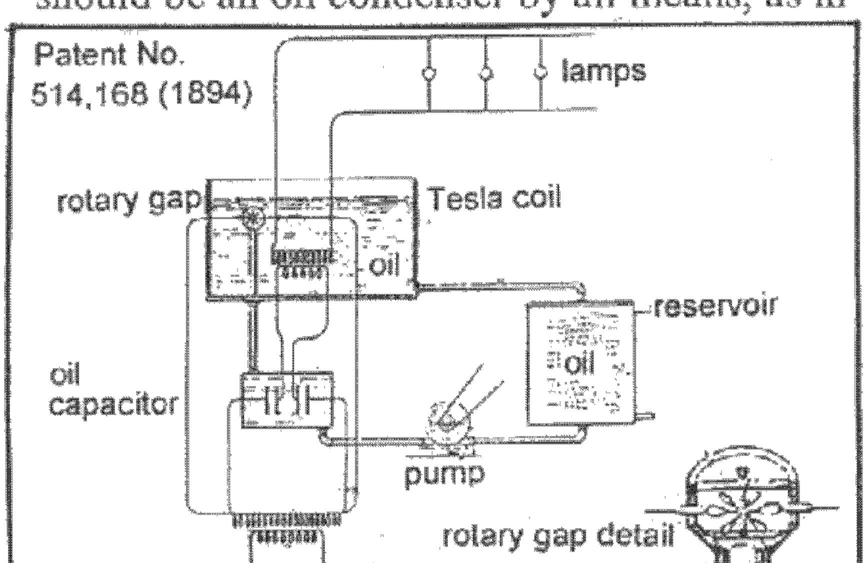

Tesla's oil-coil patent

My oil coil, a simplification of Tesla's London oil coil produces the whitest, hottest, loudest arcs of any I have built. I Have become a believer in oil magic. Chapter 2 tells how to build an oil coil.

coils that do real work

Did Tesla invent his coil just to throw sparks? One might think so given the direction Tesla coil experimenting has taken over the decades. No doubt Tesla enjoyed his coronas, sparks, and streamers as much as we do, but Tesla's effort was directed at the development of high-efficiency lighting plants and other practical equipment in which arcs and coronas could be destructive forces. This will be our effort, too.

Engineering for pyrotechnics teaches us some electrical-engineering lessons on the controlling of awesome potentials, but, in designing systems that can do practical tasks, we honor some important engineering criteria that in designing to throw sparks we customarily throw to the winds. These are:

1. **low current draw**
2. **long duty cycle**
3. **low maintenance**
4. **quiet operation**
5. **emission-free operation**
6. **compactness**
7. **safety**

These criteria amount to a big challenge to the builder. In disruptive-discharge systems, the spark gap alone presents some tremendous engineering problems.

build a Tesla lighting plant

Tesla invented the resonant transformer that came to be called the Tesla coil for the specific purpose of producing high-voltage, high-frequency currents to stimulate with unprecedented efficiencies appliances that would produce light. Let's pick up where that tradition leaves off.

With an eye to the above seven criteria, we will construct a high efficiency oil-coil lighting appliance that can run from batteries. Chapter 5 tells you how to build a solid-state pulse generator for a Tesla lighting plant that can drive fluorescent tubes from d.c. sources 12 volts and up.

in praise of tamer coils

It is amazing what can be done with relatively low power Tesla coils. It is the tamer coil that we can put to practical tasks.

Low power Tesla coils also provide some unexpected benefits for the experimenter. Try a variable transformer between the wall socket and that big supply transformer on your monster spark-thrower. Tone down that noisy, scary machine to a quieter more modest level and you might discover a whole new world of experimental delights. Experimenters often refer to the variable transformer as a "variac." Variac is a brand, as are Powerstat, Superior, Ohmite, and Staco (shown).

Have you observed the phenomenon of the sparkless spark gap? Neither had I until I tamed it down with a variac (see Chap 4). High-frequency electrotherapy is done with a coil of restrained output. Many adventurous experiments can be conducted quietly and safely. For low-power experimenting, I built what I call my "dream' power supply around a small Powerstat variable transformer (270 VA), a small open-core neon transformer (5 kv, 20 MA), and a bridge rectifier. The outputs are 0-140 volts a.c., 0-5,000 volts a.c., and 0-90 volts filtered d.c.

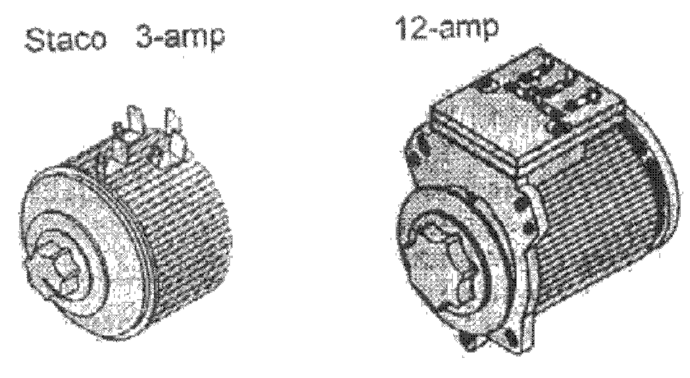

variable transformers (variacs)

Practical Tesla coils are smaller and tamer than what many of us experimenters have become accustomed to. All over the world there are big macho machines sitting (mostly idle) in basements and garages, and these monsters by themselves can generate enough Tesla power to light up the neighborhood. But, for local purposes in the home, farm, or small workplace, what is appropriate is the modestly scaled coil.

2. Build an Oil Coil

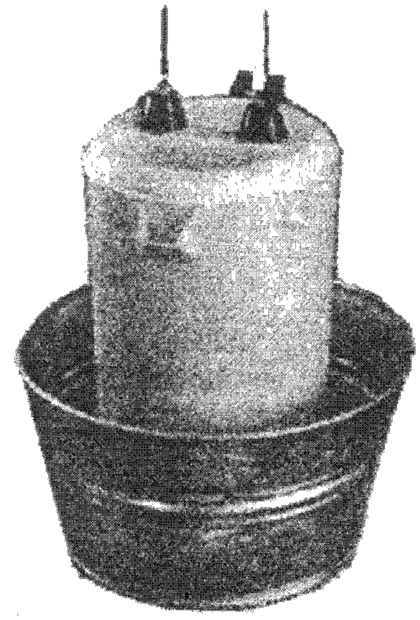

If you still want to build a veritable flame-thrower, if you are an unrepentant builder of macho coils where value is measured in coronal and streamer pyrotechnics, read the following but ignore my recipe and instead challenge yourself by building Tesla's 1892 London oil coil shown on page 3. My recipe oil coil is simpler and easier to construct. It is designed for failure-free constant running in the spark mode with only 5kv input from a transformer, and, in the solid-state mode, will probably never see much over 50 kv at the final output terminal. The recipe oil coil is resonant at 460 kc unloaded and, loaded with a four-foot florescent tube, at 115 kc.

Since both Tesla's original and my recipe coil are oil-insulated, what is the difference that enables Tesla's to produce grander voltages? Tesla's original has more layers, and the width is narrower. This engineering respects the hazard of inter-layer arcing. The narrowness allows less voltage increase per layer and hence less voltage difference between layers. If you are building for pyrotechnics, you should also imitate the best you can the thick insulation Tesla recommends in primary and secondary conductors. If you go all the way with Tesla's original and provide two coils, note that they are wound in opposite directions. Evidently, this provides for push-pull action. If you succeed with this coil, it should be dynamite. Please send me photos.

Tesla's patents show a multilayer coil, not necessarily in oil. It has a single flat-spiral primary and two secondaries.

preparations

Whichever oil coil you make, the recipe or the London, the insulation of off-the-shelf wire will be inadequate layer-to-layer, so for the oil coil you will wind on between each layer the material Tesla recommends, which is oil-soaked muslin. Muslin means cotton, as in bed sheets, so you will need to scrounge an old bed sheet (or, in desperation, a new one, if she's not looking) and rip the thing, in the direction that it wants to rip evenly, into strips of about two inches in width. Make yourself a stash of these in lengths of about ten inches or so.

The oil can be automotive oil; that's the cheapest liquid insulation you'll find. Tesla used boiled-out linseed oil. I considered using one of the new synthetic motor oils since these are supposed to take more heat. That's a consideration since a tiny trapped bubble can heat up to 4000 degrees F. (Tesla notes, though, that such bubbles will be driven off by high-frequency currents. That is, if they are free to move.) Synthetics being expensive, I settled for a synthetic-fossil blend, which was Quaker State 4x4 at $3.95. Regular motor oil (nondetergent) is cheaper and should work just fine. Get three or four quarts.

My oil coil is wound on a homemade spool the core of which is a 7-inch plastic pipe 1 1/2 inches in outside diameter. The spool-ends, which you can improvise any way you want (they need not be perfect circles), I cut out of 3/16-inch ply using a circle cutter (fully extended) on a drill press. With the same cutter, I made a 1 1/2-inch hole in the center into which I fitted the core pipe.

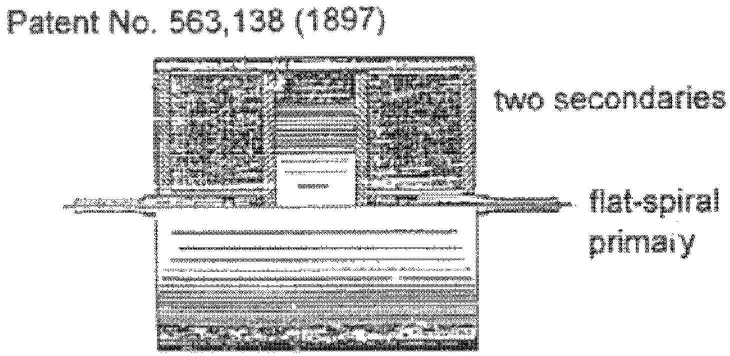

Tesla's multilayer coil

The job is best done on a winding jig like the one shown; in fact, it's hard to imagine doing the job without such a tool. The ideal way to wind an oil coil, Tesla's way, is with the whole coil and winding apparatus submerged in oil, thus assuring the complete displacement of air. You may have the capability of contriving such an arrangement, I didn't, but there are other ways of removing air. The process of winding on wire is itself air-removing.

Needless to say, even unsubmerged, winding an oil coil can be a messy job. Put a tray under the coil to contain the profusions of oil that will be dripping off as you wind. Put newspaper everywhere. Keep a roll of paper towels handy and a big metal wastebasket. By all means get yourself a box of latex surgical gloves (a drugstore item) and be prepared to change them frequently. Promptly dispose of all oil-soaked material in a safe manner that won't be prone to spontaneous combustion.

winding and wrapping

Put some oil in the tray beneath the coil and in it lay several of your bed-sheet strips. Wind some oil-soaked strips onto the bare spool core and onto this wind the first layer of wire, the primary coil.

