

TESLA COIL DRIVER

In this chapter, we deal with the design of the Tesla coil driver. We will discuss several components or subsystems, including the power supply, the controller, the gate driver and inverter, the fast fuse, and the current sense resistors. There are a limited number of interconnections between the subsystems. We will try to be clear as to what the inputs and outputs are for each subsystem.

1 Power Supply

Initially, I used a circuit without a step-up transformer. I could then apply up to ± 170 V to the inverter. This allowed for testing of many concepts, but the voltage was not high enough to make significant sparks. I then decided to use a step-up transformer ahead of the inverter. It appeared that a factor of ten ($10(\pm 170) = \pm 1700$ V) would give me enough range. So I bought a 240/2400 transformer from T&R Electric. This is an oil filled transformer, sometimes called a pole peg. Some older distribution systems used a 2400 V delta three phase configuration. These are typically being upgraded to 7200 V wye, so the 240/2400 V transformers should be available on the used market for some time. I got one rated at 5 kVA, so the rated input current is $5000/240 = 20.8$ A. This is a very conservative rating, for continuous operation in a 40°C ambient. I would expect it to run four hours at 30 A without a problem, and 10 minutes at 40 A. The high side rating is 2400 V and 2.08 A.

The power circuit is shown in the figure Fig. 1. The part ahead of the 240/2400 V transformer T1 is very common with classical Tesla coils. There is a switch SW1, which is actually a 50 A circuit breaker back at the main electrical service entrance. Then SW2 is a 40 A breaker in a metal box near the transformer. Inside the metal box is a power line filter which hopefully keeps some of the trash from getting back to the utility (or affecting other circuits in my lab). There are two inductors consisting of four turns of heavy wire on iron powder toroids, about $80 \mu\text{H}$ as measured on an Escort ELC-120 LCR meter. Then I used $0.039 \mu\text{F}$, 1600 V polyethylene capacitors on both sides of the inductors. These capacitors have a self resonant frequency of about 5 MHz. Most of my testing was done at frequencies between 120 kHz and 240 kHz, so the capacitors should help for the fundamental and several harmonics. But really high frequency RFI will go right on by. Electrical wiring is very lossy for frequencies above a few MHz, so I would not expect a problem, and indeed have not experienced anything.

I added a $6 \mu\text{F}$, 370 VAC motor run capacitor to provide more filtering for the low frequency noise. There is also a MOV (Metal Oxide Varistor), the 20K43IU, to clip any really nasty pulses. The circle marked PL is a red panel light that comes on when the 40 A breaker is turned on. This is one of many safety features that could be added.




Figure 1: Tesla Coil Power Supply

The transformer output is rectified by a full-wave bridge rectifier made from 6 kV diodes I found in my junk box. The current rating is unknown, but almost certainly inadequate to charge the capacitor bank if full transformer voltage is abruptly supplied to discharged capacitors. The variac is always turned to zero before SW2 is turned off, and is slowly brought to the desired setting after SW2 is turned on.

Current through a rectifier into a capacitor bank tends to be very peaked unless a smoothing inductor is used. I used powdered iron toroids, the Amidon T520-26, with 200 turns of 20 ga enameled wire on each one, for the inductor. One toroid had an inductance of 14 mH. Four in series were used to get a total of 56 mH.

The capacitor bank was formed of 16 electrolytic capacitors rated at 1400 μF and 450 V. Four were placed in series to get a string rated at 350 μF and 1800 V. Two strings were then paralleled to get a rated capacitor C_3 of 700 μF and 1800 V. Two more strings were added to get another capacitor C_3 for the negative supply. The figure shows the measured value of the capacitors, 740 μF rather than 700 μF .

Finally, resistors R_1 were added to discharge the capacitor bank. These are 35 k Ω , 100 W resistors, two in series on each side of the supply. At full voltage of ± 1700 V, the power dissipation on each side is

$$P = \frac{V^2}{R} = \frac{(1700)^2}{70000} = 41 \text{ W} \quad (1)$$

or about 20 W per resistor. Normally, I do not exceed ± 1200 V, or a power of about 10 W per resistor. The resistors are mounted at an angle in free air, so air can flow by convection through the interior of the resistors. No evidence of overheating (discoloration, odor, smoke) has been noticed.

The nominal rated voltage difference between $V-$ and $V+$ is $2400\sqrt{2} = 3394$ V, or about ± 1700 V. It appears the maximum voltage available under load at the inverter output will be somewhat less than this amount, in the range of ± 1500 to ± 1600 V. It turns out that ± 1200 V is adequate to get interesting sparks, so most of my testing is done at that level or below.

Power was run into and out of the capacitor bank with wire made by the Alpha Wire Corporation, No. 1638, 18 ga, with 0.090 inch rubber insulation. There are three wires, one each for $V+$, $V-$, and neutral. The transformer provides electrical isolation, so this neutral is not the same as the utility neutral. The neutral is connected to a buried ground plane through a current sense resistor (shown on a later figure) so it is not precisely at ground potential. The neutral must be kept electrically isolated from earth and from any connection that might yield a ground loop.

2 Gate Driver And Inverter

There have been several generations of driver and inverter. One early generation used two power MOSFETs, the IRF450, rated at 500 V and 13 A, in a half bridge configuration. The 240 VAC line and a variac were used to drive the rectifier without a step-up transformer. The maximum rms voltage at the output of the MOSFETs was about 170 V in this configuration. (Note that a square wave of amplitude ± 170 V has an rms value of 170 V). I could certainly measure input impedance of the Tesla coil under nonspark conditions with this voltage, but could not do much in the way of producing sparks.

A following generation had a 3:1 high frequency autotransformer at the MOSFET output, which raised the voltage to $3(170) = 510$ V. This allowed me to produce sparks up to 10 inches long. Most of the time, I had to put a small bump on the toroid to get breakout. I could now measure input impedance under sparking conditions, but I did not have enough voltage range to see how impedance varied with voltage. I could not look at spark length as a function of voltage either. The leakage inductance of the autotransformer, only a few μH , also caused some ringing problems with stray capacitances in the circuit.

The last generation is shown in Fig. 2. I needed more voltage (and power) than was readily available with MOSFETs so I decided to try IGBTs (Insulated Gate Bipolar Transistors) in series. There are four Harris HGTG18N120BND IGBTs in series in each leg. These are nominally rated at 1200 V and 18 A. The 1200 V is a rating that we *never* want to exceed. However, the device will switch up to 100 A until losses cause it to overheat. In a Tesla coil application where we operate in single-shot mode or in a low duty cycle, say, 10 ms on and 1 sec off, the device should function with rms currents up to 50 or 60 A. At this writing, the inverter has withstood 30 A rms for short periods without difficulty.

When one string of four IGBTs is on, the other string sees the entire voltage, plus to minus, of the capacitor bank. If we were operating at the capacitor rating of ± 1700 V, the four IGBTs must withstand $(2)(1700) = 3400$ V, or 950 V per device. Voltages will never divide evenly, and there will always be some overshoot or ringing due to stray inductance, so this is about the limit of a safe design.

Signals to turn on a string of IGBTs arrive from the controller by way of the fast fuse (which includes the CD4073, a three-input AND gate). The $42\ \Omega$ resistor at the left was

selected experimentally to get a fast rise time with minimal overshoot.

There are differences among components which will cause the IGBTs in a string to turn on and off at slightly different times. This causes the voltages to not be shared equally. Even a few ns makes a significant difference. At one time I tried to deal with the turn-on and turn-off times separately. I put 1N4151 diodes across the 1 k Ω pots to allow all IGBTs to be turned on as fast as possible. The 1 k Ω pots then allowed the turn-off times to be adjusted between IGBTs. I later removed them when I could not see much benefit. Without the diodes, the 1 k Ω pots slow down both the turn-on and turn-off when the resistance is increased.

The 4420 is a non-inverting driver for gates of MOSFETs and IGBTs. It can supply up to 6 A while charging a capacitive load. Each gate is charged through a pulse transformer. I used a ferrite toroid with 5 turns of 20 ga. magnet wire for the primary and 10 turns for the secondary. I ordered the Philips TX22/14/6.4-3F3 (former part number 846XT250-3F3) but I think I received the TX22/14/6.4-3E2A. Most ferrite toroids I have seen have no markings on them whatsoever. You have to keep the toroids in a container with the description. However, Philips coats the different ferrite materials with different colored epoxy paints. The 3F3 material is painted orange while the 3E2A material is painted white, and I got white toroids. The materials are not vastly different, and I suspect either one would work in this application.

The numbers 22, 14, and 6.4 in the part number refer to the outside diameter in mm, the inside diameter, and the height or thickness of the toroid, respectively. The windings need to be physically separated enough to withstand up to 1700 V, so I put the primary on one side of the toroid and the secondary on the other side. The leakage inductance would be smaller if the windings are on top on one another, but then I would need to place additional insulation between windings. This can be a challenge in relatively small toroids, so I opted for air insulation between totally separate windings.

