Solid State Tesla Coils

How to Make Them, How They Work!

Last updated 7/5/2005

Table of Contents:

Preliminary Cautions, Notes, Warnings, etc.

1. The subject matter in this file assumes you already have some electronic project knowledge and skills. You ought to have some skills in building things, as well as electronic knowledge to the point of knowing simple AC circuit analysis, series and parallel resonant circuits, some amplifier and oscillator theory, and a bit of radio stuff. I hope you know what Armstrong, Hartley and Colpitts oscilators are and how they work. If not, you may still be able to build the stuff mentioned below and burn your house down, but knowing how the stuff below works will help if you don't quite get things working. This is not the place to learn the basics of electronics - spend a week or two reading library books if you have to.

2. Things may not work right the first time. You may blow out some things. Be prepared to have spare parts or to have to buy more should something go wrong.

3. The stuff mentioned below will create high voltage. Although the high frequencies will usually protect you from electric shock, there will be plenty of ways to get burned and a few ways to start fires. Also, Tesla coils easily produce vivid corona that usually produces ozone that can corrode your lungs. And that's not all that can go wrong! Go here for more info.

4. There is no warranty of correctness nor completeness of this information.

Theory - Low Impedance Primary Feed

I actually made working solid-state Tesla coils this way as long ago as 1978 and did not know then the whole reason why this worked.

Loose coupling from the primary to the secondary is part of this trick!

Here is what happens:

The secondary is effectively a parallel resonant circuit, with stray capacitance being in parallel with the coil's inductance. This is an oversimplification, since not all stray capacitance is in parallel with all the coil's inductance. But there is a parallel resonant mode!

But at frequencies a little higher than that at which the secondary is parallel-resonant, the capacitance is a little more conductive than the inductance, and the whole thing looks capacitive. This net capacitance will series-resonate with the "leakage" inductance to give you a mechanism for impressive voltage gain!

Here is an equivalent circuit for a loosely-coupled transformer with stray capacitance across the secondary:

```
   ---------+-----UUUUUUU------+----------+
            |    leakage ind.  |          |
            )                  )          |
  primary   )      secondary   )        -----  stray
 inductance )      inductance  )        -----  capacitance
            )                  )          |
            )                  )          |
            |                  |          |
   ---------+------------------+----------+
```

The above is a model for a loosely coupled transformer with stray capacitance across the secondary. It can be messy and tricky determining values for all three inductors to simulate an actual transformer. The leakage and secondary inductors in this model form a voltage divider with a ratio equal to the coupling ratio. The primary inductor in parallel with the series combination of the other two inductors has an inductance equal to the actual primary inductance with the secondary open. If you can determine all this mathematically so far without getting messed up, you can check your work by confirming that the leakage inductor in parallel with the primary inductor equals the primary inductance when the secondary is shorted - at any frequency high enough where resistance is low compared to reactances of all inductances, but frequency is low enough for stray capacitance to not significantly affect anything.

But this model is for a non-isolating transformer with unity turns ratio. To complete the model, cascade it with an ideal transformer having the actual turns ratio. The actual effective turns ratio to use for this is not the actual ratio of the number of turns, but the square root of the ratio of the inductances of the two coils. The square root of the inductance ratio is usually less than the ratio of actual number of turns in most Tesla coils. So the voltage gain will be the coupling factor times the square root of the inductance ratio, unless you get resonant effects.

Now, the big trick is a resonant effect. At a frequency at which the stray capacitance across the secondary is a little more conductive than the secondary's inductance, this net capacitive effect will series-resonate with the leakage inductance. This frequency is a little higher than the secondary's natural resonant frequency, but that may be lowered a little since having the primary nearby adds stray capacitance. In my experience this frequency is somewhere around that at which the length of the wire used to wind the seondary is a quarter wavelength.

The series-resonant effect leads to low input impedance and very impressive voltage gain. The voltage gain from primary to secondary can easily be ten and sometimes 20 times the actual ratio of the number of turns in each coil.

Another thing to realize is that at resonance, the output voltage lags the input voltage by 90 degrees. In addition, the input impedance of the effective series resonant circuit (as seen by the primary) can easily be much less than the primary's inductance, and it is likely that you can afford to not power-factor-correct (add a capacitor in parallel with) the primary.

