Phase-Locked-Loop *SSTC Driver

Construction and operating manual

Rev 2.0

23 May 2006

This document is available online at
http://www.scopeboy.com/tesla/drsstc/

WARNINGS

1: This is an advanced project! I advise you not to attempt it unless you're already experienced in electronics, have an oscilloscope, and know how to interpret the green squiggles that appear on its screen. You should have some knowledge of high frequency power electronics and high energy/high voltage safety procedures too.

2: I disclaim all responsibility for anything that happens to you in the process of building and operating any system based on my driver circuits. I sincerely believe that all of this information is accurate, but I don't guarantee it. Even if the information was accurate, you could still make a mistake and blow yourself to hell. All I can say is, if you take both your eyes out with flying pieces of silicon, don't come crying to me! In case you didn't catch that, I repeat, there is NO WARRANTY WHATSOEVER!!!!

3: I don't guarantee to provide free support for this project to hobbyists. I may provide it, but then again I may not if I'm too busy on something else. If you are a corporation, feel free to hire me as a consultant, and be prepared to pay more than I earn in my day job....

Steve Conner

Introduction

The PLL driver was designed to drive all sorts of power electronic devices that use resonant H-bridge or half-bridge converters. It is mainly targeted at-

*Solid-state Tesla coils (single or dual resonant, CW or pulsed, MOSFET or IGBT)
*High-frequency induction heaters using MOSFETs or IGBTs

The main feature of the PLL driver is its self-tuning action. It seeks and locks to the resonant frequency of the load to ensure maximum power throughput and near-zero current switching of the converter. Hence, it works best with loads that do have a strong resonance.

The standard driver board includes overcurrent protection for the switching devices, and a peak hold rectifier circuit that can be used with an external meter to read peak current in the tank circuit. A phase angle controller is available as an add-on board. When you are running your system off a 50/60Hz AC supply, this allows you to control the converter DC link voltage using a triac or thyristors in the rectifier. The voltage is not regulated, though. This feature is mainly intended to allow you to tune and check a Tesla coil at low voltage without lugging a variac around. The phase angle controller can also be used to interrupt the gate drive, which can be useful to control power on CW-SSTCs and induction heaters.

The standard driver board can supply a gate drive transformer directly with 15V drive at up to 9A peak. (depending on what gate driver ICs you fit.) A set of gate drive amplifier boards with a matching DC-DC converter is currently under development. This will boost the gate drive to upwards of 30A peak per device at +24/-12V, to drive the largest IGBT modules as fast as you dare to.

For CW applications, the average gate drive current is limited by heating in the driver chips, heating of the 15V regulator, and the power supply transformer capacity. It is possible to drive four IRFP460s or similar at ~150kHz continuously, but not much more, without added heatsinks.

To build a working system, you will need to add:

A H-bridge or half-bridge, with associated bypass capacitors
A power supply for the H-bridge
A gate drive transformer with wiring
A small 50Hz transformer to supply around 15V AC to the board
At least one current transformer or Rogowski coil for feedback and overcurrent sensing. Some configurations use two CTs, or one CT and one Rogowski coil.
A remote controller (ie an interrupter in the case of a DRSSTC)

To build a system based on large brick IGBTs you will also need:

HF power supply unit
Two or four gate drive booster boards. Heck, why not eight, and run two H-bridges in parallel :-)

Not necessary but HIGHLY recommended- A shielded metal enclosure and mains EMI filter

Theory of operation

Driver board

The heart of the driver is an oscillator that runs all the time, at approximately the resonant frequency of the load. You set the oscillator manually using the tuning trimmer, and the optional fine tune control on the remote handset if implemented. The oscillator actually runs at twice the resonant frequency and is divided by two internally, but that doesn’t really matter.

Normally, the oscillator free runs and the gate drivers produce no output. All switches in the H-bridge are off. When you command the driver to start by connecting a voltage to the INT+ and INT- terminals (which are optoisolated) a number of things happen-

The oscillator is switched through to the gate drivers, and the H-bridge fires up. The enable/disable is synchronised to the oscillator to make sure that the H-bridge is never asked to switch a part cycle.

