

PLL tuned HF Tesla Coil for Plasma Loudspeaker¹

Bill Slade

Ver. 1.2, March 5, 2007



¹This work has not been peer reviewed, so any assertions made in this paper should be thoroughly checked by the reader. Although the author has tried to avoid errors, they may exist. Readers are encouraged to contact the author with possible errors, comments and questions. The author assumes no responsibility for the consequences readers may suffer attempting to construct this device. In short, you build this device at your own risk! This paper is intended to fulfill no other purpose than to be an educational exposition of an electrical engineering problem. This is a work in progress. Last modified: 2007- 03- 05. This work is licensed under the Creative Commons Attribution- Noncommercial- Share Alike 2.5 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by- nc- sa/2.5/> or send a letter to Creative Commons, 543 Howard Street, 5th Floor, San Francisco, California, 94105, USA.

1. Project summary:

This project was undertaken with the goal of developing an efficient continuous wave (CW) solid state Tesla coil operating above 3MHz. The attempt to construct this device was inspired by the work of Dan McCauley at EasternVoltResearch [1] and experiments carried out by the keen experimenters at 4hv.org. In conjunction with this, it was also desired to design and construct a suitable modulator for producing high quality sound from the plasma discharge. Some preliminary specifications for the Tesla coil and RF stages are:

1. Power amplifier (PA) output (typ): 200- 250 Watt
2. PA supply: 100- 120VDC
3. PA efficiency: > 90%
4. Operating frequency: 4.5 MHz
5. PA input signal level: 5V (TTL-level)
6. VCO freq range: 4.0- 5.0MHz
7. PLL loop bandwidth: >50kHz
8. BW of drain low-pass network: >15kHz
9. Drain LP network attenuation at 4.5MHz: > 100dB (conducted limits should be below 56dBuV)
10. Measured RF electric field at 300m: < 15uV/m (based on FCC part 18 for ISM devices on non- ISM bands).

Preliminary specifications for the modulator are:

1. Drain supply voltage: 100- 120V
2. Output PWM duty cycle: $0 < d < 90\%$
3. Clock frequency: 200kHz
4. Output voltage: 100V
5. Clock ripple on power supply: < 0.5V
6. Effective audio band: 20Hz – 15kHz

The overall system outline should enhance flexibility, so that various experiments can be tried in the interest of studying the Tesla coil behaviour and optimising output. Figure 1 shows the proposed self tuning Tesla coil system.

This project progressed in three stages. The first was to develop a suitable class- E (or similar) power amplifier and matching scheme to the Tesla coil resonator. As will be seen later, it is a challenge to maintain the efficient class- E like amplifier behaviour over different levels of plasma power in the Tesla coil as a result in shifts in capacitive and resistive loading caused by variations in discharge size. Certain precautions need to be observed to reduce undesired (and illegal) interference and to enhance safety.

The second task was to develop a special phase- locked- loop (PLL) that was capable of generating quadrature (at 0 and 90 degrees) VCO output signals. The PLL uses a pickup “antenna” near the Tesla coil resonator to sample the phase of the electric field near the top of the coil for the “self tuning” mechanism to work properly. Developing a stable PLL loop requires a bit of thought and care to detail. Some aspects of PLL design will be discussed as well as precautions that need to be taken to prevent interference from the PA stages.

The modulator is based on the ubiquitous TL494 switch mode power supply chip operating as a pulse-width modulator. It operates at high speed (important to get high audio quality and to reduce the size of transformer). Details on important design considerations are given to assist experimenters. As with the Tesla coil, some care needs to be taken to reduce electromagnetic interference, given the high clock speeds and sharp switching transitions encountered in this subsystem.

Now the disclaimer. I do not want to see any news reports about someone killing or injuring him/herself building musical Tesla coils in the backyard shed! **As always, anyone who experiments with these types of circuits does so at his or her own risk. There are high voltages involved as well as dangerous levels of continuous wave RF power which can cause intensely painful and possibly life-threatening burns. Be sure that you know what you are doing. The author attempts to point out possible safety hazards in the course of this document, but there is no guarantee that these warnings are definitive. It is up to you to use common sense. If you are not sure about a safety issue, ask someone who knows. Do not attempt something if you are unsure. Always be aware that the death of yourself and/or someone else is a likely penalty for your mistakes! Know the risks!**

Be aware that you are also responsible for keeping any electromagnetic interference within legal limits. This document describes the construction of high power RF circuits capable of causing harmful unintentional interference by RF coupling into power supplies, nearby cables and metal objects. Be sure to use adequate decoupling in circuits and use RF tight enclosures for system components.

Finally, use appropriately rated fuses and/or circuit breakers in your power supplies. Short circuits will result in fire and possible burns to yourself or others as well as possible destruction of valuable property. When you use an oscilloscope, voltmeter, ammeter or any other test instrument, make absolutely sure you know what you are doing!

2. Plasma Generation/CW Tesla Coil

The initial step is to develop a plasma source. For this, we rely on a modern variation on a very old technology: the Tesla Coil.

2.1 The resonator

The resonator consists of about 25 m of PVC insulated 1.5mm² copper wire close wound on a 90mm diameter PVC coil form (a 35cm long piece of PVC drain pipe). It is important that the form material have low loss at the operating frequency (approximately 4.5 MHz). PVC pipe is a good performer in this service as long as the coil windings do not get too hot (lower than 50 or 60° C). I ended up with about 85 turns on the form, forming a solenoidal secondary length of about 25 cm. (This will depend on the thickness of the insulation on the winding wire.)

At the bottom of the secondary winding, 10 turns of 1.5mm diameter enameled copper wire were close wound on top of the secondary winding. The ends can be

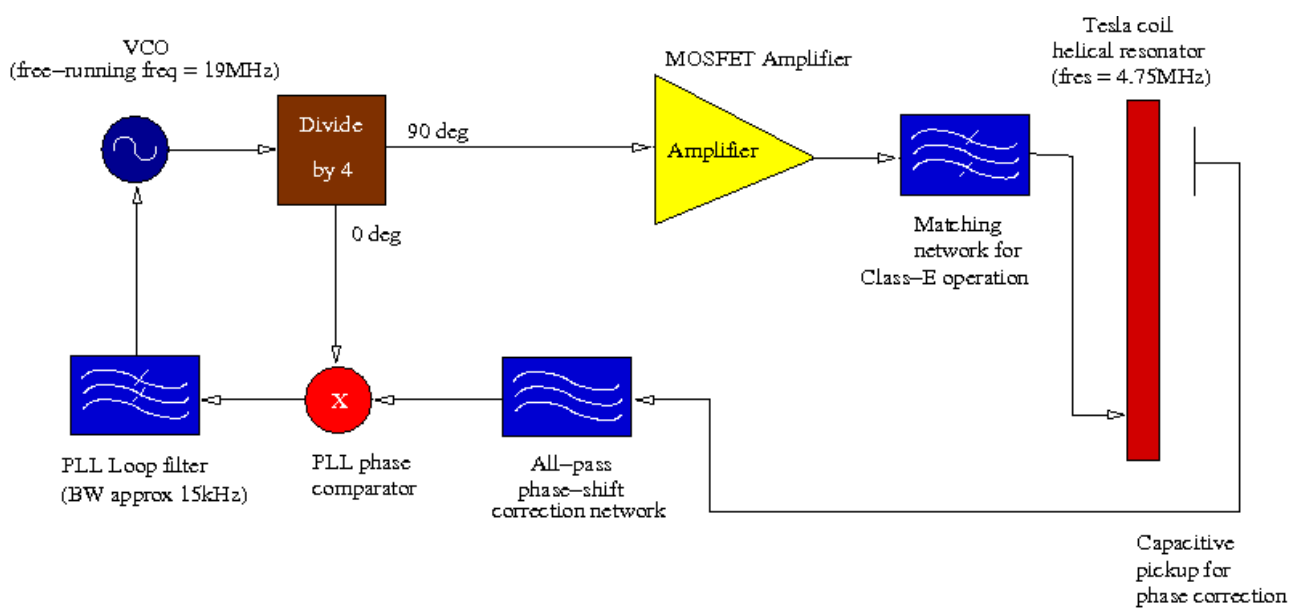


Figure 1: Proposed PLL based self-tuning Tesla coil for operation between 4 and 5 MHz.

held in place by feeding the wire ends through holes in the bottom of the form or held in place with electrical tape. Figure 2 shows a photo of the completed coil resonator.



Figure 2: Picture of energised Tesla coil resonator in its "cage" with two standard fluorescent tubes.

It is worth mentioning that there is nothing special about the resonant frequency of the resonator. The original a priori design criterion for the resonant frequency was that it be around 3-5 MHz. It turned out that a standard package of 25m of wire wound on the 90mm pipe should give a frequency in that range.

2.1.1 Modeling the resonator with NEC

Further study of the open resonator were carried out using the NEC solver [2]. This software is usually used for solving antenna problems, but any open structure is

conveniently solved. Even resonator loaded and unloaded Q factors can be computed as well as resonant frequencies. See Appendix A for some background theory on the NEC package as well as where to find the software and its user manuals. In the meantime, let us present some results which help us gauge the feedpoint impedance of the primary fed resonator under various conditions of loading. This will also help us to estimate the coupling between the primary and the secondary windings.

Using the NEC package means assembling a geometric model of our coil. We neglect the plastic of the coil form and the wire insulation and assume all air dielectric. The winding diameter and pitch can be chosen and we get something that looks like the cross-section picture in Figure 3.