A working self-oscillating low-impedance-primary-feed Tesla coil typically includes a third winding for feedback. Ideally, this would be wound close to (better just below than just above) the low end of the secondary coil. At the frequency at which the primary impedance is both low and a purely resistive load (close enough to maximum voltage gain), the voltage coming out from the feedback coil will be close to 90 degrees lagging the voltage applied to the primary. Strictly theoretically speaking, having the primary appear purly resistive means that the effective series resonant circuit has to look a bit capacitive to balance the primary's inductance - and this occurs at a frequency slightly lower than that at which the voltage phase shift is 90 degrees, which means the phase shift is a little less.

So, if you have a power amplifier and something to get a leading phase shift (net, including whatever phase shift the amplifier has), and feed all this with the feedback winding, you probably have a working solid-state Tesla coil. If the amplifier is inverting or you interchange the two leads of the feedback winding, then have the amplifier and any phase shift network between its input and the feedback winding add up to a 90 degree or slighjtly higher lagging phase shift.

Please beware that the amplifier input and any phase shift device could well present a significant load on the secondary and detract from the voltage gain. Also note that the feedback winding is not tightly coupled to anything and leakage inductance may cause you some phase shift unless the feedback winding is lightly loaded. It is recommended to have the feedback winding consist of one turn; use more turns only if this winding sees truly high impedance.

Please also note that the amplifier used in a low-impedance-primary-feed scheme should have low output impedance, whether the amplifier is running clean or distorted. Otherwise the whole system may oscillate at a different frequency where the primary's impedance is higher and the voltage gain from primary to secondary is not really high.

A web page by Richie Burnett with solid state Tesla coil theory, apparently mainly for the low impedance primary feed method. Please keep in mind that voltage across a resonating secondary lags the secondary winding current by 90 degrees.

Theory - High Impedance Primary Feed

High Impedance Primary Feed involves feeding the primary at a frequency at which the primary impedance peaks. If you have "critical coupling", or "undercoupling", you have one main primary impedance peak. If you have "overcoupling", then you have a double peak in primary impedance as a function of frequency. In any case, the voltage gain will not greatly exceed the product of coupling coefficient times the square root of the ratio of secondary inductance to primary inductance. This means the voltage gain will almost certainly be less than the turns ratio. You will get decent secondary voltage only with some rather high primary voltage.

High impedance primary feed is therefore impractical for most solid-state Tesla coils. This is better-suited to vacuum tube Tesla coils. It is a bit difficult or at least tricky to get a vacuum

tube circuit to deliver enough power for really good results, at least for tubes other than high-power radio transmitter types.
Circuits to do this include the Armstrong, Hartley, and Colpitts oscillators. These oscillators have a capacitor in parallel with the primary and oscillation occurs when the primary/capacitor combination is a parallel resonant circuit. The secondary will detune the primary resonance somewhat, usually splitting the parallel resonant mode into modes of somewhat lower and somewhat higher frequencies where the primary circuit acts as a parallel resonant circuit. Feedback is normally taken from the primary and not the secondary in the Hartley and Colpitts circuits.
At the lower parallel resonant mode, the secondary voltage typically lags the primary voltage by only a few degrees. At the upper mode, the secondary voltage lags the primary voltage by nearly 180 degrees. You may want to consider this, since any feedback extension of the primary will pick up plenty from the secondary. Chances are, a Hartley or Armstrong circuit will oscillate at the lower resonant mode.
If the coupling is loose enough and the primary and secondary resonant frequencies (when separated from each other) are equal, (critical coupling), the resonant mode will not be split but there will be a single impedance peak at the resonant frequency. The primary impedance will not be particularly high and the voltage gain may be a little impressive. The secondary voltage will lag the primary voltage by 90 degrees.