My primary consists of just one layer of #18 very well insulated high-voltage 5kv test wire, very flexible and easily wound. (I have found many uses for this wire in high-voltage work.) In my next coil, I may try multiple primary layers, and maybe, in respect to Tesla's principle of having primary and secondary of equal weights of copper, use a heavier conductor. So far my impression is that the character of the primary is not so critical, but this impression may change with more experimentation. A design alternative could have the primary as a separate coil which is inserted into the hollow, oil-filled core.

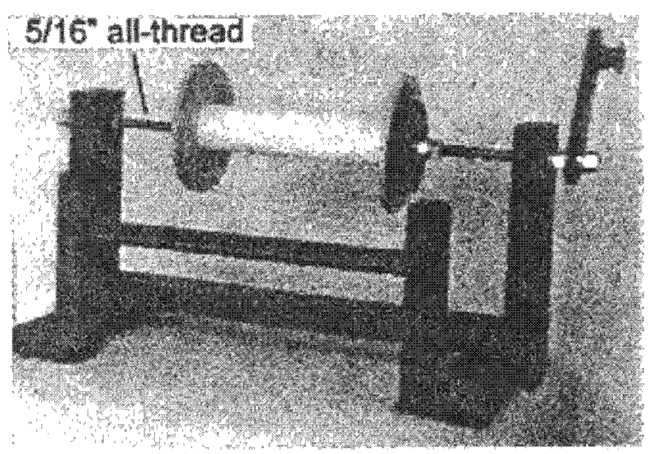

author's winding jig

Drill the spool end and pass the wire through and wind on a primary layer or two. On top of the primary winding, wrap two or three thicknesses of your oil-soaked muslin strips. The breakdown between primary and secondary is a real concern, so especially if your primary wire is weakly insulated, lay on the oiled strips.

As you are winding on the wire, make sure the coil is constantly saturated with oil so that it is wet and dripping. Baste the whole length of the coil while winding by dipping your latex-gloved fingers into the oil tray and dripping the oil onto wire and muslin. When you rest between layers, wrap the coil loosely in plastic sheet to keep dirt out.

My secondary conductor is plastic insulated, #20 solid, a fat wire as Tesla secondaries go but defensible since it's tough and resistant to breaking as you stretch it taut in the act of winding, and it is also capable of carrying lots of current. You may want to use a smaller gauge wire, but a plastic insulation, rather than enamel, is appropriate for a multi-layer coil. The #20 gives 15 turns per inch, 105 turns per layer. Drill a little hole in the spool-end to let the wire in; you'll drill another at the end of the final layer to let the wire out. As you wind the wire on, keep it taut and it will squeeze out the oil as it progresses down the length of the form.

You'll be putting on six layers (total turns: about 630). Between each layer, you'll be putting on more oiled muslin, and, since voltage will build layer-to-layer, each successive secondary layer gets progressively more thickness of insulative cloth between. The final layer needs no cloth over it.

containment

Promptly remove the coil from the winding jig and slowly submerge it in oil. Use some appropriate insulative container. Mine is a plastic one-gallon pitcher from our old friends at Rubbermaid. I have considered, in the interest of a better appearance, a tiny aquarium of hexagonal cross section called a Minihex from Island Aqarium. Tesla's container was a sealed wooden box with an outer lining of zinc sheet.

Run the wire leads out of the container and supply some appropriate primary and secondary terminals to connect to.

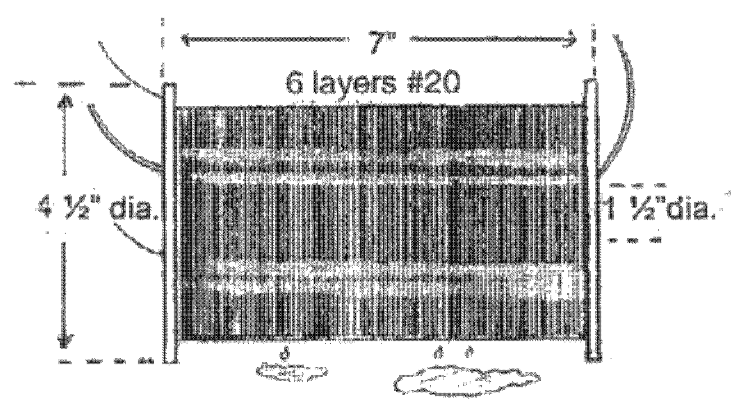

recipe oil coil

Hazard note: Particularly if you are going to push this coil to the limit, and since your insulative container may be of some fusible material, take the precaution of setting the whole thing in a metal container that has enough capacity to serve as a containment birm in the case of meltdown.

This said, I will now confess that I have put into the primary system of the recipe oil coil as much as 12,000 volts at 30 MA for brief periods and with gap closed down without incident. After I exhaust my curiosity with a series of experiments, I intend to run it wide open to determine at what point, if any, I can make this oil coil fail.

3.An Alternative Fluorescent

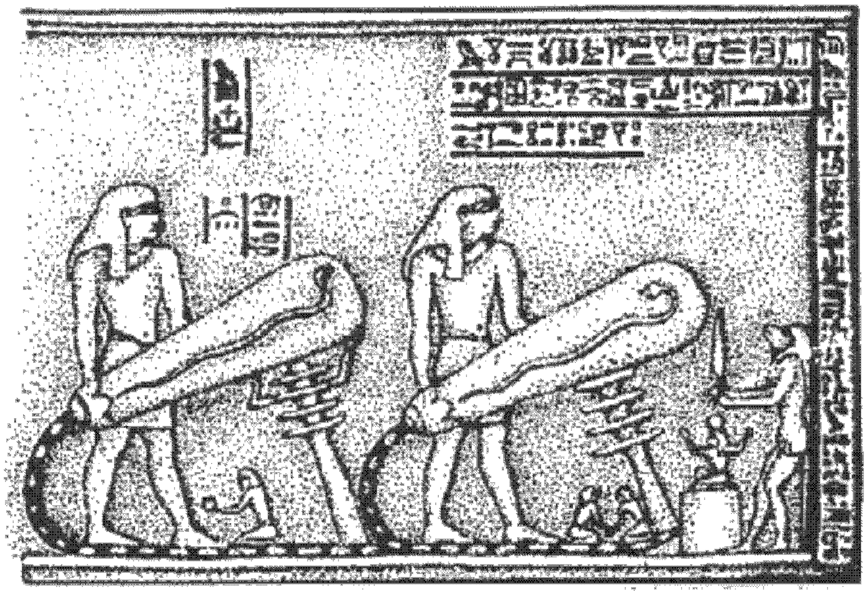

ancient Egyptian wall carving

Tesla-coilers know that, under the influence of high-potential, high-frequency currents, a fluorescent tube just wants to light up; there is a natural compatibility. Hold a fluorescent tube in the vicinity of a Tesla coil terminal, and it lights up wirelessly. A fluorescent tube forgotten in some corner of the lab is sometimes unexpectedly seen flickering spontaneously like it wants to get in on the fun.

You have no Tesla coil? You need only take a fluorescent tube into a darkened room to get a demonstration of its electrostatic receptivity. Just rub the tube briskly with a piece of fur or cloth, the hair on the back of your arm, your cat, and it will produce little flashes of light, so naturally. Fluorescent tubes just want to light up under the slightest electrostatic disturbance, whether directly wired to a source or not.

conventional fluorescent

Compare the standard fluorescent fixture wired into 120-volt 60-cycle a.c. The apparatus needed to light up the tube is surprisingly complex. By comparison, it's like the tube has to be clobbered into activity.

A starting device is needed to provide a jolt at two to five times the line voltage, up to 600 volts, to awaken the tube. A clever little automatic thermal starting switch is required as well as step-up apparatus. The tube also requires a steady electric flow generated by the burning of cathode filaments. These erode slowly as they burn at up to 950 degrees F and, of course, make the tube emit some heat. Like the filaments in an Edison incandescent, it's the on-and-off especially that fatigues them. One start equals three hours of steady burning. One thousand starts kills them. Our Tesla system suggested here has no use for these cathodes except as cold electrodes, so it's ok if tubes are burnt out in the cathode department.

The starting jolt establishes an arc through the interior argon gas and mercury vapor, and the resulting ultraviolet light stimulates to luminescence the phosphors that coat the inside surface of the tube. **Caution:** the phosphors within those fragile florescent tubes are toxic; they make cuts resulting from the implosion of these fragile lamps particularly painful and persistent.

Conventional full-spectrum fluorescents, for growing or for environmental eye-ease, are painted with a mixture of different phosphorescent powders that produce a more natural sun-like emission.

"ballasts"

Once conducting, the tube has almost no resistance and would present a dead short to the circuit were it not for the so-called ballast transformer, which is not a transformer in this function but a choke coil, and was called that in the early literature. The current-limiting choke's impedance eats up some current itself, adding 10 to 20 percent to the wattage. A tube's rating does not include the ballast, and for a tube rated 40 you can add about 8 more watts. A variation, the instant-start Slimline, has a ten-pound ballast that consumes 11 watts. The ballast choke can also be used to supply the starting jolt. When the starting switch opens, the choke's field collapses, inducing a high-voltage surge. There may be a step-down transformer winding on the ballast used for a preheat circuit. Heavy ballasts require that florescent-tube fixtures be of heavy, cumbersome construction, although the tube itself is feather-weight and easy to support.

There is a small capacitor in the starting switch to suppress radio interference when it goes on and off, and another suppression capacitor across the line, and yet another may be used to put two lamps out of phase.

You cannot dim these fluorescents. Special ones that dim require yet more ancillary hardware.

fluorescent sensitivities

Whether you drive them conventionally or alternatively, fluorescent tubes are sensitive to both temperature and humidity.

Humidity increase means it takes more power to drive a tube, and particularly humidity affects a tube's willingness to start up. In the conventional mode, the voltage necessary to start can vary from 250 to 750 volts depending on humidity. Moisture on the surface degrades the glass' dielectric quality so energy is not concentrated at the electrodes where it belongs. I have gotten brighter lights and easier start-ups in experiments with Jim in drier California than with tubes similarly driven in my own lab up in rainy Oregon. To fight humidity, in manufacturing fluorescent tubes are coated with silicone. This surface makes moisture bead up instead of forming a conductive sheen. Since this coating may be worn off of older tubes, I spray mine with silicone lubricant.

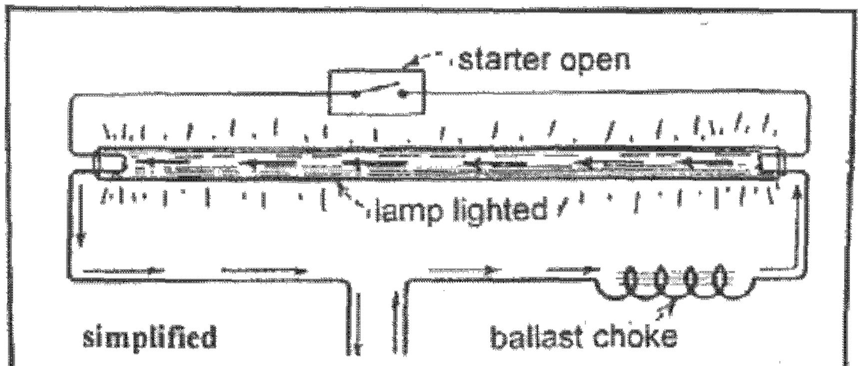

conventional fluorescent circuit

Temperature affects the performance of these cranky fluorescents. A tube designed for indoor use will fail completely outside in the cold. Outdoor tubes wear a jacket of vinyl tubing. Such a sheath may offer protection against humidity, too, as well as breakage.