Figure 3: Schematic of resonator geometry. Note that drawing is not to scale.

The input file to NEC is given by:

```

CM *****
CM ** Model of a coupled HF Tesla coil resonator **
CM ** Compute S11 **
CM ** primary excited w/1.0. **
CM ** Dimensions are: **
CM ** **
CM ** Bill Slade **
CM ** 17- 3- 2006 **
CM *****
CE
GH 1 2000 0.00258 0.232 0.045 0.045 0.045 0.045 0.0001
GM 0 0 0. 0. 0. 0. 0. 0.01 1.0
GH 3 200 0.00258 0.01295 0.048 0.048 0.048 0.048 0.0001
  
```

```

GM 0 0 0. 0. 0. 0. 0. 0.01 3.0
GA 4 10 0.01 0. 90. 0.0001
GM 0 0 0. 0. -90. 0.045 0. 0. 4.0
GW 5 15 0.048 0. 0.01 0.068 0. 0.01 0.0001
GW 5 15 0.068 0. 0.01 0.068 0. 0. 0.0001
GW 7 15 0.0476445857 0.00583039067 0.02295 0.068 0.0023370052 0.01518
0.0001
GW 7 15 0.068 0.0023370052 0.01518 0.068 0.0023370052 0.0 0.0001
GE 1
GN 1 0 0 0 0. 0. 0. 0. 0.
FR 0 30 0 0 4.70 0.005
EX 0 0 2269 1 1.0 0.
LD 0 0 2268 2268 50. 0. 0.
XQ 0
PT - 2
EN

```

The frequencies of the fundamental modes of the inductively coupled Tesla coil resonator are found where the imaginary part of the admittance (the blue curve) experiences a zero crossing in Figure 4. This happens first around 4.71MHz and corresponds to a “parallel” resonance. That is, the point where the feedpoint admittance becomes very small (the impedance becomes very large, like we expect for a parallel LCR circuit). The second resonance occurs near 4.85 Mhz and behaves like a “series” resonance, i.e. where the admittance is a maximum and the coupled resonator looks like a series LCR circuit.

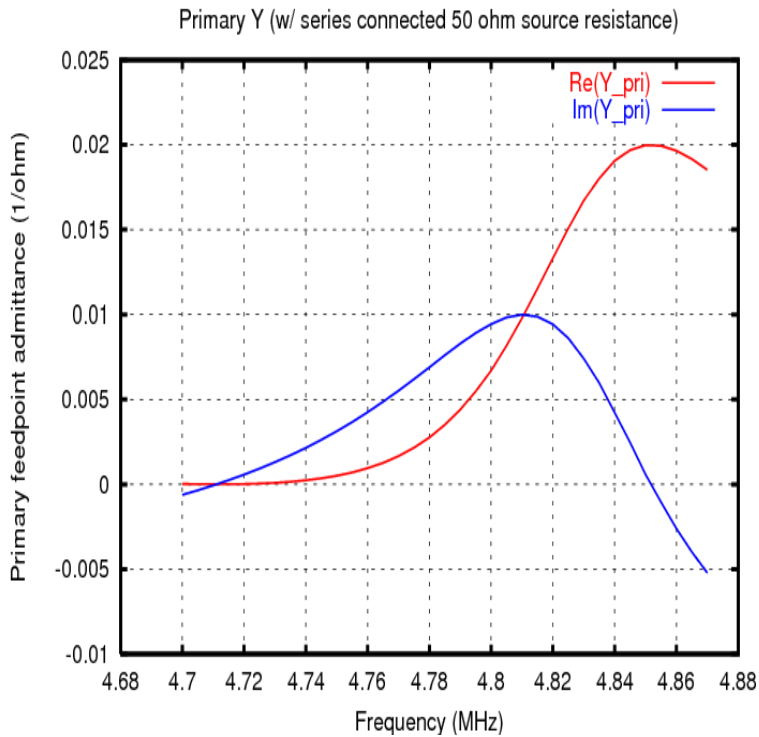


Figure 4: The primary feedpoint admittance of a lossless TC resonator. (The primary is in series with a 50 ohm resistor.)

We can simulate this behaviour using the circuit model in Fig. 5.

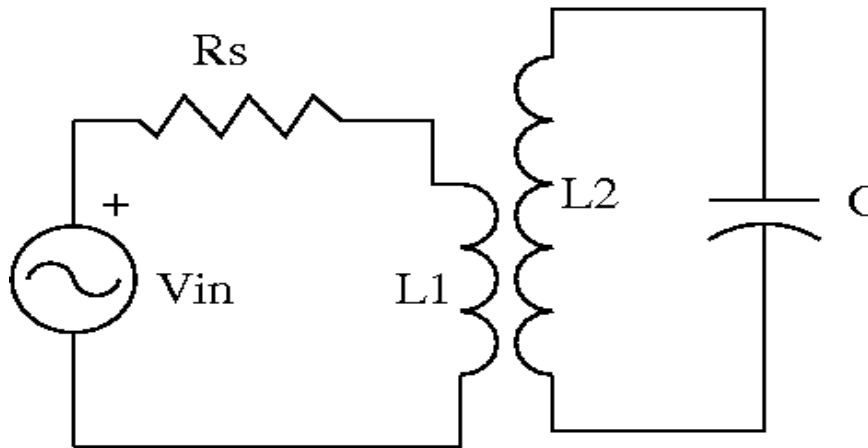


Figure 5: Circuit model of the coupled resonator. If $R_s=50$, $L_1=10\mu H$, $L_2=196.8\mu H$ and $C=5.805p$. Coupling coefficient $k = 0.24$.

Assembling the circuit equations in the frequency domain:

$$\begin{Bmatrix} V_i \\ 0 \end{Bmatrix} = \begin{bmatrix} R_s + j\omega L_1 & -j\omega k \sqrt{L_1 L_2} \\ -j\omega k \sqrt{L_1 L_2} & j\omega L_2 + 1/j\omega C \end{bmatrix} \begin{Bmatrix} I_1 \\ I_2 \end{Bmatrix},$$

where the currents I are the loop currents on both sides of the transformer.

To find the admittance seen by the voltage source, solve for I_1 :

$$I_1 = V_i \frac{j\omega L_2 + 1/j\omega C}{(R_s + j\omega L_1)(j\omega L_2 + 1/j\omega C) + \omega^2 k^2 L_1 L_2}$$

This can be modified slightly for more attractive presentation as well as solving for the admittance:

$$Y_i = \frac{1 - \omega^2 L_2 C}{(R_s + j\omega L_1)(1 - \omega^2 L_2 C) + jk^2 \omega^3 L_1 L_2 C}$$

The plot of the admittance Y_i versus frequency for this circuit model is seen in Figure 6. Note from the formula that the parallel resonance occurs when

$$\omega = \frac{1}{\sqrt{L_2 C}} ,$$

i.e., the admittance vanishes (input impedance tends to infinity).




Figure 6: Plot of admittance seen by the voltage source in the circuit model of Fig. 5.

The rolloff above resonance is a bit sharper for the circuit model than for the NEC model, because we are assuming in the circuit model a perfect LC circuit. The NEC model contains a small resistance to account for slight radiation effects. However, this model will provide a good starting point for circuit simulations of the Class- E

amplifier based Tesla coil.

An important part of the modeling is to estimate the loading provided by the arc. For the purpose of developing a matching network between the amplifier and the coil resonator, I assume a resonator Q factor between 10 and 20. This is very difficult to measure in practice, because the coil must be operating under high power conditions. Furthermore, the characteristics of the arc depend on power delivered to the coil and are also, by nature, time varying as a result of air movement and plasma instabilities. An in-depth discussion of discharge properties is beyond the scope of this paper, but is an active area of research.

2.2 Class E Amplifier and Matching

In this section, we touch on some of the important aspects of class-E switching amplifier design. Starting from a simplified version of the topology found in [3], we construct the basic output matching/filter network (Figure 7).

Starting from the left of Figure 7, DC power is supplied to the circuit by the supply VDD. The inductor RFC must have a sufficiently high value to block the RF from going back into the power supply as well as to look like a current source at RF to the switching/filter circuit. If the inductive reactance is several hundred ohms at the working frequency, this is sufficient. (At 4.5 MHz, 17 – 25uH should be enough.) Do not make this inductor too big, otherwise it could resonate (with possibly disastrous consequences). Also, if you wish to audio modulate the supply (to make a "singing Tesla coil"), you do not want the audio rolloff to be too noticeable in the treble range.

