SOLID STATE TESLA COIL

Jump to most recent progress

Introduction

As a power electronics engineer, I frequently work with large semiconductors in power supplies and motor drives, etc. These often switch thousands of watts at several hundreds of kilohertz. Modern power transistors offer an increasingly viable alternative to the Vacuum Tube Tesla Coils, as performance improves and prices continue to fall.

Whilst testing a switch mode power supply for a customer, the TC resonator at the end of the bench caught my eye, and curiosity got the better of me. I could not resist the temptation to see what would happen if I replaced the high frequency transformer in the supply with a primary coil feeding the resonator. The worst thing that could happen was that the power transistors would fail catastrophically , and after all the supply wasn't mine anyway ;-)

It actually worked surprisingly well (for a few seconds), and I decided to design my own solid-state mini coil.

Design

The first design was based around two IRF740 MOSFET devices made by International Rectifier. The two switching devices are connected in a half bridge configuration as shown in the schematic below. These devices are very close to the theoretical "ideal switch". They can switch 400volts at 10amps in around 50 nanoseconds and are reasonably priced.

The half bridge is fed from the 240Vrms mains supply, and the MOSFET devices are turned on alternately at roughly 250kHz. The high voltage square-wave output from the transistors is fed into a 25 turn primary which is tightly coupled to the bottom portion of the resonator. At resonance the base current of the resonator is sinusoidal, and a sinusoidal current flows in the primary coil also.

Simplified circuit:

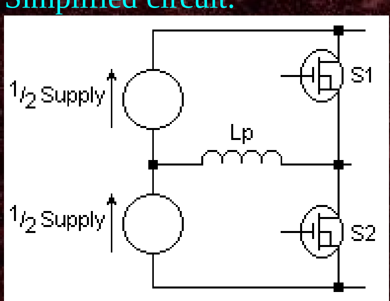

Primary voltage waveform (Red square wave,)
and primary current waveform (Green sine wave,)

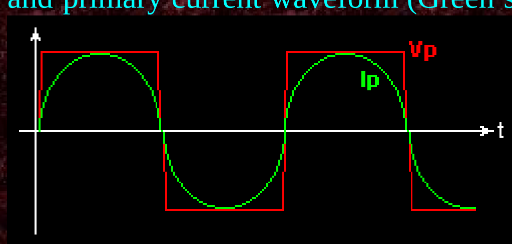

If the Tesla Coil is driven at its resonant frequency then the switching transitions of S1 and S2 occur when the current (Ip) passes through zero. This means that switching losses in the MOSFETs are practically eliminated, and heating is due to conduction losses only. (This is technique is explained as "soft-switching" or "ZCS" in many Power Electronics papers.)

An advantage of the primary feed method is
that it provides the necessary voltage

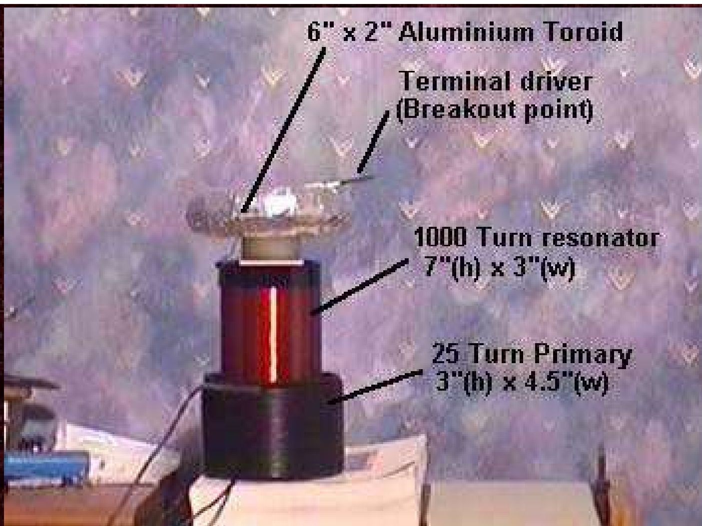

transformation required to match the output impedance of the inverter to the resonator. This negates the need to employ a separate high frequency matching transformer or the use of elevated supply rails to get the required drive voltage.

A significant disadvantage of the primary feed method is that very tight coupling is required (k>0.35) in order to get good power transfer. This makes insulating the primary from the secondary somewhat challenging as the power level is increased.

The drive electronics is based around the TL494 PWM controller IC made by Texas Instruments. This IC is fairly "long in the tooth" but it is well behaved and is also easy to obtain. The IC contains an internal sawtooth generator and the necessary comparators and latches to produce the drive signals required for each MOSFET in the half bridge. The IC generates two complementary drive signals with a short dead time between transitions to ensure that one MOSFET has had time to turn off before the opposing device is turned on. Without this precaution the conduction times of both devices can overlap shorting the mains supply with interesting *(read expensive,)* consequences.

Click here to view original schematic…

The two outputs from the TL494 are boosted in current by push-pull stages and are used to drive the primary of a small ferrite transformer. This transformer serves to isolate the sensitive low voltage control circuitry from the high power MOSFET side, whilst coupling the drive signals to the gates of the two MOSFETs. (Power semiconductors usually fail short-circuit. Without this isolating transformer such a failure would almost certainly lead to damage of the control circuitry too.) ***(Please note that the schematic linked opposite is not a finished design, and contains some errors relating to reverse recovery of the MOSFET body diodes. It is presented here only as a reference to show the progression of the design !***

The latest schematic can be found further down this page.

This isolation transformer has two secondary windings wound in opposite directions to drive the gates of each MOSFET. This serves two functions. Firstly, it ensures that when one MOSFET is turned on by a positive gate voltage, the opposing device is held firmly in the off state by a negative gate voltage. (This negative bias is useful to prevent spurious turn-on due to the Miller capacitance from drain to gate of the MOSFET.) Secondly, the two isolated secondary windings allow the high-side (top) MOSFET to be driven without the need for complicated floating or bootstrap power supplies.

Operating modes

The overall arrangement is very flexible because the oscillator is continuously running. The MOSFETs merely chop up the supply voltage as instructed by the oscillator and feed the RF to the primary coil. This means that the supply voltage to the half bridge can be DC or virtually any desired waveform you choose to throw at it.

The effect of varying the supply voltage is to Amplitude Modulate the RF applied to the TC primary winding.

I tried 4 different supply schemes which gave different RF envelopes and radically different spark characteristics:

Half wave rectification,

RF envelope:

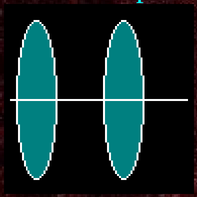

This was achieved by inserting a diode in series with the mains supply to the MOSFET half bridge so that only positive half-cycles resulted in current flow. (This is necessary anyway in order to prevent shorting of the negative supply cycles by the MOSFET body diodes !) The RF envelope consisted of rounded bursts of RF lasting 10ms with 10ms gaps in between.

Spark appearance:

Sparks were roughly 6 inches long, very straight and "sword-like" in character. The absence of branching in the streamers struck me as being very odd. Apparently this appearance is common in Vacuum Tube TCs also.

The sound was like a muffled 50Hz buzz but still quite loud.

Power is estimated to be around 160 watts in the picture opposite.

Full wave rectification,

RF Envelope:

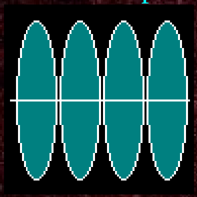

This was achieved by using a full wave bridge rectifier between the mains line and the MOSFET bridge. This ensures that there is current flowing through the inverter during the entire supply cycle. The power drawn from the mains line roughly doubled as expected and the RF envelope assumed the classic full-wave rectified shape. This implies a considerable increase in the average RF energy applied to the TC.

Spark appearance:

Sparks became noticeably fatter and more bushy, but there was no increase in length. The picture opposite clearly shows the greater "fullness" of the discharge including wispy branches leading off from the main feature.

The tone of the sound changed to twice the pitch (100Hz) and became distinctly more "full-throated" and hissy.

Power is estimated to be around 300 watts.

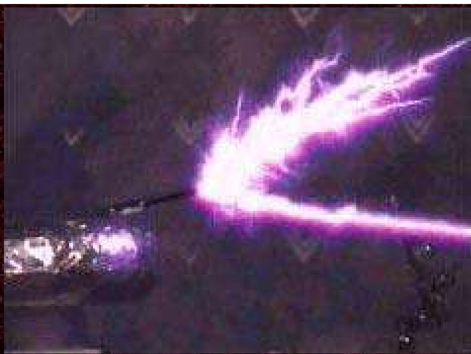

These two pictures show the ability of the coil to produce a lot of corona from points. Notice how the discharge often divides into two jets of corona right at the breakout points.

Smoothed DC,

RF Envelope:

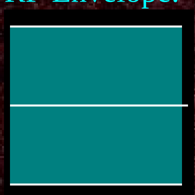

This was achieved by using a full wave bridge rectifier and a large high voltage reservoir capacitor ahead of the MOSFET bridge. This provides a constant supply of around 350VDC to the inverter. Power draw increased again due to the sustained high voltage, and the RF envelope was that of a continuous-wave source. During this test some warming of the MOSFET heatsink was noted due to the high average current.

The discharge from the breakout point became very bushy. It looked and sounded like a jet of burning gas, and spread out like a cone from the discharge point.

All of the buzzing was gone to leave a pure hissing sound. This test produced a lot of ozone really quickly and also overheated the thin wire at the base of the secondary coil blistering the varnish.

Power is estimated at 420W in the picture shown opposite.

Phase angle controller,

RF Envelope:

A phase angle controller (similar to a commercial light dimmer) was connected ahead of the MOSFET inverter in order to interrupt the supply to the half-bridge. The phase angle controller was set to turn on exactly at the peak of the mains supply cycles and remain on until the end of each half-cycle. This leads to a very sharp rise in the voltage applied to the inverter, rising from 0 to around 350 volts in a matter of microseconds. This sudden application of power results in a sharp rise in the RF envelope and an interesting effect on the spark characteristic.

Spark appearance:

The discharge from the breakout point became branched like a conventional spark gap TC. The sparks were about 6 inches long, distinctly spidery and danced about frantically.

The sound was considerably sharper and more raspy, no doubt due to the rapid rise of the RF envelope. It was similar to the sound from a conventional 100BPS synchronous TC, but sounded slightly deeper and more fuller.

RMS Power level was thought to be around the 180 watt mark, although this measurement may not be particularly accurate.

Another picture showing a peculiar branching in the discharge. The arc to the lower right of the picture is striking a piece of metal which was not earthed.

Average RF power

Unlike a conventional damped-wave Tesla Coil, the solid state Tesla coil is capable of producing considerable amounts of sustained RF power. This leads to a few unusual things:

Firstly, the base of the secondary became very hot due to the high RMS current flowing through the fine wire. Maybe skin effect plays some part in this also. This is particularly noticeable if the system is run in CW mode for any length of time.

There is visibly more current in ground strikes than found with my spark gap TC.

Sparks to ground appear like pale ghostly white flames and arch upwards with the heat like the arc from a Jacobs Ladder. Anything flammable catches fire instantly in the arc.

I noticed that I got tiny RF burns if I touched anything metallic in the vicinity of the running coil even at fairly low power levels.

At one time I forgot to put the breakout point on the solid state coil, and an unused resonator about 2 feet from the solid state coil, (but quite close to me,) sprang to life with a firey crown of corona. Boy did that surprise me !!!

Conclusion

Please note that building a solid state tesla coil **IS NOT EASY**. In fact both the design and construction present significantly different and far more complex challenges than those encountered in conventional tesla coil work. A very fast (100MHz) oscilloscope and a large bag of MOSFETs are essential. The biggest problem with this type of design is that a blown MOSFET is often the first sign you get about an underlying problem, so diagnosing the cause of blown semiconductors can sometimes be difficult. Despite these difficulties, the Solid State coil is a beautiful thing when it works properly.

I think that a small solid state (or Vacuum tube) Tesla Coil is the method of choice for "up-close" analysis of spark behaviour and demonstrations. When compared to a conventional coil it has the following characteristics:-

- Less audible noise. Not so much crash and bang, more hum and hiss.
- Less RF hash radiated. The SSTC is very clean due to its continuous electronic source of RF, and causes no TVI.
- Higher average RF power. The SSTC seems to produce a much stronger RF field than a similarly rated spark gap TC.
- Great for corona displays and lighting neon tubes at a distance without wires.
- Wide variety of spark characteristics from forked lightning to flames by modulating the RF generator in different ways.
- Does not shock so much as burn, but such RF burns are reported to be very nasty.
- Ideal for research because the system is under full electronic control. (Maybe one could adjust the burst rate fast enough to play the national anthem ?)

Significant downsides to the solid state approach are as follows:-

- Requires that the designer have a good working knowledge of power electronics,
- Careful attention must be given to layout, and screening,
- Suitable semiconductors are moderately expensive,
- Semiconductors are still quite fragile in such an application, and are not forgiving of any mistakes.

Application notes and design examples provided by device manufacturers help greatly with design and layout tips, and power semiconductors are consistently becoming faster, more robust, and cheaper, so the future looks promising for the Solid State Tesla Coil.

Recent developments (18" sparks)

Late last year I tackled the reliability problems with the original SSTC design. I also modified the circuit to form a full H-bridge configuration in search of some longer sparks.