Blackening of fluorescent tubes results from mercury migration and streaking.

Dimming can occur slowly over time due to phosphor decay.

How will the application of novel Tesla

currents affect these idiosyncrasies? The higher voltages should facilitate starting, but humidity and temperature will still be factors. As to the affect on mercury streaking and phosphor decay, only usage over time will tell.

Fluorescent tubes even if driven by Tesla currents, are prone to flickerings and pulsings, often inexplicable, as are the standing wave patterns sometimes seen marching down the length of the tube.

the Tesla alternative

Early in the era of the Edison incandescent, Tesla said, "The more we progress in the study of electric and magnetic phenomena the more we become convinced that the present methods will be short-lived. For the production of light, at least, such heavy machinery would seem to be unnecessary. The energy required is very small, and if light can be obtained as efficiently as theoretically appears possible, the apparatus need have a very small output."

Tesla was an optimist in believing that the old methods would be "short lived." The old system persists, and (have you noticed?) it's even getting worse. Those energy consuming Edison filament bulbs, with which we still largely light our homes, now have a shorter life than ever and cost more. The power delivery system is more expensive to the consumer and as fragile as it ever was.

"I have made the discovery," Tesla announced in an 1891 patent, "that an electrical current of excessively small period and very high potential may be utilized economically and practicably to great advantage in the production of light." Pure Tesla lighting is high voltage and high frequency, and, ideally, resonant as well. Strides have been made in commercial-world lighting, but they exist only in a few specialized areas. Some of these advances do show a respect for Tesla's system.

electronic neon

High voltage is basic to neon. There is also a "cold cathode" (no heated filament) form of industrial fluorescent lighting that uses high voltage, up to 15,000 volts, the same as in neon and the same step-up transformers as in neon. Both neon and cold-cathode fluorescent systems are typically low frequency (60 cycles) and, of course, nonresonant.

Ventex 22kc 9kv 6.5"x1.9"x1 ¼"

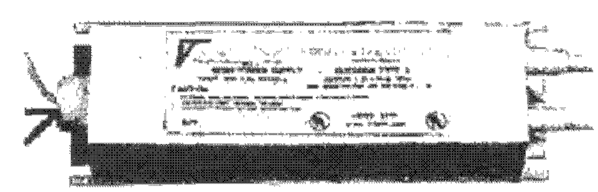

electronic neon transformer

The recently introduced electronic neon transformers are the closest specialized lighting has come to Tesla technology. These systems are high voltage and high frequency, running typically at 22 kilocycles and 3,000 to 5,000 volts. Most operate from a 120 -volt a.c. source, but there are some 12-volt d.c. versions as well. My experiments with both the 120- and 12-volt versions suggest the possibility of using these for fluorescent-tube drivers. Can they serve as little off-the-shelf Tesla-style lighting plants? It may be worth pursuing.

Jim turned an electronic neon transformer into a hot arc thrower that may have potential as a practical torch, perhaps a cutting tool. We burned out a couple of transformers doing stuff like this. The electronic neon transformer is a delicate item.

Tesla's high-frequency criterion is honored today in the compact fluorescent, which uses self-contained electronic pulsing circuits at high frequency (but low voltage) to produce a more efficient fluorescent. These are useful especially to off-the-grid home-power folks who draw from solar-charged battery power with tight economies. Some compact fluorescents work from 12-volt d.c. sources. These little units are relatively expensive.

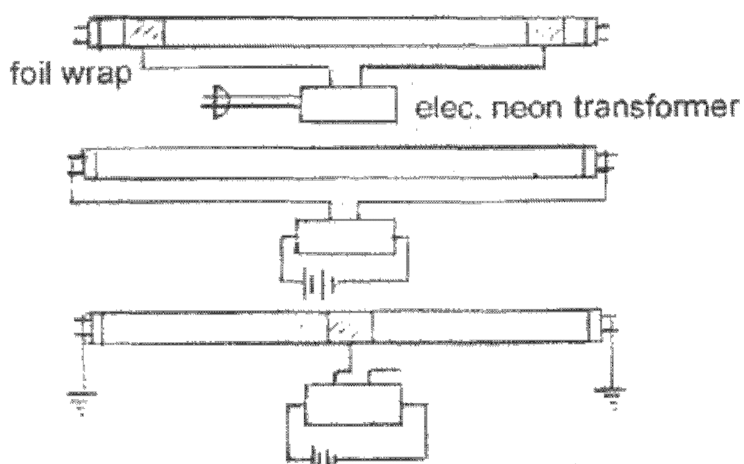

electronic-transformer experiments

Tesla invented the old 60-cycle system, and the Edison lamp became a fixture in it. Tesla reinvented his system, introducing patents for a system of high-potential high-frequency lighting for which the power supply became the Tesla coil. But, with the specialized exceptions above, the old system has become institutionalized, refuses to budge, and Tesla's alternative has been cast out. Such social inertia does not alter the scientific reality that Tesla currents so effectively stimulate gas and phosphorescent lamps that the phenomenon may belong in another world of efficiencies.

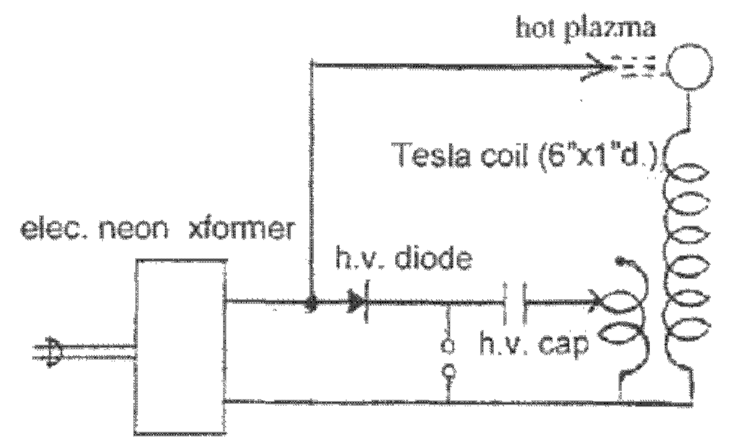

Jim's torch

With the oil coil we will drive conventional off-the-shelf fluorescent tubes with currents that meet Tesla's criteria: high voltage, high frequency, and resonant.

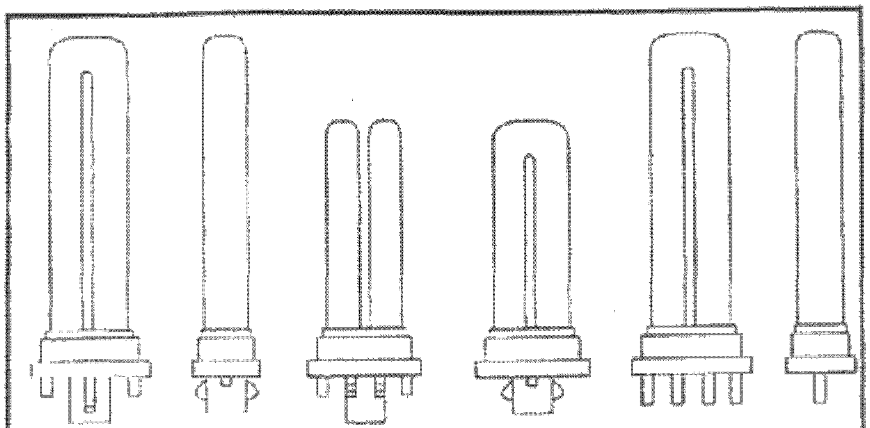
compact fluorescents

4. Spark vs. Solid State

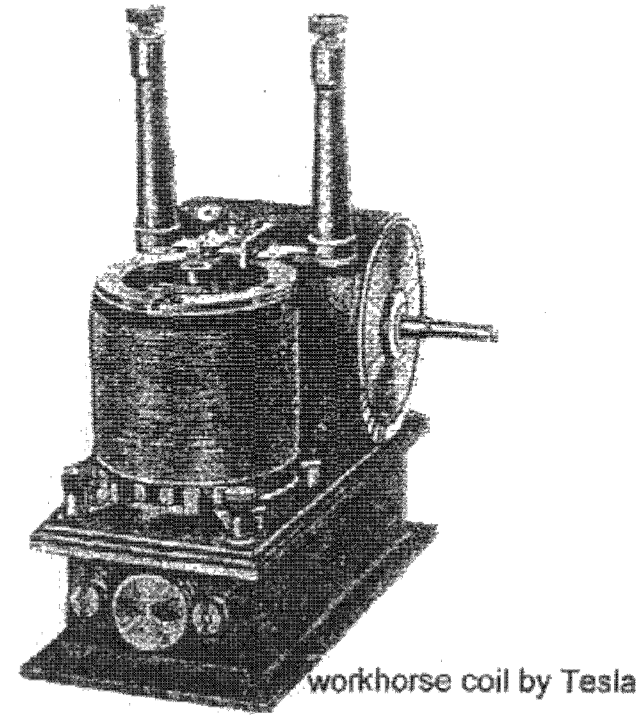
workhorse coil by Tesla

Experimenters defend the spark gap as the premier shock-excitation switch. They say no vacuum tube or semiconductor can match disruptive discharge's ability to drive the primary by dumping great quantities of current with great suddenness and in the process to withstand the extreme electrical stresses. I wonder how the spark gap might have developed had the project not been abandoned by official electrical engineering circa 1920. Might we now have off-the-shelf high-voltage control circuits of great simplicity, power, speed, and precision?

Still, the ratings of the commercial power field-effect transistors (FETs) we will be using in the next chapter are impressive for products of the effete world of semiconductors. In the series are ratings up to 500 volts and 25 amperes. Today there are integrated circuits and transistors that compete in rapid on-off switching very well with spark. These have been invented and mass-produced to meet the needs of high-speed computing and are now inexpensively available to us off-the-shelf to bend to our own perverse purposes. These new tools hold promise that an ideal solid-state Tesla-coil pulsing circuit can be designed from scratch.

the ideal pulse

My solid-state mentor, Silicon Valley inventor Jim Campos, pondered the question on his duel-trace scope: what kind of pulse does the secondary want from the primary to perform at its peak? A secondary coil has a particular resonant frequency. The primary ideally must match that vibration beat-for-beat. That's the ideal, but spark experimenters are accustomed to primaries that vibrate at some harmonic fraction of the secondary's resonant frequency. As a matter of fact, so was Tesla. But, for ideal results, the frequency must be precisely controllable over a wide range and thus capable of matching the secondary's.