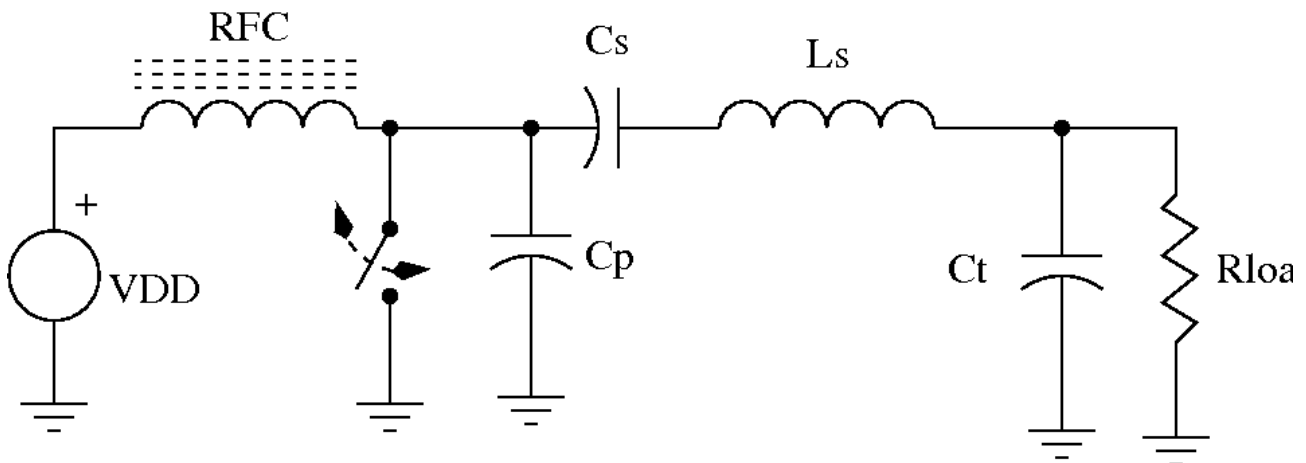


Figure 7: Topology of the class-E switching amplifier.

The switch shown usually takes the form of a power MOSFET or IGBT. In the megahertz range, MOSFETs work well. (I use a 2SK2698. IRFP460 may be used, but it has a higher gate capacitance and is therefore more difficult to drive.) The switch is made to open and close only when the voltage across the switch is close to zero. This greatly reduces the power dissipated in the transistor during on-off switching transients (when the transistor is in its linear region, i.e. looks like a resistor). This means that the timing of the on and off switching must be carefully controlled. This is accomplished through the careful design of the RLC circuit to the right of

the switch.

Using the rule of thumb provided on the Class- E forum [4], we design L_s to have a reactance of about $X_{L_2} = 3 * 0.56 * V_{DD}/I_{DD_{max}}$ for operating into a load $R_{load} \approx 50$ ohms. (For a more systematic set of design rules, see [5].) From there we will use SPICE simulations to optimise its value. The values of C_s and C_p should be chosen to resonate (including the transistor drain capacitance) with the coil L_2 near the operating frequency. Given our choice of inductor, we will have created a low Q resonant circuit (assuming a 50 ohm load resistance). The value for C_p will ultimately have to be tweaked, since real transistor drain capacitances will vary slightly. Generally, smaller values of C_p will mean higher instantaneous drain voltages. Care should be taken to ensure that the transistor maximum tolerable drain voltage is not exceeded. Use slightly larger values than those indicated by simulation and work downwards in capacitance until desired performance is obtained. C_t is a shunt load tuning capacitor. An air-variable capacitor is ideal for this service (30- 800pF). This is used to achieve a good match and to tune for class-E operation for load variations (this will be important for good Tesla coil operation).

In order to ensure that we have found some good starting values for the matching network, a time- domain SPICE simulation is carried out on the following netlist:

```
VIN      100      0      DC      115.0
LDD      100      1      250u IC=0

VGG      201      0      SIN(0.0 15.0 4.69meg 0 0)
RGG      201      200     4.7

XQ1      1        200      0      STW14NM50

RDUM     1        0      100000
CDD      1        0      270p
LS       1        2      2.5u
CP       3        0      720p
RLOAD    3        0      50
```

```
.OPTIONS NOPAGE RELTOL=1.0E-8
.END
```

Notice that we have replaced the idealised switch with a model for the 2SK2698 MOSFET whose gate is driven by a time- domain sinusoidal source.

Using these values, we perform a simulation that yields reasonable Class- E drain voltage waveform seen in Figure 8.

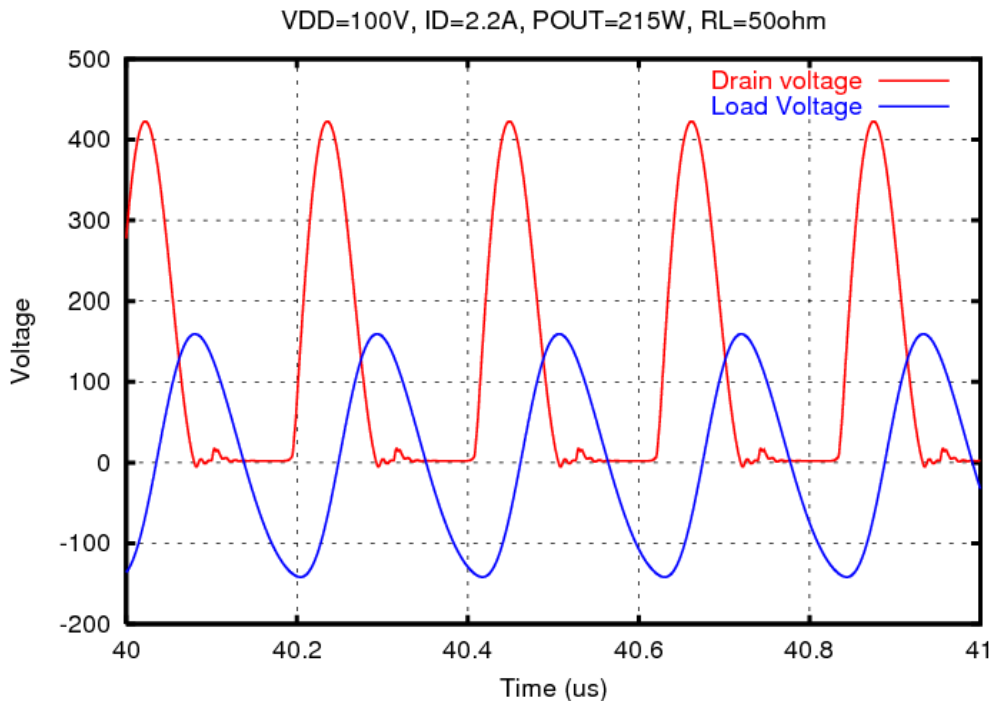


Figure 8: SPICE Simulated waveforms for the simple class- E output stage shown in Figure 7. The switch is replaced with a model of the 2SK2698 MOSFET.

Notice the phase lag of the load voltage with respect to the drain voltage as well as the slight distortion of the output voltage waveform on the load. Since we are developing this circuit for a Tesla coil (instead of a radio transmitter), we do not worry so much about this distortion. A radio transmitter would need additional filtering to ensure low harmonic output.

Now that we have an idea of what the amplifier looks like, let us replace the 50 ohm load with the circuit model of the Tesla coil resonator in Figure 9.

Table 1 gives the values of the components that yield good simulation results.

Component	Value	Notes
VDD	100V	DC power supply
RFC	50uH	RF choke
SW		Use 2SK2698 switched w/12V sinusoid
Cp	270pF	This is addition to the drain capacitance of the 2SK2698 (about 600pF)
Cs	2700pF	
Ls	2.1uH	
Ct	270pF	This will depend heavily on TC loading
Lp	1uH	primary winding
Lr	17.68uH	secondary

Component	Value	Notes
K	0.3	coupling coefficient
Cr	67.87p	
Rr	10000 ohm	for setting unloaded Q (here about 20)

Table 1: Matching circuit component values.

Note that the coupling coefficient is somewhat different than that predicted by the resonator modeling. This is because the performance of the coil depends strongly on the loading from the arc (the loaded resonator Q). The load that appears at the




Figure 9: Circuit model of matching network for class-E Tesla coil.

feedpoint depends on a combination of the Q and the coupling. These are basically determined using a decidedly unscientific trial-and-error process. We can be satisfied that the coupling coefficient is within 20% of that predicted by the resonator modeling.

The last row in the table is the equivalent resistance of the secondary. We have set it such that the resonator Q is about 20 (simulating the effect of the arc loading). Note that this is a much oversimplified model of the Tesla coil resonator (for example, the actual voltage magnification is not modeled, nor does this represent a good model of the arc discharge). It will suffice, however, to put some starting values on the matching circuit components.

A Spice simulation yields the results in Figure 10.

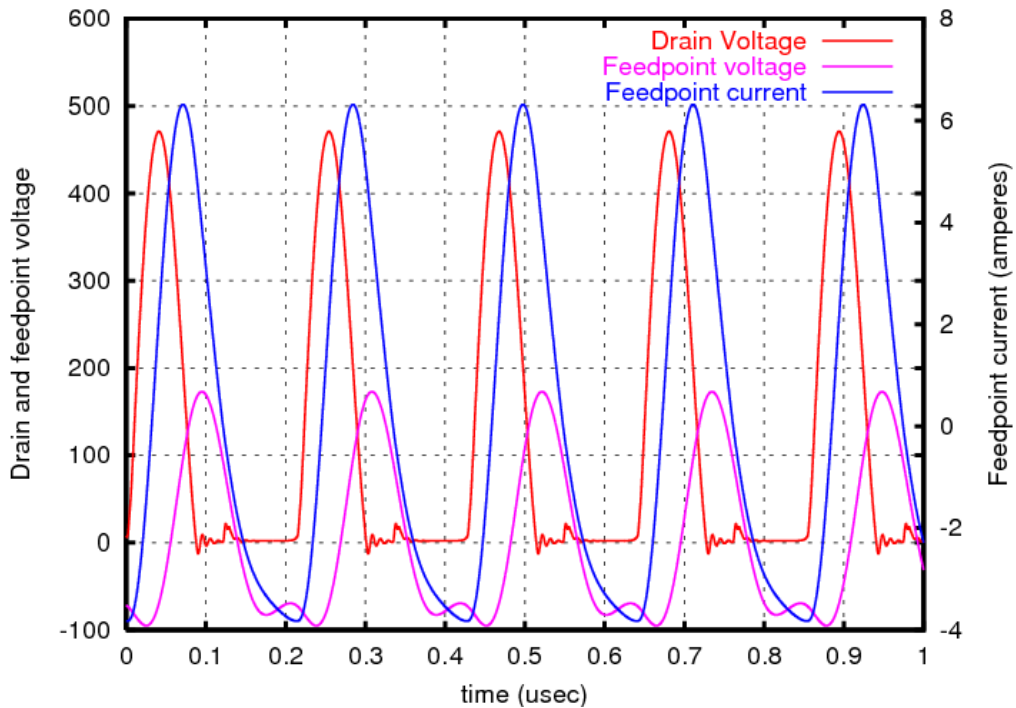


Figure 10: Time domain waveforms for Tesla coil driven near Class- E operation.