The resonator pictured here is 3.5" x 16" and is topped with a 6"x1.5" toroid. Sparks look

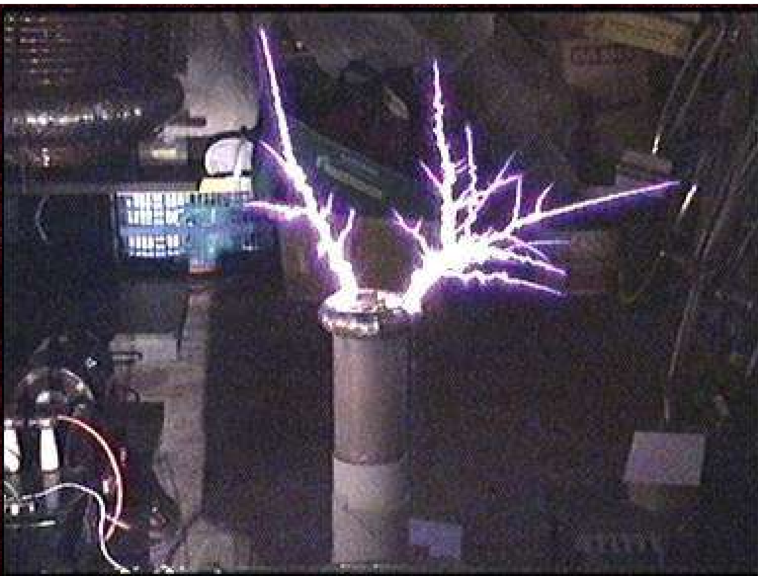

similar to those from a conventional spark gap Tesla Coil, although they are somewhat thicker and hotter.

The driver runs directly from the 240V 50Hz AC mains which is then half wave rectified. Current draw is approximately 5 amps RMS.

It uses four STW15NB50 MOSFET devices connected in a H-bridge arrangement. This drives the two ends of the primary coil in opposition (anti-phase) and effectively doubles the voltage swing that can be developed across the primary winding. (This really helped achieve a good spark length.)

There are currently no smoothing or energy storage capacitors in use here. The primary consists of 19 turns of wire and is link coupled over the bottom third of the resonator. (k estimated at around 0.40) The RF peak envelope power has been measured at 4800 watts, so the RMS power input should be about 1200 watts or so. Peak RF current in the primary is 22 amps, and the resonator appears like a constant current sink once the breakout potential has been reached.

The resonant frequency is nominally 350kHz, but the frequency of the driver is dynamically swept throughout the mains supply cycle in an attempt to maintain correct tuning as the sparks grow.

This dynamic tuning is of some importance to achieving long sparks. Without some automatic adjustment of the oscillator the growing sparks "snub" themselves out as they detune the resonator and limit the terminal voltage. Dynamic tuning is achieved by feeding a small portion of the supply voltage into the frequency determining part of the driver circuit. This causes a progressive drop in the drive frequency as the supply voltage increases and the sparks propagate. It is very crude but definitely makes an improvement.

The four MOSFETs are only slightly warm after a 3 minute run, and I have run it for 30 minutes continuously to check reliability. After the longer run, the heatsink was quite warm, and both primary and secondary displayed noticeable heating.

Average spark length is around 14 inches with occasional hits out at 18" or so. The sparks generate a loud thudding humming sound, and appear like inch thick flames where they contact the toroid.

I have also seen several brilliant white balls emitted from the toroid during operation. (See frame sequence opposite.) These are thought to be balls of burning Aluminium which come from the surface of the foil covered toroid, although it really surprised me when it first happened !

The surface of the toroid is covered with small 1/8th inch foil bumps to promote breakout. **If a smooth toroid is fitted without any breakout**

points, there are severe flashovers which instantly burn up the plastic primary form: A significant problem with using such a tight coupling.

I have no plans to increase the spark length of this particular design, as it is intended as a compact portable unit. It will be packaged neatly and used for demonstrations at Teslathons etc. However, I may try to build a bigger solid state system in the future, as the flexibility and absence of TV and radio interference is appealing to me.

Several members of the Tesla List have also pointed out that this coil represents a good platform for investigating the little understood areas of spark loading and impedance matching to corona.

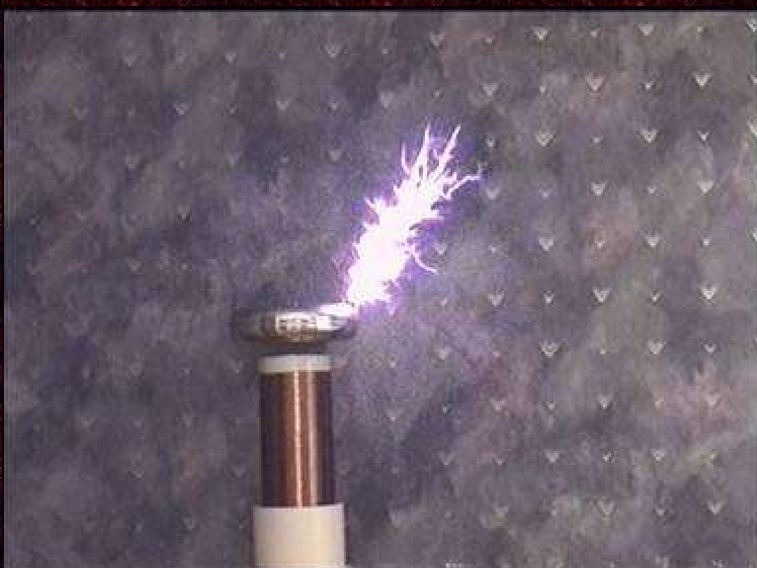

The picture opposite shows a fierce 9 inch flame produced by running the Solid State driver from a continuous smoothed DC supply. (CW mode)

This causes noticeable heating in the driver, the primary winding, and the lower portion of the Tesla resonator.

Input power was measured at 1500 watts. Despite this low power the end of the breakout point (small terminal driver) became melted into a ball. The tip continued to glow for seconds after the power was switched off.

The discharge was very hot in this mode, and "heat-shimmer" was visible above the corona. Sound was a rushing, hissing, crackling noise.

An unusual phenomena frequently observed above power arcs. In these two frames the coil is arcing over a distance of 6 inches to a grounded wire.

Brief smudges of pale yellow light are seen rising above the main flame-like arc. I have been informed that this is due to combustion of trace gases in the air.

H-bridge driver schematics

The latest schematics for the control electronics and power electronics can be downloaded by clicking the two links below. These schematics have been checked for obvious mistakes, and are believed to be error free. Significant improvements have been made from the original design in the following areas:

1. Re-configuration of the MOSFET gate drivers to reduce the dead-time at switching transitions. The old design had a large 5% dead-time at switching instants to allow one MOSFET to turn off before the other is turned on. This dead-time was found to be excessive, and caused the body diodes of the MOSFETs to conduct heavily due to the free-wheeling current.

2. Isolation of the MOSFET body diodes, using series a Schottky diode and parallel fast recovery diodes. This eliminates problems due to the slow reverse recovery characteristics of the body diode. This modification combined with the one explained above, dramatically reduces MOSFET mortality rates !

3. Dynamic tuning. The carrier signal generated by the TL494 is "Frequency modulated" by the HV supply voltage. The frequency is actually swept down by a few percent as the voltage increases in an attempt to track the resonator frequency as the sparks grow. This is quite rough, as is does not take into account that detuning only happens above breakout, etc. However it has been found to be very effective, most likely because the loaded resonator Q is low, and the tuning range is actually quite broad during sparking conditions ?

Click here to view the schematic for the control electronics:

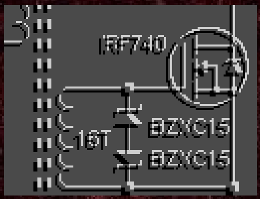

Click here to view the schematic for the power electronics:

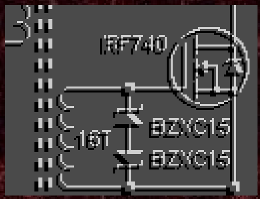

Pictures of the driver board

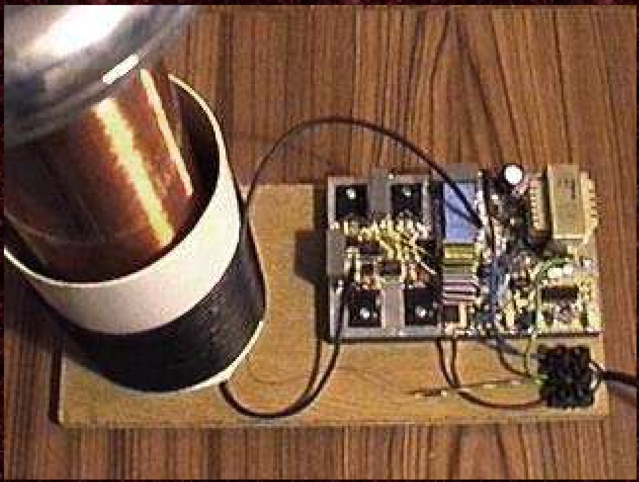

The picture opposite shows my Solid State Driver board connected to the Tesla Coil primary winding.

The PCB measures 6" x 4" and is mounted directly on top of a large Aluminium heatsink for cooling the power semiconductors.

The small transformer at the top right of the PCB provides a 15 volt supply to power the control electronics. Everything else runs directly off the 240V mains supply.

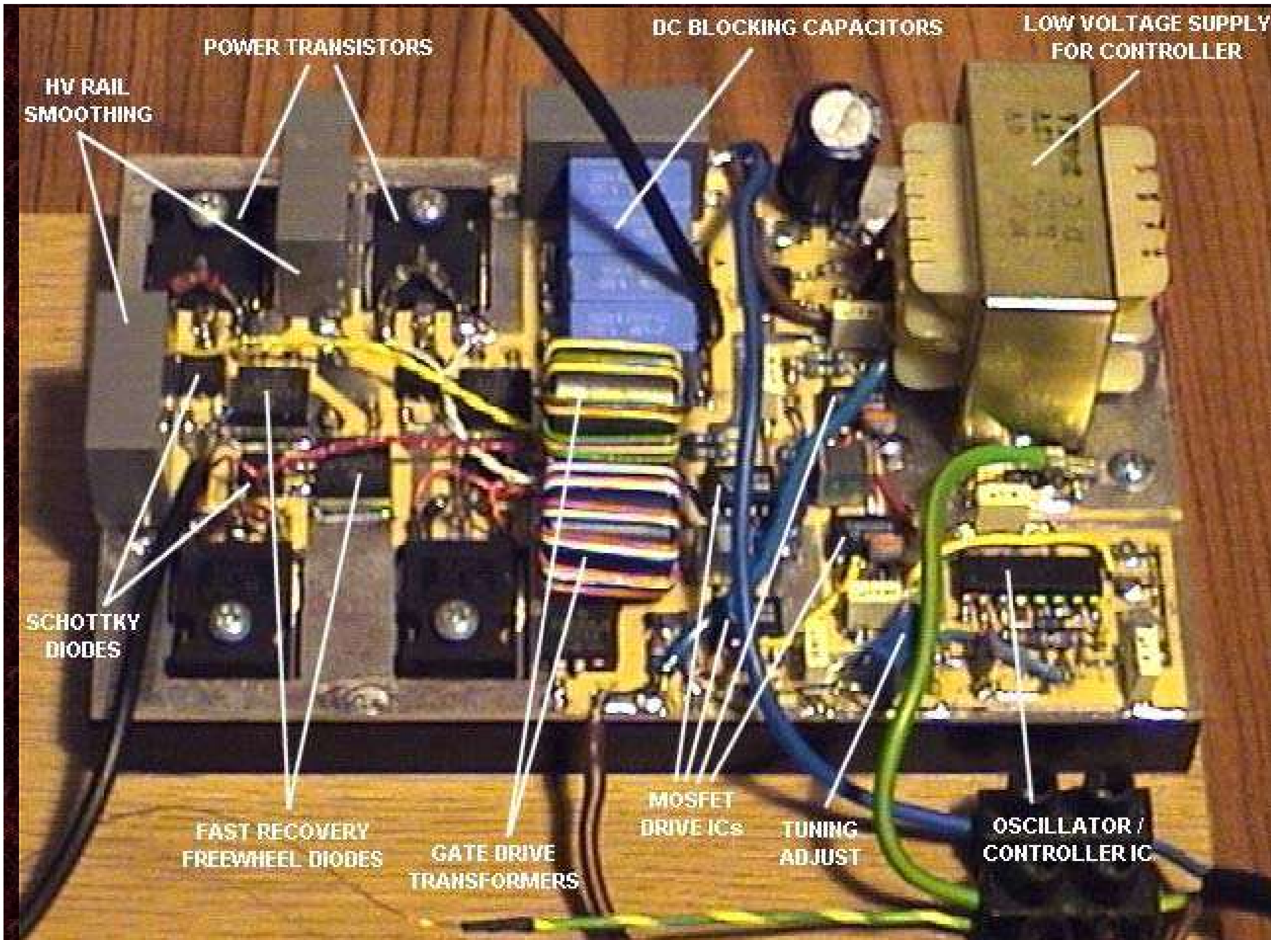

As a warning to anyone contemplating building a similar system, the development of this project was **VERY** expensive, I spent the equivalent of around 600 dollars on various semiconductors and had to borrow some sophisticated test gear to debug the design. However, now that it works well I think it is beautiful ! Careful attention must be paid to layout, wiring lengths, heatsinking and shielding to ensure reliable operation. Some design and construction tips can be found by clicking the link below.

Burning steel and singing arcs

Click here to see new pictures of a steel breakout point burning away, and music coming from a spark !!!

Solid State Tesla Coil theory

If you have bothered to read this far you probably want to know more about SSTC operation. But be warned, this is where it gets a bit more heavy...

Click here to read my in depth SOLID STATE DRIVER THEORY pages. (Recommended reading if you are going to build your own driver.)

or click on this link to see some reasons why MOSFET devices fail in solid state TC duty, if you have built your own driver but have problems!