The primary pulse should have an extremely sharp, almost vertical rise time. ("Slam it like a brick," Jim said.) Spark does fulfill that, but the pulse should not fall off in its own good time, as the spark pulse generally does, but be sustained for a particular period. Hence, pulse width, too, must be precisely controllable over a wide range.

The duel-trace scope told this story: The pulse must be long enough to take advantage of the excursion phase in the secondary, the upswing, but not so long as to interfere with the downswing. Such interference causes damping and disturbs the clear, steady ringing of the secondary in resonance.

If the primary pulses the secondary at exactly its resonant frequency beat-for-beat, the decay typical of intermittent pulsing is absent. The wave is not damped but is sustained as an undamped continuous vibration.

The scope showed that the ideal period for sustaining the pulse is one quarter of the wavelength. For example, if the frequency is 250 kilocycles, the total wave duration is 4 microseconds, thus the pulse duration must be 1 microsecond.

Sharp rise, quarter-wave duration, sharp fall: a square wave. That's the ideal pulse.

chopping battery power

Before we can put battery power to work for any alternating system like a Tesla coil, the battery direct current has to be interrupted into a series of pulses, or "chopped." (The power inverter that converts battery d.c. to household a.c. for RV's, boats, and home power is a sophisticated type of chopper.)

Another example: a battery Tesla coil seen in the old literature uses a model-T spark coil for its high-voltage supply transformer. As in the laboratory induction coil, the model-T coil did its chopping by means of a doorbell-like magnetic interrupter. Later in the evolution of the automobile, chopping would be done by the ignition points, and later by a solid-state module. This kind of chopping is at low frequencies, down in the audio range, and at relatively low voltages. In high-voltage, high-frequency Tesla-coiling, we might construe

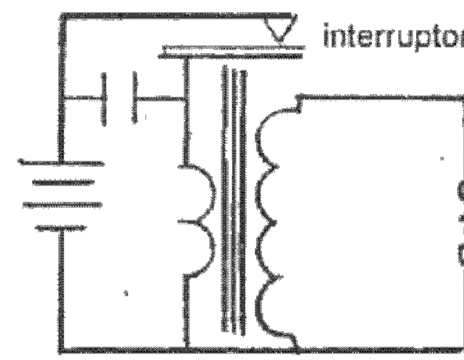

Model-T spark coil

the spark gap as a sort of chopper, be it a static gap or a rotary gap. Tesla's evolutions of the rotary spark gap, his "circuit controllers," might be seen as consummate mechanical choppers. These had frequency limitations which solid-state choppers promise to surpass.

solid-state chopping

A solid-state battery-current chopper seen among Tesla-coilers uses the 555 Timer chip in an astable mode to drive power transistors like the 3055. Both the 555 and the 3055 are widely available, even from Radio Shack. The battery d.c., chopped into a proper pulse, drives an auto ignition coil that serves as the high-voltage transformer for a small spark-gap Tesla coil. Mine had a chrome-plated hot-rod ignition coil. The solid-state driven ignition coil itself makes a nice laboratory induction coil or a high-voltage power supply for many experiments. The solid-state pulse circuit by itself makes a neat little audio tone generator. The little battery Tesla unit served as my medical coil for many years. The circuit shown was inspired by Jim, by Bob Beck, and by Walt Noon.

Flybacks

Most solid-state Tesla coils in the literature and on the Internet are TV flyback transformers. The flyback is a multilayer coil. It generates the high voltage for the TV picture tube and is a little Tesla coil, demonstrating that when pushed by necessity conventional engineering does resort to Tesla technology. The flyback is a 12-volt device. A limit of the flyback is that it cannot be pushed to resonance without burning out.

The same chopper that drives the ignition coil can also drive a flyback. However, the most frequently seen flyback Tesla coil is driven by an inefficient inverter circuit.

The 555 can be controlled for frequency by the turn of a pot, and the same goes for pulse

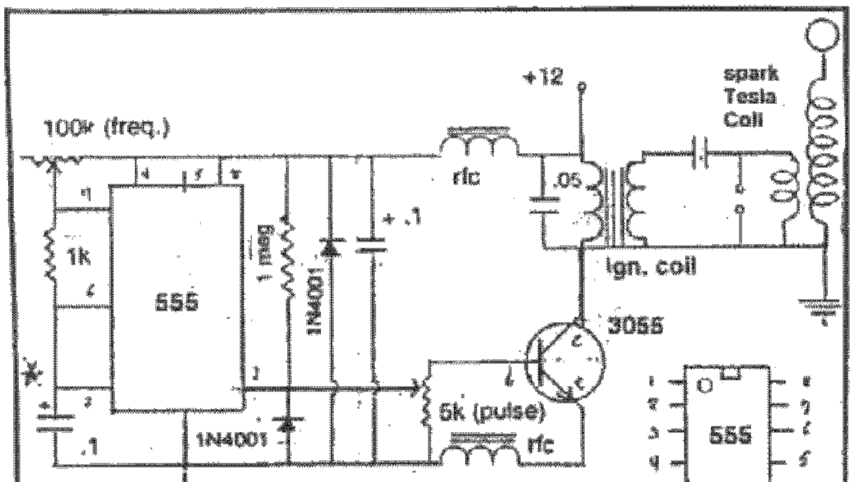

555 drives ignition coil

width. To drive an ignition coil, the frequency band is set down in the audio range. (This circuit plus a loudspeaker makes a simple tone generator you can have some fun with.) The capacitor marked * sets the frequency range as do the values of the pots. Audio is in the lower range of the amazing 555, which can go up to one million cycles (or two million with the new low-power 555).

This ignition-coil system is a hybrid, a wedding of audio-frequency solid state and spark. The manufacture of faster transistors that can handle higher frequencies and of transistors that can handle higher power opened the way to a Tesla coil that could be pulsed directly by solid state.

redeeming spark

Tesla intended his coils to perform utilitarian tasks: lighting, power transmission, radio transmitting, electrochemistry, electrotherapy, ozone production. The apparatus, he said, should not be "liable to derangement" and should be "capable of long use without attention or adjustment." And we might add, and he would agree, that it should not be noisy, polluting, or wasteful of current. Tesla did not plan on hex Schmitt triggers or power FETs. The apparatus was to be driven by disruptive discharge, by sophisticated rotary and series gaps that were "capable of effecting an extremely rapid interruption of the circuit." Tesla's patents show eight high-speed rotary "circuit controllers," all were mercury-actuated.

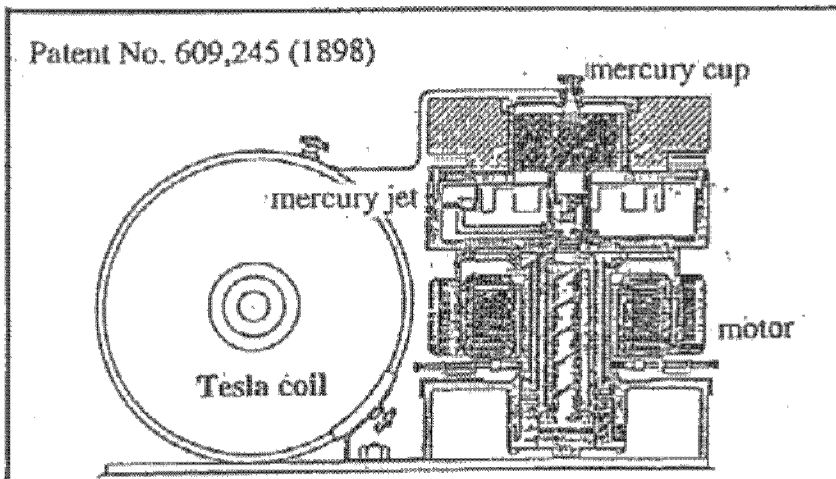

circuit controller

In the same London lecture as the oil coil, Tesla introduces a series gap: "The dielectric strength of a given width of airspace is greater when a great many small air gaps are used instead of one." The principle evolved into the airtight series gaps that were commercially manufactured for spark radio transmitters. An engineering rule of thumb: a gap should be under 1/100th of an inch, 1000 to 1200 volts per gap. My effort with the series airtight is

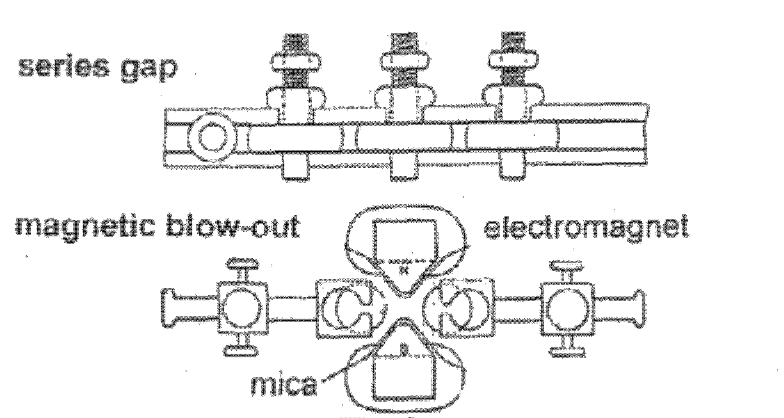

twoTesla gaps

the spark-plug gap shown. Use nonresistor plugs, like marine, small-engine, or the big Champion 509's for tractors. The atmosphere within an airtight can be altered to advantage. Add helium, alcohol vapor, or, by some burning process, remove the oxygen from the entrapped air so destructive ozone won't form.

Also in the London lecture is Tesla's magnetic blow-out: "The intense magnetic

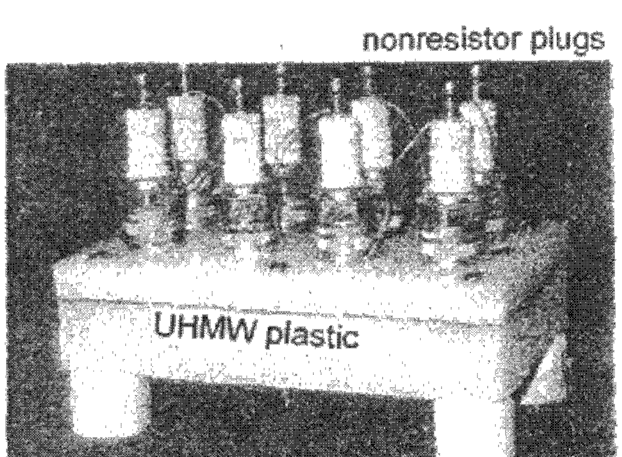

airtight spark-plug series gap

field then serves to blow out the arc as soon as it is formed, and the fundamental discharges occur in quicker succession." (quenching)

It was with my spark-plug series gap that I discovered the strange phenomenon of the sparkless spark gap. This happens when a Tesla coil (with the help of a variac) is driven slightly below the sparking level but still performs. And it does so with unusual solid-state-like smoothness. Later I found that Tesla had observed it, too: "With many spark gaps in series…I was able to produce oscillations without even a spark being visible between the knobs." Does the sparkless phenomenon suggest possibilities for redeeming spark?