The simulation results show well timed drain switching transitions. The switch opens and closes when the voltage is close to zero, which is the condition for class-E operation. We also see that the Tesla coil primary voltage and current is rich in harmonics. This is not a problem, because we want to make sparks, not a radio transmitter. This is why the entire coil should be well shielded in an RF tight enclosure (Faraday cage).

In practical Tesla coil service, it is difficult to maintain good class-E operation, but by careful tuning, efficient (although not perfect) switching behaviour can be attained. Figure 11 shows a typically attainable power amplifier drain waveform.




Figure 11: Scope shot of power MOSFET drain waveform during operation of the Tesla coil.

The waveform has the tendency to change as the Tesla coil discharge power varies or the primary/secondary windings heat up and expand, slightly changing the coupling. The parallel variable capacitor can be used to compensate for these changes. Of course this is a manual adjustment that is somewhat difficult to automate. Some improvement to primary/secondary design to reduce thermal expansion would undoubtedly help.

2.3 The PLL and oscillator chain

Since we use a phase locked loop to auto-tune the Tesla coil, it is worth reviewing some basic theory behind the PLL [6]. Most basic frequency multiplying (integer- N) PLLs have a block diagram that looks something like Figure 12.

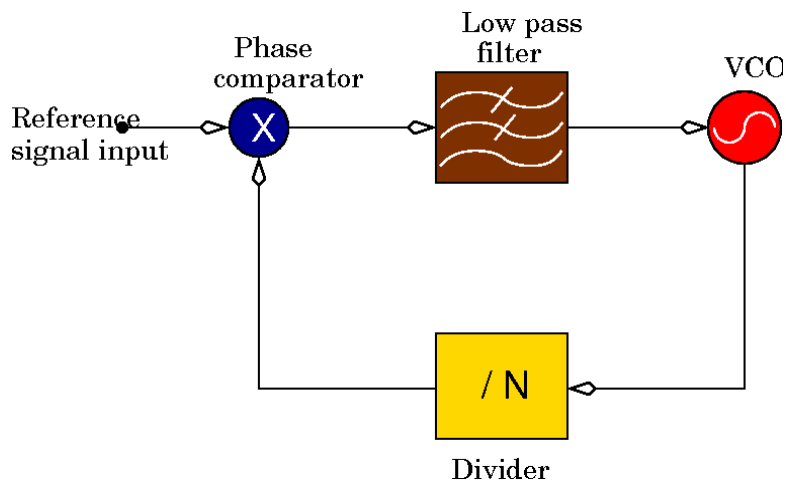


Figure 12: Generic phase- locked loop.

A reference signal is applied to the input to the phase comparator. If you were making a frequency synthesizer, this might be a signal from a highly stable crystal

oscillator. In our case, it will be a signal proportional to the voltage at the top of the Tesla coil resonator. The phase comparator in our circuit is essentially a balanced mixer like those found in most radio circuits. The reference signal is "mixed" with the frequency divided signal from the voltage controlled oscillator producing both low and high frequency components. Only the low frequency components pass through the low-pass filter. These low frequency components steer the voltage controlled oscillator such that the frequency divided VCO output and the reference signal are always exactly 90 degrees out of phase, simultaneously keeping the VCO frequency exactly N times the reference frequency.

If the loop is close to the locked state, we can perform a linear analysis of the phase behavior. First, let us list the pertinent variables in Table 2.

Variable	Label	Notes
Reference phase	ϕ_{ref}	
Phase detector output voltage	V_{ϕ}	
Phase detector gain	K_{ϕ}	units of volts/radian
Low pass filter transfer function	$H(s)$	
VCO input voltage	V_{tune}	
VCO tuning gain	K_{VCO}	units of Hz/volt
Divider ratio	N	
Divider phase	ϕ_N	

Table 2: Variables associated with linear PLL analysis.

We can treat the system in Figure 11 as a linear feedback system with the VCO tuning voltage as the output and the reference phase as the input. Remembering that the phase comparator (mixer) gives us the difference between the reference phase and the divider phase. Hence,

$$V_{\phi} = K_{\phi}(\phi_{ref} - \phi_N) \ .$$

The tuning voltage (in the frequency domain) of the VCO is then

$$V_{tune} = H(s)K_{\phi}(\phi_{ref} - \phi_N)$$

The tuning voltage, in reality, does not change the phase of the VCO, but its frequency. In terms of phase, this means

$$\frac{d\phi_{VCO}}{dt} = K_{VCO} V_{tune} \ .$$

In the frequency domain, the VCO phase can be written as

$$\phi_{vco} = \frac{K_{VCO} V_{tune}}{s} ,$$

where we substitute s for d/dt .

Remembering that the divider also "divides" the VCO phase (an exercise to the reader: convince yourself of this), we can close the loop and get

$$V_{tune} = H(s) K_{\phi} \left(\phi_{ref} - \frac{K_{VCO} V_{tune} / N}{s} \right) .$$

After some rearrangement, we can put the VCO tuning voltage in terms of the reference phase:

$$\frac{V_{tune}}{\phi_{ref}} = \frac{s H(s) K_{\phi}}{s + \frac{K_{phi} K_{VCO}}{N} H(s)} .$$

Putting this in terms of reference frequency (by setting $\omega_{ref} = s \phi_{ref}$) yields the classic PLL closed-loop transfer function

$$Y(s) = \frac{\omega_{VCO}}{N \omega_{ref}} = \frac{\xi H(s)}{s + \xi H(s)} ,$$

where

$$\xi = \frac{K_{VCO} K_{\phi}}{N} .$$

The variable ξ is usually set by the PLL chip and VCO (although we may be able to change the divide ratio).

The key to designing the PLL lies in finding a good filter transfer function $H(s)$ that yields the tracking speed we need while simultaneously remaining stable. This can be done by ensuring that the poles of the transfer function all lie well on the left-hand-side of the complex plane. Alternatively, at the point where the magnitude of the transfer function lies on the edge of the unit disk in the s -plane, there should be a good phase margin away from 180 degrees (the condition for oscillation instability).

2.3.1 A design example

Let us start with a simple example: a loop low pass filter that is just a simple RC circuit where

$$H(s) = \frac{1}{1 + s\tau}$$

and τ is the time constant of the RC circuit. The transfer function can now be written as

$$Y(s) = \frac{\xi/(1+s\tau)}{s + \xi/(1+s\tau)} .$$

This can be rearranged into

$$Y(s) = \frac{\xi/\tau}{s^2 + s/\tau + \xi/\tau} .$$

we see we have created a second-order PLL using a first-order RC low pass filter. The poles of $Y(s)$ lie at

$$s = -\frac{1}{2\tau} \left(1 \pm \sqrt{1 - 4\xi\tau} \right) .$$

If $\xi\tau > 1/4$, then the system is said to be underdamped. If there is a sudden change in reference frequency, the VCO output frequency will exhibit damped oscillatory behaviour. According to the linear theory, these oscillations will damp out over time. Keep in mind, however, that this linear model is an approximation of a nonlinear system. If the oscillations are large, this model breaks down, because it cannot account for the nonlinear instability that can occur if the frequency excursions are too large. For this reason, if the linear theory predicts highly underdamped solutions, expect instability to appear.

If $\xi\tau < 1/4$, then the PLL is overdamped. Here stability is not generally a problem. However, slow locking and poor tracking can be problems of the loop is overly damped.

The ideal situation is when $\xi\tau = 1/4$. This is the critically damped situation, which exhibits the fastest locking and tracking times for a given loop bandwidth. This allows us to produce the following design rule for the second order PLL:

$$\frac{K_{VCO} K_{\phi}}{N} = \frac{1}{4RC} .$$

2.3.2 A "real" PLL design example

Let us consider a somewhat different filter topology. By adding a second resistor and capacitor as in Figure 13, we can achieve more control and better stability over the frequency response of the loop (than with the simple filter in the preceding section). This is the famous "lead-lag" filter often used in PLLs.




Figure 13: 3rd order filter used in the Tesla coil.