Photo from Cambridge 2001 by Mark Hales.

Future developments

Here are a few thoughts for future developments on the SSTC theme:

1. It is planned to modify this existing rig to operate from smoothed DC later this year, for the purpose of investigating corona impedance and loading issues.
2. More investigation into decreasing the un-loaded resonator base impedance, in order to get more power into the resonator without requiring very a high coupling coefficient. (Possibly winding a physically larger secondary to be operated above a ground plane.)
3. Dynamic tuning based on sensing of the resonator base current. Essentially the resonator is made to be the frequency determining part of the oscillator, so the driver frequency "perfectly" tracks the resonant frequency during streamer growth. This will also ensure that switching transitions occur at zero current, resulting in reducing switching losses.
4. Development of a twin SSTC. This should provide longer sparks between towers, without increasing voltage stresses across each tower.
5. Improve the heatsinking of the present design, as there are still some thermal issues when run times are long.

Credits and Links

Many thanks to John Freau, Alan Sharp and Paul Nicholson for providing information, tips and suggestions for my CW coil work.

Here is a link to **Alan Sharp's Web Page** which contains excellent information about Solid State Tesla Coil design and construction. **International Rectifier's** web site also contains many valuable application notes covering topics such as MOSFETs drive circuits, etc. Here is a link to **John Freau's Web page** which contains some information about Vacuum Tube Tesla coils. (Vacuum tube coils are very similar to solid state coils in their method of operation.)

Also be sure to check out the dedicated section of SSTC related links on my main links page. Countless people have contributed ideas to my SSTC work, and many others have built solid state Tesla coils based on information presented here. There is an ever growing amount of information available about SSTC stuff on these sites.

Back to home page

SOLID STATE TESLA COIL

DRIVER CIRCUIT THEORY
(Page 1)

Abstract

This paper discusses many of the factors that influence the design of a solid state resonant coil driver. It assumes that the reader has basic knowledge of circuit theory and is primarily aimed at those with some experience in the field of power electronics. Simple circuit theory is used to analyse the nature of the load presented by a Tesla coil. This model is then used to investigate implications on the design of the drive electronics. Finally the operation of the drive electronics is explored under a variety of operating conditions. This paper attempts to concentrate more on conclusions and their implications, rather than getting bogged down in complex maths and circuit theory. It aims to highlight the shortfalls in classical inverter designs when being applied to this unusual application.

Introduction

When designing a Solid State Tesla Coil driver, the designer can consult many power electronics books, technical papers, and manufacturers data sheets. These often highlight good design practice, and may include valuable tips for layout and construction. However, the designer should be aware that the solid state Tesla coil driver has a distinct difference when compared to most classical inverter designs. This difference is in the nature of the load, and is worth looking at in some detail since it can make the difference between a robust design that performs well, and one that fails prematurely.

Inverters normally supply power to loads that are a combination of resistance and inductance, (e.g. Electric motors, transformers, etc.) This combination of inductance and resistance gives rise to a load current which has a triangular wave shape. Most inverters produce some sort of square-wave output voltage that is applied across the load. The inductance of the load "integrates" the applied voltage and the current rises and falls during the high and low portions of the output voltage waveform. See diagram opposite.

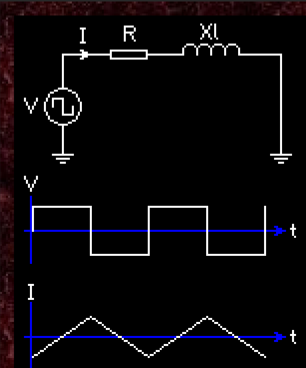

Classic power electronics theory is based around this idea of applying a square voltage waveform to an inductive load resulting in a triangular current flowing from the inverter through the load. (Note that the load current is at its maximum when the voltage changes polarity, so the active components in the inverter have to switch a significant current.)

In contrast, the inverter used to drive a Tesla coil is supplying power to a load that contains resistive, inductive and capacitive components. (The Tesla coil resonator.)

This combination of R, L and C results in a resonant condition which responds most favourably to a particular frequency, (its natural resonant frequency.) If this load is driven with a square voltage waveform at its resonant frequency, it will result in a **sinusoidal current** flow.

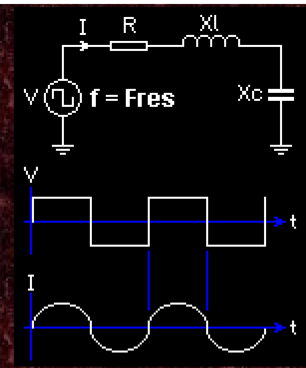

The square voltage waveform contains a fundamental frequency and all of the odd harmonics of this frequency, but the resonant load only "sees" the fundamental frequency. The load current contains only the fundamental frequency from the square wave that was applied, so it is sinusoidal in shape.

(Note that the sinusoidal current passes through zero at the same time as the voltage changes polarity, so the inverter does not have to switch an appreciable current in this case.)

With the exception of some modern "soft-switching" schemes, this resonant load condition is not discussed in many power electronics books. However, it is worth looking into the nature of the resonant load in more detail. After all, we seek to maximise the power that we can deliver to this resonant load, whilst also ensuring that our inverter will have a long and happy life driving this load. Building an inverter based on traditional power electronics theory does not necessarily ensure either of these requirements are met.

The sections that follow describe the nature of the load that a Tesla coil resonator presents to the driver under various conditions, and the implications that this has on the design of the driver circuit. Some thought is also given to the effects of spark breakout and various drive methods for coupling power into the coil effectively.

The behaviour of the load

Such a high-Q resonant load as that presented by a Tesla coil is rarely encountered in either the power electronics or RF engineering fields. High Q-resonators essentially store up energy over time. This is the exact opposite of most power supply and RF systems where the aim is to transfer power quickly with the minimum of losses. (The storing-up of energy in a high-Q resonator implies high circulating currents and therefore increased losses, which are usually undesirable.) However, in our application it is this "ring-up" over several cycles that gives rise to the high output voltages and sparks that we desire.

The load presented by the Tesla coil, **when fed from its base**, can be modelled as a Resistance ***R*** in series with inductive ***Xl*** and capacitive ***Xc*** components. This is familiar as the series RLC circuit shown below:

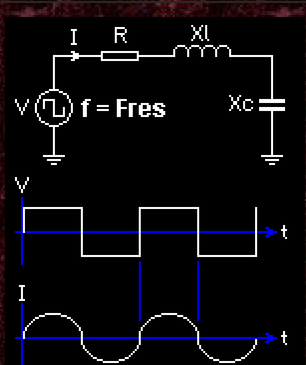

The resistive part ***R*** is the sum of all of the resistive losses for the resonator. For example, DC resistance, skin effect, proximity effect, ground-plane losses, coil former losses, radiation resistance, and corona losses all contribute towards this resistive component.

The inductive part ***Xl*** is due to the inductance of the Tesla coil resonator. The reactance ***Xl*** depends on the inductance of the coil and increases as the applied frequency increases:
Xl = 2 pi f L

The capacitive part ***Xc*** is due to the self-capacitance of the Tesla coil, and the

	capacitance of the terminal connected to the top of the resonator. The reactance ***Xc*** depends on the capacitance and its magnitude decreases as the applied frequency increases: ***Xc = - 1 / 2 pi f C***

Since all three elements are connected in series, then the total impedance presented to the driver is the sum of the three components.

Zload = R + j Xl + j Xc Where: ***j*** is the imaginary number ***sqrt(-1)***
(Used to keep the mathematicians happy ;-)

Note that real power can only be dissipated in ***R***, since ***Xl*** and ***Xc*** are purely reactive.

(Please do not get deterred from reading this because of the algebra. You do not need to understand all of the maths to understand this paper. I have tried to keep the maths to an absolute minimum !)

Driving the coil at exactly Fres

At a particular frequency the values of the inductive reactance ***Xl*** and capacitive reactance ***Xc*** are equal and opposite. This is known as the resonant frequency and can be found as follows:

Xl	=	***- Xc***
2 pi f L	=	***1 / 2 pi f C***

Some rearranging to find ***f*** confirms that the resonant frequency depends on the values of ***L*** and ***C*** as expected.

Fres = 1 / 2 pi sqrt (L C)

Since the reactances ***Xl*** and ***Xc*** are equal and opposite when the coil is driven at ***Fres***, they cancel each other out and the load impedance appears like a pure resistance.

Zload = R + ~~j Xl + j Xc~~

Zload = R

There are three important consequences of this:

1. The impedance ***Zload*** at resonance is the lowest impedance that the Tesla coil can ever present to the driver. Therefore the coil will draw maximum current from the driver when driven at ***Fres***. The impedance is always higher at any other operating frequency due to the addition of the reactive components.
.
2. If maximum current flows through the load at the resonant frequency, then application of Ohm's law shows that maximum voltage must also be developed when driven at ***Fres***.
.
3. The pure resistive load at resonance implies no phase shifts. It is reactive components that give rise to phase shifts and since these cancel at ***Fres*** then the load current drawn from the driver must be in phase with the applied voltage.
.

The resonator presents a very well behaved load to the driver provided it is driven at **exactly** the correct resonant frequency. However, this is difficult to ensure under all conditions. Therefore we should look at what happens if we drive the resonator at a frequency that deviates significantly from its natural resonant frequency.

So how much does the driver frequency need to deviate from the resonant frequency before something interesting happens…?

Q factor and bandwidth

In the case of a Tesla coil resonator we usually seek to minimise the winding resistance, and hence maximise the unloaded Q factor of the resonator. This keeps power dissipation due to I^2R losses to a minimum, and also maximises the voltage gain that we get due to resonant rise.

Q = Xl / R

For a given inductance and operating frequency, ***Xl*** remains fixed, so maximising ***Q*** implies minimising ***R***.
Since the load impedance equals ***R*** at resonance, this also ensures that the resonator presents a nice low impedance to the driver at resonance.

In order to see the effect of driving the resonator at frequencies above and below its natural resonant frequency, we need to look at the formula for the load impedance again because ***Xl*** and ***Xc*** will no longer cancel each other out:

Zload	=	***R***	+	***jXl***	+	***jXc***
Zload	=	***R***	+	***j 2 pi f L***	-	***j 1 / 2 pi f C***

The ***Xl*** term is proportional to the applied frequency, and the ***Xc*** term is inversely proportional to the applied frequency. (The resistive part ***R*** is just plain old ***R*** regardless of the frequency.)

As we move away from the natural resonant frequency, ***Xl*** and ***Xc*** will no longer cancel each other out, and we have a small net reactive contribution. Minimising ***R*** has the side effect that it makes even small amounts of ***Xl*** or ***Xc*** become a significant part of the total load impedance.
(EE's will draw some vector diagrams at this point, just to be sure ;-)

In summary the load impedance becomes very sensitive to the smallest of changes in *f* around the natural resonant frequency.

Driving the coil well away from Fres

We can quickly see that driving the coil at a **very high frequency**, (for example ***f*** = Infinity,) would cause the inductive reactance ***Xl*** to be very large and the inductive reactance ***Xc*** to be insignificant. In this case the load impedance would be very high because it is dominated by the large ***jXl*** component. Also, because the ***jXl*** term is very large, it swamps the resistive part ***R*** and the load is considered to be highly reactive.
At frequencies far above the natural resonant frequency of the coil, the load appears almost like a pure inductor. There is a 90 degree lag between the applied voltage and the current that flows through the load.

Similarly, if we consider driving the coil at a very low frequency, (for example ***f*** = 0 Hz,) then we can see that the opposite situation occurs. The capacitive reactance ***Xc*** becomes very large and the inductive reactance ***Xl*** becomes insignificant. In this case the load impedance is very high because it is dominated by the large ***jXc*** component. Also, because the ***jXc*** term is very large, it swamps the resistive part ***R*** and the load is highly reactive again.
At frequencies far below the natural resonant frequency of the coil, the load appears almost like a pure

capacitor. The current that flows in the load leads the applied voltage by 90 degrees.

Since both of these situations result in a high load impedance then little current flows. Therefore there is minimal danger to the driver circuit when it is tuned a long way from the resonant frequency of the Tesla coil.

If driving the Tesla coil at the exact resonant frequency results in a pure resistive load, and driving the resonator far away from resonance results in a high impedance (open circuit) condition, then what happens at frequencies in-between ?

Problems associated with being "out-of-tune" are most troublesome when the inverter attempts to drive the resonator at a frequency that is only slightly away from the natural resonant frequency of the coil. This is very important when you consider that the small bandwidth of a high Q resonator demands very precise tuning : A difficult thing to guarantee in practice given the influence of corona loading and the surrounding environment !

Driving the coil just above Fres

If we consider a resonator being driven at a frequency ***f*** just above its natural resonant frequency, we can see that there will be a small net inductive reactance. This net reactance is small compared to the resistive part of the load, so the total magnitude of ***Zload*** is still low and considerable current flows. However, there is enough inductive reactance to cause a significant phase shift in the current that flows through the load.

When driven at the exact resonant frequency, we saw that the load current was sinusoidal and was precisely in phase with the applied voltage waveform. This is no longer the case when the coil is driven slightly above its resonant frequency. A substantial current still flows, but the current drawn by the load starts to lag behind the applied voltage waveform by a short time as shown below:

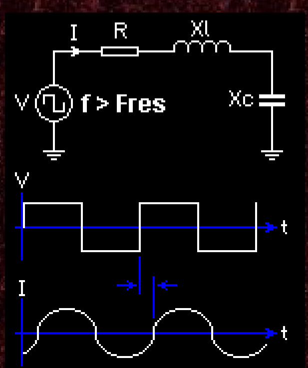

Essentially the load current does not change direction at the same time as the applied voltage, but continues to "free-wheel" in the same direction for some time.