The bimetallic Chaffee gap is another inspiration. Commercially developed circa 1918 for aviation spark radio transmitting, the Chaffee gap, which had electrodes of copper and aluminum, quenched so capably that it

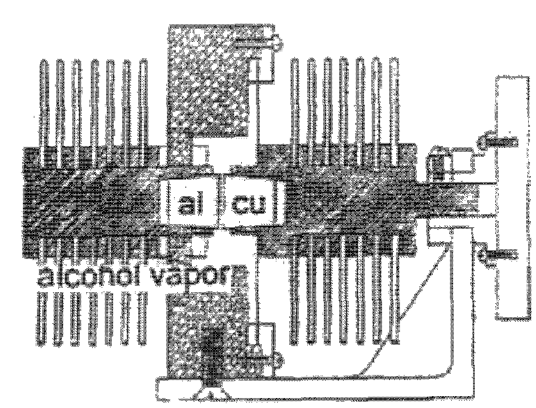

Chaffee gap

enabled modulated voice transmission with great clarity. Says a radio manual of the period, "The practically instantaneous reestablishment of the high initial gap resistance when the current becomes zero, is due probably to the formation of an insulating oxide film on the aluminum." Sounds like semiconductor talk to me.

5. Build a Solid-State Pulse Generator

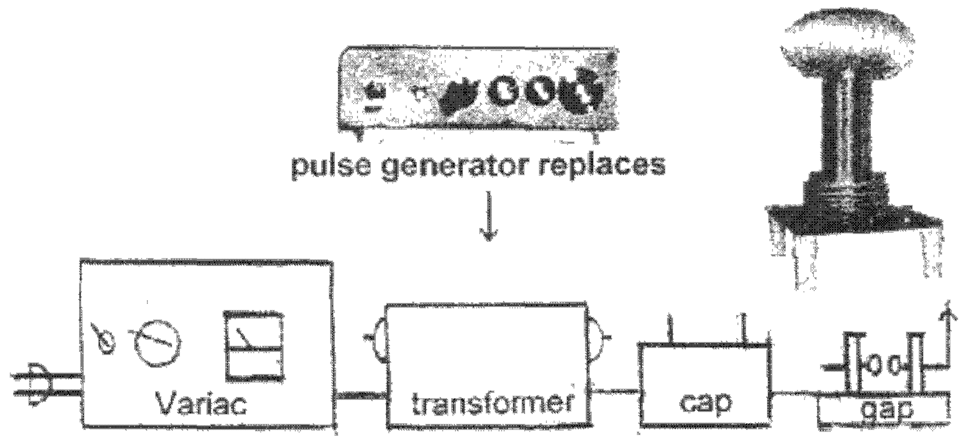

This Chapter tells you how to build an experimental solid-state pulse generator that will drive directly, beat for beat, any closely coupled Tesla coil that is resonant at 50 to 300 kilocycles. The oil coil of Chapter 2 (when loaded) meets these specs just fine.

The pulse generator weighs only under half a pound but replaces the conventional disruptive-discharge Tesla coil's ponderous transformer, capacitors, spark gap, and variac. It need not be plugged into the fragile grid and can run from batteries, if necessary, for it is powered by 12 to 48 volts of direct current.

Schmitt trigger or 555

The front end of such a Tesla-coil pulse unit can again be the old 555 workhorse. We tried the CMOS (the so-called low-power) version but it did not prove to be superior to regular. Also very appropriate is the 4584 hex Schmitt trigger chip. The specs told Jim that these might be the superior slam-it-with-a-brick entries from the progressively more speedy slam-bang world of digital. Jim worked out all the details of this circuit in a series of experiments conducted over three years. Countless components and values were tried and rejected in oscilloscope drudgeries. An objective was to simplify and minimize the number of parts which now amount to about a dozen.

The 555 and 4584 can be wired into themselves to become high-frequency square-wave pulse-formers whose frequency and pulse width can be precisely controlled.

The upper frequency limit of the 4584, according to the specs, is 2 megacycles. The 555: one mc. The 4429 chip is called a FET driver. The final output power transistor of choice is the IRF 640, rated 200 volts and 18 amperes, but other power FETs in the IRF series (and in other formats) could be used here if sufficiently speedy.

The full designation for the hex Schmitt trigger is CD4584. Some suppliers know it only by its replacement-part number, which is NTE 23740. A similar package is the CD40106. Radio Shack carries neither of these chips but does carry the 555. Radio Shack also carries, from the IRF power-FET series, the IRF 510 rated 100 volts and just 5.6 amperes.

only 12 volts in?

It was not among our expectations that our experimental solid-state pulse generator would drive a Tesla coil with an input of under 50 volts or so. But we were delighted to discover that with a mere 12 volts of battery power we

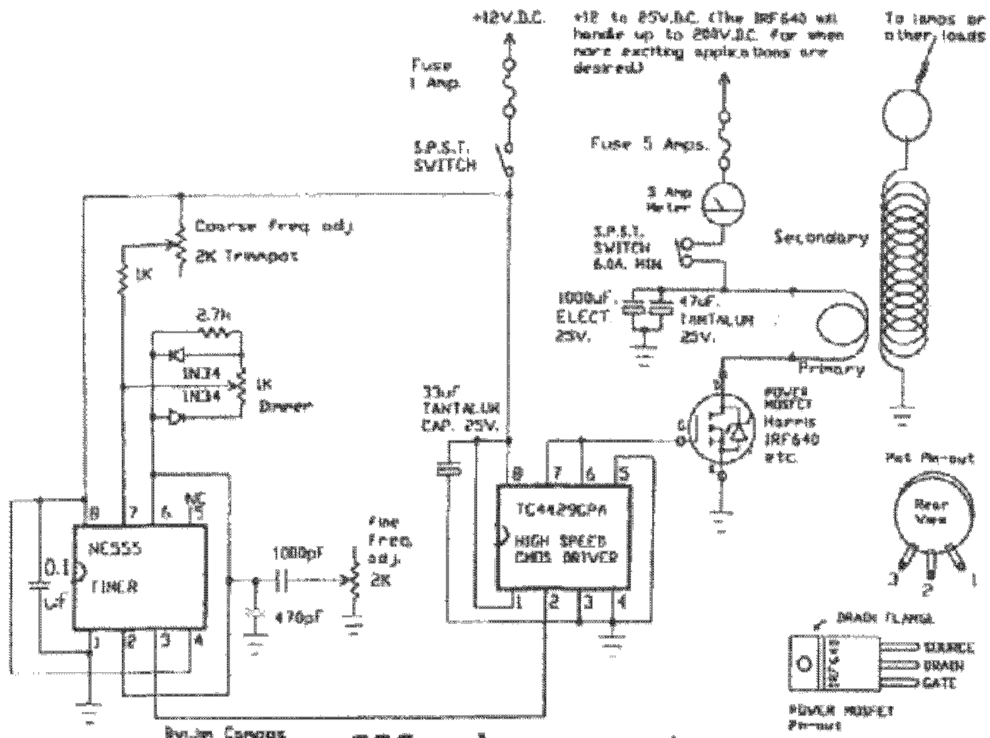

555 pulse generator

could get results. In fact, with just 12 volts input we could coax outputs from the Tesla-coil terminal of over 2000 volts. This was sufficient for lighting our four-foot fluorescent test lamp to practical levels. So here is clear testimony to the effectiveness of a closely coupled system pulsed one-for-one with the ideal waveform.

Series battery hook-ups of 24, 36, and 48 volts (possible in some home-power situations) produce superior results and efficiencies. The 12- to 2000-volt step-up itself represents a magnification of 166. Operating draw for the 12-volt lamp driver was 450MA to 1.5 amps, depending on the brightness desired. That's 5.4 to 18 watts.

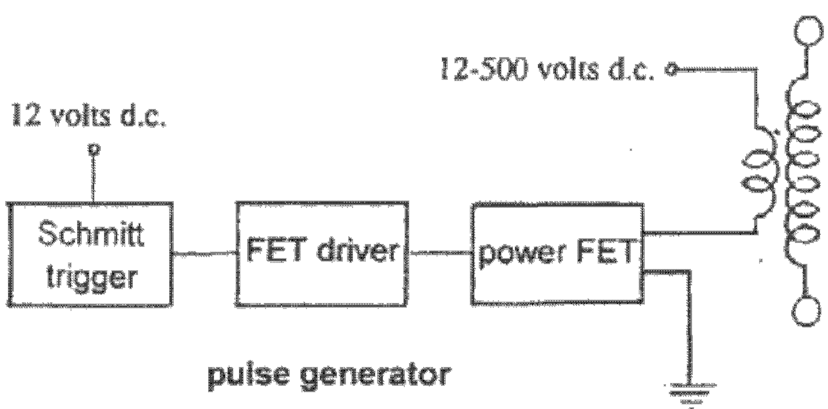

direct pulsing

fighting heat

To keep the on-time at a minimum, the pulse must have a very sharp rise time and fall time, a sharply square wave. Any “ramping” will extend the on-time, hence more heat. We fixed a heating problem that plagued us in the FET and FET driver as well when we changed FET drivers from a slower, amplifying 7667, which we used through many experiments, to the fast switching 4429.

When the circuit is in the “on” state conducting through the primary, its condition quickly becomes, after passing the 1/4-wave mark, a dead short. Momentarily, currents in excess of 18 amperes may be coursing through the power FET. At one megacycle and a quarter-wave pulse width, this short circuit would exist for nine percent of the time. The higher the frequency, the more of these short-circuit episodes in a given period of time. Heat builds and can only be dissipated at a limited rate.

We did a series of experiments to determine a safe operating frequency range in which this device could run cool and problem-free indefinitely. Thus the pulse generator now operates comfortably driving Tesla coils at 300 kc and below

Another good reason for a lower frequency of operation is that this level control by pulse-width control diminishes as the pulse width shrinks at higher frequencies.

If the circuit is pushed beyond the lighting-plant application here, you may want to put in additional FETS. The FET driver has a second output that would come in handy. In high-voltage applications the FETs, as well as other components in that circuit should match the voltage applied.

alternatives

There are other ways to drive power transistors besides the 555/4584-4429 combination, and you may want to explore them. Many months of experimentation went into the development of the recipe circuit, however, and any new direction may require a similar investment. A possibility that has Jim's interest substitutes for the 555/4584 a quad-D flip-flop, which, in his scheme, would automatically provide for any given frequency a pulse width of quarter-wave duration.

An alternative to the 4429 of our recipe circuit may well be the 4080 full-bridge FET driver. The circuit shown is from a typical application schematic in a reference source.