The gate needs at least 10 V to turn on, and 12 to 15 V is better, to assure some noise immunity. The gate voltage should not exceed ± 20 V on a repetitive basis, and the absolute maximum is ± 30 V. When the first pulse is applied, the capacitor C_7 is discharged, and its voltage does not change instantaneously. If 15 V is applied to the left capacitor terminal, the right side of the capacitor will also try to have 15 V on it. However, the two back-to-back 1N4737 zeners will clamp the voltage at about 8 V. There will then be 7 V across the impedance of the 4420, the resistance of circuit traces, and C_7 . The pulse transformer will put out about double the input voltage or 16 V. If anything else is happening, the bidirectional zener, P6KE18, clamps the gate voltage at 18 V.

While current is flowing out the pulse transformer secondary to the IGBT gate and the bidirectional zener, it is also flowing into C_7 . With a pulse train of approximately 50% duty cycle, C_7 will charge to 7.5 V after the first few cycles. Once C_7 is charged, voltage excursions of +15 to 0 at the 4420 output will be seen as a pulse train of ± 7.5 V at the pulse transformer input. The gate would then see ± 15 V. The bidirectional zener and the 1N4737 zeners will not conduct after the first few cycles.

There are many tradeoffs in the gate driver circuit. C_7 and the leakage inductance form a

resonant circuit which is excited by the first pulse. As energy trades back and forth between C_7 and the inductance, the gate voltage can drop below the threshold value, causing the IGBT to try to turn off when it should be on. This effect seems to be helped by keeping the leakage inductance small, that is, by having as few turns as possible on the pulse transformer. On the other hand, the time that a pulse transformer can hold a voltage is proportional to the number of turns, so for low frequency operation, one needs many turns.

One should not use ceramic capacitors for C_7 . Some of them change capacitance value with temperature, and the initial pulse will heat the capacitor enough to cause a thermal resonance in the gate voltage in the 3 to 10 kHz range. A polyethylene capacitor rated at 50 V was much better than a ceramic capacitor.

If it were not for the initial transient, design of the driver would be very straightforward. I could not find a book or paper which discussed the design of this particular configuration. Modeling would be difficult because the current and voltage output of the 4420 under overload conditions the first pulse would be strongly dependent on the exact construction of the printed circuit board. The leakage inductance of the pulse transformer would vary with the specific toroid and also with the hand winding technique. I did tests with various values of C_7 and R_9 , and basically quit when I got something to work. There could easily be much better designs than this one.

Circuit board layout of the drivers and IGBTs is also critical to success. The system needs to be as compact as possible, consistent with adequate heat sinking and voltage separations. The three dimensional aspects are difficult to present in two dimensional drawings, so I will attempt a verbal description. All the components shown in Fig. 2 are mounted on a double sided PC board, approximately 6 by 6 inches. Four IGBTs are mounted along one edge, and their 4420 drivers toward the center. The other drivers are located the other side of center and the other IGBTs on the opposite edge. Drivers and IGBTs are on the ground side of the PC board, with their pins extending through to the trace side. The ground foil is only on the interior of the board, under and around the drivers. The IGBTs are mounted on separate heat sinks and placed on the board where the ground foil has been removed. The heat sinks are finned black aluminum, approximately 1 by 1.6 by 2 inches outside dimensions. There are two pins about an inch apart that go through the board along with the IGBT pins to provide some mechanical stability.

Conventional two-sided board practices are used in the interior of the board. One gets to ground by simply going through a via. The board becomes a single sided board when we get to the outside, where the heat sinks and IGBTs are located. The board itself will easily handle the maximum of ± 1700 V in this application, but holes through the board could compromise this integrity. So the board in the heat sink region is treated as a surface mount board. Components like C_7 , R_9 , D_2 , and the pulse transformer are soldered to pads on the board rather than placed in holes drilled through the board and then soldered. There are obviously many ways to lay out a board, and all designers have their own preferences, but this particular layout seems like a reasonable approach.

The IGBTs are mounted on their heat sinks without an insulating pad. This maximizes the heat transfer, but with the disadvantage that the heat sinks are electrically charged. I think one set of IGBTs was destroyed due to a discharge from a heat sink to the grounded case. I now set the PC board, oriented with heat sinks down, on a 1/16 inch thick layer of polyethylene lying on the bottom of the case.

Next to the MOSFET symbol in Fig. 2 is an upward directed diode. This is built into these particular IGBTs. These diodes are an important part of the circuit. They allow for current to continue flowing after one set of IGBTs turns off and before the other set turns on. They also allow for energy stored in the coil to flow back into the capacitors when the controller is turned off and all gate pulses cease. The waveforms of voltage and current look almost the same after the controller is turned off, except for a 180° phase shift in voltage. While the gate pulses are applied, power flows from IGBTs to Tesla coil. When pulses are removed, power flows from coil to IGBTs.

During inverter operation, voltages divide across the IGBTs according to the integral capacitances from drain to source. Adding some capacitance helps to swamp out individual differences among IGBTs. The energy stored on these capacitors has to be dumped through the IGBTs when they are turned on, so the capacitors need to be kept small. At one time, I put 270 pF across each IGBT, but eventually removed them, except for the case where one IGBT takes more than its share of voltage. This means that its integral capacitance is smaller than the others. In such a case, the voltage across the IGBT that is taking more than its share can be reduced by adding some capacitance in parallel to that IGBT, say 50 or 100 pF.

A capacitor rating of 1000 V should be adequate. A rating of 1000 V is less than the IGBT rating of 1200 V, however, the capacitors have more derating built in than the IGBTs. The IGBTs will be failing quickly at 1300 V, while 1000 V capacitors would probably survive in this application at 2000 V or more.

Static voltage balance during long off periods is achieved by resistors R_{11} . Without them, the drain-to-source capacitances will charge to random values from the small (and different) leakage currents of the IGBTs. A 430 k Ω resistor has a much greater current than the IGBT leakage current, so static voltage balance is quite good. Each resistor carries 1/8 of the total supply voltage during off periods. During inverter operation, each resistor carries 1/4 of the total supply voltage while its IGBT is off, and 0 while its IGBT is on. The power dissipation for a total voltage of 3400 V is




Figure 2: Gate Drive and IGBT Circuit

$$P = \frac{V^2}{R} = \frac{(3400/8)^2}{430000} = 0.42 \text{ W} \quad (2)$$

during the off period, and double this amount, or 0.84 W during inverter operation. I use 2 W resistors and have not noticed any problem with them.

Shown to the right of the 430 k Ω resistors is a long string of 1 M Ω resistors and 180 pF capacitors. This is a voltage divider which outputs one tenth of the input voltage when a X10 scope probe is attached. My scope can measure up to ± 200 V at the tip of its X10 probe, so with this divider I can measure up to ± 2000 V at the input to the Tesla coil.

The divider is a resistive divider at dc and low frequencies, and a capacitive divider at higher frequencies. At 160 kHz, for example, the capacitive reactance of a 180 pF capacitor is about 5500 Ω , much less than the 1 M Ω resistors. If an rms voltage of 1700 V was applied to this capacitor chain, the current would be about 18 mA and the total reactive power would be about 31 VA. In many applications this would be too great a load for the source, but for this situation it should be acceptable.

I bought a X100 probe before I built the divider, hoping the divider would be unnecessary. However, when I received the probe and looked at the specifications, I realized it would not work because it was rated for only 450 Vac.

3 Current Sense Resistors

In the metal case containing the inverter, there are several other components besides those shown in Fig. 2. These include the current sense resistors, a filter for the current waveform to the scope, a small current transformer for supplying current information to the controller, and some high frequency capacitors for voltage and current support. These are shown in Fig. 3. The switches SW1-SW8 represent the 8 IGBTs in the inverter.

When the IGBT string (SW1-SW4 or SW5-SW8) turns on, the circuit needs to be able to supply tens of amperes in a time of 2 or 3 μ s, if the Tesla coil is operating in the range of 200 kHz. There is no way that electrolytic capacitors shown in Fig. 1 can do this, so additional high frequency capacitors are installed close to the IGBTs to support the voltage. I used two strings of capacitors on each side, one string made of 15 polyester capacitors rated at 0.22 μ F and 1250 V, connected in parallel, and the other string made of 15 polyethylene capacitors rated at 0.039 μ F and 1600 V, likewise connected in parallel. The result was approximately 3.9 μ F on each leg of the power supply. The polyester capacitors will be operated above rating at high variac settings but this should be within the safety factor of the capacitors. Polyester capacitors are more lossy than polyethylene capacitors, so continuous operation at 1250 V and a high ripple factor could cause these capacitors to overheat. Like several other components, we are hopeful that a low duty cycle will keep these capacitors from overheating.