It eventually changes direction shortly after the voltage changed polarity. It is the net inductance of the Tesla coil at this operating frequency that keeps the current free-wheeling for a short time.

(It may help to imagine fluid flowing in a pipe, which keeps flowing for a short time under its own momentum when the pump is turned off.)

This has some implications on the design of the inverter if it is to withstand this condition without damage. (See DRIVER CIRCUIT BEHAVIOUR sections later.)

Driving the coil just below Fres

If we consider a resonator being driven at a frequency ***f*** just below its natural resonant frequency, we can see that there will be a small net capacitive reactance. This reactance is small compared to the resistive part of the load, so the total magnitude of ***Zload*** is still low and considerable current flows. However, there is enough capacitive reactance to cause a significant phase shift in the current that flows through the load.

When driven at the exact resonant frequency, we saw that the load current was sinusoidal and was precisely in phase with the applied voltage waveform. This is no longer the case when the resonator is driven slightly below its resonant frequency. A substantial current still flows, but the current drawn by the load starts to lead the applied voltage by a short time as shown below:

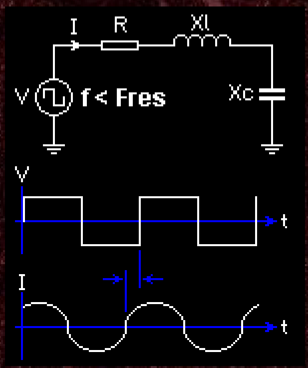

Essentially the load current changes direction slightly before the applied voltage changes polarity !

At first this behaviour seems odd, but it is due to the net capacitance in the Tesla coil at this operating frequency.

(It may be helpful to imagine a container being "over-filled" with fluid. When the container becomes full there is a back-pressure, and fluid starts to flow back out of the container. This analogy is pretty weak, but it helps me to get my head round what is happening here.)

This property of load current reversal before the applied voltage is somewhat unique and is a direct consequence of charge storage in the capacitance of the Tesla coil. Most inverters for motor drives or power supply applications drive power into loads which are highly inductive. For this reason the behaviour discussed here is rarely encountered elsewhere in the Power Electronics field.

This strange behaviour has some important implications on the design of the driver. Specifically it must be able to withstand this leading power factor condition without damage. (See DRIVER CIRCUIT BEHAVIOUR sections later.)

Impedance plot

The impedance plot below summarises how the impedance of a base-fed Tesla resonator changes with tuning. It shows how the **magnitude** of the load impedance varies as the frequency of the driver is swept from very low to very high.

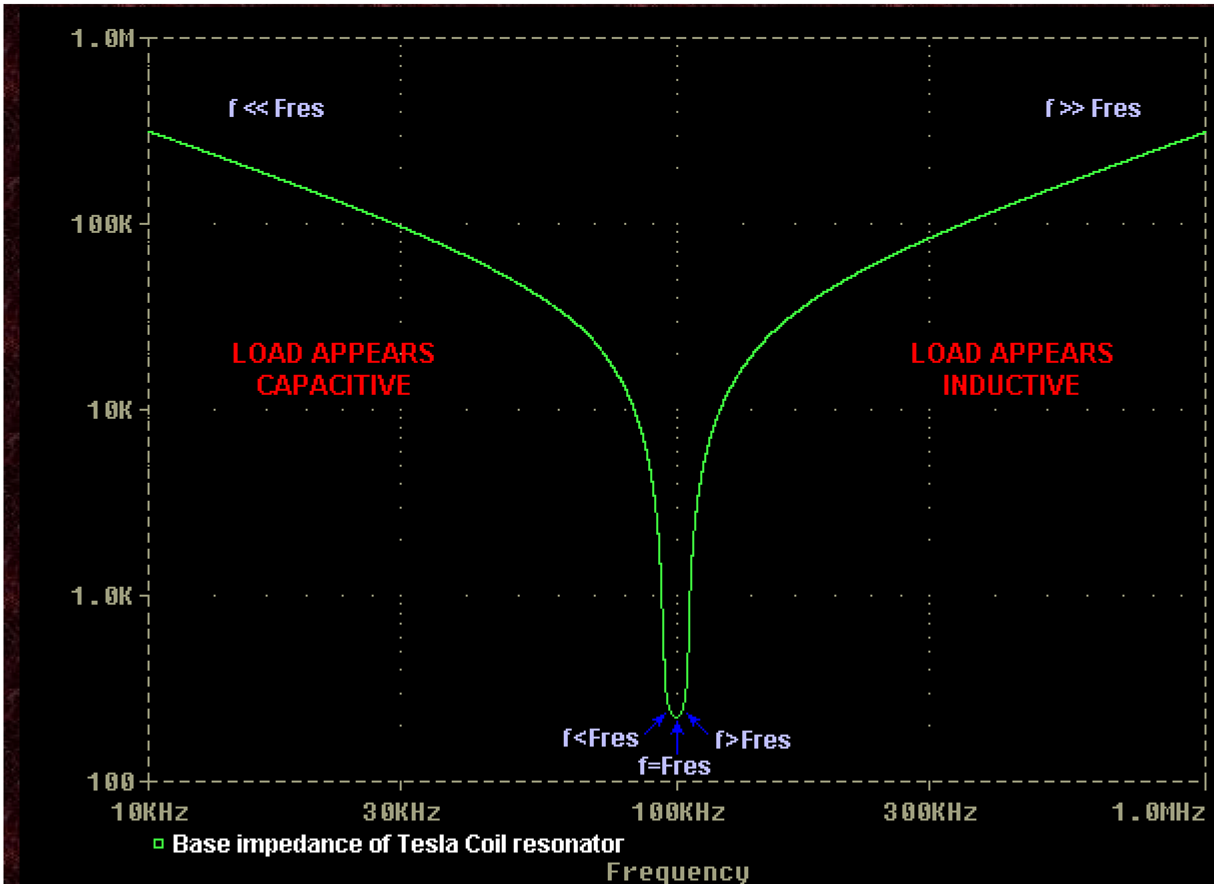

Either side of the resonant frequency the impedance presented at the base of the coil is very high and little current is drawn. Close to the resonant frequency the impedance presented to the driver is low and considerable current is drawn. The width and depth of the "impedance dip" depend on the Q factor of the resonant circuit.

Higher Q-factors are desirable because they result in maximum voltage gain at the top of the resonator. They also produces a narrow impedance dip which is very deep. This characteristic offers good loading of the driver provided it can be accurately tuned to the resonant frequency of the coil.

The plot below shows the current drawn (top trace,) and the phase of this current relative to the applied voltage over the same range of frequencies.

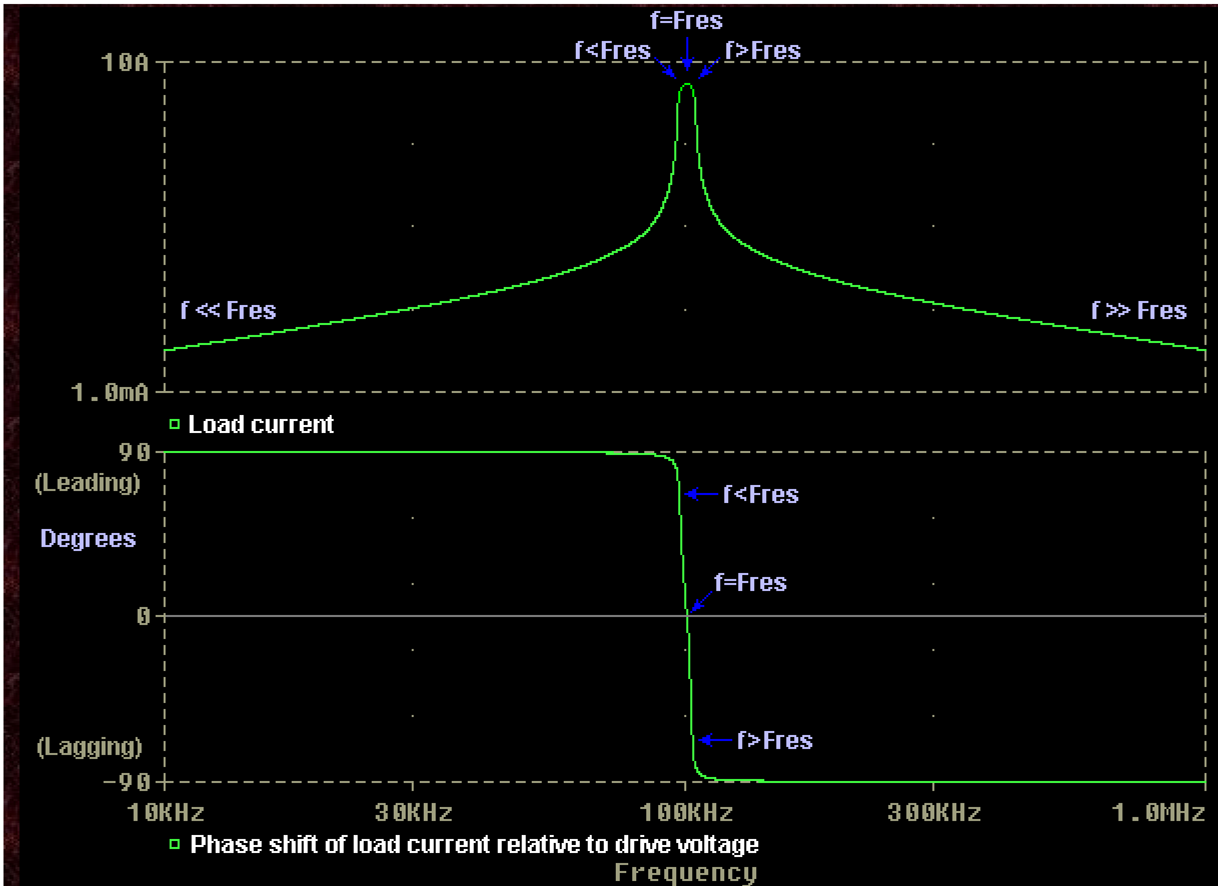

Notice how considerable current flows for a narrow band of frequencies around the resonant frequency of the coil. Also notice that the phase of the current is shifted dramatically for small deviations either side of the resonant frequency. It is this phase shift that is the cause of many headaches in solid state coiling.

Higher resonant modes

At this point it should be mentioned that a practical Tesla resonator exhibits other resonant modes in addition to the "quarter wave" or fundamental mode that we have discussed so far. In this respect it behaves much like a transmission line or guitar string. The 1/4 wave mode that we have investigated represents the lowest frequency peak in the response of the resonator, and results in a single voltage maxima at the end of the resonator. Driving the resonator at this fundamental frequency is most useful to us since we desire maximum voltage rise at the top of the resonator only. We can also conveniently model it as a series RLC circuit.

When driven at approximately **twice** this frequency the Tesla resonator develops a voltage maxima at the centre of its length and a minima at the far end. This is known as the "half wave" mode. It is not of great use to us since maximum voltage is developed part way along the length of the coil.

When driven at higher frequencies still, the resonator displays a number of smaller voltage maxima and minima distributed along its length. These resonant peaks occur at approximate multiples of the fundamental frequency. However, we should resist the temptation to call them harmonics. They are not true harmonics since they do not appear at exact multiples of the fundamental frequency.

These higher resonant modes are of little interest to us since we drive our coil at the fundamental frequency, but a plot is shown below for completeness. The graph shows how the base impedance of a resonator varies over a wide frequency range. Therefore the dips in the graph represent the frequencies at which the impedance is minimum. Maximum current will be drawn from the driver at these resonant frequencies.

(The data for this plot was kindly provided by Paul Nicholson.)

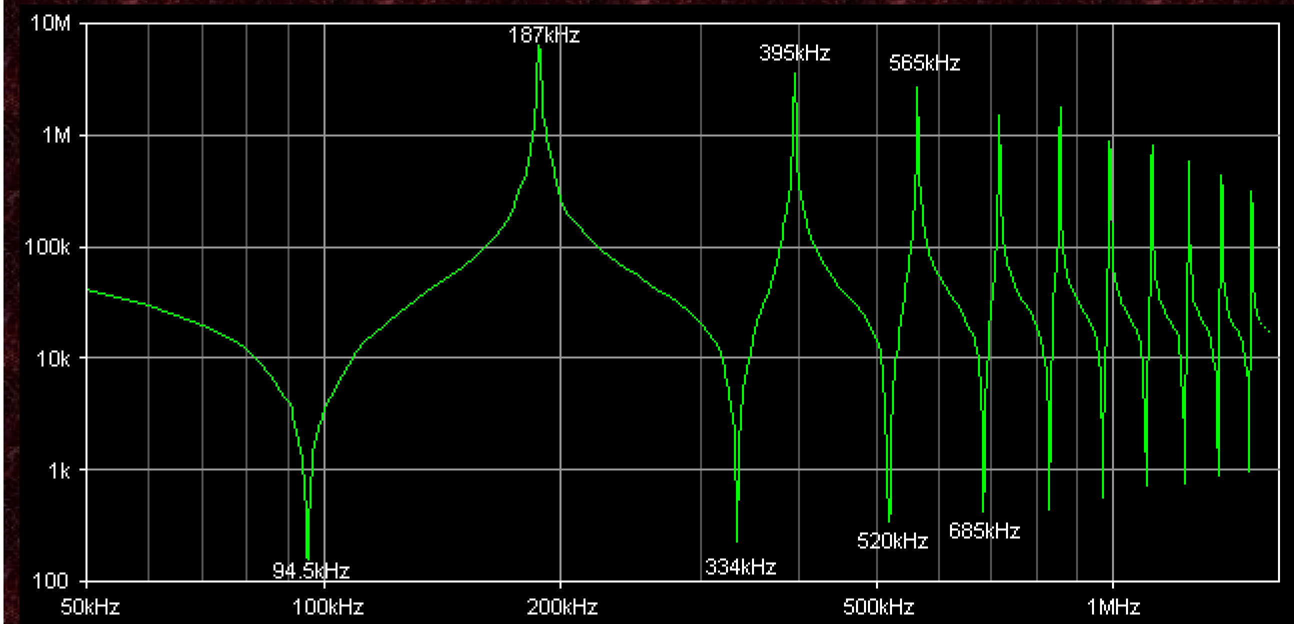

Note: Although we often drive our resonators from a square wave voltage source which is rich in odd harmonics, the current drawn by the resonator is sinusoidal as it responds to the fundamental frequency mostly. The harmonics of the squarewave do not appear to excite higher resonant modes present in the resonator. This is thought to be due to the fact that the higher modes of the resonator do not occur at exact multiples of the driver frequency.