Theory - End Feed Without a Primary

In this scheme, you have basically a series circuit consisting of:

1. The inductance of the coil.
2. The capacitance from the coil to earth ground.

The terminals of this series circuit are the "low" end of the coil and ground. This is an oversimplification, since capacitance is distributed from different parts of the coil to ground and not all of the capacitance is in series with all of the inductance. But still, you can consider the coil and capacitance from it to ground to make a nice series resonant circuit at a particular frequency.
The coil becomes a helical radio antenna - but excessively compacted and with little true radiating ability.
If you have a source of AC voltage at the resonant frequency applied to the "low" end of the coil with respect to ground, you get an end-feed Tesla coil. The current flowing from the AC voltage source is the ratio of applied voltage to the resistance of the whole resonant circuit. The resistance will differ from the wire's resistance for a few reasons:

1. Some current gets capacitively coupled to ground before flowing through all the resistance. Ideally, this makes the resonant-feed resistance half the coil's DC resistance.

2. The "skin effect" makes the coil's resistance higher at high frequencies.

3. Losses such as RF radiation, corona discharge, dielectric losses in any insulating materials in the coil or nearby nonconductors in the high voltage AC electric field, eddy current losses in conductors in the coil's magnetic field, etc. will show up as extra resistance.

But the voltage at the "high" end of the coil is (at least by theory) the coil's inductance times 2 times Pi times the current times the frequency. Derate this a bit since some current gets capacitively shunted to ground before flowing through all the inductance. Maybe only as an oversimplification, multiply everything above by 2/Pi which is approx. .62.

My experience seems to indicate that the series-resonant frequency will be close to the frequency at which the length of the wire used to wind the coil is 1/4 wavelength long. Divide 1/4 of the speed of light by the wire length, and you get the frequency. I tried a bit of theory, and got a theoretical result of wire length being 5/8 wave instead of 1/4 wave, so I wonder where I went wrong and will try in the future to reconcile all this.

But as for theory for making a working end-feed Tesla coil? Voltage gain will be not too far from the actual "Q" of the coil. This means that if you are doing well, voltage gain may be a few hundred. This means you need at least around a hundred volts and maybe around a thousand volts of high frequency AC applied to the "low" end of the coil with respect to ground. In a self-oscillating scheme, you need feedback. You could use a feedback coil, which will have an output voltage in phase with the high voltage output (or opposite). This is 90 degrees lagging (or opposite this, which is 90 degrees leading) the curent flowing through the main coil; and that current is ideally in phase with the end feed voltage. If you come up with some sort of power amplifier that delivers preferably a few hundred watts (at a decent voltage of hundreds of volts) with a way to get a 90 degree phase shift at the coil's resonant frequency, you're in business. Simply have one turn of wire wound around or just below the low end of the coil to feed into this amplifier.

Another trick for self-oscillating schemes would be to take feedback from an extra resistor added in series with the whole resonant circuit. If you can sense the voltage across a resistor added in series with the "low" end of the coil, this is ideal. But neither end of such a series resistor is at ground, so this makes things a bit trickier - maybe use a transformer to couple this to the input of a power amplifier that drives the "low" end of the coil.

Another option is to put a low value resistor between the "ground" connection of whatever power amplifier is used and actual ground. Most of the resonant current will flow through this resistor; probably enough to make things work by connecting the "high" end of this resistor to the power amplifier's input.

Please note that with a current-sense scheme, there is no phase shift. Current fed into the "low" end of the coil is in phase with the voltage fed to it, and you need the power amplifier to have zero (or maybe 180 degrees with a transformer with a reversed winding) phase shift.

Theory - Energy Storage in the Primary

In this scheme, the primary has a capacitor in parallel with it, forming a resonant circuit having the same resonant frequency as the secondary. A DC voltage is applied to the primary and current increases, which stores energy in the primary's inductance. The current source is then shut off, and the stored energy becomes a strong oscillation in the primary circuit (primary coil and capacitor). This energy is resonantly coupled into the secondary.

Variations of this scheme can include schemes to suddenly charge the capacitor in parallel with the primary by ways other than storing energy in the primary's inductance - including variations of the usual spark gap method.

Typically, the primary and secondary are "overcoupled" to each other. This means that the

energy stored in the primary resonant circuit will transfer to the resonant secondary during a few cycles of the resonant frequency, then back to the primary, and repeating back and forth. The waveform of primary voltage, primary current, and secondary voltage will all usually be some sort of modulated sinewave.
If the "Q"'s of the resonant items are reduced, you can get "critical coupling" - where energy transfers to the secondary, but damps out before bouncing back to the primary. This is usually not advantageous, since maybe half the stored energy will be lost before the secondary voltage peaks. Overcoupling will normally be better since this results in a majority (preferably approaching all) of the stored energy will be in the secondary at some point. Undercoupling is a situation more lossy than critical coupling and is clearly not advantageous.