The new filter transfer function is

$$H(s) = \frac{1 + s\tau_2}{s^2\tau_1\tau_2 + s\tau_3 + 1} ,$$

where $\tau_1 = R_1C_1$, $\tau_2 = R_2C_2$ and $\tau_3 = R_1(C_1 + C_2) + R_2C_2$. This is a bit of a mess to insert into the PLL transfer function, so let us present a practical example. For the NE564 PLL used in the Tesla coil, we shall use the parameters in Table 3 (optimised by trial- and- error substitution using readily available capacitor and resistor values into the transfer function until we get a curve that gives good, not excessively peaked response). It is also important to keep the rolloff well above the audio range if we want to use drain modulation for creating good quality audio from the arc. This is because as the modulator changes the arc power, the PLL needs to track the shift in resonant frequency of the Tesla coil. This improves the linearity of the audio response. Note also that $R_1 = 1.3K$ is fixed internally by the NE564 chip.

Filter element	Value
R_1	1.3K Ω
C_1	4.7nF
R_2	100 Ω
C_2	47nF
K_{VCO}	1.26×10^8 rad/s/V
K_ϕ	0.477 V/rad
N	4

Table 3: Filter component values and PLL parameters.

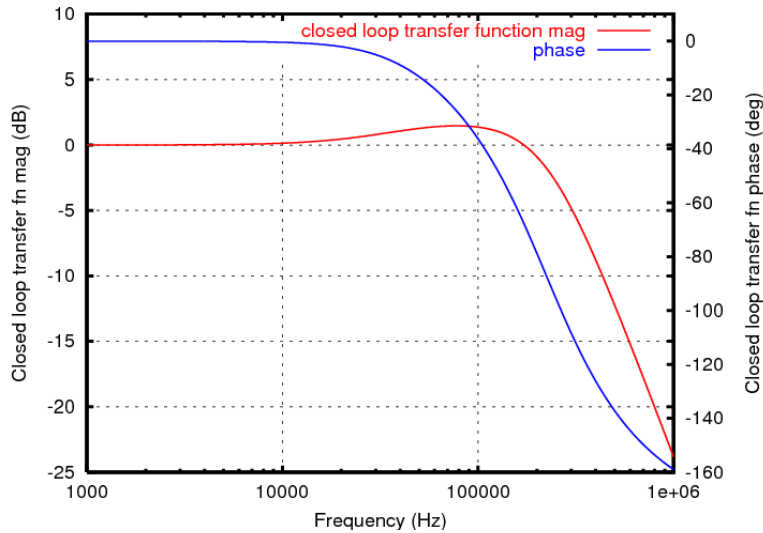


Figure 14: Closed loop response curves for 3rd order PLL.

Figure 14 is the closed-loop response of our PLL. It exhibits -3dB rolloff around 300kHz, which is more than suitable to allow the PLL to track audio variations in the coil loading while still attenuating the VCO signal well in the PLL loop. Figure 15 shows the PLL filter response by itself). Here we see the effect of the two filter poles (that lie near 4400Hz and 790kHz) and a zero at 34kHz.

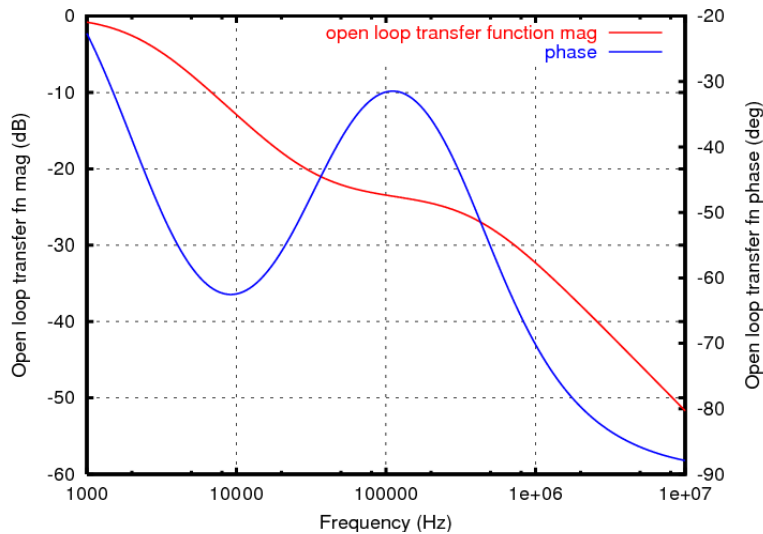


Figure 15: Loop filter transfer function.

The stability of the PLL can be measured by the amount of phase margin there is in the denominator of the PLL transfer function. Generally speaking, we encounter instability if the denominator approaches zero, viz.

$$\frac{\xi H(s)}{s} \rightarrow -1$$

We can plot the left-hand-side of this expression versus frequency and look where the magnitude approaches unity (0 dB). The corresponding distance between the 0dB point on the phase plot and ± 180 degrees gives the phase margin (in Figure

16). It should generally be more than 45° for good performance.

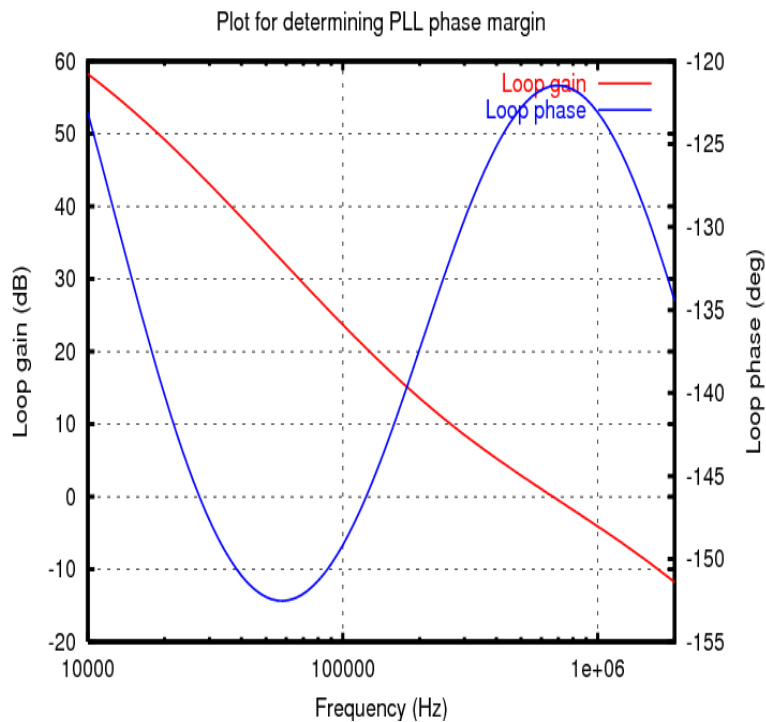


Figure 16: Plot for computing phase margin.

Figure 16 shows that we have nearly 60° phase margin (seen at $f \approx 700\text{kHz}$, where the red loop gain curve crosses the 0dB grid), so the PLL should remain stable.

2.4 Notes on implementation and circuit diagrams

2.4.1 Quadrature generation

The PLL oscillator generates a signal around 22.5MHz. By using two D flip-flops in a circular shift register configuration, we can divide the frequency by four and have 4.5 MHz signals available at 0, 90, 180 and 270 degrees. The circuit is in Figure 17.




Figure 17: Quadrature generator/frequency divider using two D flip-flops.

The 74F74 TTL flip-flop is used, but any chip capable of working above 25MHz should be fine.

2.4.2 The RF generator and Tesla coil apparatus

Figure 18 shows the full schematic of the Tesla coil system (without modulator or power supply).