Consider the data in the graph above. The quarter wave resonant frequency of the Tesla coil is 94.5kHz. If this is driven from a squarewave voltage source at 94.5kHz then the most dominant harmonic of the drive waveform is the 3rd harmonic at 283.5kHz. This is some distance away from the 3/4 wave resonant mode of the Tesla coil at 334kHz. Therefore little current is drawn from the driver at its 3rd harmonic frequency. The next most dominant harmonic of the drive signal is the 5th harmonic at 472.5kHz. However, this is even further from the coil's 5/4 resonant mode at 520kHz.

For a squarewave drive, only odd harmonics are present, and the amplitudes of these harmonics diminish in proporion to 1/f. Therefore it is unlikely that any high order harmonic that happens to coincide with a high order resonant mode would result in significant current flow at that frequency. For this reason we will only concern ourselves with the 1/4 wave mode for most of the following discussions. It will be modelled using a simple series resonant RLC circuit.

Now we have a good understanding of the behaviour of the Tesla coil as a load when driven at frequencies around its natural resonant frequency. We will go on to look at the implications that this has on the drive electronics, and then how to design a driver that satisfies these criteria.

Click here for the next section on **DRIVER CIRCUIT BEHAVIOUR.**

Back to home page

SOLID STATE TESLA COIL

DRIVER CIRCUIT THEORY
(Page 2)

DRIVER CIRCUIT BEHAVIOUR

This is the detail of how to design a good Solid State coil driver. It assumes that you have read the previous section about the load presented by a Tesla coil to the driver.

There are many possible configurations for the solid state driver, but the following sections describe the behaviour of a simple half-bridge switching arrangement feeding the base of a Tesla resonator. In later sections other drive methods and more complex topologies are discussed.

We will start with the simplest analysis using "ideal" components. The circuit is then developed to investigate the effects of real-world components and behaviour under adverse conditions.

Effects of resonant load on the driver at exactly Fres

The circuit opposite shows a half-bridge arrangement consisting of two switches, driving the base of a Tesla resonator at precisely its natural resonant frequency.

The Tesla coil is modelled as a series LCR circuit as explained in the previous sections. Each switch is closed in turn for half of the total switching cycle.

The mid-point of the half-bridge develops a square voltage waveform which is fed to the base of the resonator.

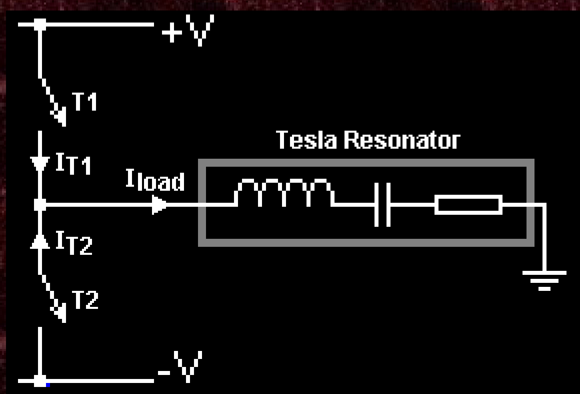

In practice solid state devices such as MOSFETs or IGBTs are used for the switches T1 and T2 to achieve the required switching speed. The arrow on each switch in the circuit diagram serves as a reminder that these devices usually only control conduction in one direction.

The operation is described step by step below:

1. T2 is turned off, T1 is turned on.

2. Sinusoidal current flows down through T1, and flows from left to right through the load.

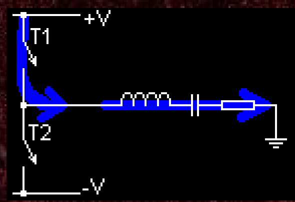

3. Sinusoidal load current passes through zero.	
4. T1 is turned off, T2 is turned on. 5. Sinusoidal current flows down through T2, and flows from right to left through the load.	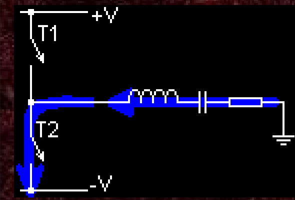
6. Sinusoidal load current passes through zero.	
7. END OF SWITCHING CYCLE. (Go back to step 1.)	

On a timing diagram the waveforms look like this:

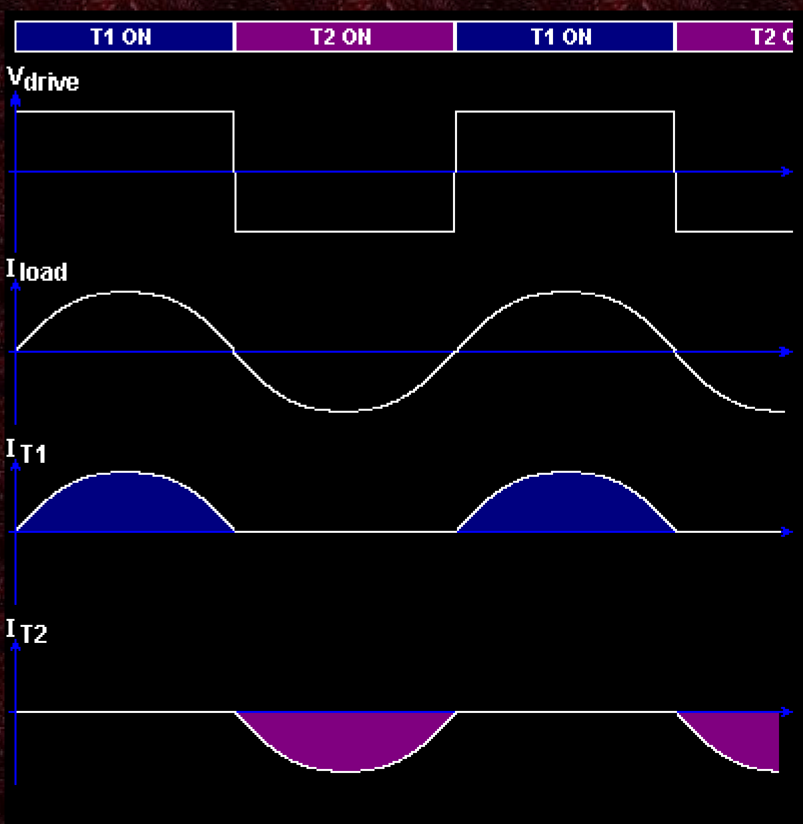

The coloured blocks at the top indicate when each switch is closed and can conduct.

The V*drive* plot shows the voltage developed at the mid-point of the half-bridge which is used to feed the resonator.

The I*load* plot shows the sinusoidal load current drawn by the tesla resonator. This is in-phase with the applied voltage because we are driving the coil at its resonant frequency.

The I*T1* plot shows that the positive half-cycles of the load current are carried by switch T1 during its on-time.

Likewise, the I*T2* plot shows the negative half-cycles of the load current being conducted by switch T2 during its on-time.

Although this analysis is simplified slightly, it is very close to the real behaviour of a practical driver circuit, and a few things are worth noting from the results. Notice how the current flowing through the load (and through the switches) passes through zero at the instant when the switches change state. This only happens when the **driver is perfectly in tune with the resonator**, and has a number of advantages:

i. **Greatly reduced switching losses.**
In other inverter applications there is usually significant current flow through the switches that must be interrupted at the switching instants. Therefore switching losses form a large proportion of the total losses in conventional inverter designs. However, when the load is resonant at the drive frequency there is no current flowing at the switching instant so no power is dissipated during the switching transitions. This reduces heating in the switches.
.

ii. **Greatly reduced spikes and ringing.**

There is no abrupt switching of large currents, so the maximum rate of change of current (di/dt) is simply the slope of the sinusoidal load current.
This can be contrasted with hard switched applications where currents of many amps are switched in tens of nanoseconds. (di/dt effects on stray circuit inductances are reduced markedly due to this Zero Current Switching behaviour.)
.

iii. **Improved load sharing of paralleled switches.**
When several switches are paralleled for greater current handling, the load current must be shared evenly between devices. The temperature coefficient of MOSFETs assists static balancing of the load current. However dynamic load sharing can be problematic at high switching frequencies, due to mismatches in the switching speeds of the paralleled devices. In this particular case the current through the switches is near zero for some time around the switching instants, so matching of switching times is less critical. The load current increases slowly at the start of each conduction period so the first devices to turn-on do not get hit with the full load current immediately. Similarly the last devices to turn-off only have to handle the tail of the current as it falls to zero.
.

iv. **Reduced avalanche stresses for series connected switches.**
When switches are series connected for greater voltage handling, avalanching occurs in the last devices to turn on, and the first devices to turn off. Under the conditions being investigated here little current flows through the switches at the switching instant. Therefore the avalanche energy of series connected devices is reduced dramatically. (Although manufacturers now specify avalanche ratings, this mode of operating is stressful to the devices. Series connection of silicon switches can be troublesome and minimising avalanche energy will improve reliability.)
.

Now that we have got excited about how well this simple half-bridge driver would work in an ideal world with ideal components and with perfect tuning, it is time to take a look at what can go wrong in practice !

The requirement for "free-wheel" diodes

Unfortunately real life is never so simple. Two factors act to complicate issues slightly:

1. It is not possible to guarantee that the driver will always be in perfect tune with the Tesla resonator at all times, and under all conditions.

2. In reality it is not possible to open one switch and close the opposing switch at **exactly the same instant in time**.

The second point is a serious issue since simultaneous conduction of both switches T1 and T2 shorts the supply. This leads to a very high current pulse down the leg of the half-bridge and is very damaging to silicon switches. Such a condition is known as shoot-through (or cross-conduction) and must be eliminated. The shoot-through problem is easily overcome by including a short delay between opening one switch and closing the opposing switch. This delay is called "dead-time" since no switches are conducting, and ensures that both switches cannot be closed simultaneously.

The introduction of a dead-time where no switches are conducting, necessitates an enhancement to the circuit so that the inductive load current can still be supported while both switches are open. This function is provided by a pair of "Free-wheel" diodes. They are connected across each switch and provide alternative conduction paths for the load current to flow in either direction during the dead-time.

(If no path is provided for the free-wheeling load current when both switches are open, then a large voltage pulse is developed. This is often referred to as "back-EMF" and is encountered with many inductive loads such as motors and solenoids, when the current is interrupted abruptly.) The action of the free-wheel diodes can be thought of as clamping the back EMF to the supply rails when each switch opens.

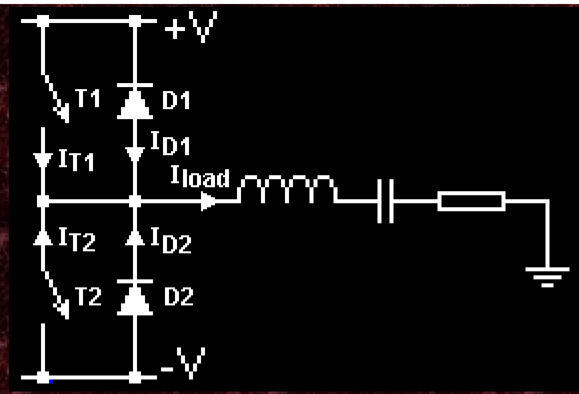

The introduction of the two free-wheel diodes means that there are now 4 possible states in the switching cycle. Not only can the load current be carried by switch T1 or T2, but it can also be carried by diode D1 or D2. Now things get a little more complicated.

The way in which the load current is transferred from one device to another depends on the tuning of the driver with respect to the resonator…

Effects of being "out of tune" on the driver

In earlier sections covering the behaviour of the Tesla coil as a load, we saw that the load current does not necessarily change direction at the same time as the applied voltage. In fact it was found to lead or lag behind the applied voltage depending on the tuning of the driver. This determines the direction of the load current at the instant of the dead-time:

- If the driver is tuned to a frequency **above the resonant frequency** of the Tesla coil, then the load changes direction **after the dead time**.
  In this case the free-wheel diodes sustain the load current during the dead time.
  .
- If the driver is accurately tuned **to the resonant frequency** of the Tesla coil, then the load current will be **near zero during the deadtime**.
  In this case little current flows through the free-wheel diodes.
  .
- If the driver is tuned to a frequency **below the resonant frequency** of the Tesla coil, the load current changes direction **before the dead time**.
  In this case, it is the free-wheel diodes which allow the load current to change direction before the dead-time. (Remember that semiconductor switches can only control conduction in one direction.)
  .

This is of interest to us as it determines how much current flows in each component, and what happens when the load current is passed from one component to another. Most of the stresses that a solid-state driver experiences in this application are a result of large currents being commutated between switches and free-wheel diodes. This can cause over-voltage transients or current surges. Both are undesirable and lead to device failures.