Any of these schemes generally result in Tesla coils oscillating intermittently like the usual spark gap coils, as opposed to the usual solid state coils which oscillate continuously. It is typically difficult to use switching semiconductors to store enough energy in the primary's inductance to charge up the secondary's capacitance to impressive voltages. The peak primary current and/or peak primary voltage will be big numbers!

Dipole Vs. Monopole Coils (These are mainly monopoles)

All of the actual Tesla coils mentioned below are monopoles, which have one high voltage end and the other end is grounded. A dipole has both ends having high voltage out-of-phase with each other, and the middle having little AC potential with respect to ground.
Dipoles can generate more end-to-end voltage for a given amount of energy/power and corona trouble, and can make really long sparks between electrodes connected to each end. Dipoles will usually make less near-field electric field several feet away and further away than monopoles.

Monopoles are better if you want to light neon lamps and fluorescent tubes a several feet away - and the end feed ones really do this since there is no primary coil shielding part of the secondary.

Primary fed dipoles work just like the monopole versions.

A dipole version of the end-fed coil is fed in the middle. You need a break in the wire in the middle of the coil, which is where the coil is fed. I recommend using a transformer to feed the coil to avoid strange ground currents mucking things up should the coil be not perfectly symmetric or not in a perfectly symmetric environment.

Please note that a given coil will resonate at a higher frequency, theoreticaly twice as high, when used as a dipole than when used as a monopole.

Actual Circuits - Low Impedance Primary Feed

More will come soon, but for now I have:

1. The 1978 Circuit!

This one resembles a Hartley oscillator. I did not know the whole story then! It is not an optimum circuit, but it usually worked for me.

```
             _________  |  HV out
          C /         ) (
      Q1  |/          ) (
+---------|     B+    ) (
|      B  |\     +____) (
|         E |         ) (
|          Gnd        ) (
|     R1              ) (
|______VVVVVV_________) (
     |         |         |
     |___||____|         Gnd
         ||  C1
```

The secondary is a large oatmeal box covered with one layer of magnet wire around 32 or 33 Gauge. This is nearly 6 inches (15 cm) in diameter by nearly 10 inches (25 cm) long. Cardboard is usually not good for winding high-Q, really high impedance coils, but oatmeal boxes are made of exceptionally good cardboard for this.

I have done this with other secondaries, using PVC pipe and unusually good cardboard:

1. A cardboard one 4.2 inches (10.6 cm) in outside diameter and about 12 inches (30.5 cm) long, wound with 32 gauge wire.

2. One on 2-inch PVC pipe, which is 2.5 inches (6.3 cm) in outside diameter, and this one was maybe 16 inches (40.5 cm) long, and wound with 33 gauge magnet wire. I remember this one not working as well as the larger diameter ones, but still giving impressive results.

The primary is four turns of really thick copper wire or 1/4 inch copper tubing. The coil diameter was nearly twice the secondary diameter, or about 10 inches. The coil height was about 5-6 inches, with the upper half being the primary winding and the lower half being the feedback winding. The bottom of this coil was level with the low end of the secondary.

Q1 was a 2N3055, 2N5629-2N5631, or 2N6029-6031 power transistor in a decent size heatsink. The 2N6029-6031's are PNP so you need to reverse the polarity of the power supply for these. 2N3055's are not as good as the others for not burning out at high power.

R1 and C1 - I don't remember too well - I think R1 was around a kilo-ohm or two, and C1 was maybe around .02 uF. Your mileage will probably vary anyway.