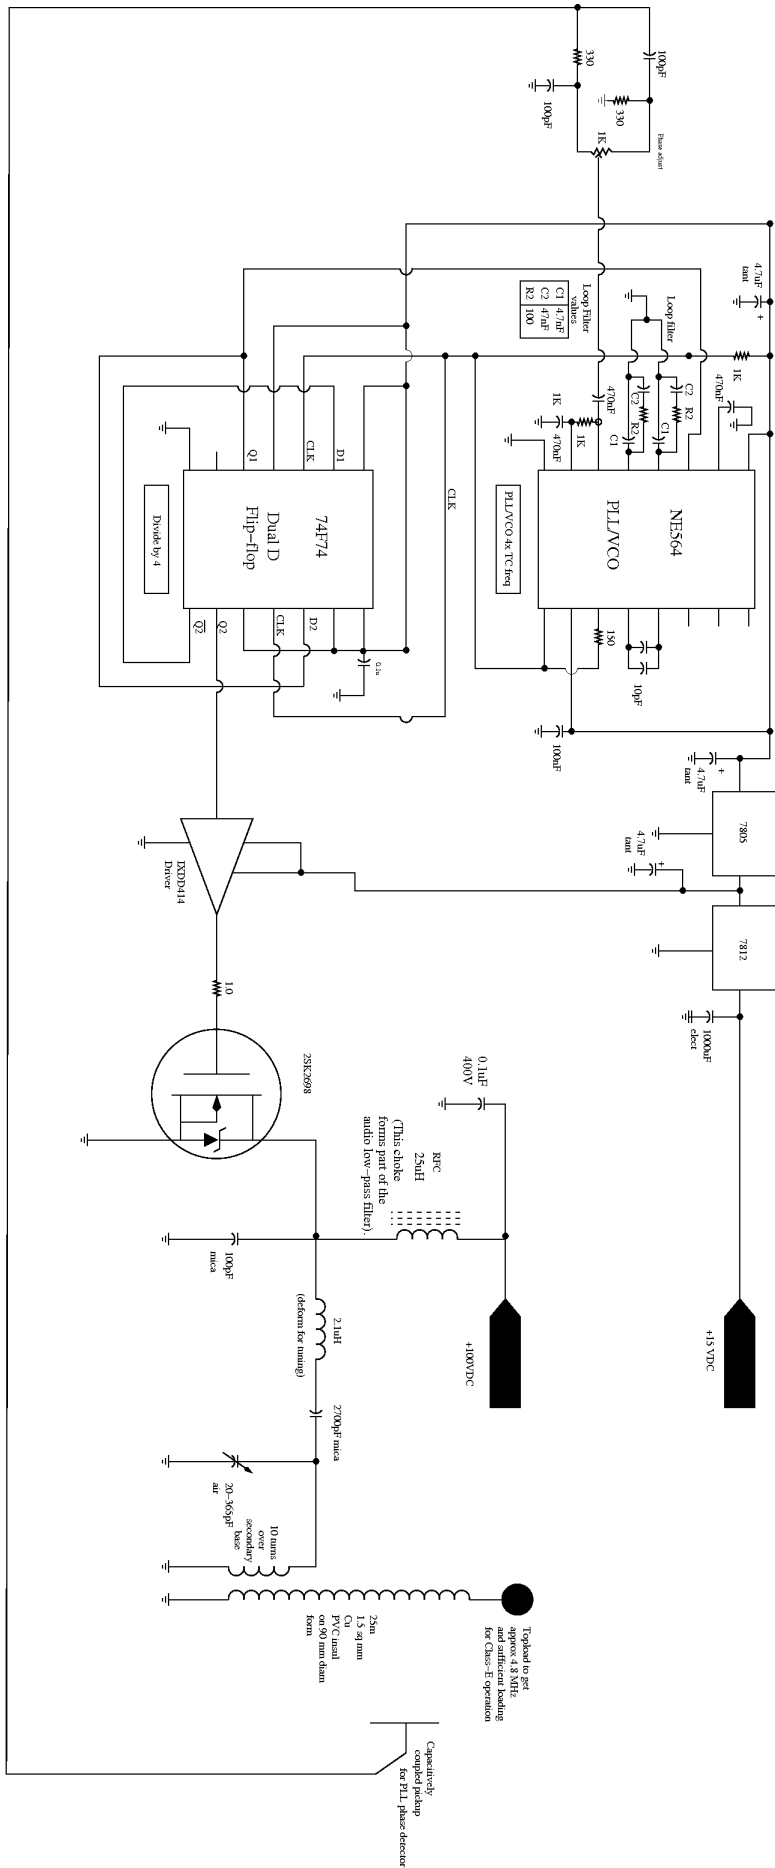


Figure 18: Schematic: PLL tuned class-E Tesla coil.

Number	Part number	Description	Supplier
1	NE564	PLL/VCO system	Digikey
1	74F74	Fast TTL D flip- flop	Digikey
1	IXDD414	MOSFET Gate Driver	Digikey
1	2SK2698	15A/600V N-Ch power MOSFET	Digikey
1	7805	Voltage regulator	Digikey
1	7812	Voltage regulator	Digikey
2	330R	Resistor 1/4W 1206 chip	Digikey
2	1K0	Resistor 1/4W 1206 chip	Digikey
2	100R	Resistor 1/4W 1206 chip	Digikey
1	150R	Resistor 1/4W 1206 chip	Digikey
1	1K0 pot	variable resistor, 1/4 W	Digikey
1	1.0R	Resistor, 2W, axial	Digikey
1	0.47uF	Capacitor, poly, 400V	Digikey
2	100pF	Capacitor, ceramic NPO, 603 chip	Digikey
1	100pF	Capacitor, silver- mica, 500V	Digikey
2	4.7nF	Capacitor, poly., 100V	Digikey
2	47nF	Capacitor, poly., 100V	Digikey
2	100nF	Capacitor, poly., 100V	Digikey
3	4.7uF	Capacitor, tant., 16V	Digikey
2	470nF	Capacitor, poly., 100V	Digikey
2	10pF	Capacitor, ceramic, NPO,603 chip	Digikey
1	2700pF	Capacitor, silver- mica, 500V	Digikey
1	70- 1000pF	Capacitor, air- variable, 250V	RF Parts
1	100uF	Capacitor, electrolytic, 50V	Digikey
1	1000uF	Capacitor, electrolytic, 50V	Digikey
2		Iron pwdr core Amidon mix 2	Reichelt Electronics
20m	1.5 mm ² Cu wire	enameled wire for coils	Reichelt Electronics
25m	1.5mm ² Cu wire	Insulated wire for secondary	Local hardware
45cm	9cm diam PVC pipe	winding form	Local hardware
1	Faraday cage		See text
1	enclosure	for PCB	Local supplier
4	BNC connectors	For enclosure and Faraday cage	Local supplier
	assrt'd hardware		Local hardware
1	PCB		Local supplier
1	Heat sink		Digikey

Table 4: Bill of Materials, PLL tuned Tesla coil.

2.5

3. The modulator

Figure 19 shows the schematic of the audio modulator.

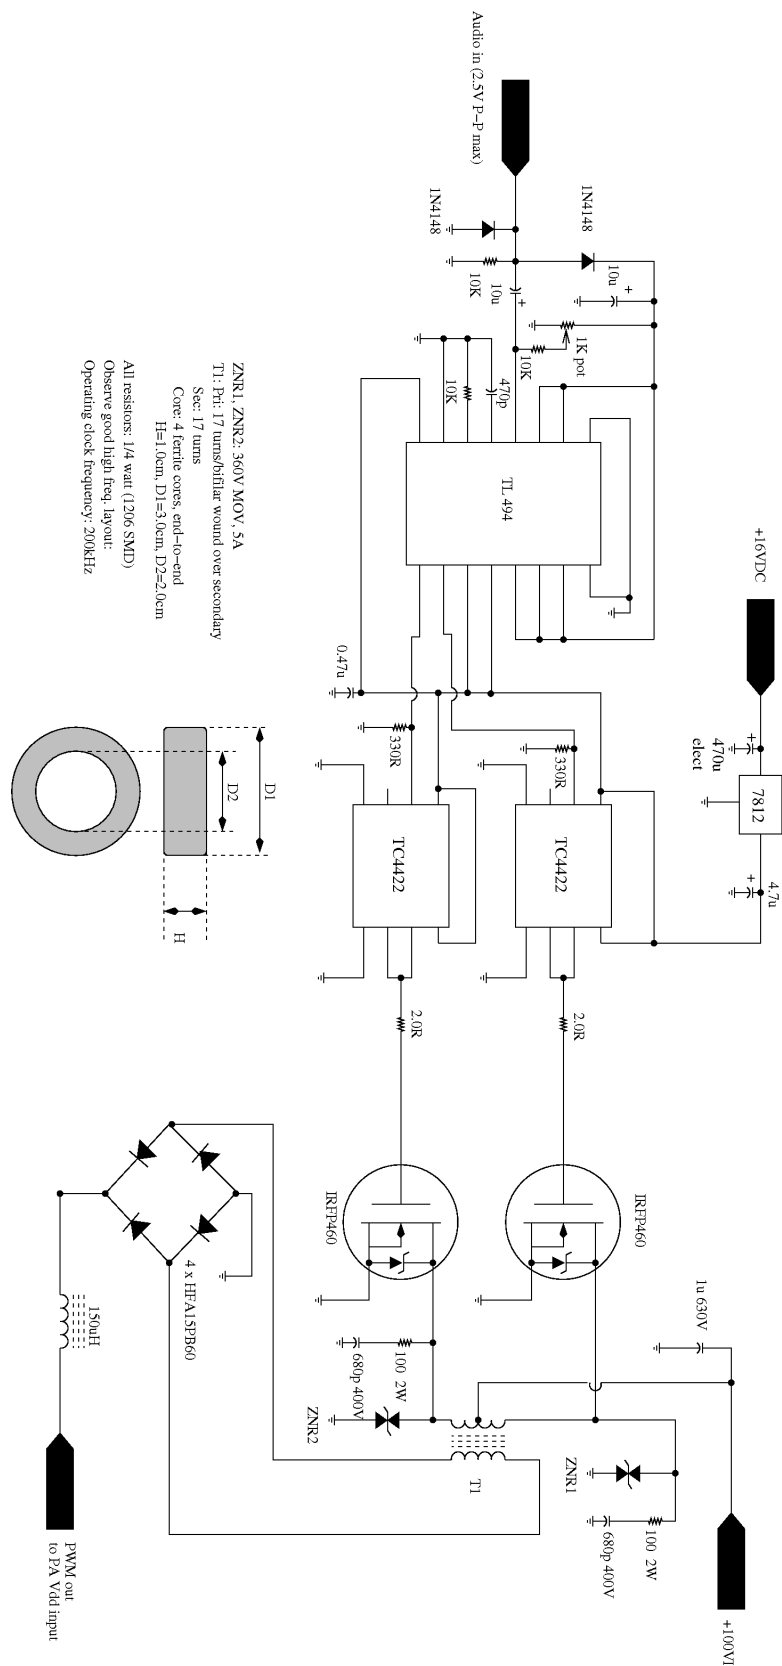



Figure 19: The audio modulator schematic.