The switcher provides a square wave voltage to the Tesla coil input. A square wave can

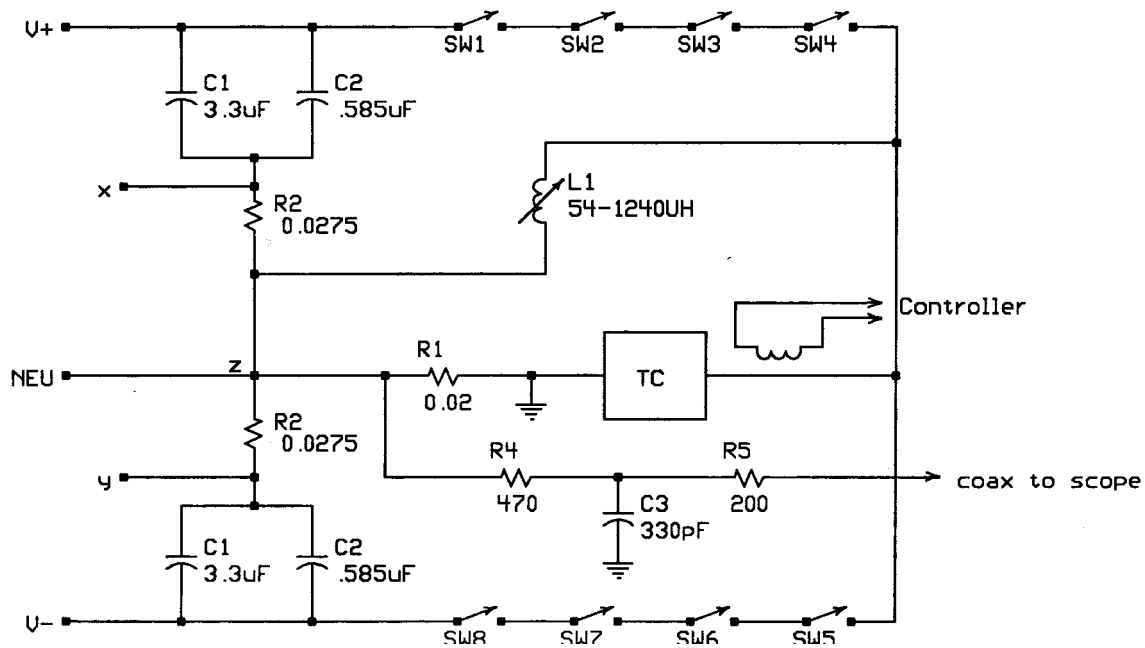


Figure 3: Support Functions for IGBT Circuit

always be composed into a fundamental and a series of odd harmonics of sine waves. The magnifier is not resonant at exactly three (or five, or seven) times the fundamental frequency, so the harmonics always face a very high surge impedance. The current will build up at the resonant frequency but not at the harmonics. This means that a square wave of applied voltage will produce only a sine wave of current. This sine wave will be in phase with the voltage at resonance, will lag above resonance, and lead below resonance.

IGBT switching at resonance means the IGBTs turn on and off when the current through them is small. This is a desirable feature, in that it reduces the losses in the IGBTs. It cannot be too small, however.

It is often not emphasized in power electronics books, but power MOSFETs and IGBTs in this circuit configuration essentially require an inductive load. With a lagging current, when one string of IGBTs is turned off, current will continue to flow, charging and discharging the drain-to-source capacitances, until the voltage across the opposite string has been reduced to zero. If the voltage reaches zero while the devices are off, current will start to flow through the diodes of the other string. When a gate pulse arrives at those IGBTs, the devices already have zero voltage across them, so the turn-on process works very smoothly. On the other hand, a leading current will result in each string turning on with full voltage across them. Very large transient currents will flow, with steep rise times. Stray inductances will cause substantial ringing in voltage due to these transient currents. A leading current is a disaster on its way to happen in an inverter circuit like this.

We need to be able to tune the Tesla coil across resonance. One way to assure the IGBTs are always operating into a lagging load, even if the Tesla coil happens to be drawing a leading current, is to add an additional inductance in parallel with the Tesla coil. This is shown as L_1 in Fig. 3. This inductor current also provides for relatively soft switching of the IGBTs when the Tesla coil current is very small.

The difference between hard and soft switching is shown in Fig. 4. We have two plots of an (approximately) square wave of voltage applied to a Tesla coil, with the resulting (approximately) sinusoidal current. In the Dec 19 plot, the voltage takes about 300 ns to make the transition. There is little ripple and the current waveform is reasonably smooth. In the Sep 30 plot, however, the voltage transition lasts only about 100 ns. This means that the voltage waveform contains more high frequency components, which can excite various resonances in the Tesla coil itself, the transmission line, and the IGBT driver. There is some ripple on the voltage and more on the current. The Dec 19 plot is obviously the preferred method of operation.


There are two necessary conditions to get soft switching. The current flowing through one bank of IGBTs must be sufficient at the time of turn-off, and the dead band (the time lapse before the other bank of IGBTs turn on) must be adequate. The controller chip (discussed later) can have a dead band as low as 80 ns, but is set for about 400 ns in the Dec 19 plot.

Assume in Fig. 3 that string SW5-SW8 have been on, the Tesla coil current is small, and inductor current is flowing to the left through SW5-SW8. The Tesla coil voltage is $V-$. IGBTs SW5-SW8 are now turned off. Current continues to flow to the left, charging the drain-to-source capacitances of SW5-SW8. A similar current flows to the left in IGBTs SW1-SW4, which has the effect of discharging their drain-to-source capacitances. An (approximately) constant current flowing into (approximately) constant capacitances will cause the voltage to slew at an (approximately) constant rate.

If the current is too small, the voltage will not have reached $V+$ when SW1-SW4 turn on. The remainder of the transition will be very rapid, with associated ripple. On the other hand, if the current is too large, the entire transition will also be rapid, with similar results. Actually, the Sep 30 plot looks like a case where the inductor current is larger than necessary. When an IGBT string is turned on by the gates during a voltage transition, the remaining transition will be nearly instantaneous.

There is therefore a range of inductance values that will yield acceptable results. For these IGBTs operating at 150 kHz, 122 μH and 210 μH worked. A value of 317 μH would give even nicer waveforms but tuning was very critical. Any small adjustment in phase would yield really ugly waveforms.

Note that the peak inductor current is the critical factor, not the Tesla coil impedance. At higher frequencies, the peak current in a given inductor will decrease, requiring the use of a smaller inductor. If 200 μH works at 150 kHz, then we would expect a 300 kHz coil to need a 100 μH inductor.



HP54645D System A.02.08 13:07:19 Sat Sep 30, 2000




Figure 4: Soft switching (Dec 19), harder switching (Sep 30).

A coil wound on a ferrite toroid would be an attractive solution to the need for an inductor except for two problems: The number of turns would have to be large in order to avoid saturation, and we would need to build several such inductors, swapping them in and out of the circuit until we found one with approximately the proper inductance. I decided to build a variable inductor instead. I wound 140 turns of 22 gauge yellow insulated hookup wire on a length of 3/4" CPVC pipe. The pipe O.D. is approximately 0.9". The winding length is about 8.4". I then placed three pieces of Amidon R33-050-300 ferrite rods inside a length of 1/2" CPVC pipe. The rods are 0.5" O.D. and 3" long. The O.D. of the rods and the I.D. of the pipe are the same, so the rods fit snugly. Inductance is then varied by sliding the smaller pipe inside the larger one. Minimum inductance is about 54 μH and the maximum is about 1240 μH .

Returning to our discussion of Fig. 3, the current to the Tesla coil is measured as the voltage across R_1 , a 0.02 Ω resistor. R_1 is fabricated from 11 surface mount resistors, 0.22 Ω and 2 W each. They are mounted in parallel across a gap in two sections of 2 inch wide copper foil. Where possible, the ground path is made of wide conductors to lower the stray inductance of the circuit.

When a voltage pulse train of resonant frequency is applied to the Tesla coil, current will build up until the coil contains enough energy to drive the spark. The current amplitude at spark remains relatively constant as input voltage is increased further. A ± 1200 V input will have about the same maximum current as a ± 900 V input. A given value of current will be reached more quickly with a higher voltage of course. A spark that occurs 1.9 ms into the pulse train with ± 900 V applied might occur 1.2 ms after start with ± 1200 V applied. A current of about 30 A (rms) just before spark was observed frequently. The power dissipated in R_1 with this current in it is

$$P = I^2 R = (30)^2 (0.02) = 18 \text{ W} \quad (3)$$

which is within the range of a 22 W resistor. The most critical limit in the circuit is that of the IGBTs, which is 100 A peak in short bursts. If the peak current is 100 A in a sinusoidal waveform, then the rms current limit is 70 A. At 70 A, the power dissipation in R_1 is

$$P = (70)^2 (0.02) = 98 \text{ W} \quad (4)$$

which is a little over four times the steady state rating. For the intermittent operation used here, R_1 should last indefinitely, at least until the IGBTs blow. If the IGBTs fail in a shorted state, then the stored energy in the capacitor bank is available to turn the 0.22 Ω resistors, traces on the printed circuit board, and anything else in the way, into fuses. An IGBT that fails shorted will eventually be blown apart and become an open circuit. A metal cover over the inverter is important at high power testing to keep shrapnel from flying around the room. Hearing protection is not a bad idea either.