Supply Voltage - DC of no more than 40 volts. Anything higher usually burned out the transistor. Lower voltages 25-40 volts sometimes did this anyway. One thing to try is unfiltered DC, which will give you the same peak voltage but half the transistor abuse. Maybe try halfwave rectified unfiltered DC to cut the transistor heating to 1/4 of that you would get with filtered DC, but some transformers don't like substantial current in a halfwave scheme since net DC will flow through the transformer secondary and can saturate the transformer core. You should have a capacitor of several microfarads (maybe a hundred or two) from the transistor emitter to the primary center tap in order to power this whole thing with what looks like an ideal low impedance supply, at least at the oscillation frequency. This capacitor should be able to conduct an amp or two of high frequency current without heating. Tantalum types and axial lead electrolytics generally work better than "radial" lead electrolytics.

This circuit idles safely with no oscillation and it is safe to short or load down the secondary's high voltage output.

RESULTS: I got peak voltages around 50 kilovolts or a little more with a 40 volt (peak if unfiltered) DC supply and the large oatmeal box secondary, with sparks over 2 inches. With filtered DC power, this was at currents of at least a couple milliamps. Corona effects were very spectacular. The 4 inch secondary delivered at least 45 kilovolts peak and the 2.5 inch secondary delivered at least 35 kilovolts peak.

The "Mos-Zilla" circuit!

The Mos-Zilla is an overgrown CMOS buffer modified into a "linear" amplifier, with a resistance coupled gain stage at its input. This simple circuit is not really good with the linearity so it will distort somewhat in audio applications. My suspicions are that despite the distortion, it will not sound good as a guitar amp.
(Corrected 4:40 AM GMT 6/15/98 - C6 was missing.)

```
                  B+             B+             B+       B+
                   +              +              +        +
feedback           |              |              |        |
winding            >              |              >        |
+-UUUUUU-+         >R1            |              >R5      |
|        |         >        ___   |              >_____   |
>       Gnd        |       |  G||-+S             |    G||-+S
> R8               |       |   ||Q2              |     ||Q4
>                  |       |   ||-+D             |     ||-+D
|  C4              |       |      |              |        |
+-||-+       +-----+---||--+-VVV--+-----||-------+-VVV----+--||---+
|    |       |     |   C1  |  R4  |     C2       |  R7    | + C6  ) Primary
|   Gnd   R2 >  ||-+D      |   ||-+D             |     ||-+       )
|            >  || Q1      |   ||Q3              |     ||Q5       )
+_VVV___||___>__||-+S      |___||-+S             |_____||-+       |
  R9 |  ||   | G   |          G   |              >    G   |      Gnd
     |  C3   >     |              |              >R6      |
  C5===      >R3   |              |              >        |
     |       >     |              |              |        |
     |       |     |             Gnd            Gnd      Gnd
    Gnd     Gnd   Gnd
```

The secondary used for this consisted of approx. 900 turns of 33 gauge magnet wire covering approx. 8 inches (20 cm) of 6-inch plastic pipe having an outside diameter of about 6.5 inches (16.5 cm). The plastic of this particular pipe was not PVC, but I expect similar results with PVC.

The primary was, for best results with no step-up transformer, 2 turns of any thick wire. I wound this on a 30 pound potato salad bucket, which is roughly 9 inches square (23 cm) with rounded corners.
For better results still, I used a ferrite core step-up transformer with a 4 to 1 turns ratio feeding a primary winding of 6 turns. The transformer consisted of a pair of Ferroxcube 4229 3B7 pot core halves with no air gap; the primary was two turns of a quadruple strand of 18 gauge hookup wire; and the secondary was eight turns of a double strand of 18 gauge hookup wire. You can probably make a decent substitute using a flyback transformer core. You may need 3 and 12 turns instead of 2 and 8 for better performance of your transformer. The secondary should be wound over the primary or vice versa; don't put the windings over different parts of the core.

The feedback winding is a single turn around the bottom end of the secondary. You may do better with two turns - but feedback winding loading becomes a greater concern with more than one turn.

Q1 is a power MOSFET, IRF510 or IRFZ10.

(I actually did this with a P-channel IRF9Z10, which required building the whole first amplification stage "upside down".) Q2 is a power MOSFET, IRF9Z34 (P-channel).

Q3 is a power MOSFET, IRFZ34.

Q4 is four IRF9Z34 power MOSFETs in parallel.

Q5 is four IRFZ34 power MOSFETs in parallel.

You may get away with quite a range of substitutions/changes for these.