We will now look at what happens in our half-bridge when it is not tuned perfectly to the resonant frequency of the Tesla coil. Let us start by looking at the case which is easiest to analyse…

More to come shortly… This page is currently under development…

Effects of resonant load on the driver just above Fres

Steps in the switching cycle,
Voltage and current waveforms,
Free wheel current."
Stray inductance.
ZVS transition.
Forward recovery problems.
Voltage overshoot at device turn-off.

Effects of resonant load on the driver just below Fres

Steps in the switching cycle,
Voltage and current waveforms,
Free-wheel diode behaviour.
Reverse recovery problems.
Isolating the MOSFET body diodes.
Current overshoot at device turn-on.
Effects on current monitoring.

Beware the MOSFET body-diode !

Those familiar with MOSFETs will know that the fabrication process results in a built-in anti-parallel diode between the source and drain terminals of the device. This is often referred to as the "body-diode." Referring to any MOSFET data sheet will reveal specifications for this intrinsic diode.

At first it appears that the internal body-diodes are a bonus since they provide the desired free-wheel diode function for free. This is often the case in many power electronics applications where they provide the function of the free-wheel diodes with ease. Sadly it is not the case in this application. The MOSFET body-diode is a side effect of the fabrication process and is not a particularly good diode. The same design criteria for good MOSFET characteristics do not produce the best body-diode characteristics. The design of a MOSFET is always a compromise, and it is the characteristics of the body-diode that suffer.

When compared to discrete high speed diodes, the body-diode's reverse recovery time is very long. This means that the diode takes a long time to turn off when the current flowing through it changes direction. As explained previously, this leads to a shoot-through condition when the opposing switch is turned-on. For this reason the body-diodes are clearly not suitable for free-wheel diode duty in this application and should be isolated.

The body-diode is isolated by means of a Schottky Barrier diode connected in series

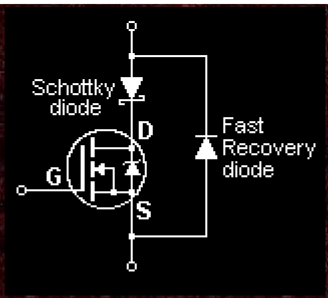

with the MOSFET drain lead. (Schottky diodes operate due to majority carrier conduction, and therefore do not exhibit any significant reverse recovery time. Essentially, they turn off immediately when the current tries to change direction.) This prevents current from flowing through the MOSFET body-diode and forward biasing it. If the body-diode is never forward biased it does not exhibit a reverse recovery problem.

An external fast recovery diode is then connected across the pair to provide the necessary path for the free-wheel current.

The use of external free-wheel diodes gives the designer greater choice in the characteristics of this critical component. It also removes a source of dissipation from the MOSFETs, since the free-wheel current no longer enters onto the MOSFET die.

Although dedicated fast-recovery diodes are much faster than the MOSFET body-diodes, they still have a finite reverse recovery time. This is typically several tens of nanoseconds. Any remaining reverse-recovery problems can usually be solved by slowing down the turn-on of the switches. This allows longer for the free-wheel diodes to recover, and reduces the peak reverse recovery current. (The recovery is also softer and results in less radiated interference too!)

Beware excessive dead-time !

The amount of dead-time between opening one switch and closing the opposing switch should be just sufficient to prevent simultaneous conduction of both switches, and no longer.

When the driver is accurately in tune with the Tesla resonator there is little load current flowing at the switching instants. Therefore little current flows in the free-wheel diodes during the dead-time, and the conduction time is short.

(See red shaded areas on plot opposite.)

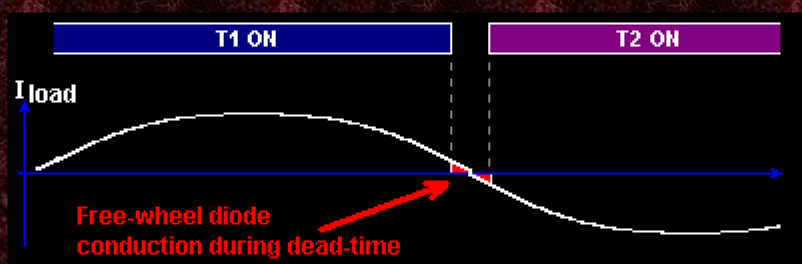

If the dead-time is increased beyond this minimal value, the load current is forced to flow through the free-wheel diodes for longer. It should also be noted that the load current increases either side of the zero crossing, so excessive dead-time also makes the diodes see a higher peak current.

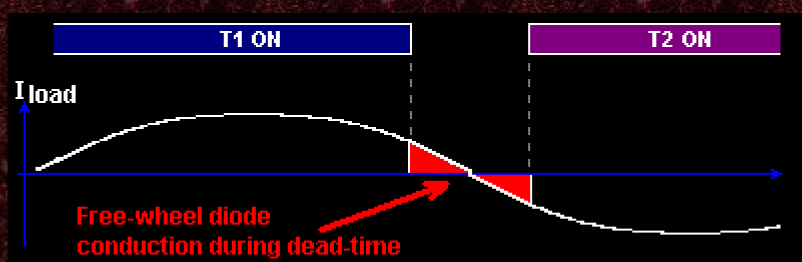

Higher currents conducted for longer periods lead to higher dissipation in the free-wheel diodes. However, more importantly, the harder the free-wheel diodes are driven into forward conduction, the harder they are to turn off. (The reverse recovery time becomes longer.)

For this reason the driver should **be accurately tuned to minimise the current flow during the deadtime**, and the **dead-time should be kept short to minimise diode conduction time.**

Snubbers

Over-voltage snubbing,
Over-current snubbing,
Damping of parasitic resonances,
Snubber energy,
Snubber effectiveness,
Lossless snubbing and ZVS transition,
Diode clamping loop,

Resonator impedance

Components of resistance.
DC resistance.
Skin effect.
Ground plane losses.
Radiation resistance.
Proximity effect.
Impedance changes due to spark loading.
Reflected load theory.
Impedance inversion.
The constant current sink model.
Resonator detuning.
Frequency tracking (dynamic tuning).

Drive methods

Base feed method.

This is the drive method that has been discussed in all of the previous sections and is by-far the simplest arrangement. The base of the resonator is connected directly to the output of the inverter with a single wire. The return current is capacitively coupled to the surroundings and ultimately back to ground. The base-fed resonator is modelled as a series resonant LCR circuit, and its behaviour has already been described in considerable detail.

In practice the base feed method does not give good results unless a relatively high drive voltage is used. This is because the base impedance at resonance can still be tens or hundreds of ohms, and rises further during spark breakout. Therefore the resonator does not draw much current unless a voltage of several hundreds or thousands of volts is used.

The following list summarises some of the properties of the base-feed approach:

Advantages:	**Disadvantages:**
Very simple to analyse using circuit theory,	Requires high drive voltage for good sparks,
Simple to set up. (No adjustments to be made.)	Not very flexible in terms of impedance matching,
No primary winding in close proximity to the resonator. This allows the coil to resonate	Resonator is directly connected to the HV inverter, possibly allowing 50Hz or DC

freely, and also reduces flashover problems.	components into the discharge.

(The author knows of one individual who has obtained excellent spark performance from a base-fed solid state Tesla coil. However, his design uses a supply voltage of 1500 volts, in order to force many kilowatts of power into the resonator. This approach appears relatively simple on the surface, although the construction of a robust, high power, high frequency inverter to operate from a 1.5kv supply is not for the beginner !)

For this reason one of the methods described later can be used to provide a better impedance match between the low impedance of the driver and the medium-to-high impedance of the Tesla resonator. This allows the driver to operate from a lower supply voltage such as a battery, or directly from the AC mains.

Primary coil drive method.

The Tesla coil can be driven by a primary coil that is coupled to the resonator. Energy is then transferred from the primary coil to the resonator (secondary coil) by means of the magnetic coupling between them. When this drive method is employed, we also benefit from some "step-up" of the drive voltage due to transformer action between the two coils.

We will see that this drive approach is very flexible. It has a number of advantages over the base-feed method, although it is slightly more complex to analyse…

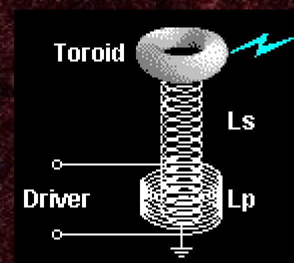

When the Tesla resonator is driven via a primary coil we will see that the resonant nature of the Tesla coil is coupled back to the primary winding. We will start with the model opposite that shows a primary coil Lp, coupled to a secondary coil Ls. The resonator is represented by the usual series resonant LCR circuit, and consists of Ls, Cs, and Rs. The degree of magnetic coupling between the two coils is represented by the factor k.

There are clearly several variables in this arrangement. However, for a given resonator, Ls, Cs and Rs are fixed. Therefore the designer is left with the choice of Lp and k.

As k is increased the primary coil "sees" more of the secondary winding. We can imagine the secondary winding Ls being divided into two parts called Ls1 and Ls2. [***Ls1+Ls2 = Ls***]
Ls1 is the portion of the secondary that is coupled to the primary coil. Ls2 is the remaining portion of the secondary inductance that is not coupled to Lp and is free to resonate with Cs.

(Ls2 is actually the leakage inductance of the original loosely coupled two coil system.)

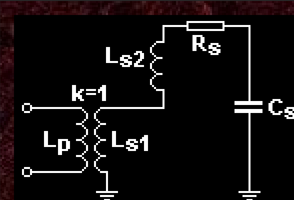

Ls1 = k² Ls

Ls2 = (1-k²) Ls

Notice the similarity of the above circuit to that of the magnifier arrangement. We now have a closely coupled transformer (consisting of Lp and Ls1,) which is base feeding a free resonator Ls2.

*(The coupled primary approach is electrically equivalent to base feeding a smaller resonator through a step-up transformer. In fact it is possible to do a transformation between these two equivalent drive methods. See the equations for **Ls1** and **Ls2**.)*

The transformation ratio of the transformer is clearly equal to ***sqrt (Lp / Ls1)***, so impedances are transformed by ***(Lp / Ls1)***. Therefore we can "reflect" the resonant load

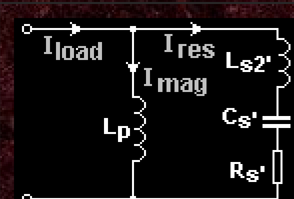

Ls2' = (Lp / k²) - Lp

Cs' = k² Ls Cs / Lp

(Ls2, Rs and Cs) from the secondary side to the primary side of the transformer. (***Ls2', Cs' and Rs'*** are the equivalent component values at the primary side.)	***Rs' = Rs Lp / (k^2 Ls)***

This analysis shows us that the series resonant nature of the Tesla coil is reflected back to the primary due to the magnetic coupling between the coils. The series resonant LCR model of the secondary appears as if it were actually connected in parallel across the primary winding, although its impedances are transformed down due to the turns-ratio between the two coils.

From the last schematic above, we can see that the load current **(Iload)** supplied by the driver is now the sum of two components. These are **Imag** and **Ires**. We will consider the significance of each of these components below:

Imag.

This is the **magnetising current** that flows through the primary inductor regardless of whether it is magnetically coupled to anything else. It can be likened to the off-load current drawn by any transformer and is inversely proportional to both Frequency and the winding Inductance (Lp).

The current waveform is triangular in nature, and ramps up and down during the high and low portions of the drive voltage respectively. It lags the drive voltage by 90 degrees due to the inductance.

The magnetising current does not perform any transfer of real power, and merely represents current sloshing back and forth between the primary winding and the drive circuit.

We would like to minimise **Imag** to reduce its contribution to the total load current, and the best way to do this is to use a large Lp value. I.e. Specify a generous number of primary turns.

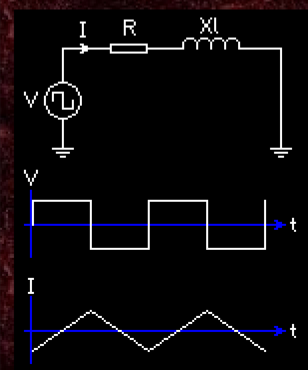

Ires.

This is the current that couples to the secondary and **contributes to the resonant action.** Its magnitude is equal to the secondary base-current multiplied by the transformation ratio ***sqrt (Ls1 / Lp)***.

When driven at the correct resonant frequency this current component is sinusoidal in shape and is in phase with the applied voltage.

We seek to maximise **Ires** because this is the current which pumps the resonator and gives us the desired resonant voltage rise. The best ways to maximise **Ires** are:

- Reduce the effective resistance of the secondary coil as far as possible, (i.e. high Q.)
- Reduce the primary inductance **Lp** to get the maximum step-up ratio.
- Increase k to get maximum coupling and therefore maximise the step-up ratio.

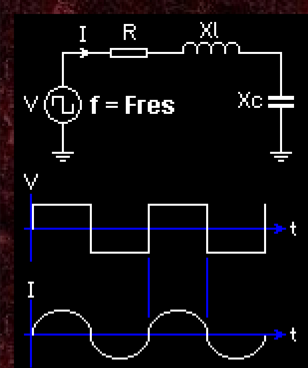

Effect of changing k

As the degree of coupling is increased, two things happen:

1. More of the secondary winding is coupled to the primary so Ls1 increases. This increases the voltage gain due to transformer action. The turns ratio effectively increases because the primary sees a greater number of the secondary turns.)

2. Less of the secondary winding is left free to resonate with Cs. This decrease in the remaining resonant inductance Ls2 causes the resonant frequency to rise as the coupling is increased !