A current transformer is mounted on the 'hot' lead of the Tesla coil, opposite R_1 . I made it from two stacked ferrite toroids, the TX22/14/6.4-3E2A mentioned earlier. The toroids fit

over a straight 14 ga wire connecting the inverter to the transmission line. This wire forms a single turn primary. There are 25 turns of 20 ga magnet wire wound on the toroids. If the two ends of this winding were shorted, and if everything were properly designed, a current of 25 A to the Tesla coil should result in a 1 A current in this winding.

We already have a means of measuring Tesla coil current, with R_1 . What is needed is a voltage waveform at the controller that looks just like the current waveform, except cleaner and with sharper transitions. That is, we want to convert the sine wave in Fig. 4 into a square wave. That current wave has a peak value on the order of 100 mV, far too low to clean up with diodes. We need a square wave with amplitude of a volt or two.

A current transformer is always operated with a short across the secondary for two reasons. One is linearity, which we are not particularly interested in here. The other is that if the secondary is open circuited, the secondary voltage can rise to extremely high values, enough to damage the insulation in current transformers built for 60 Hz service. The second reason is used to advantage here. The secondary is loaded with two chains of two diodes each, the chains pointed in opposite directions. As the Tesla coil current passes through zero, the secondary voltage also goes through zero but very rapidly. Two diodes in series have a very high impedance when the voltage across each diode is less than 0.5 V or so. Then while the Tesla coil current is still near zero, the secondary voltage will be clamped at two diode drops, say 1.4 V. The secondary voltage is not a perfect flat top square wave, but will reach 1 V for a Tesla coil current of hundreds of mA, and then rise slowly to 1.4 V or a little more for Tesla coil currents of tens of A. This approach works surprisingly well for the purpose.

There are two additional resistors in Fig. 3, labeled R_2 , that are used for overcurrent protection discussed in the next section. These resistors are made of eight $0.22\ \Omega$ 2 W resistors in parallel. When SW1-4 are on, current flows to the left through R_1 and up through the R_2 attached to x . When SW5-8 are on, current flows to the right through R_1 and up through the lower R_2 . With these current directions, point x will be negative with respect to ground while point y will be positive.

As mentioned, good sparks are obtained with a Tesla coil current on the order of 30 A, while the IGBTs are rated for 100 A in short bursts. We will select 70 A as the overcurrent set point. This will hopefully be high enough to prevent nuisance trips and low enough to protect the IGBTs. At 70 A, the voltage at y will be $(70)(0.0475) = 3.32$ V, and a -3.32 V at x .

4 IGBT Overcurrent Protection

After one explosion of my inverter, I decided to add overcurrent protection. It is well known in the power electronics business that silicon is a faster acting fuse than any traditional device sold as a fuse. That is, a ten dollar IGBT will always act to protect a ten cent fuse. At least, it seems that way. My only hope seemed to be to remove the gate pulses to the IGBTs as

rapidly as possible.

The circuit used is shown in Fig. 5. Gate pulses from the controller (discussed in the next section) are brought to the inverter board through three input AND gates, the 4073. Current through IGBTs 1-4 is measured separately from current through IGBTs 5-8. If either current exceeds a pre-set limit, one input to each three input AND gate will go low, interrupting all gate pulses immediately.

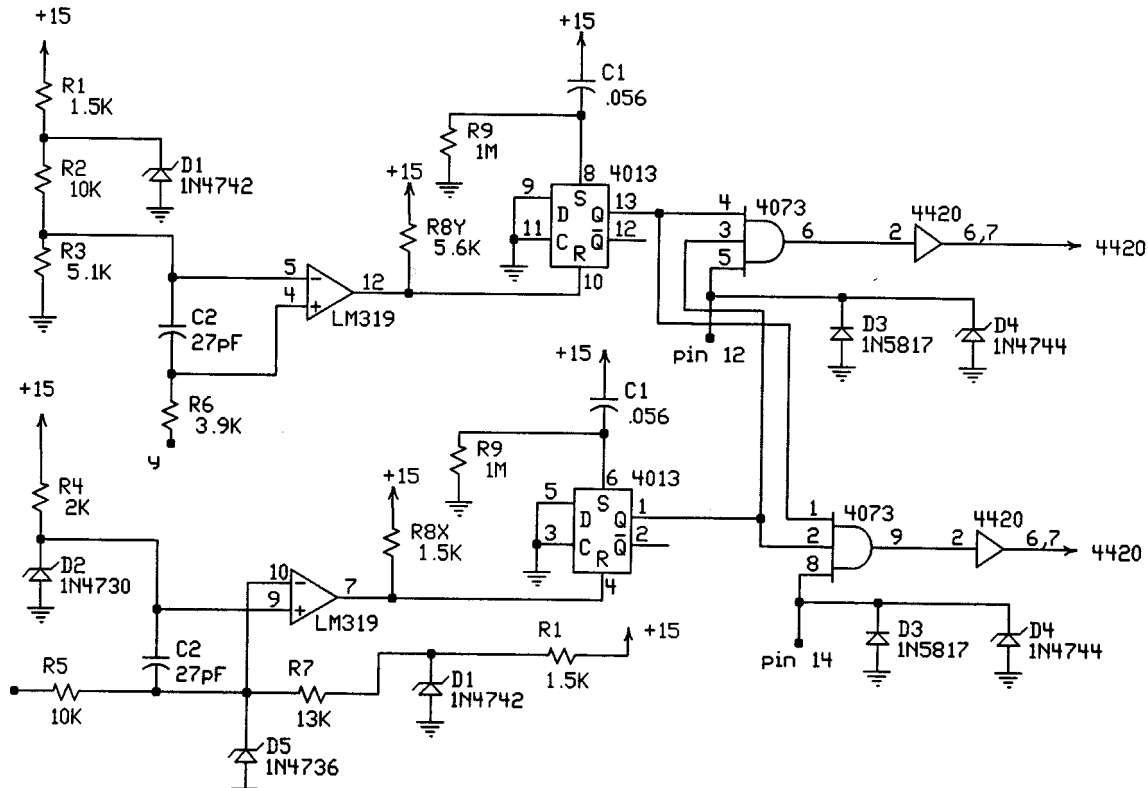


Figure 5: IGBT Overcurrent Protection

Starting at the top left of Fig. 5, we see a voltage divider chain, R_1 , R_2 , and R_3 . Zener D_1 sets the top of R_2 at 12 V. During fault events and even during normal operation, the +15V line is likely to be noisy, so this adds some noise suppression. Pin 5 of the LM319 comparator will be held at about 1/3 of 12 V, or about 4 V. When point y exceeds 4 V, the comparator output goes high, resetting the 4013 and causing its output at pin 13 to go low.

The negative going voltage at x causes more of a problem for a system that uses only +15 V and ground. Instead of setting the inverting input at a fixed value and waiting for the noninverting input to go higher than the inverting input, we set the noninverting input and wait for the inverting input to go lower. D_2 is a nominal 3.9 V zener, so pin 9 is held at this value. Pin 10 is held to a value not over 6.8 V, the rating of D_5 . As x goes below zero, the

voltage at pin 10 drops, making D_5 effectively into an open circuit. The voltage on pin 10 is then determined by voltage division across R_7 and R_5 .

Resistors R_{8Y} and R_{8X} are pullup resistors required by this comparator. C_2 provides a small amount of noise suppression. C_1 and R_9 provide an effective short pulse at startup to make sure the 4013 flipflops wake up with output Q high. D_3 and D_4 help to keep what should be a circuit trace with a digital signal of 0 to +15 V between -0.3 and about +15 V in a noisy environment.

When this circuit acts, the +15 V supply must be turned off and back on to reset the fuse action.

5 Tesla Coil Controller

The controller described in Fig. 6 has taken an amazing amount of time to get to its present state. I have worked on it intermittently for ten years, visiting with many students and faculty along the way, before getting something that basically does the job. It still seems that it should not be that hard. I fully expect some reader to glance at the circuit, get the basic concept, and redesign to achieve something far superior, in a short period of time. But that is the way of science.