NOT SHOWN - Put a back-to-back series pair of 15 volt zener diodes from gate to source of each MOSFET. Parallel banks of MOSFETs need only one diode pair per parallel bank, not per MOSFET.

R1 is a 16 ohm power resistor, which must dissipate a few watts.

R2 is a 100K resistor, as little as 1/4 watt is OK.

R3 is a 150K resistor, as little as 1/4 watt is OK.

R4 is a 100 ohm resistor, as little as 1/2 watt is OK.

R5 and R6 are 8 ohm power resistors that must dissipate a few watts.

R7 is a 50 ohm resistor that should probably be at least 1 watt.

Please note that resistor values are not critical - deviating by 10 or maybe even 25 percent should be OK.

R8 is 2.2K preferably 1 watt, and R9 is 3.3K, preferably at least 1/2 watt. You should experiment with these. R8, R9, C4, and C5 (which is in parallel with the gate-source capacitance of Q1) combine to form a phase shift network that should have a phase shift of slightly less than 90 degrees lagging at the oscillation frequency with as little loss or feedback winding loading as practical.

C1 is a 10 uF nonpolarized capacitor, type not critical. I used a few 2uF 50V polypropylene ones in parallel. You can probably get away with as little as 4 uF.

C2 is similar to C1, preferably larger. I got away with a direct connection in place of a capacitor, but I do not guarantee not needing a capacitor here for the output stage to bias itself well enough to give enough gain for the whole thing to start oscillating.

C3 is 2 uF nonpolarized, but you can probably get away with much less maybe like .1 uF.

C4 is 270 pF, but you will have to experiment here.

C5 was zero (no capacitor at all, open circuit not a short) but you may get better results with a little capacitance here and a smaller C4 than best results with C5 being zero. The gate-source capacitance of Q1 can be significant here.

C6 is an electrolytic that does not overheat with a few, maybe quite a few amps of high frequency AC flowing through, and a voltage rating of at least 25 volts. I recommend any or all of the following to minimize capacitor overheating:

1. Use a large value of at least a few thousand microfarads.
2. Use one of large physical size.
3. Use an axial lead electrolytic - these have less resistance than "radial" lead types.

Supply voltage - Use at least 8-10 volts, but no more than 15 volts. Use lower voltages in this range until you reliably get this thing oscillating. If you use higher voltages with no oscillation, the complimentary pairs of MOSFETs will overheat.

RESULTS: With the transformer, 6 turn primary and a 14 volt supply, I got 50 KV peak at a couple milliamps. This produced 2-inch sparks. This is close to the maximum possible with a secondary wound with 33 gauge magnet wire and no silicone rubber or anything else to stop corona from forming on the top turn. With no transformer and a 2 turn primary, I got around 35 KV peak.

When the output was 50 KV peak, the current flowing through the low end of the secondary was roughly .4 amp. The voltage being applied to the primary was roughly a 25 volt square wave, the fundamental component of which was a sinewave of roughly 20 volts RMS. The primary current was roughly 7.5 amps.

One more monstrous variation, which I did not yet try: Make a duplicate of the power stage (the portion of the circuit that has Q4 and Q5) and connect that extra stage's input to the output of the stage with Q4 and Q5. Now you have two power stages whose outputs are out of phase with each other - a bridged amplifier! Connect the primary to these two outputs. You may need to increase the number of primary turns (by up to 40 percent) to avoid excessive primary current and increased MOSFET heating.

NOW HERE - links to actual circuits by a few others who have achieved solid state Tesla coils with this method.

Jump down to The Helpful Hints, where I now have these links!

Actual Circuits - End Feed

This will come soon.

But for now, there is a Tesla coil project mentioned in the September 1991 issue of Radio Electronics (a magazine now called "Electronics Now"). Please note this one is not self-oscillating, but uses an independent oscillator and amplifier. Approaching the coil while it is running will detune it.

Design Example (untested) - Energy Storage in the Primary

Secondary - 3300 turns of 34 gauge wire covering approx. 26 inches (66 cm) of a 6" (inside diameter nominal) PVC pipe, approx. 6.5 inches (16.5 cm) O.D.

Primary - 4 turns of heavy wire or thin copper tubing 16 inches (40.5 cm) in diameter and 16 inches long.