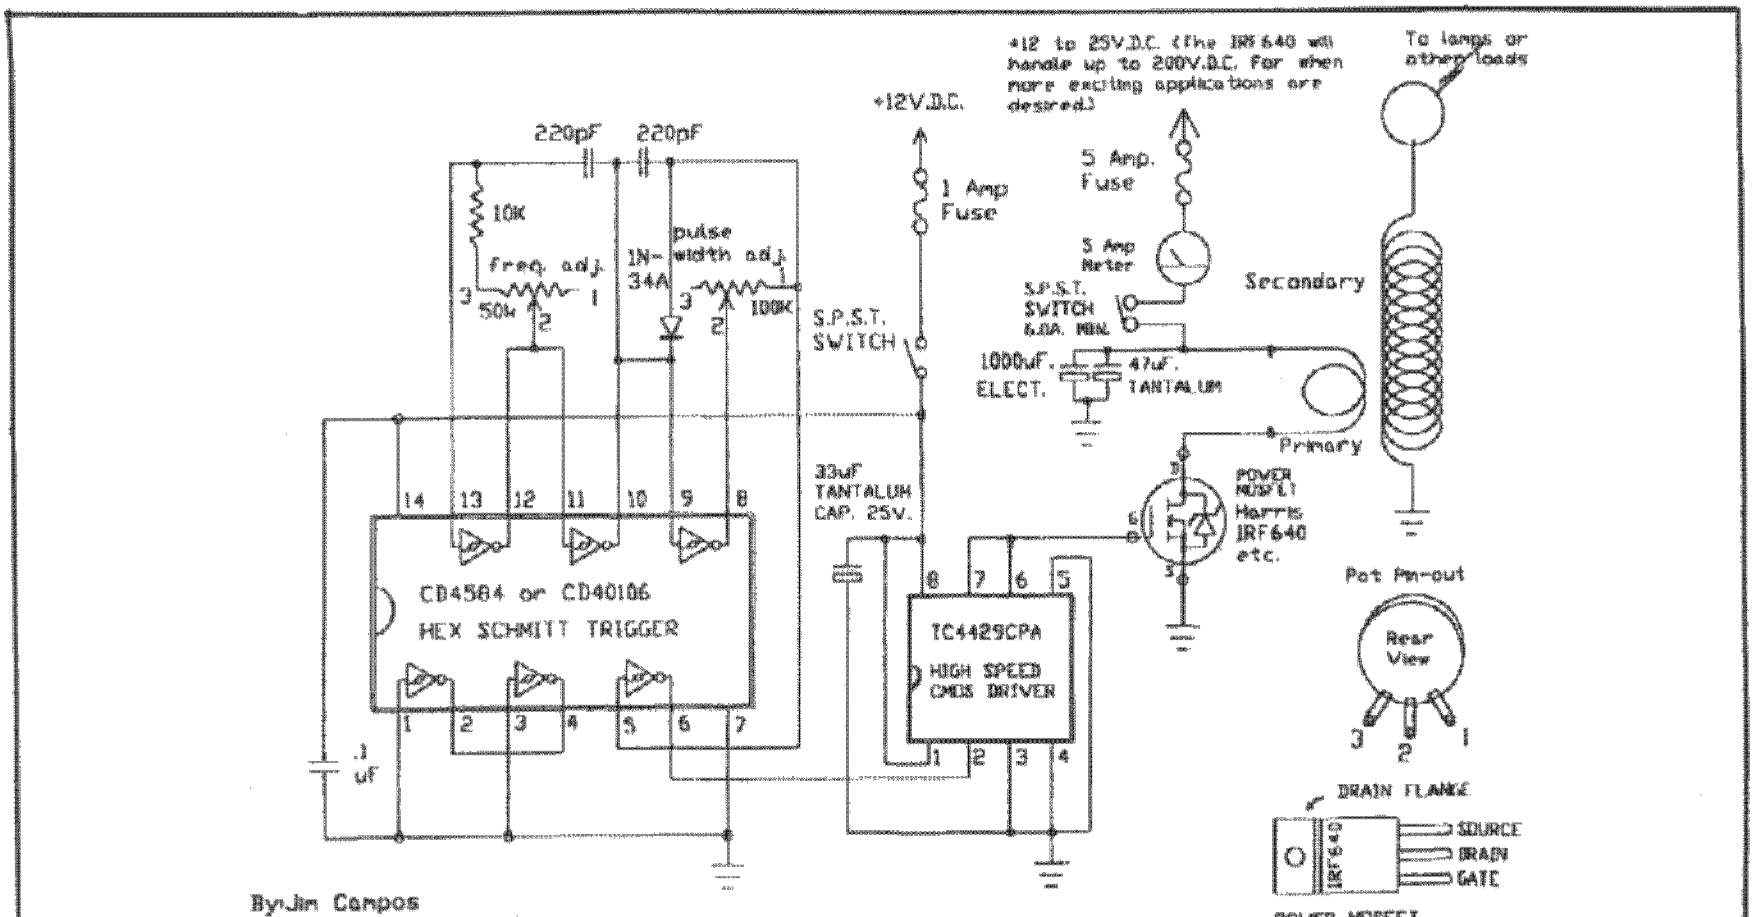

hex Schmitt trigger pulse generator

The full ordering number for this chip is HIP4080AIP. It is more expensive than the 4429 and harder to find. An internet posting on the 4080 by a pyrotechnically oriented Tesla-coiler referred to “sheets of flame.”

recipe construction

Locate the big 1000 mfd. electrolytic close to the battery-supply terminal. Battery current will fill this big capacitor and it will hold plenty, ready for release. Right next to it and in parallel put the 100 mfd. tantalum. This faster capacitor will release the energy in quick spurts. **Hazard note:** both of these capacitors are polarized, one lead is plus, the other minus. Connect them backwards and you get an explosion. Particularly dangerous is the tantalum, which becomes a little fragmentation bomb. Make sure the voltage rating of these caps matches your input voltage.

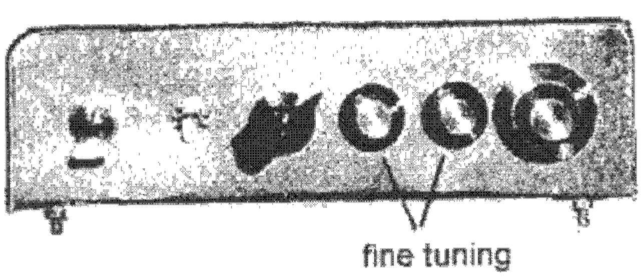

experimental pulse generator

In the interest of expedient current delivery, use heavy wire, #18 or bigger (stranded, insulated) throughout the business end of the circuit: from the supply terminal to the heavy-duty 6-amp power switch, to the source and drain of the FET, and through to the primary.

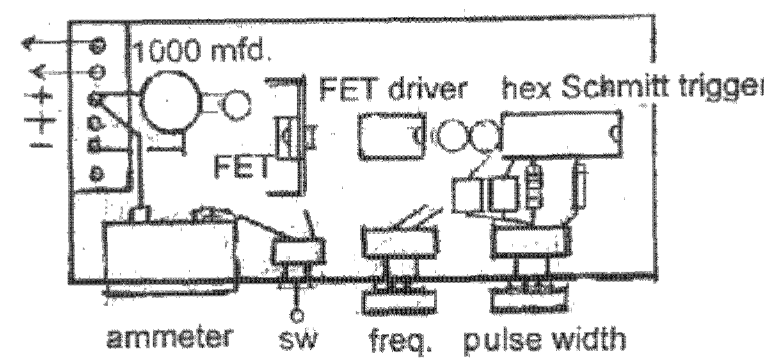

recipe pulse generator suggested layout

Make the wires to the unit from the battery heavy, like #14 or bigger.

You may want to breadboard the circuit and get the feel of it before hardwiring. Lots of breadboard work has already been done, though. I built the first hardwired version, and

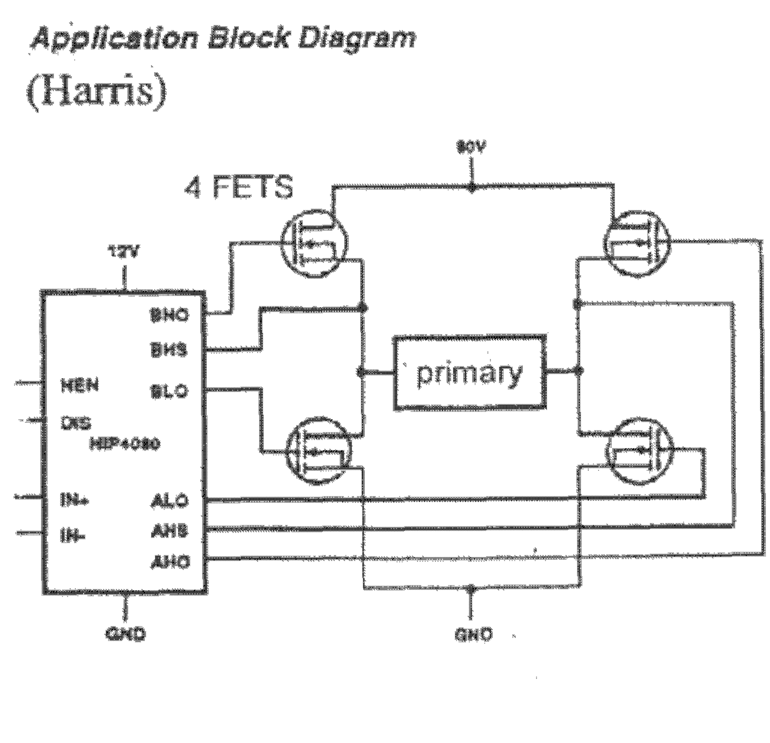

full-bridge FET driver alternative

it did work, but Jim took it over and redid the wiring at the chips so it now looks like jeweler's work. Stray capacitance can distort the pulse, so keep wires short. The chips and allied components are jammed tightly together to facilitate short connections. The little tantalum buffer caps should be snug up against their respective chips.

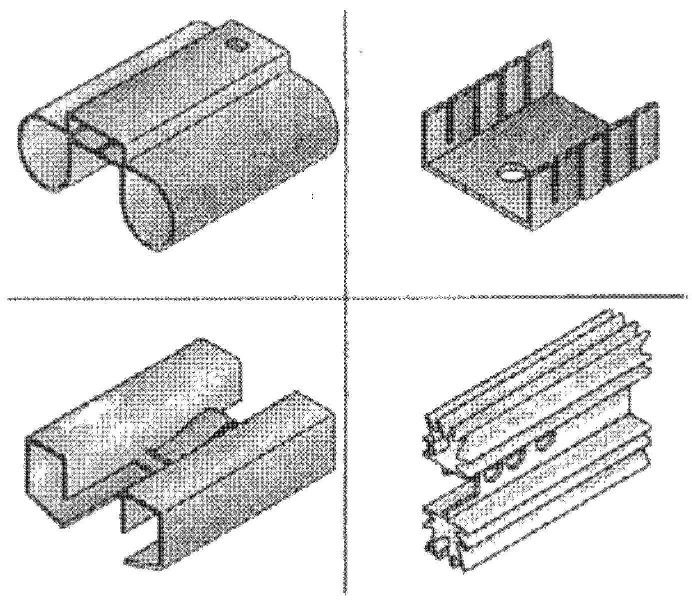

heat sinks

Parts Guide

Hard-to-find parts are specified here. See schematic for such common items as pots, resistors, etc.