The heart of the circuit is the Motorola 34066. It puts out complementary gate drive pulses on pins 12 and 14 at a frequency proportional to the current into pin 3. The gate pulses have a minimum dead band of about 80 ns (so both strings of solid state switches cannot be on at the same time). The dead band can be increased by the resistor R_{DT} . Initially, I used $R_{DT} = 0$ but eventually I put in a 1 k Ω pot for R_{DT} . I would adjust this pot to get minimum transition noise in the current flow to the Tesla coil, which probably occurred at 200-300 ns. A dead band of up to 400 ns will allow for somewhat more variation in other parameters.

The 34066 is turned on by an ENABLE signal at pin 9. The ENABLE signal is supplied in one of two ways: a pulse from a momentary contact switch SW1 or a pulse from an external pulse generator to the port marked BNC. The switch SW1 sets the 4013 flipflop so the 34066 is on until power is removed. The BNC port allows the controller to be turned on and off for specified periods as determined by the pulse generator. With adequate voltage applied, a spark from the Tesla coil would occur in typically 0.5 to 3 ms, so the on period was typically set at 4 to 6 ms. The off period would then be set at 1 second, so we would have 1 pps (pulse per second). On occasion, I would adjust that up to 10 pps, to observe the effect.

It is important that the 4013 flipflop U1A between switch SW1 and the 34066 always wakes up with Q in the low state. Otherwise the controller may start operation immediately after turn on, with possible negative consequences. C_6 keeps the Set low and C_5 and R_{46} keep the Reset of the 4013 high during turn on, so the chip always starts in the right mode.

One of the 34066 outputs (pin 12) is used for a clock pulse into U1B, the other half of the 4013. The data pin is always high, so Q is set high every clock pulse if ENABLE is high. If ENABLE is low, Reset is always high, which holds Q low.

There are two modes of operation, manual or automatic, depending on SW2. In manual, the current into the 34066 pin 3 is determined by the settings of R_{49} and R_{51} . Each mode has two submodes, depending on whether ENABLE is on all the time or just when the pulse generator sends a pulse. Normally in manual, SW1 is also flipped, so we have continuous operation. Current and voltage to the Tesla coil are monitored on a scope, and R_{51} is adjusted to give current in phase with the voltage. This is a resonant point, so current is maximum when it is in phase with the voltage.

The manual mode is not suitable for long term operation. Minor differences in power supply voltage, heating effects in R_{49} and R_{51} , drift in the 34066, the movement of a person watching the sparks, or the change in capacitance due to the space charge of a spark, will cause the frequency to shift by a few hundred hertz. For a high Q Tesla coil, this small shift in frequency can reduce the current by a factor of two, or more. We therefore need a circuit that will monitor the current and automatically adjust the input to pin 3 of the 34066 to keep the Tesla coil current in phase with the voltage. It was this task that turned out to be a challenge.

An important task is to get a voltage waveform into the controller that is in phase with the current. At first I used two Rogowski coils next to the ground lead from the IGBTs. These consisted of many turns of 30 ga wire wrapped around a rectangular block of polyethylene. Each block was 0.5 inch thick by 1 inch wide. The winding length of the coil was about 2 inches. The ground lead is a three inch wide piece of copper foil at this point of the circuit. One coil was pressed against the foil from the bottom, while the other coil was just laid on the foil from above. The two coils were connected in series. Measured inductance of the coils in place was about 400 μH . Since the polyethylene is nonmagnetic, one did not have to worry about saturation and other magnetic effects.

The current signal from the Rogowski coils was brought into the controller enclosure and onto the printed circuit board at the point marked i1, top left corner of Fig. 6. Noise was filtered by R_1 , R_2 , and C_1 , then the signal was limited by the two diodes D_1 and D_2 . Above some minimum current, the voltage across these two diodes would start to look like a square wave, except with sloped sides and rounded shoulders.

The problem with this technique was that the signal looked too much like the current. During buildup, the signal would look sinusoidal, with a peak value well under one diode drop. As the current increased, the slope of the sides of the signal would also increase, changing the apparent phase detected by the following circuit by an excessive amount. The Rogowski coils also had their own resonant frequencies in the range of a few hundred kHz, which would tend to amplify the noise at those frequencies.

We therefore went to a small current transformer that was described earlier. This current transformer, when terminated by diodes, acts more like a comparator in that a few mV swing

on the input would result in maximum swing on the output. It presents a fairly good square wave to port i1 starting with the first pulse to the Tesla coil, many pulses before any current is visible on a scope.

The next step is to convert this small analog signal (plus and minus one diode drop) to a clean digital signal, +15 to zero, a nontrivial task. The solution is clever, even elegant. I wish I could say it was original, but I saw it in an application note. Two pins of the 4046 phase-lock loop are used, one for input and the other for output. There is no feedback. The 4046 converts the analog into a digital signal very nicely. It has an internal biasing network so it will work if a capacitor C3 is used to allow a dc offset. It was my experience that the duty cycle of the resulting square wave was not close enough to 50%, so I added an external bias network R17, R18, and R19. R18 needs to be adjusted so the output at pin 2 of the 4046 is near a 50% duty cycle. This is only done once, in initial testing of the circuit.

The current signal is passed through an LM319 comparator and a 4049 inverter to provide additional buffering. R_{14} and C_{12} provide an adjustable amount of phase delay. This delay allows the operator to tune through resonance a significant amount. Without the adjustable delay, one might just be able to tune to resonance, or even not quite to resonance. There is also the option of using another 4049 inverter to get another 180° of delay. One could accomplish the same function by reversing the leads from the current transformer.

U1B, the 4013 flipflop, now has the voltage waveform (actually the gate signal to IGBTs 1-4) applied to the Clock and the current waveform applied to the Reset. Q is set high by the voltage and low by the current. While Q is high, the switch U2C, a 4066, is turned on, causing the capacitor C_{11} to discharge a little through resistor R_{20} . When Q is low, C_{11} charges through the resistors R_9 and R_{35} . One looks at Tesla coil voltage and current on the scope and adjusts R_{20} until the current is approximately in phase with the voltage and has acceptable smoothness. If the current looks like the Sep 30 plot in Fig. 4, I would look for ways to smooth it out.

At resonance, the duty cycle on Q should be on the order of 0.3 to 0.5. As the frequency drifts below resonance, the duty cycle decreases, R_{20} is not connected as much, the voltage on C_{11} increases, increasing the current to pin 3 of the 34066 and thus increasing the frequency. The opposite happens when the frequency drifts above resonance. If the duty cycle at resonance is 0.4, things need to work over a duty cycle range of say 0.2 to 0.6. If we start off with a resonance duty cycle of 0.1 or 0.9, we just do not have the tuning range that we need.

The voltage on C_{11} is buffered through U6, a LM324 op amp in a voltage follower configuration. This is an old, slow op amp, but great speed is not needed in this application. A precision op amp would probably be better, but a quick look failed to turn up a single supply op amp that met all the other requirements. I did not want to add a -15 V supply to the printed circuit board just for the op amp if this single supply version would do the job, and it does reasonably well.

There are some patches around the op amp that I threw in to get it to work in a minimum fashion. One challenge was that the firing of the IGBTs would shift under load. For example,

the controller might be carefully tuned so current was exactly in phase with voltage at an applied voltage of say ± 100 V. The variac setting would be increased so the applied voltage became ± 200 V. The current would now be lagging the voltage by perhaps 10° . The controller was sending out a gate drive signal at the same phase point with respect to the current as before, but the lag between the gate signal and the zero crossing of the voltage waveform to the Tesla coil was changing with current. What was needed was for the current into C_{11} to decrease slightly as Tesla coil current increased, similar to a small adjustment on R_{20} . But it needed to be dynamic, during the growth of Tesla coil current under spark conditions.

The patch was the two op amps at the top of the figure. The voltage across diodes D_1 and D_2 will increase slightly with current. Diode D_{13} pulls off the negative peak, so the junction between D_{13} and R_3 becomes slightly more negative as Tesla coil current increases. The pin 7 output of the inverting op amp will become slightly more positive.

D_4 is a 7.5 V zener that limits the allowable voltage to the op amp and to the 34066.

In one-shot mode, ENABLE is off until the spark is desired. The ENABLE signal is ran through D_{10} , R_{34} , an LM 319, and a 4049 inverter to switch U2D, which forces the voltage on C_{11} to be that set on R_{11} when the switch is on. When ENABLE is turned on, the switch turns off in a time determined by the time constant of R_{34} and C_4 . This causes the circuit to always start hunting for resonance from the right starting point. This has to be retuned for each different coil or other adjustment which changes the resonance frequency.