Resonant Frequency - 44 KHz.

Inductances (predicted): Secondary = .41 henry, primary = 4.4 uH.

Primary Capacitor: 2.8 uF.

Voltage Gain (optimistic!): Should be the square root of the ratio of inductances. Since the secondary's capacitance is distributed and not all current flows through all the inductance, I think effectively half of each are storing energy and this increases the voltage gain by 41 percent. 1.41 times SQR (.41/.0000044) is 430. I say hope for 350 or maybe almost 400.

To get 100KV peak with a voltage gain of 350, you need a peak primary voltage of 286 volts. This will store .115 joule in the 2.8 uF capacitor. When this much energy is stored in the 4.4 uH inductance, the current is 229 amps.

Lets say you have a 12 volt DC supply voltage and you are using a bank of fourty IRF730 or twenty IRF740 power MOSFETs (resistance around .025-.028 ohm for the whole bank). With no resistance at all, 12 volts should build up the current in a 4.4 uH inductor to 229 amps in 84 microseconds. With .025 ohm, this takes 115 microseconds. With .03 ohm, this takes 126 microseconds.

You can switch four or five IRF740's, and twicce as many IRF730's with a 555 timer IC. My experience is that National Semiconductor' LM555 works best.

Since the resonanting primary will have voltages swinging hundreds of volts in both directions, you will need diodes in series with the drains of the MOSFETs. Otherwise the resonance will stop with the first negative voltage swing since power MOSFETs have internal diodes connecting drain to source. Diodes often do not parallel well, so you need one big diode, or one in series with each MOSFET drain, or one in series with each of whatever banks of MOSFETs are in parallel. You may need high speed "fast recovery" diodes. The voltage drop of the diodes will increase the energy buildup time a little unless you increase the supply voltage accordingly.

Now what if you make this thing a dipole instead of a monopole? The resonant frequency is doubled and the amount of energy stored to produce a given peak secondary voltage (end to end) is cut by 75 percent. But adding a pair of electrodes to make a spark gap will add capacitance, so the resonant frequency will be less than doubled and the stored energy requirement will be more than 1/4 that needed for the monopole version. The primary capacitor in the above example would probably have to be around 1 microfarad, and the required primary current would be around 140 amps (for 100KV peak secondary voltage). With 12 volts applied to the 4.4 uH primary inductance, the primary current will build up to 140 amps in 51 microseconds - a little longer with resistance and the diode drop.

Helpful hints for building solid state Tesla coils!

1. I have had bad luck making push-pull amplifiers with power MOSFETs. I usually get strange and nasty parasitic oscillations. If not, I get nasty ringing when a MOSFET cuts off - such as at the end of a half cycle of a squarewave. If anyone solves this, please e-mail me!

UPDATE 7/13/99: Someone did e-mail me with some hints to make MOSFETs behave better in general, including in push-pull circuits. He suggested values usually used in audio power amplifiers, and advised that they may have to be changed for best results in solid state Tesla coils. I have yet to try this. Doing both A1 and A2 below is important; B is less important:

A1. Put a resistor and a capacitor in series, and connect this RC series combo from the MOSFET drain to ground. Suggested values are 10 ohms for the resistor and .01 uF for the capacitor. You probably want a decent capacitor and a resistor with a decent power rating, as considerable current may flow through this!

A2. Wind 25 turns of wire around a 10 ohm 1 watt or 2 watt resistor and connect the ends of the wire to the resistor leads to make a resistor-inductor parallel combination. Put this in series with whatever connects to the MOSFET drain, and close to the MOSFET drain.

B. Put a ferrite bead in series with the gate of the MOSFET.

UPDATE 7/5/2005: I tried some web searching in this area, and found some trend of successful solid state Tesla coils with the "low impedance primary feed method" using multiple MOSFETs to use a single winding primary fed by either a "half bridge" or an "H-bridge" rather than "push-pull" with a 2-winding primary.

Examples:

This one by Carl Willis

One page earlier than the above one in the site of Carl Willis, with links.

Jan Wagner's page on a self-resonant (self-oscillating) solid state Tesla coil.

Jan Wagner's page on MOSFET gate driving tips.