This behaviour is summarised in the drive impedance plot shown below:

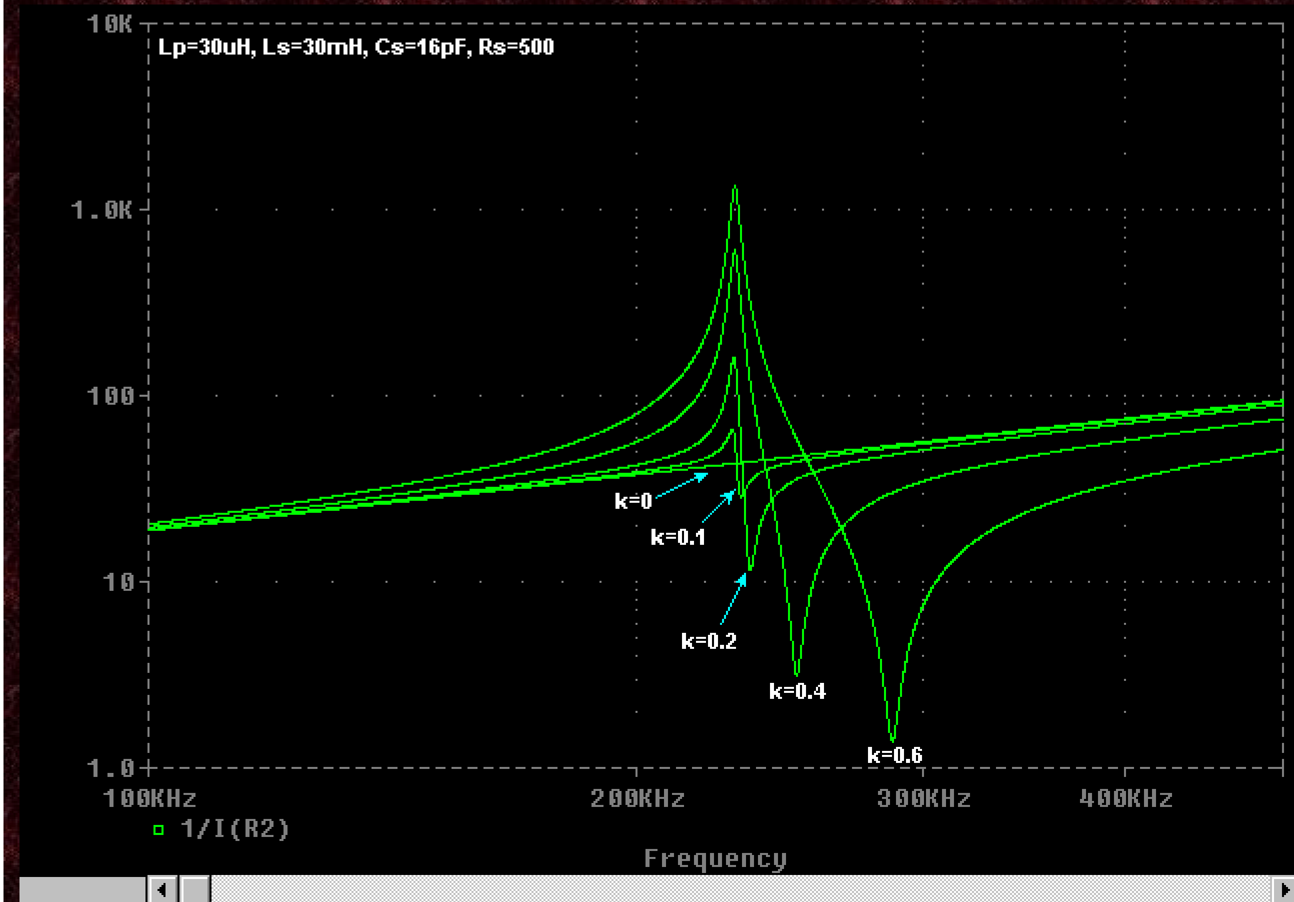

Graph showing typical variation in primary impedance with different degrees of coupling.

Notice how all of the lines lie around the same straight line labelled k=0. This is the impedance of the primary coil on its own, with no coupling to the seconadry. The impedance plot begins to deviate from this straight line as the coupling is increased and the primary begins to "see" more of the resonator.

The points of interest on this graph are the impedance minima. The minimum point for each line represents the frequency at which maximum current is drawn from the driver and the resonant component **Ires** is in phase with the applied voltage. This is also the point which gives the maximum voltage gain at the top of the resonator, and therefore best spark performance.

Notice that the resonant frequency with minimum primary impedance changes as the coupling is varied. This is because more of the resonator is coupled to the primary as k is increased, leaving less of the secondary coil free to resonate.

The resonant part of the secondary was earlier found to be: ***Ls2 = (1 - k²) Ls***
Since frequency is proportional to ***1 / sqrt (L C)*** then it can be shown that the resonant frequency seen at the primary is equal to the free-standing (un-coupled) resonant frequency of the secondary multiplied by the following factor:

1 / sqrt (1 - k²)

If k=0, the resonant frequency is equal to the natural resonant frequency of the resonator, but there is no coupling. We effectively have two separate circuits. As k is increased towards unity, ***Ls2*** decreases to zero, and the resonant frequency seen at the primary heads towards infinity! At k=1 we basically have a conventional transformer with no resonant behaviour.

This shows that it is important to remember to re-tune after adjusting the coupling when using a solid state driver and a primary coil to drive the resonator!

Note: *An impedance maximum point also exists in the frequency response above. This is because the reflected resonator impedance is slightly capacitive at this frequency and it cancels out the inductive reactance of the primary coil. At this frequency a parallel resonant circuit is formed between the primary inductance and the net capacitive reactance reflected back from the secondary. It is of little practical importance to us since it represents a high impedance to the driver. Therefore it draws negligible current from the driver. It is mentioned here for completeness.*

Effect of changing primary inductance Lp

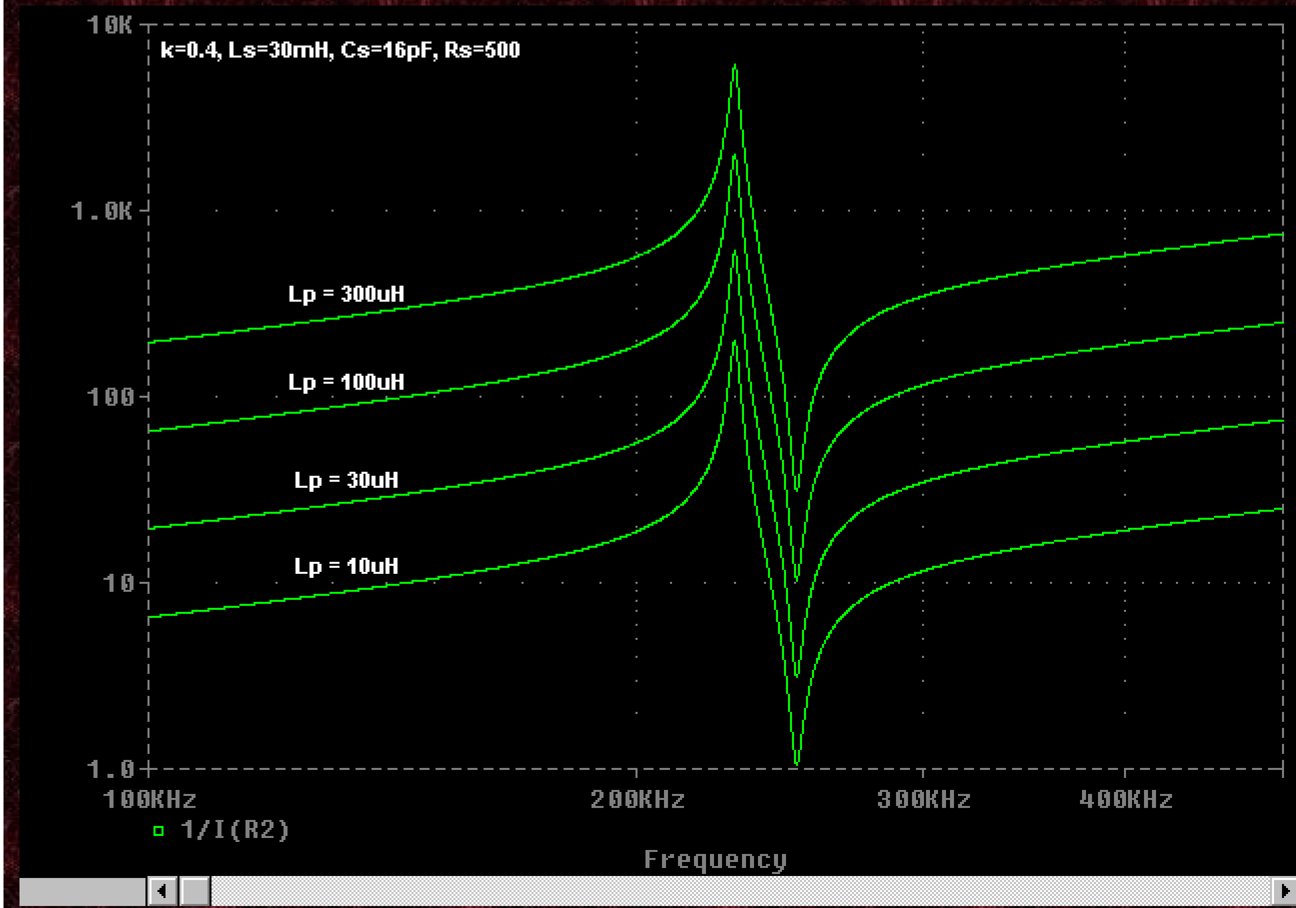

Graph showing typical variation in primary impedance with different numbers of primary turns.

Notice how increasing or decreasing the primary inductance (Lp) effects the whole primary impedance in the same way. Both the magnetising current and the resonator current components change as a direction function of Lp. However, the positions of the resonant peaks are not changed provided the coupling remains constant. Altering Lp provides an easy way to increase or decrease the total load current without changing the resonant frequency.

For example, halving the number of primary turns reduces the primary inductance to approximately one quarter of its previous value. In an ideal world the impedance presented to the driver would reduce to one quarter of its previous value and there would be a corresponding increase in the current drawn from the driver.

Note: *In practice the impedance does not fall quite as far as one quarter of its previous value. Therefore the current does not quite quadruple. We will see later that this is due to a strange "negative resistance" property of the spark discharge from the top of the coil, when we attempt to drive more power into it!*

Summary on effects of Lp and k.

The choice of Lp and k are very important. In fact I think that one of the main challenges in solid state Tesla coil work is in juggling these two parameters to get maximum power transfer into the resonator.

If Lp is too small, the magnetising current becomes unacceptably high. Although this does not contribute to the supply current of the driver, it does contribute to heating in the switches. It also increases the current ripple seen by the supply reservoir capacitors.

If Lp is too great, then there is little step-up due to transformer action. As a result, the Tesla resonator is not well matched to the driver, and little power is drawn from the driver.

If k is too small, then the load current is dominated by the magnetising current, and little of the current supplied by the driver actually contributes to resonance in the secondary winding.

If k is too great, then there is a risk of electrical flashover. This is due to the close proximity demanded to achieve this degree of pri-sec coupling, and the large sparks that result !

It is my opinion that Lp should be made only sufficiently large to obtain an acceptable magnetising current, then k should be maximised as far as possible to get the minimum possible load impedance at resonance. This is presently my recipe for effective power transfer using this particular drive method. Good power transfer can only be achieved with reduced coupling if the drive voltage is increased. Consider that there is less "step-up" due to transformer action as the coupling is reduced. Therefore you have to drive the primary with a correspondingly higher voltage to make up for this, and acheive the same performance.

The author's present solid state Tesla coil system makes use of a 18 turn 20uH primary inductor driven with +/- 340v at 350kHz. This is tightly coupled to the secondary winding with a coupling coefficient of approximately 0.55. This provides good power throughput, with an acceptably low magnetising current. However, such a high coupling coefficient demands close positioning of the primary and secondary, and tends to promote flashovers. The distance between the primary and secondary windings is less than half of an inch, and sparks of 18" are common from the toroid, so insulating between the two coils is a challange.

One may think that such a high coupling coefficient produces most of the voltage gain by transformer action, and little voltage rise due to resonance in the un-coupled portion of the secondary. This is actually true, but we must realise that spark breakout loads the secondary winding ruining the Q, so resonant rise is not good after spark breakout. The base impedance of a physically small resonator seems relatively high to start with, and this increases further following spark breakout, therefore a high coupling coefficient is required to obtain maximum transformer action, and achieve good matching under sparking conditions.

More to come shortly… This page is currently under development…

Back to home page

LINKS TO OTHER SITES

Information about UK events and much more.	 Locate other Tesla coil enthusiasts in UK.	Pupman Tesla Coil Mailing List. (USA)	Terry's Tesla Coil web server. (USA)

UK Quick links... US Quick links...

Safety Documents:

Pupman safety sheet	Thorough safety checklist with lots of useful tips. Check it out, you only live once !
Terry Fritz's Safety page	Highlights many of the dangers of working with Tesla Coils.
Bert Hickman's safety page	Serious safety tips, backed up with some funny pictures ;-)
Herb's Tesla safety page	Short summary of the dangers associated with Tesla Coils.

UK Sources:

RS Components	Electronic equipment, components, materials, and tools.
Tunewell Transformers	Manufacturer of high quality Neon Sign Transformers.

HV Capacitors:

Norfolk Capacitors	UK Manufacturer of HV pulse capacitors.
Hivolt Capacitors	UK Manufacturer of various high voltage capacitors.
Cambridge Capacitors	UK Manufacturer of Metalised Polypropylene PFC capacitors.
Arcotronics	Manufacturer of various capacitors including Motor Run and PFC.
Aerovox	High quality electrolytic and pulse capacitors.
Wima	More capacitors, data sheets and applications notes.
Murata	High power RF capacitors, HV capacitors and flyback transformers.