Actually life is even more complicated. I would occasionally see a problem with hunting or oscillation about the resonant frequency. R_{20} would be set correctly, but current would drift from a 20° lag with respect to the applied voltage to a 20° lead and back. In this high Q circuit, this is too much phase shift. The IGBTs would go into a hard switching mode, stray inductances would ring, and the current waveform would get really ugly. To fix this problem, I added the compensating capacitor C_{15} . Now the current would stay in phase with the voltage to within a few degrees. But a new problem popped up. With ENABLE off, pin 3 of the 34066 shows two pn junctions in series to the outside world. With R_{41} in place, the voltage of pin 3 is about one volt. When ENABLE turns on, pin 3 wants to accept current at 1.4 volts. So the pin 3 voltage needs to change by 0.4 volts quickly. However, the voltage across the compensating capacitor C_{15} cannot change instantly. There will therefore be a few cycles where the pin 3 voltage is too low, such that the output frequency will be too low as well.

The patch for this problem is the RC time constant formed by R_{34} and C_4 . When ENABLE goes high, pin 12 of the 4049 stays high for 10 to 20 cycles. R_{11} controls frequency while the 34066 is settling down. Hopefully, the current to the Tesla coil will build up enough during this time that when pin 12 of the 4049 goes low, frequency control is immediately assumed by a pulse train of about the right duty cycle to the 4066 switch controlling R_{20} .

Another patch is the combination of R_{34} , R_{21} , C_2 , C_9 , and the three switches U9A, U9B, and U9C connected to pin 16 of the 34066. This is the one shot control input of the 34066. Frequency and pulse width output of the 34066 are controlled by external components attached

to pins 1 and 2 (oscillator) or pin 16 (one shot). In our case, we select values so the oscillator input does the controlling and the one shot input just floats along, doing basically nothing. The only exception to this is the first two pulses out of the 34066. The IGBT gate drive circuits have capacitors that need to be charged by the first few pulses. One effect of starting operation with a discharged capacitor is that with a standard length gate pulse, the first bank of IGBTs will still be on when the second bank turns on. This forms a short across the main power supply. I lost at least one inverter to the high currents of this short. My fix was to use the one shot input to make the first pulse to each bank of IGBTs shorter than the following pulses.

The patch works in the following manner: When ENABLE goes high, C_2 charges through R_{34} . While the voltage on C_2 is still below the transition voltage of the 4049 inverter U4A, pin 2 of the 4049 will be high and the 4066 switch U9C will be on. This switch connects pin 16 of the 34066 to two variable resistor strings through two other switches, U9A and U9B. One switch is controlled by pin 12 of the 34066 and the other is controlled by pin 14. For the first pulse at pin 12, switches U9C and U9A are on, and the resistor string R_{36} and R_{37} are connected to pin 16 of the 34066. The parallel combination of the resistor string and C_9 determines the discharge time for the one shot function of the 34066. R_{37} is adjusted so the one shot turns off the pin 12 pulse at the desired time, instead of the longer time of the 34066 oscillator function, controlled by R_{45} and C_7 . R_{34} is adjusted so that after the first pulse at pins 12 and 14, switch U9C is turned off. The one shot time constant is now determined by C_9 and R_{21} , which is longer than the oscillator time constant, so this part of the circuit is no longer doing anything.

This patch actually has two functions: As mentioned earlier, slightly shorter pulses allow the gate drive capacitors to charge up without the ugly feature of having both sides of the IGBT inverter on at the same time. And the shorter pulses also improve symmetry of the inductor current. In steady state operation, the IGBTs switch when the inductor current is at a peak. The voltage across the variable inductor is approximately a square wave, so the inductor current is basically a saw tooth wave. This saw tooth should be symmetric about zero. The current for the first pulse starts at zero rather than a negative peak, so if the first pulse is the same length as the following ones, the inductor current will rise to about double the steady-state peak. The current waveform will then drift down with time as the losses in the circuit cause a change toward symmetry. The IGBT switching current is too high for half the IGBTs and too low for the other half during the first few cycles. This can easily cause ringing and IGBT destruction. The current waveform is helped immensely by shortening the first pulse until the peak inductor current is about equal to the long term peak (rather than double). R_{39} is then adjusted so the peak inductor currents during the first few cycles are not drastically different from the steady-state peaks. It is not essential that the peaks be exactly those of steady state during the first few cycles. Peak values between 70% and 130% of the long term peaks seem to work in an acceptable fashion.

It should be mentioned that driving a digital chip like the 4049 with an analog signal from an RC network is considered poor design. The RC voltage goes through the transition voltage

of the inverter slowly, so there will be an interval when the inverter output is neither high nor low. This confuses any following chips. While the control voltage to the 4066 switch is at mid range, the resistance through the switch will be in the $k\Omega$ range, which adds to the resistance of the attached string (i.e. R_{36} and R_{37}), and changes the one shot timing. A better way of doing it is with a comparitor, similar to what is shown in the lower right corner of the TCC601 schematic. But poor designs sometime work, as is the case here, and saved me from building yet another PCB with another chip on it.

Yet another patch to deal with a peculiarity of the 34066 is R_{33} , D_9 , and the pulse transformer connected to pin 10, the FAULT input. The 34066 starts immediately when ENABLE is applied, but may put out another hundred pulses after ENABLE is removed. Residual energy stored in the Tesla coil will ring down over a similar time period. If there are no gate pulses on the IGBTs, and if energy is adequate, the IGBT diodes will conduct on the half cycle as ringing occurs. This is not a problem. Power is flowing from the Tesla coil to the power supply rather than the other way around. An increase in voltage on the power supply capacitors can be noted.

If gate pulses continue during the ENABLE off period, when the frequency control has been disabled, we have the possibility of gate pulses getting out of phase such that one bank of IGBTs will turn on while the diodes of the other bank are conducting. This is not as bad as the gates of both sides being on at the same time, but still results in hard switching and massive ringing within the inverter and power supply. This ringing can destroy IGBTs being operated close to their rating. We therefore want the 34066 to turn off when ENABLE goes low.

The 34066 has a FAULT input where a 1 V pulse will shut down the output pulse train within 70 ns. So a positive pulse is needed from the negative going edge of the ENABLE pulse. I first considered using an unused inverter in the 4049. This meant that the FAULT input would be high before ENABLE turned on, and for a few ns afterwards. This would very probably mean that the 34066 would never function.

My next choice was to use a pulse transformer. I had a good supply of the 204XT250-3F3 ferrite toroids on hand, so I wound two windings of 30 ga. wire on a toroid, one about 60 turns and the other about 20 turns. The winding with 60 turns is connected to R_{33} . When ENABLE goes high, a negative pulse appears at the FAULT input, limited by the forward voltage drop across D_9 . When ENABLE goes low, a positive pulse appears at the FAULT input, limited by the zener voltage of D_9 . The 1N4730 is a 3.9 V zener. Voltage limits on the FAULT input are -1 to $+6$ V, so we should be within these limits. No effort was made to do a detailed design of the pulse transformer. Current has to be low, 1 mA in this case, to avoid overloading the 4071. Plenty of turns are needed for this low current. A step down ratio of say three to one would be appropriate, so 1 mA in the primary would result in approximately 3 mA in the secondary. My first attempt worked, with a peak voltage of about 2 V at the FAULT input, so I did not proceed further.

One other patch, R_{43} and R_{44} between pins 1 and 12 of the 34066, attempts to fix a very

subtle problem. C_{11} needs to be small in order to have a high slew rate when a spark occurs. A spark increases the effective size of the Tesla coil top load, which increases the effective secondary capacitance, which lowers the operating frequency. The controller needs to hold frequency at resonance within perhaps 100 Hz (better than 100 parts per million) before spark, and then slew at rates up to 100 Hz per μs when the spark occurs. This controller will do this, but at the expense of significant ripple on C_{11} . The current into pin 3 of the 34066, that controls frequency, will also have this ripple present.

There will be a double frequency ripple at pin 1, one hump controlling gate pulse length at pin 12 and the following hump the length at pin 14. The ripple on pin 3 will cause one of the gate pulses to be longer than the other by a few percent. The Tesla coil sees a slightly non symmetrical waveform with the time spent at $V+$ somewhat different from the time at $V-$. This is not a particular problem to the Tesla coil, but the inductor across the Tesla coil will integrate this non symmetry over the duration of the ENABLE pulse. Instead of an inductor current varying symmetrically from $+1$ to -1 A, it might vary from $+0.6$ to -1.4 A. We might have not enough current to properly switch the IGBTs at one extreme, and too much at the other extreme. The current waveform might show more noise on the positive going edge than the negative going edge, or vice versa.

A value of $1.5\text{ M}\Omega$ for $R_{43} + R_{44}$ was found empirically to balance the current into the current mirror of the 34066 and give equal length output pulses. R_{44} , a $500\text{ k}\Omega$ pot, allows a small adjustment if it seems necessary. One would look at the inductor with a current probe (Hall effect so it shows the total current, not just the ac component), and adjust the pot until the inductor current is symmetrical about zero.