Number	Part number	Description	Supplier
1	TL494	PWM/SPMS controller	Digikey
2	TC4422	Power MOSFET gate driver	Digikey
2	IRFP460	Power MOSFETs	Digikey
1	7812	12V/1A Voltage regulator	Digikey
2	1N4148	Silicon general purpose diode	Digikey
2	360V MOV	5A/360V MOV for clamping drain	Digikey
4	HFA15PB60	600V/15A ultra- fast diodes	Digikey
4	Ferrite cores	See schematic & text	Digikey
3	10K	Resistor, 1/4W, 1206 chip	Digikey
2	330R	Resistor, 1/4W, 1206 chip	Digikey
2	2.0R	Resistor, 1/4W, 1206 chip	Digikey
1	1K pot	Potentiometer, 1/4W	Digikey
1	470pF	Capacitor, ceramic, NPO, 603 chip	Digikey
2	10uF	Capacitor, elect., 50V	Digikey
1	0.47uF	Capacitor, poly, 50V	Digikey
1	4.7uF	Capacitor, tant., 16V	Digikey
1	470uF	Capacitor, elect., 50V	Digikey
1	0.22uF	Capacitor, poly, 630V	local supplier
1	1.0uF	Capacitor, poly, 630V	local supplier
1	PCB	FR-4, double sided	local supplier
		1.0mm ² enameled wire	Reichelt
		Asst'd. hardware	local supplier

Table 5: Bill of materials, modulator.

The modulator operates at a clock frequency of about 200kHz. This helps to keep the transformer small and the sound quality good. The transformer is made of four ferrite ring cores stacked end- to- end to improve the power handling (we need to control up to 200- 300 watts), so a conservative approach to sizing the transformer core is good. Seventeen turns of 1mm² enameled wire are tightly wound as a secondary on the four cores. Cover the cores with a layer of PVC electricians' tape first, to prevent scoring the enamel on the wire (and possible short circuits). A layer of electrical tape is also put on top of the secondary before winding a bifilar centre- tapped primary (each primary winding has 17 turns). See Figure 20 for the finished transformer. When connecting the primary to the MOSFET drains, be sure the transformer senses are correct, or you will see fire and smoke on power- up!! Design your circuit board to be double sided, with one side as a ground plane to enhance performance and to avoid RFI problems.



*Figure 20: Modulator prototype board being tested.
Finished transformer seen at bottom.*

It is important that the MOSFETS are mounted on a suitable heatsink. Also, the fast rectifier diodes will need heatsinking for sustained high power operation. However, this modulator is highly efficient (over 95% at 250W output power), so large heatsinks are not necessary. Mounting the modulator PCB (and hence, heatsinking the MOSFETS and ultra- fast rectifiers to the box wall) in a die cast aluminium box (also providing good RFI shielding) may be sufficient.

This modulator gives reasonable audio performance in an efficient manner. However, the use of a dedicated filter is recommended to between the modulator and the RF amplifier drain input. Two things are worth noting about the existing circuit in Figure 19:

- The capacitor that follows the bridge rectifier will cause high current peaks through the diodes in the bridge, particularly when pulse duty cycle is less than 50%. This reduces efficiency and can stress the diodes. A series inductor here would reduce stress on the diode bridge.
- The simple capacitor/inductor L-network may not provide enough RF decoupling to keep conducted RF into the power supply as small as possible.

Figure 21 shows how the PA stage is connected to the modulator via a L-C-L T-network. Figure 22 shows the response of this filter. It is slightly peaked and has a cutoff frequency of about 40kHz. At 4.5MHz, this filter should provide around 100dB of attenuation. Conducted RF should be on the order of millivolts on the DC supply side, if good layout and enclosures are used. For the 25uH inductor, use RF grade powdered iron cores. Low-loss, high permeability ferrite should be used for the 150uH inductor.

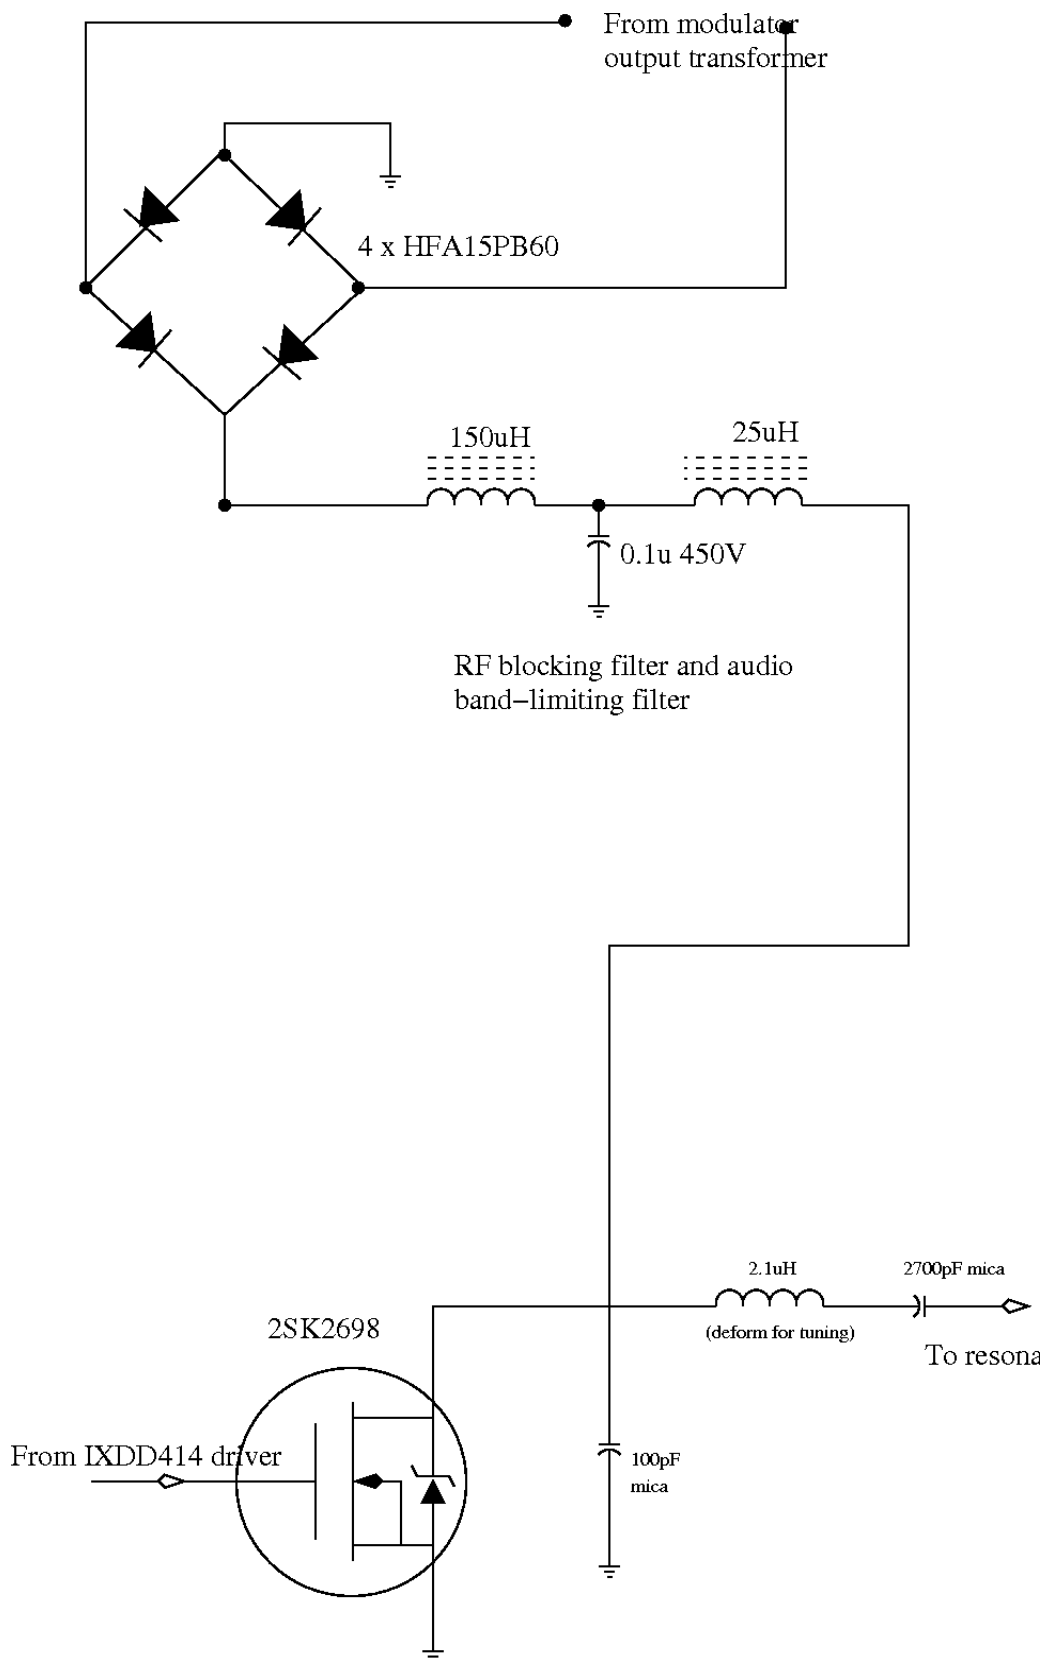


Figure 21: Recommended implementation of the RF blocking filter between RF power amplifier and audio modulator.