Jan Wagner's page on solid state Tesla coils with a "half bridge".

Jan Wagner's page on Tesla coils and high voltage stuff in general.

2. PVC pipes are generally good for winding secondary coils. But PVC pipes larger than 4 inches (approx. 10 cm) (inside diameter) are not easy to get since they are not normally used to build nor repair homes. (This is in the USA at least.)

From Robert Eastman (kodak@flash.net):

If one is partial to PVC pipe, for large coils a single large piece of pipe is not the best approach. Rather, make a skeletal coil form from numerous (typically 8 to 12) smaller pipes fastened at their ends to rings of the desired form diameter. I use rings made from PVC sheet (avail from plastic supply houses), or from masonite boiled in beeswax. Not only is this approach cheaper than using a single large pipe, but the form is much less unwieldy -- even for very large forms. Secondly, the dielectric losses are considerably reduced -- as is surface leakage -- resulting in a higher coil Q.

3. Many cardboard tubes, including most thicker ones such as carpet tubes, are too conductive to work well for Tesla coils.

From Robert Eastman (kodak@flash.net):

1) If one is using a cardboard tube as the secondary-coil form, it helps greatly to first boil the tube in beeswax before winding the coil about it. This step drives out any residual moisture from the cardboard, and keeps it out. Secondly, the wax will saturate the cardboard, making it a much better insulator. Beeswax is preferable because of its very low dielectric losses at the frequencies found in TCs (100's of kHz, typically).

(A couple things about this beeswax stuff from Don:

Don't actually make the wax boil. You do need the wax to get well above the boiling point of water to boil out every last trace of water. But if you make the wax itself boil or even get close to this, you produce extremely flammable wax vapors and probably also cause some chemical breakdown of the wax. If you see "smoke" coming from the wax, it is hot enough or even a little too hot.

I don't think beeswax has magically low dielectric losses compared to plastic tubing. However, a waxed thin cardboard tube may well have substantially lower dielectric losses than a heavy thick plastic tube simply by having much less material mass.)

4. If you short or nearly short the top end of the secondary, you may cause any oscillation to fail. Avoid doing so unless the circuitry is known to safely idle with no oscillation.

5. Corona can load down the high voltage output. Due to the corona's capacitance (this effect requires only a few picofarads!), the corona can conduct significant current - often over a milliamp. The corona is capable of burning combustible objects. You should put a "bead" of silicone rubber over the top turn of the secondary if you anticipate peak voltages of more than maybe 50 kilovolts - depending on the gauge of the wire, etc. (thicker wire will be corona-free to slightly higher voltages.) Any corona at the top turn may burn the coil tubing. This can leave carbonized spots (usually known as "carbon tracks") which can add to corona problems!

6. The end of the top turn can form a big piece of corona. This can load down the high voltage. This may even develop a bit of jet propulsion which can make the top lead fly around. You may be able to get around this by bending the top lead downward and towards the center of the coil, but with the tip pointing up but below the level of the top turn. Better still is to attach a corona-resistant electrode of some sort to the top of the coil.

7. There is some chance that drawing a spark or even severe corona from the top end of the coil may discharge its capacitance severely and you may not get a continuously arcing spark. I usually got a continuously arcing spark. If you get trouble due to sparks discharging the coil's capacitance, it may pay to use an unfiltered supply to give a pulsating output. But there should

still be some significant capacitance across the power supply connections of the primary circuitry, as close to the output amplifier stage (or power oscillator stage if used instead) for the circuitry to work well.

8. I repeat, have a couple hundred microfarads of capacitance across the supply rails as close to the power transistor(s) as possible. This capacitance may well be conducting substantial high frequency current and should be of a type good for this. This may be a parallel bank of tantalum capacitors.

9. Connect the low end of the secondary to the nearest good ground. Connect the primary circuit's ground to this point, and not too far from the secondary. Connect all nearby substantial metal objects to this point also, and not too far from the secondary. You want ground currents from objects exposed to the secondary's electric field to return to the low end of the secondary in as short a path as feasible.

Back up to Don's Tesla coil main page.

Back up to Don's home page.

Written by Don Klipstein.

Please read my Copyright and authorship info.
Please read my Disclaimer.