Hex Schmitt trigger CD4584 (Harris), also CD40106 or NTE23740, from Mouser (44 cents). Or 555 timer from Mouser or Radio Shack. **FET driver** TC4429CPA (Telecom) from Digikey ($2.63). **Power FET** IRF 640, (200 volts, 18 amps (Harris), from Mouser $1.41), or lower power IRF 510 from Radio Shack. **Tantalum caps**: 68uf, 25-volt ($5) or 33uf, 25 volt ($1.52) from Mouser. **Ammeter,** 0-5 amps, Mouser ($18.85). **Mouser (800)346-6873** and **Digikey (800)344-4539.**

6. Tesla Lighting Plant

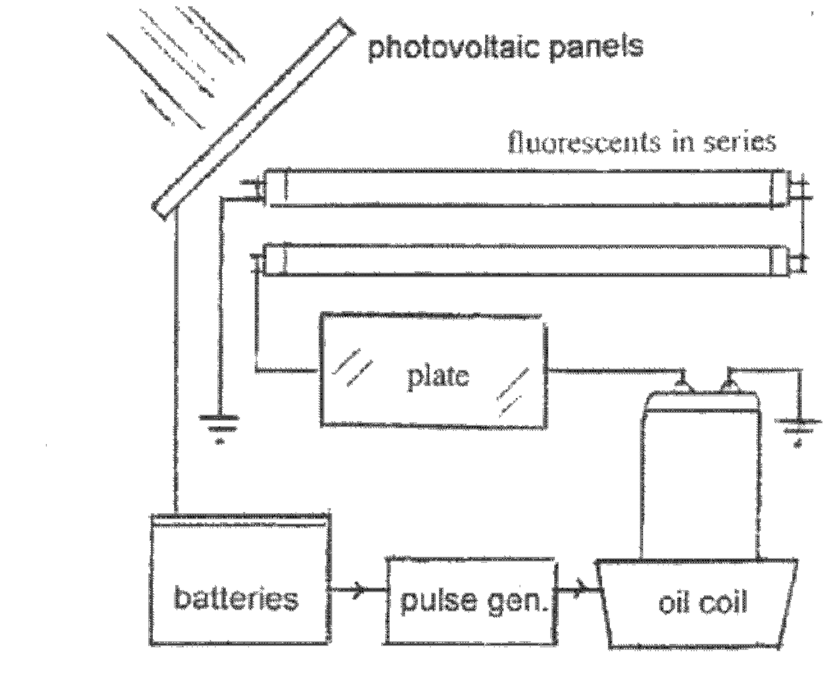

If we use the solid-state pulse generator to drive a suitable Tesla coil, can we drive one or more fluorescent tubes at current-draw efficiencies superior to conventional fluorescents?

I have done some experimenting suggesting that the answer to this question is yes. Experimenting also suggests that a design like this battery electronic Tesla coil could function coolly over long duty cycles, that it could function quietly and emission-free and require little maintenance. Properly installed, it should be safer in terms of shock hazard than conventional 120 VAC devices and safer probably in respect to fire hazard as well. Although the prototype illustrated is far from compact, it's obvious that this criterion could be fulfilled, too. Yes, experimenting does suggest these potentialities for such a system: a Tesla Lighting Plant.

experiments

I soldered the two connecting wires directly to the fluorescent tube's end-pins, shorting each pair out. If you don't want to solder, try conventional fluorescent-tube sockets, which probably can be adapted for the purpose and will also serve to support the tube. Otherwise the tube can be lightly suspended in any well insulated fashion. Another way "into" the tube that I experimented with was via bands of aluminum foil wrapped around the glass, a capacitive connection, but a direct connection to the pins works better.

My "plate" was just aluminum foil glued on to some 1/8-inch plastic sheet. The plate functions the same as a terminal capacitor (the sphere or toroid) on a Tesla coil, raising up the voltage (and also incidentally lowering the frequency). The tube itself contributes some to the terminal capacity, but for best results this should be augmented by the plate. Like a flywheel in a mechanical system, some inertial load is needed for the Tesla coil to work on. In a practical appliance, the plate could be incorporated as a reflector, a shade, or a diffusor. The plate is a cheap way to kick up the voltage, but too much terminal capacity could drive the frequency down to where the system cannot be tuned.

For current-draw measurements I used a big laboratory ammeter set to the 0-3 amp scale. An ammeter is essential for tuning this system to efficiency and would be incorporated into any ultimate commercial device as a panel meter (0-5 amp is commonly available).

Exact voltage measurements, necessary to compute wattage, were made on a digital multimeter (DMM). The same DMM has a frequency-counter feature that was put to use, and I also used it for a crude light meter by plugging into it a cadmium sulfide photoresistor and setting the DMM for resistance measurement.

experimental limits

Considerably more time has been devoted to the development of the pulse generator and Tesla coil than in developing the ideal way to connect the system to florescent tubes for peak performance and amp economy. I suspect

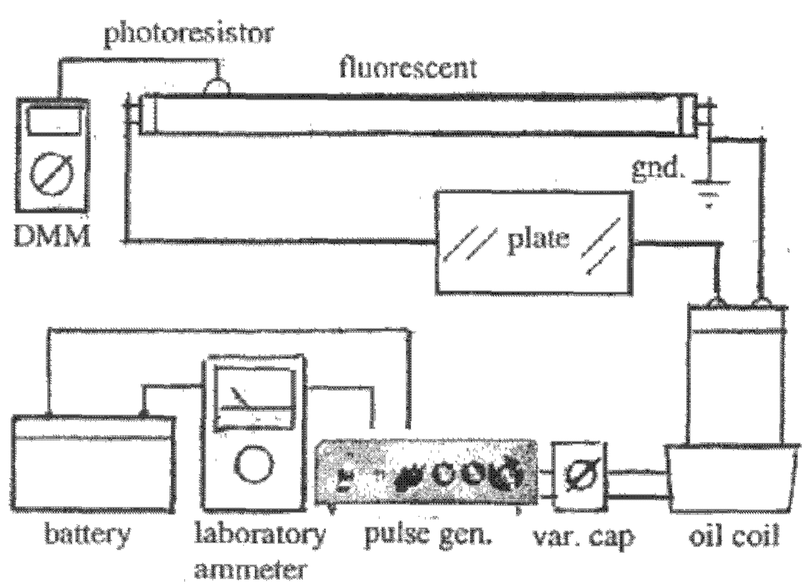

experimental set-up

there are discoveries ahead, and you may make some in your own experimenting. For example, in the interest of that fly-wheel effect, we have put a high-voltage capacitor across the Tesla coil's output terminals and succeeded in driving the tube to a brilliance more than twice that of the conventional. Current draw was excessive in this experiment, but it was suggestive. What would variable capacity do? Also, experiment with various capacities placed across the primary coil in order to better tune that circuit.

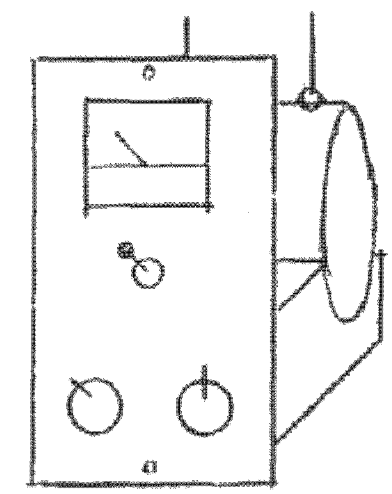

Tesla lighting plant commercial concept

Considerably more experimenting has been done with a single tube at 12 volts than with multiple tubes at higher inputs.

Finally, experimenting has been limited to lighting, while I suspect that this electronic Tesla coil may have other practical applications, such as in ozone production or even heating or motive power.

current economies

The pulse generator's internal chip electronics consumes a negligible 22 milliamperes. The ammeter measures the current of interest which is through the FET-Tesla-coil primary circuit. On the pulse generator you manipulate the frequency and pulse-width controls to produce the desired level of light at the minimum current draw. Check for excessive heat in the FETs and chips as you experiment. If current draw is held down to one ampere or less, the FETs will run cool and unstressed and long running times are possible. If the tube does not light up completely but glows at the ends, wait, and it will usually snap on bright. Touching the tube sometimes helps.

In tuning, you'll discover that an increase in light output does not necessarily require a commensurate increase in current consumption. In conventional lighting, an increase in light output always means more current, but here that rule is suspended. In fact, as you adjust frequency, there are settings where you'll see the amps dip as the lamp brightens. This is the dip familiar to the radioman tuning his transmitter, the dip on his milliammeter that tells him that the rig is operating on the money. The resonant dip.

A typical operating frequency in these experiments was 115 kc.

At 12 volts and one tube it is possible to produce low-level practical room illumination, light suitable for walking around, drawing 450 milliamps or 5.4 watts. Turning it up toward adequate reading light level by advancing the pulse width, it's pulling 700 MA or 8.4 watts. At 1 amp the light level is about half that of a conventional tube while pulling one-third the power, 12 watts.

At 24 volts a single fluorescent tube starts up more easily and burns more brightly. Connect multiple tubes in series. It is in multiple-tube lighting that the advantages of this sort of system become more dramatic. In con-conventional fluorescents, a second tube means, inescapably, twice the wattage is used, a third tube, three times, etc. But in the Tesla system, this equation dissolves. One amp input lights three tubes all to the same level as one amp lit one tube. The single tube lit at one amp and 12 volts equaled 12 watts. The three tubes at 24 volts are together pulling one amp, which equals 24 watts, but now that's only 8 watts per tube. Projected over a number of tubes this efficiency could add up to a considerable conservation of battery power.

Fluorescents lit Tesla-style can be adjusted for light level with the pulse-width control. Conventional fluorescents operate at a fixed output, are overly bright for some situations and tastes and cannot be mellowed out. Light levels can be attenuated for particular needs resulting in current economies.

Fluorescent tubes lit Tesla-style are healthier. The conventional fluorescent flickering at 60 cycles is irritating to the nervous system and can cause eyestrain and headaches. At 100 kc no flicker is perceptible.

space heating?

Wouldn't it be great to invent a Tesla-style battery powered space heater that pulls just a few amps? I don't have the answer, but I suspect that it would lie in the realm of vacuum technology. Tesla patented some high-potential, high-frequency lamps that contained solid elements in vacuum globes. Solid elements like little blocks of carbon burn brightly and at high voltages reach such temperatures as to cause the solid element to fuse or vaporize. Could this phenomenon be translated into space heating?

growing

The only serious indoor growing operation I have observed close-up utilized an array of standard, cheap cool-white four-foot fluorescent tubes. The grower, who could have afforded the pricier grow lights or arc lamps, claimed that the old cool-white was his lamp of choice. I'll vouch for the quality of his crop. His lamps, of course, are the same cool-whites that abound, in a technically burnt-out condition, in the dumpsters of office and apartment buildings everywhere waiting to get a new life by feeding on high-potential, high-frequency Tesla currents.