6 Initial Tuning

People not building a controller board can skip this section. I include it as a personal check list. I lay out the circuit to fit on a double sided printed circuit board, 4.5 by 6 inches, using a software package (HiWire in my case). I etch and drill my own boards, but am thinking about using the experts to do this. I would certainly get a better board to work with.

After soldering on the components, but before mounting the chips in the sockets, I apply $+15$ and ground to the board. This would detect some solder bridges and some gross errors in component values. For the board TCC601, I measured 11 mA in the $+15\text{ V}$ line, which seemed reasonable. I then inserted the chips and applied $+15\text{ V}$ again. I measured 39 mA, which still seemed reasonable.

I then put the board into Manual Mode, and push the Enable button, so the board should be operating at fixed frequency. I then check that there are complementary signals at pins 12 and 14 of the 34066, that R_{DT} changes dead time, and that R_{49} and R_{51} will adjust frequency.

The variable inductor and the Tesla coil are connected to the IGBT inverter, and the variac is adjusted to some low voltage like 6% of rated. In Manual Mode the Enable button

is pushed and R_{49}/R_{51} are tuned for resonance. Heating effects on the board will probably prevent one from holding the exact resonant frequency, but we should be able to tune through resonance. With the frequency near resonance, R_{18} is adjusted so pin 2 of the 4046 has a 50% duty cycle.

With the circuit still operating, the duty cycle at pin 13 of the 4013 is checked. At resonance, the duty cycle appeared to be about 30%, and would vary from perhaps 10% to 50% as R_{49}/R_{51} are varied across resonance. If this general condition was not observed, I would try adjusting R_{14} , and failing that, would reverse the input current leads at port i1 or add (or subtract) an inverter stage between the 4046 and the 4013.

With the circuit still operating, measure the voltage at the wiper of R_{51} (2.09 V for a coil operating at 177 kHz). Set R_{11} so the wiper has the same voltage. This should start the circuit at the resonant frequency in Automatic Mode, so the current will build up during the first few cycles.

Turn the board off, switch to Automatic Mode, and apply +15 V pulses to the port marked BNC from a pulse generator in single-shot mode. Adjust R_{37} and R_{39} so the first pulse of both pin 12 and pin 14 is shorter than the following pulses. More precise adjustments will be made later.

Check that pin 12 of the 4049 goes low about 60 μ s after Enable is turned on. This allows current to build up for 10 to 15 cycles under control of R_{11} before control is changed to the Tesla coil current as monitored at port i1.

Check that pin 12 of the 34066 goes low (output pulses stop) when Enable goes low. This confirms that the fault transformer is connected correctly.

In Automatic Mode and with the variac set at about 6% we want to adjust the circuit so current will build up to its steady state value while a short pulse (1 to 5 ms) is applied to the BNC input. Set R_{20} to midrange (say 25 k Ω). Current should build up for a few cycles under control of R_{11} , and then may decline after that due to improper adjustments. Adjust R_8 until it appears current is continuing to build after R_{11} loses control. Look at the Tesla coil voltage and current and continue to adjust R_8 until current and voltage are approximately in phase. R_8 is set only once (for each coil). Phase adjustments during regular operation are made with R_{20} .

R_8 is a part of two LM324 opamp stages which help the controller maintain the proper phase relationship as the TC current builds (or declines during the discharge), and as applied voltage is changed. When the current waveform is a good sinusoid and is in phase with the voltage at steady state, the current waveform should also be a good sinusoid when the current is at 20% of the steady state value, even though the phase between voltage and current will be somewhat different. There will be some interaction between R_8 and R_{20} in producing this phase tracking. It is a bit tedious, but one could examine this interaction by starting with R_{20} at a different value, say 15 k Ω or 35 k Ω , adjusting R_8 for resonance, and comparing the current waveforms at 100% and 20% for this new setting.

For the next step, we remove the Tesla coil entirely, and drive only the variable inductor. First we look at pin 16 of the 34066 and adjust R_{34} until U9C is only on for the first two triangular pulses. This may require some adjustment of R_{37} and R_{39} . The waveform on pin 16 should show a rapid rise to 5.1 V and a linear decline to 3.6 V, followed by a constant voltage of 3.6 V, for the first two pulses. Later pulses should show a rapid rise to 5.1 V and a linear decline to perhaps 4.5 V. If pin 16 does not drop to 3.6 V, then this pin does not control the turnoff of gate pulses.

Now we hook a current probe on a lead to the inductor. Adjust R_{37} and R_{39} until the inductor current is somewhat symmetric. The first pin 14 pulse should terminate when the inductor current is about half the steady state peak-to-peak value. The first pin 12 pulse should terminate when the inductor current is about the same negative value as the previous positive value. The inductor current will not be perfectly symmetric at this point, so adjustments should continue until the inductor current is as close to being symmetric as possible. The first few peaks should be within ± 20 or $\pm 30\%$ of the steady state peaks.

7 Waveform Plots

Perhaps a few plots of voltage and current will help understanding. In the following figures, A1 and A2 refer to the two analog channels of the HP 54645D oscilloscope. A1 is the voltage waveform applied to the Tesla coil, measured at the output of the inverter. There is a X10 voltage divider in addition to the probe, so a scale of 20 V/div is actually 200 V/div. A2 is the current, as represented by the voltage across a 0.02 Ω resistor. A scale of 100 mV/div would be 5 A/div.

The plots at the bottom of the figures are five of the sixteen digital channels of the scope. Channel 0 is the clock input of the 4013, an inverted output of the 34066. Channel 1 is the output of the 4046. Channel 2 is the Reset input to the 4013. Channel 3 is the output of the 4013 and the input to the 4066 that controls frequency by controlling voltage on the LM324 op amp. Channel 4 is pin 12 of the 4049. While it is high, frequency is controlled by R_{11} .

Fig. 7 shows the various waveforms for the first 180 μ s. The first voltage pulse is shortened, as mentioned earlier. The current waveform is mostly switching noise for the first few cycles, although we start to see some current signal by the third or fourth cycle. The 4046 in Channel 1 appears to synchronize with the voltage immediately. So the controller has a decent digital signal for current even when the analog current waveform is still small and noisy during the first few cycles.

When Channel 4 goes low the controller is allowed to dynamically control frequency. The circuit is properly adjusted so this transition occurs smoothly, with the current waveform becoming more sinusoidal and continuing to grow.

Fig. 8 shows the same waveforms after 1 ms of operation (except for Channel 4, which is no longer of interest to us). The voltage is a square wave with relatively slow transitions from

HP54645D System A.02.08

13:16:48 Sat Jun 30, 2001

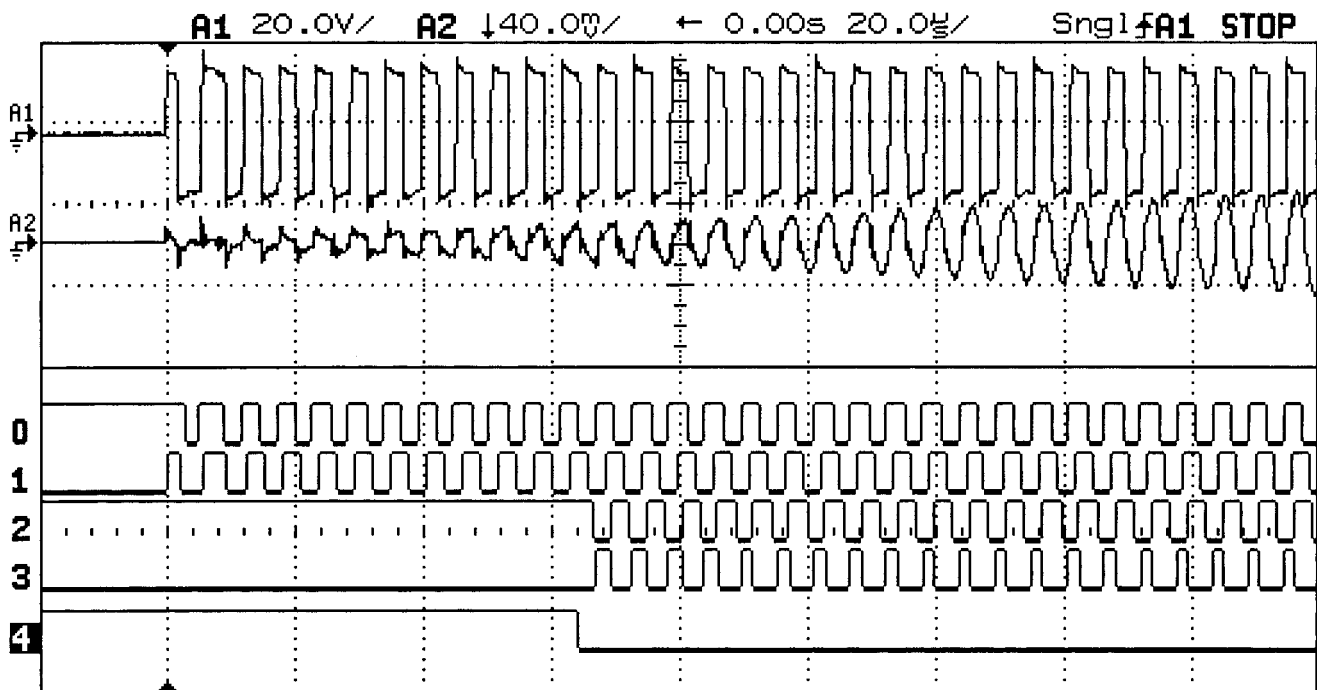


Figure 7: Controller Waveforms at Start

one rail to the other. The current is an almost perfect sine wave. If the voltage makes more rapid transitions from one rail to the other, switching noise will appear on the current. For this particular case, the current is lagging the voltage by about 7° and the duty cycle of signal on Channel 3 is about 28%, both values in a quite acceptable range.