Solid State TC:

Alan Sharp's home page	Excellent solid state coiling pages. Includes many design and construction tips.
James Pawson's TC page	Sound techincal information on a variety of SSTC designs and topics.
Gary Johnson's Tesla Coil page	8 Chapters documenting one of the largest SSTCs around. Essential reading.
Jan Wagner's Tesla Coil page	Excellent Solid State coil section. Be sure to check out the "General design notes."
www.hvguy.com	Well documented SSTC experiments with PCBs artworks, schematics and videos!

Dan McCauley's SSTCs	Pictures and schematics for several Solid State Tesla coil projects.
John Tomacic's Tesla Coil	Pictures and some documentation of a fairly large base fed SSTC.
Don Klipstein's SSTC page	More solid state stuff. Discusses many different ways of driving a TC resonator.
Marco Denicolai's page	In-depth documentation for a high-power solid state Tesla Coil supply.
Chris Hill's SMPS page	This particular page contains a good summary of switch mode inverter topologies.

Power Devices:

International Rectifier	High power MOSFET and IGBT devices, datasheets and application notes.
Intersil (formerly Harris)	More power semiconductors. Good application notes for IGBTs.
ON Semiconductor (Motorola)	More power semiconductors, data sheets and app notes.
Semikron	High power MOSFET and IGBT packs. Also driver modules.
Westcode	High current Thyristors and Diodes.

Materials:

Tufnol Plastics	Phenolic laminate range. Includes material properties and machining tips.
Tungsten Properties	More than just a hard grey metal ?
Web Elements	Click any element to see its properties. (Conductivity, weight, melting pt, etc.)

Engineering Info:

Tesla Coil Design Equations	A list of Tesla Coil related Formulae compiled by Matt Behrend.
Electronic Eng training series	An excellent reference for basic Electrical Engineering information.
Twisted Pair	Electronics abbreviations, terms and laws explained !
Pulse switching devices	John Pasley's paper on Pulse Power Switching Devices.

Search Engines:

Dogpile	Sends your request to several search engines at once. Unleash the hound !
Applegate Engineering Directory	Comprehensive directory of UK Engineering companies.

Back to home page

Why MOSFETs fail in Solid State TC duty

After emptying many pounds worth of dead MOSFETs from the trash can in my workshop, I decided that it might be worth actually compiling a list of reasons why MOSFET devices might fail in solid state tesla coil applications.

There are quite a few possible causes for device failures, ranging from obvious reasons to some mechanisms which are not so immediately obvious. When investigating a fault, this list can be used as a check-list:

- **Over-voltage,**

  MOSFETs have very little tolerance to overvoltage. Damage to devices may result even if the voltage rating is exceeded for as little as a few nanoseconds. MOSFET devices should be rated conservatively for the anticipated voltage levels, and careful attention should be paid to suppressing any voltage spikes or ringing.

- **Prolonged current overload,**

  High average current causes considerable thermal dissipation in MOSFET devices due to the relatively high on-resistance. If the current is very high and heatsinking is poor, then the device can be destroyed by excessive temperature rise. MOSFET devices can be paralleled directly to share high load currents.

- **Transient current overload,**

  Short duration, massive current overload can cause progressive damage to the device with little noticeable temperature rise prior to failure. (Also see shoot-through and reverse recovery sections below.)

- **Shoot-through, (cross conduction.)**

  If the control signals to two opposing MOSFETs overlap, then a situation can occur where both MOSFETs are switched on together. This effectively short-circuits the supply and is known as a shoot-through condition. If this occurs, the supply decoupling capacitor is discharged rapidly through both devices every time a switching transition occurs ! This results in very short but incredibly intense current pulses through both switching devices.

  The chances of shoot-through occuring are minimised by allowing a dead time between switching transitions, during which neither MOSFET is turned on. This allows time for one device to turn off before the opposite device is turned on.

- **No free-wheel current path,**

  When switching current through any inductive load (such as a Tesla Coil,) a back EMF is produced when the current is turned off. It is essential to provide a path for this current to free-wheel in the time when neither switching device is carrying the load current.

  This current is usually directed safely back to the supply rails by means of a free-wheel diode connected anti-parallel with each switching device. When MOSFETs are employed as the switching devices, the designer gets the free-wheel diode "for free" in the form of the MOSFET's intrinsic body diode. This solves one problem, but creates a whole new one...

- **Slow reverse recovery of MOSFET body diode**

  A high Q resonant circuit such as a Tesla Coil is capable of storing considerable energy in its inductance and self capacitance. Under certain tuning conditions, this causes the current to "free-wheel" through the internal body diodes of the MOSFET devices as one MOSFET turns off and the other device turns on. This behaviour is not a problem in itself, but a problem arrises due to the slow turn-off (or reverse recovery) of the internal body diode when the opposing MOSFET tries to turn on.

  MOSFET body diodes generally have a long reverse recovery time compared to the performance of the MOSFET itself. If the body diode of one MOSFET is conducting when the opposing device is switched on, then a "short circuit" occurs similar to the shoot-through condition described above.

  This problem is usually eased by the addition of two diodes surrounding each MOSFET. Firstly, a Schottky diode is connected in series with the MOSFET source. The schottky diode prevents the MOSFET body diode from ever being forward biased by the free-wheeling current. Secondly, a high speed (fast recovery) diode is connected in parallel to the MOSFET/Schottky pair so that the free-wheeling current bypasses the MOSFET and Schottky completely.

  This ensures that the MOSFET body diode is never driven into conduction. The free-wheel current is handled by the fast recovery diodes which present less of a "shoot-through" problem.

- **Excessive gate drive,**

  If the MOSFET gate is driven with too high a voltage, then the gate oxide insulation can be punctured rendering the device useless. Gate-source voltages in excess of +/- 15 volts are likely to cause damage to the gate insulation and lead to failure. Care should be taken to ensure that the gate drive signal is free from any narrow voltage spikes that could exceed the maximum allowable gate voltage.

- **Insufficient gate drive, (incomplete turn on)**

  MOSFET devices are only capable of switching large amounts of power because they are designed to dissipate minimal power when they are turned on. It is the responsibility of the designer to ensure that the MOSFET devices are turned hard on to minimise dissipation during conduction. If the device is not fully turned on then the device will have a high resistance during conduction and will dissipate considerable power as heat. A gate voltage of between 10 and 15 volts ensures full turn-on with most MOSFET devices.

- **Slow switching transitions,**

  Little energy is dissipated during the steady on and off states, but considerable energy is dissipated during the times of a transition. Therefore it is desirable to switch between states as quickly as possible to minimise power dissipation during switching. Since the MOSFET gate appears capacitive, it requires considerable current pulses in order to charge and discharge the gate in a few tens of nano-seconds. Peak gate currents can be as high as an amp.

- **Spurious oscillation,**

  MOSFETs are capable of switching large ammounts of current in incredibly short times. Their inputs are also relatively high impedance, which can lead to stability problems. Under certain conditions high voltage MOSFET devices can oscillate at very high frequencies due to stray inductance and capacitance in the surrounding circuit. (Frequencies usually in the low MHz.) This behaviour is highly undesirable since it occurs due to linear operation, and represents a high dissipation condition.

  Spurious oscillation can be prevented by minimising stray inductance and capacitance around the MOSFETs. A low impedance gate-drive circuit should also be used to prevent stray signals from coupling to the gate of the device.

- **The "Miller" effect,**

  MOSFET devices have considerable "Miller capacitance" between their gate and drain terminals. In low voltage or slow switching applications this gate-drain capacitance is rarely a concern, however it can cause problems when high voltages are switched quickly.

  A potential problem occurs when the drain voltage of the bottom device rises very quickly due to turn on of the top MOSFET. This high rate of rise of voltage couples capacitively to the gate of the MOSFET via the Miller capacitance. This can cause the gate voltage of the bottom MOSFET to rise resulting in turn on of this device as well ! A shoot-through condition exists and MOSFET failure is certain if not immediate.

  The Miller effect can be minimised by using a low impedance gate drive which clamps the gate voltage to 0 volts when in the off state. This reduces the effect of any spikes coupled from the drain. Further protection can be gained by applying a negative voltage to the gate during the off state. Eg. Applying -10 volts to the gate would require over 12 volts of noise in order to risk turning on a MOSFET that is meant to be turned off !

- **Radiated interference with controller,**

  Imagine the effect of connecting just 1pF of capacitance from the top of your sparking tesla coil to each of the sensitive points in your solid state controller. The hundreds of kilovolts of RF present would have no problem driving significant current through the tiny capacitors directly into the control circuit.

  Well this is exactly what happens in practice if the controller is not placed in a screened enclosure !

  It takes little stray capacitance to high impedance points of the control circuitry to cause abnormal operation. Bare in mind that a controller which is not operating correctly could attempt to turn on two opposing MOSFET devices at the same time. Effective RF screening of the control electronics is essential.

  It is also highly desirable to segregate power and control circuitry. Rapidly changing currents and voltages present in the power switching circuit still have the ability to radiate significant interference.

- **Conducted interference with controller,**

  Rapid switching of large currents can cause voltage dips and transient spikes on the power supply rails. If one or more supply rails are common to the power and control electronics, then interference can be conducted to the control circuitry.

Good decoupling, and star-point earthing are techniques which should be employed to reduce the effects of conducted interference. The author has also found transformer coupling to drive the MOSFETs very effective at preventing electrical noise from being conducted back to the controller.

- **Static electricity damage,**

  Antistatic handling precautions should be used to prevent gate oxide damage when installing MOSFET or IGBT devices.

- **High VSWR,**

  (I am not an RF Engineer, so I don't fully understand this one. However, Jim Lux kindly offered this excellent explanation:

  In a pulsed system, VSWR isn't as big an issue as in a CW system, although it's still an issue.

  In a CW system (or, for that matter, something where you put RF power out for many cycles of the RF, so almost any form of transmitter is in this class).

  The typical transmitter is designed for a 50 ohm resistive output impedance (at least, the designer dreams of this), then is connected via some sort of transmission line to a load. Hopefully, the load and the line are also 50 ohms, and power flows nicely down the wire. However, if the load impedance is not 50 ohms, then some amount of the power is reflected back from the impedance discontinuity. The reflected power causes several potential problems:

  1) the transmitter looks like a load and absorbs it all... If you have a high efficiency transmitter (say, running class C) the power devices (tubes, transistors, or whatever) are dissipating a small fraction of the output power. If your amp were, say, 80% efficient, and you're putting in a kilowatt, normally, the devices dissipate about 200 W, and 800W goes down the line. IF all that 800W gets reflected back, now all of a sudden, your devices are dissipating the full kW.

  *2) The combination of forward and reflected waves causes standing waves in the transmission line, where the voltage can get quite high at points 1/2 wavelength apart, resulting in breakdown or other bad things. This is essentially the result of the apparent load impedance (at the transmitter) not being what is expected. The transmission line effectively transforms the load impedance to some other value (1/4 waves are neat because Zin*Zout = Zline*Zline... see what happens if Zout goes to zero). If your transmitter is a constant current source, and the load impedance goes higher than expected, then the voltage gets higher than expected, etc. (by the way, for grins, if you have a "tough" RF power source at a few 10's of MHz, you can rig up an open parallel wire transmission line, drive it with an open circuit on the other end (or a short, it doesn't matter) and see the arcing at the voltage peaks along the line...)*

  In pulsed systems (probably more representative of switchers, tesla coils, etc.) you have a problem with the pulse propagating down the line, hitting the discontinuity in impedance, reflecting back, and summing with the next pulse being sent. Whether the reflected pulse is the same or different polarity depends on the distance, and the relative impedances. If you have several mismatches, you can get lots of pulses moving back and forth which reinforce or cancel as the case may be. (This is a real big problem on commercial power distribution, because the propagation time down the line is a significant fraction of the line frequency period, causing problems when circuit breakers open and close and for lightning strikes... the impulse from the lightning whips on down the line, the protective gap fires, shorting the line to ground, causing another impulse to propagate back, etc...)

All those cool pulse forming networks based on transmission lines work (like Blumlein, Guillemin, etc.) all make use of this idea... charge the line up, short the end, and you get a nice square pulse out the other end. Charge up a whole raft in parallel (say, though HV chokes), short one end in parallel, hook the other ends in series, and you can make a big HV pulse with a lower voltage source.

Folks talk about VSWR because that's what's easy to measure with an RF wattmeter or a bridge. In reality, what's important is the reflection coeffcient, which can be complex, by the way. VSWR can be calculated from the magnitude of the reflection coefficent VSWR = (1+mag(Gamma))/(1-mag(Gamma)), where Gamma is the reflection coefficient. (mag(Gamma) is always in the range 0 to +1).

When a device does fail...

One other thing worth mentioning here, is that MOSFET devices usually fail short-circuit, as opposed to "burning" open-circuit.

Q: Why does a MOSFET always fail short-circuit ?

A: So that the opposing MOSFET is also destroyed of course ! (It must be a Murphy law thing.)

But, seriously MOSFETs usually fail short circuit, causing the opposing device to fail as well, (also short-circuit.) This short-circuits the supply, so current limiting or suitable fuse protection should be employed to prevent damage to the supply. Electrical isolation from the control circuitry is also desirable in these circumstances. Small isolation transformers prevent damaging power from flowing back into the control circuitry when MOSFETs fail with all 3 leads shorted together.

The list provided here is by no means complete, but is meant to represent a starting point for fault finding. If you have discovered another failure mode please let me know so that I can add it to the list.

Back to home page