Florescents driven Tesla-style run cool. No cathodes are burning. The Tesla-current stimulation itself produces no heat that could be detectable.

Florescents driven Tesla-style may deliver current economies that would allow operation from solar-charged battery sources rather than noisy and potentially flammable gasoline generators.

There is scientific evidence that suggests that plants exposed to high-frequency, high-potential fields experience an accelerated rate of growth. High electric potentials at conventional 60-cycle frequencies evidently have a positive effect judging by the number of nurseries located under high-tension power lines. But high frequency may be superior.

safety

Particularly if driven by battery, the shock hazard of the system is negligible. Although the voltage at the Tesla-coil output end of the system is in the thousands, the frequency, being a hundred kc or more, does not register on the nervous system, which cuts off at about 3 kc. It is possible to touch bare terminals without sensation or harm. There is, however, a possibility of getting a high-frequency (or so-called rf) burn if your finger tip serves to complete a circuit.

For fire safety (as well as to reduce electrical losses) it is a good idea to wire up the lamps with a well-insulated conductor, such as 5 kv test wire. Poorly insulated wires that intersect can conceivably arc to one another causing a

From *H.F Apparatus*
Curtis (1920). Thanks to Rex Research
rexresearch@biogate.com

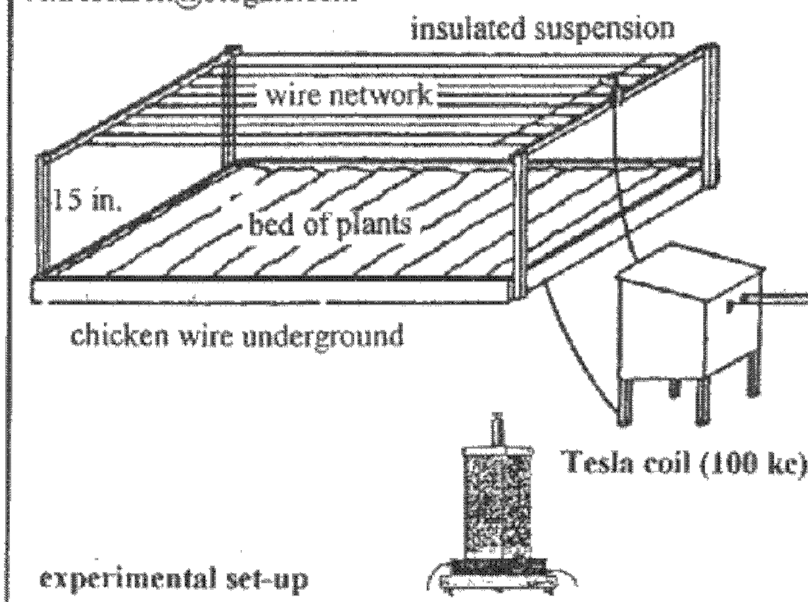

high-frequency plant-growth stimulation

fire hazard. If you are making a long run from Tesla coil to lamps, keep wires that are in parallel a good distance apart. As more tubes are added and more input voltage supplied to the system, the output voltages can reach exceedingly high values making insulative anti-arcing and anti-coronal measures even more important.

radio interference

The conventional florescent tube, flickering at 60 cycles, produces wide-band disturbances which are carried over power lines and can blot out all but strong local stations over large segments of the AM broadcast band.

The Tesla system can produce at discreet points on the band little tones and squeals from harmonics. These fade off quickly with distance from the source and are not carried over the power grid. Output at the fundamental frequency is assumed to be translated into light, the tube serving, as it were, as a dummy load. Any fundamental, or

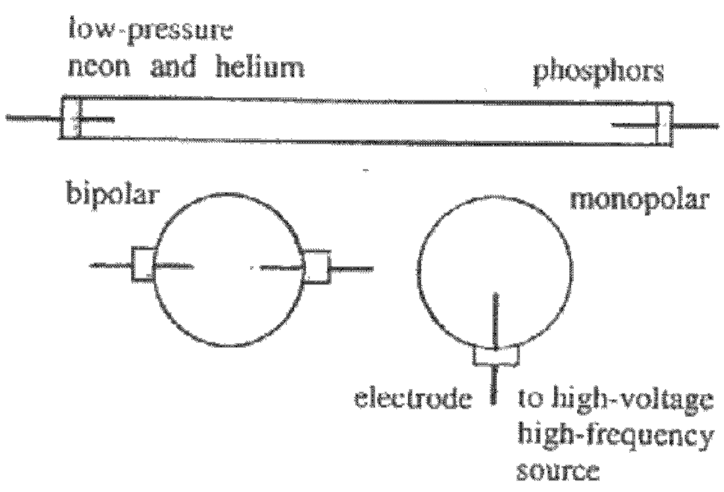

ideal Tesla lamp

any prominent harmonic, propagation that may be in evidence and unwanted could conceivably be subdued by filtration.

the ideal appliance

The ubiquitous cool-white notwithstanding, what would be the ideal lamp for Tesla currents were one to invent it from scratch? I dropped this very question upon Jim Hardesty, whose PV Scientific produces elegant Crookes tubes used to demonstrate vacuum phenomena. His suggestion: Fill a tube with a mixture of helium and neon at low pressure, say .1 mm of mercury. Vacuum pressures under 10 mm. enable maximum current flow through a tube, he said, no matter what gas is used, even air. (Normal atmospheric pressure

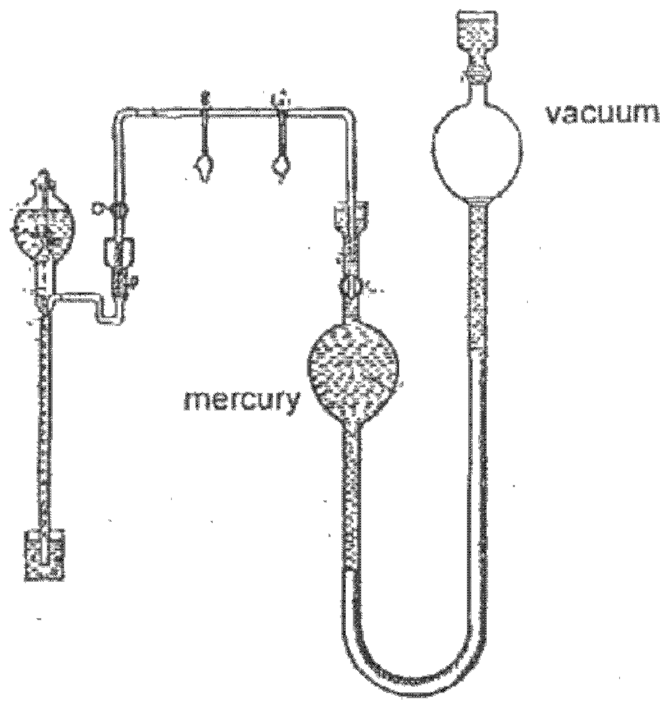

vacuum pump

is 760 mm.) One can experiment with a vacuum pump on an electrified tube to determine the optimal pressure.

A clear glass tube can be used, but, by experimenting with different phosphors for interior coatings, one can produce double or triple the amount of light for the energy put in.

7. For More Information

in print

Tesla Coil by George Trinkaus from High Voltage Press is the parent of this booklet. (See ad, page over.)

The Inventions, Researches, and Writings of NikolaTesla, by James Commerford Martin (1894) includes the London oil-coil lecture. Reprints from Omni, P.O. Box 900566, Palmdale, CA 93590 or Lindsay, P.O. Box 12, Bradley, IL 60915.

Lectures, Patents, Articles of Tesla also includes the London lecture. From Borderland Sciences, P.O. Box 6250, Eureka, CA 95502. The reference to the tesla coil as capacitor is from the ***Journal of Borderland Sciences,*** "Tesla's Broadcast Power" by Gerry Vassilatos, back issues, Vol. 52, Nos. 2 and 3.

Power MosFETS is the 1000-page Harris data manual. From Aznet, Santa Clara, CA.

Neon Signs is a 1935 text with how-to info on gases, vacuum pumps, and glass blowing. From Camden, Barrow Farm, Rode, nr. Bath, BA36PS, UK.

555 Timer IC Circuits by Forest Mims is a booklet distributed by Radio Shack.

components:

Electronic neon transformers. Some brands: Bertoni, Transfotech, Ventex, DiAnalog. From your local sign-shop wholesaler.

Variacs. Mouser carries Staco 3-amp ($71) to 12-amp ($152). Surplus: try R5D3, Portland, OR, (503)774-6560

Oil-coil containers. Rubbermaid 1-gal. pitcher from your supermarket. **Birm:** 5-quart galvanized bucket from Ace Hardware ($8.99).

High-voltage test wire, rubber insulated, 1500- to 10,000-volt from Mouser ($22 to $36 for 100 ft. spool). Also, for smaller order: Metro Electronics, 1831 J St., Sacramento, CA 95814.

Crooks tubes as well as static generators, crystal and regenerative receivers, etc. from PV Scientific, Jim Hardesty, 309 2nd, Ithica, NY 14850, www.arcsandsparks.com

Home-power suppliers publish interesting catalogs: photovoltaics, industrial batteries, power inverters, 12-volt lamps and other 12-volt devices: **Applied Power** (800)777-6609, **Sierra Solar** (916)265-8441, **Alternative Power** (608)637-2722.

$7.75

"I have made the discovery that an electrical current of small period and very high potential may be utilized economically and practically to great advantage in the production of light."

Tesla, page 11

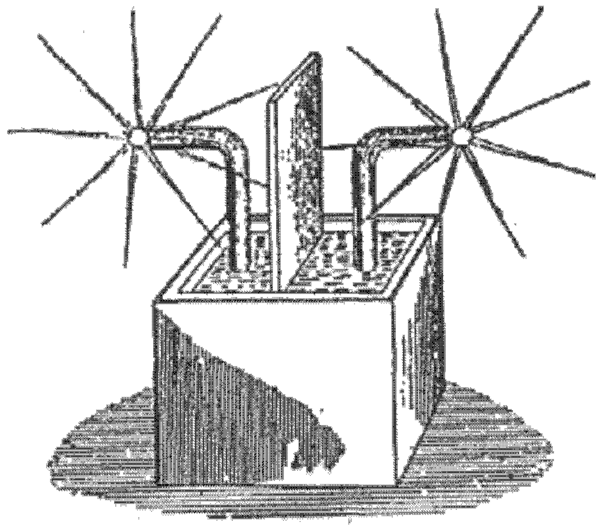

High Voltage Press
P.O. Box 1525
Portland, OR 97207
USA

toll-free 877-263-1215