The gate pulses from the 34066 have a duty cycle less than 50% because of the dead band. The inverted gate pulses in Channel 0 therefore has a duty cycle greater than 50%. This is not a problem in the tuning operation. The square wave representing current, Channel 1, does have a duty cycle close to 50%, however, since R18 has been adjusted to this effect. Channel 2 shows the Channel 1 signal after 2 inversions and a phase shift. The time delay between Channel 0 and Channel 2 is essential to the functioning of the circuit. If more lag is desired, R_{14} can be increased, although a better design might be to eliminate the analog delay and use the delay in two more inverter stages.

The frequency control signal, Channel 3, goes high when Channel 0 goes high, and goes low when Channel 2 goes high. If the current starts to lead the voltage, the Channel 2 waveform shifts to the left. The control signal (Channel 3) is not high as long. The voltage of the LM324 drifts up a little to bring the current back into phase.

Fig. 9 shows the voltage applied to the Tesla coil and the resulting current for about 4.5




Figure 8: Tesla Coil and Controller Waveforms after 1 ms

ms. The voltage is A1 at the top of the figure and the current is A2 at the bottom. Operation is well below breakout. The ENABLE pulse is actually about 2.1 ms long. Current builds up as predicted by either transmission line theory or lumped RLC analysis. Current builds nearly linearly at first, then the rate of increase slows as losses increase. Current will hit its limit when the input power is equal to the losses, in this case between 2 and 3 ms after start.

As current increases, the voltage drop across the IGBTs and the droop in the electrolytic capacitors becomes greater. The decrease in voltage applied to the Tesla coil is not great, but is most easily noticed when ENABLE is turned off. Energy stored in the Tesla coil has to be dissipated somehow. Instead of power flowing from inverter to coil, it now flows in the opposite direction. Voltage is constrained by the built-in diodes of the IGBTs. Power supply capacitors are now being charged instead of discharged. All the voltage drops in wiring and the IGBTs reverse in sign. We therefore see a small step increase in voltage when ENABLE is removed.

After about 0.6 ms, the stored energy is no longer able to force the IGBT diodes into forward conduction. Without a power supply affecting the circuit, we now have a classic RCL ring down. Both voltage and current are decreasing exponentially. During this portion of the cycle, the IGBTs are acting as capacitors.




Figure 9: Tesla Coil Voltage and Current Waveforms

Fig. 10 shows a closeup of the waveforms when ENABLE goes low. Both outputs of the 34066 go low, so Channel 0 goes high and stays there. This is the clock input to the 4013, which does not respond to constant voltage levels. The next reset pulse (Channel 2) sets the 4013 output low (Channel 3), where it remains. The controller is still measuring Tesla coil current, as seen by Channels 1 and 2, but is no longer controlling frequency.

Frequency is now being controlled by the Tesla coil. Instead of the controller and inverter trying to apply power to the Tesla coil at the right frequency and phase, the Tesla coil is applying power to the inverter. Frequency control is no longer a problem. Voltage and current will always be in the proper phase for this high Q resonant circuit during ring down.

As power flow reverses in the line between Tesla coil and inverter, either voltage or current (or both) need to shift phase accordingly. It would be an interesting quiz question to ask students in a power electronics course to predict the top two waveforms in Fig. 10. It turns out that the current never misses a beat. The voltage has a cycle that is half the normal length, such that at the end of the short cycle, voltage is out of phase with current. One is reminded of the old teaching aid for electricity where current flow is equivalent to water flow. Moving water has inertia, and any pressures or heads will be the result of the inertia in the water. I would not want to carry the analogy very far, but in this case it gives the right




Figure 10: Voltage Phase Shift at ENABLE Turnoff

answer. Current appears to have inertia and voltage follows along as best it can.

Fig. 11 shows the voltage and current waveforms later in the cycle. Stored energy is no longer able to drive the IGBT diodes into conduction. The relative phase between voltage and current has shifted about 90° . One possible explanation is that Fig. 10 is showing voltage and current of a stiff power supply while Fig. 11 is showing voltage and current of a capacitor. The stiff power supply being recharged acts like a resistor as far as phase angles are concerned. A power supply will have real power going out (normally) or in (in this case), hence must have voltage and current at least somewhat in phase. When the Tesla coil stored energy is inadequate to connect to the power supply, we just measure voltage across and current through the intrinsic capacitance of the IGBTs. A pure capacitance always has current leading the voltage by 90° . Therefore, a phase shift between Fig. 10 and Fig. 11 is not too surprising.

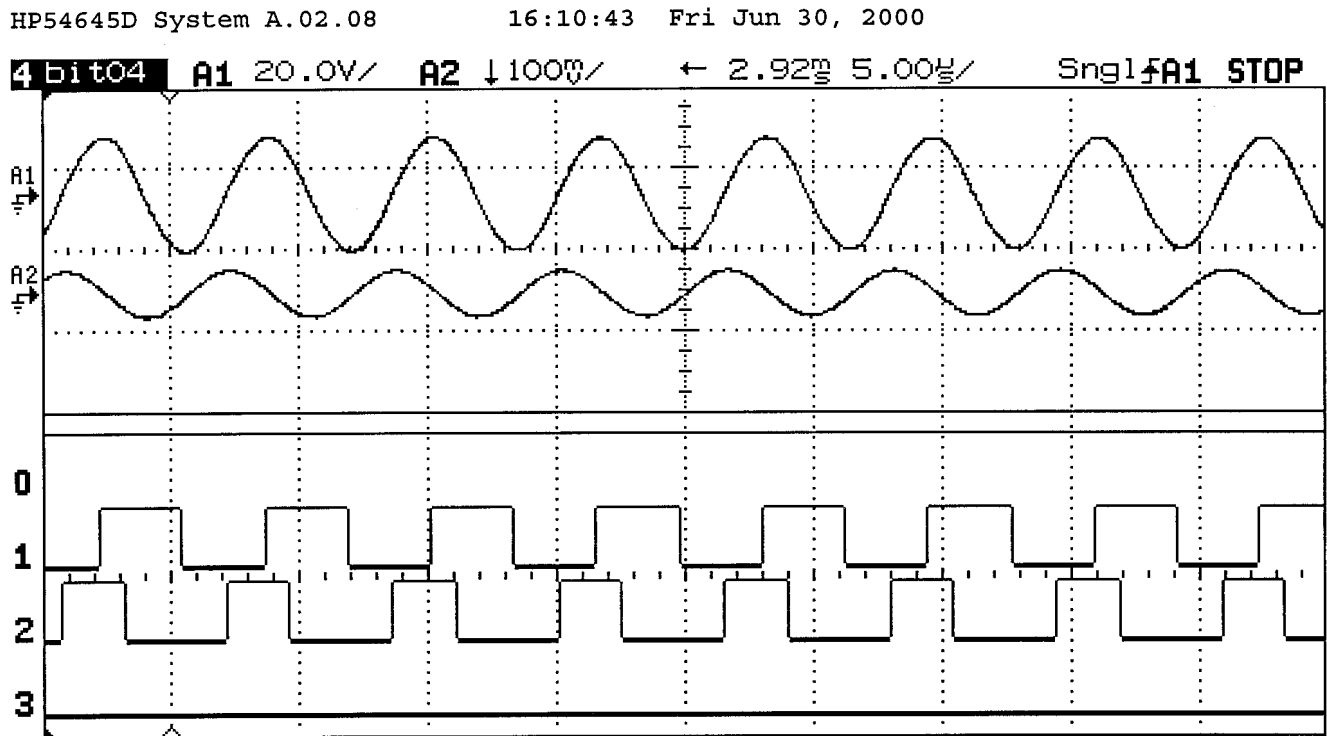


Figure 11: Voltage and Current During Ring Down