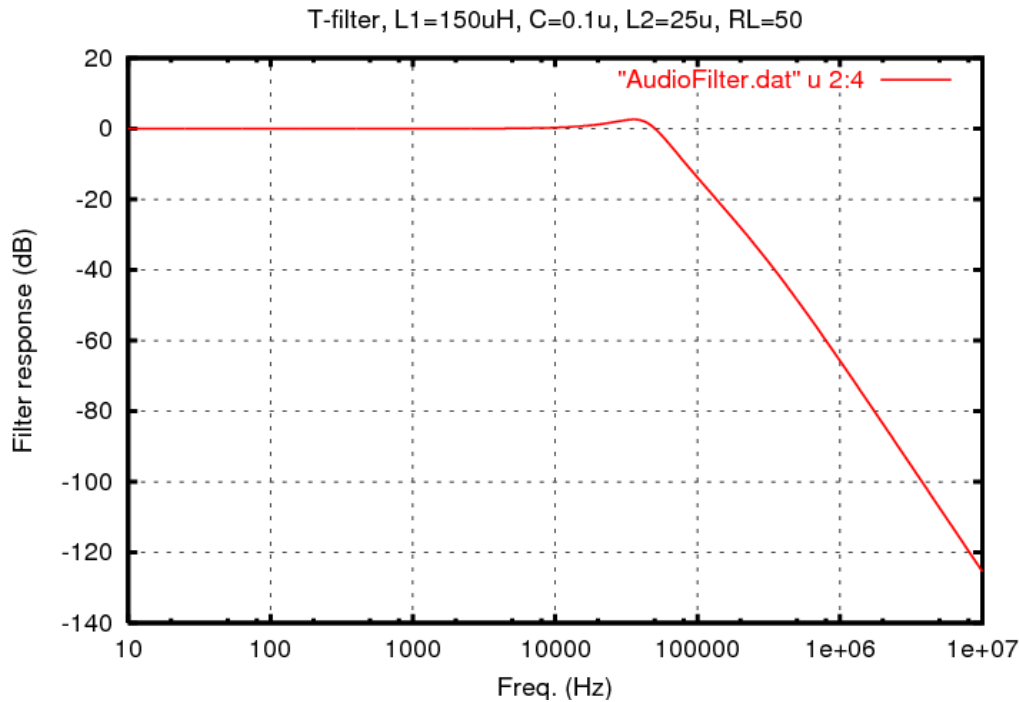


Figure 22: Frequency response of RF-blocking low-pass filter.

4. The DC power supply




Figure 23: Schematic of power supply.

The power supply is relatively straightforward. Given the voltages and currents involved, make sure that all wiring is suited for the current. (I used 1.5mm² copper wire for all power connections in this circuit.) Fuses or circuit breakers are absolutely necessary. If a power MOSFET fails and suffers a short circuit, there must be a fuse or circuit breaker present to prevent a fire! Also, since mains voltages are present on the transformer primary and high secondary voltages are present in the circuit, all connections must be insulated well and inaccessible to curious fingers or inadvertent bodily contact.

5. The Faraday cage

A very important part of this experiment was the reduction of injury risks and harmful interference. For this reason a wire-mesh enclosure like that shown in Figure 24 was constructed.



Figure 24: Coil inside of shielding cage.

The purpose of the Faraday cage is twofold. It helps prevent accidental contact with the energised coil by presenting a physical barrier to those observing the experiment. Additionally, it greatly attenuates the electromagnetic fields present in the vicinity of the coil, thereby reducing personal exposure to strong fields as well as reducing coupling into nearby metal objects that may act as inadvertent antennas.

Some care must be taken in the construction of the Faraday cage. Most important is the electrical integrity of all joints in the wire mesh. This means all joints should be tightly clamped or soldered, as seen in Figure 25.



Figure 25: Illustration of the solder joints needed for the Faraday cage.


Where cables enter the cage, the ground must be well established near the entry

point by using proper connectors soldered or clamped to the mesh (Figure 26). Otherwise, radiation from ground loops could present problems. Coaxial cables are strongly recommended for RF power applications, although closely twisted pairs will function reasonably well, particularly in balanced (differential) transmission lines.



Figure 26: Coaxial cable entry points through the mesh using BNC connectors.

Furthermore, be sure to mount the Tesla coil driving circuit in an RF-tight box (like the die cast box in Figure 27). Bypass all DC and low frequency (audio) inputs so RF coupling into the power supply is minimised. Keep all leads as short as possible and read up on electromagnetic compatibility. Then you can apply the proper precautions for preventing harmful interference. Use a field strength meter (like that seen in Figure 28) to check for possible radiation with all shielding in place.



Take readings along power cables to see if parasitic coupling is a problem. Readings should be near zero more than in RF-tight box (with lid removed for a meter or two from the Faraday cage. If testing).

there are significant readings near the power cord where it plugs into the wall, check your RF decoupling. A suitable radio receiver with an RSSI meter can also be used for this measurement. It will be far more sensitive than a simple field-strength meter and will give a good indication of any harmful radiation at large distances (10s to 100s of meters) from the Tesla coil.



Figure 28: Using a field strength meter to check shielding effectiveness.

Remember that safety is your utmost responsibility. Make sure your cables, interconnects, primary and secondary windings can handle the power and never allow anyone to come in contact the plasma or any energised circuit. Use appropriate fuses or circuit breakers.

6. Review of specifications

In this section, some of the measured performance parameters are compared with the initial "wish- list" specifications posed in the Introduction.

Tesla coil and RF stages are:

1. Power amplifier (PA) output: 200- 250 Watt: measured at 250W w/VDD=115V
2. PA supply: 100- 120VDC: Runs well at 115VDC
3. PA efficiency: > 90%: Achieved when well tuned
4. Operating frequency: 4.5 MHz: Without arc, resonant frequency is 4.53MHz. With full power arc, drops to somewhere around 4.4 MHz.
5. PA input signal level: 5V (TTL-level): Input to IXDD414.
6. VCO freq range: 4.0- 5.0MHz: Satisfied.
7. PLL loop bandwidth: >50kHz: PLL bandwidth calculated to be around 300kHz.
8. BW of drain low-pass network: >15kHz: This has not been measured yet. Filter needs to be implemented.
9. Drain LP network attenuation at 4.5MHz: > 100dB: To be tested.
10. Measured RF at 300m: < 15uV/m: Do not have instrumentation to measure this, but field strength meter shows no deviation on its most sensitive setting when more than 2m from the Faraday cage.

Preliminary specifications for the modulator are:

1. Drain supply voltage: 100- 120V: Satisfied.
2. Output PWM duty cycle: $0 < d < 90\%$: Satisfied.
3. Clock frequency: 200kHz: Satisfied.
4. Output voltage: 100V: Output voltage to 108V using 115V power supply.
5. Clock ripple on power supply: < 0.5V: Satisfied.
6. Effective audio band: 20Hz – 15kHz: Satisfied.

Be sure to follow all the Tesla coil discussions on <http://4hv.org/>!

6. Appendix A: Notes on the NEC solver

The Numerical Electromagnetics Code (NEC) is a moment method solver for wire antennas. The idea behind the model is fairly straightforward. A curved wire (see Figure 27) is assumed to have a current flowing on it as a result of an incident electric field or a voltage/current source connected to the wire. A scattered electromagnetic field results from the currents on the wire.




Figure 29: Current carrying wire is broken up into short straight segments. Electric and magnetic fields in the space surrounding the wire can then be found.

Basically, NEC solves a form of the following integral equation along the surface of the wire (whose geometry is modelled by lots of small straight line segments) for the currents $I(l)$ along the wire given an incident electric field (from a voltage source, in our case) \mathbf{E}^i .

$$-\mathbf{E}^i(l) = -\frac{j\eta}{4\pi k_0} \left(k_0^2 + \frac{\partial^2}{\partial l^2} \right) \int_C \frac{I(l') d\mathbf{l}'}{|\mathbf{r} - \mathbf{r}'|} e^{-jk_0(|\mathbf{r} - \mathbf{r}'|)} .$$

While there is a lot of complicated detail in the solution of this equation, we are interested in knowing the current that flows through the voltage source so we can use Ohm's law to compute the feedpoint impedance. NEC can do this automatically.

Those readers interested in the theory and practice of electromagnetic modeling of antennas and open wire structures are referred to the online literature [2], which consists of user manuals and theoretical exposition.

7. References

- [1] http://www.easternvoltage.com/design_resource.htm
- [2] NEC, Numerical Electromagnetics Code, [online], <http://www.nec2.org/>
- [3] J. F. Davis and D. B. Rutledge, " A low- cost class- E power amplifier with sine wave drive," *IEEE MTT- S Int. Microwave Symp. Dig.*, vol.2, pp. 1113- 1116, June 1998.
- [4] Class- E Forum, [online], <http://classe.monkeypuppet.com/>
- [5] N. O. Sokal, "Class- E RF power amplifiers," *QEX Mag.*, Jan.- Feb. 2001, pp. 9- 19.
- [6] K Shu and E. Sanchez- Sinencio, *CMOS PLL Synthesizers: Analysis and Design*, Springer, 2005.

8. Document History

1. First version v1.0, 17 Jan., 2007
2. Version 1.1, 8 Feb. 2007. Added Figure 21, 22 on RF-blocking filter/audio filter. Added supporting text. Made other small changes to wording throughout the document.
3. Version 1.2, 5 Mar. 2007: Changed schematics slightly. Added reference [1], [5].