The phase locked loop is enabled, and starts to look at the signal from the current transformer (CTSIG1 in diagrams). It adjusts the oscillator frequency to try and move the zero crossings of CTSIG1 so they occur just as the bridge is switching. The frequency range is limited, and if it can’t achieve its goal without straying out of the range, one of the tuning LEDs will light. The LOOP GAIN trimmer controls how hard the PLL tries to maintain tune, and the TARGET PHASE trimmer allows you to skew the desired phase to compensate for delays in the gate driver and H-bridge. See later for a detailed guide on setting up these adjustments.

The overcurrent detector starts monitoring CTSIG2. If it strays outside the window set by the CURRENT LIMIT trimmer, the gate drive will be stopped until the current falls back within limits. These start-stop events are synchronised to whole cycles too.

When you end the burst by removing the voltage from INT+ and INT-, the gate drive is stopped on the next whole cycle, the PLL is disabled, and the oscillator frequency is reset to its original free running value.

If the PLL gets no feedback signal on CTSIG1, it may stay more or less at its free running value, or wander to either extreme. The result depends on the LOOP GAIN and TARGET PHASE settings. It is not recommended to run the driver with no feedback unless you modify it by removing the LOOP GAIN trimmer. This disables the PLL completely.

There is an undervoltage lockout circuit that disables gate drive until supply voltages are within spec. This prevents odd behaviour on startup and shutdown. It is safe to power the driver board up/down while the DC link is hot.

Interrupter

Not written yet

Phase angle controller

Not written yet

Gate drive booster system

The gate drive booster system is used for driving IGBTs that are too big for the gate drive ICs on the main board to cope with by themselves. The system consists of a separate small driver board for each IGBT, and a HF power supply unit that supplies high-frequency power to each driver board. The system is designed for burst duty, such as in a Tesla coil application, and probably can't drive large IGBTs CW without burning out. The extra drive stage introduces a delay of 200-400ns, which can be compensated by the PLL.

The gate drive signal and HF power are fed through ferrite cored isolating transformers on the driver boards, and an undervoltage lockout signal is sent back through optoisolators. This disables the gate drive and lights the FAULT LED on the main board. The purpose of this is to stop the system from being fired up if any of the gate driver boards or the HF power supply has a fault that might damage the IGBTs through inadequate drive. A gate-to-emitter short on an IGBT should also trigger the lockout.

This protection system can't catch all possible faults. If one IGBT in an arm fails short during a burst, it's not guaranteed to shut down fast enough to save the other device. In fact it won't shut down at all unless the bad IGBT has shorted from gate to emitter as well as from collector to emitter. It's mostly designed to stop good IGBTs being fried by bad driver boards. I thought this was a reasonable compromise, since it's much simpler than the full protection used in commercial motor drives etc.

The isolating transformers are specially designed with the primary and secondary windings separated to reduce stray capacitance. They mount on a prong routed out of the PCB that sticks through the centre hole of the toroid core. The prong carries a Faraday shield to reduce the stray capacitive coupling even further.

The HF power supply unit consists of a simple self-oscillating inverter that runs at approximately 60kHz. It is a current-fed inverter and produces a sine wave voltage. The output of around 35V RMS is taken from a centre tapped winding on the oscillator transformer with the centre tap grounded to prevent it floating to high voltages if any leakage happens through the gate drivers' isolating transformers.

The inverter is supplied from the same +15VD supply in the PLL unit that feeds its own onboard gate drivers. The increased current draw means that the power supply transformer may need uprated and the regulator IC1 may need mounted on a heatsink. (TBD.)

The HF power supply unit also contains an AND circuit that combines the outputs of all the optoisolators, and signals a fault to the main board if they are not all asserted.

The gate driver board consists of two isolating transformers (one for power and one for gate drive), a rectifier, L-C filter and zener regulator for the power supply, a half-bridge of power MOSFETs to do the gate driving, a circuit to ensure the IGBT is positively turned off in the absence of a drive signal, and LEDs to show the status of the supply rails and HF gate drive output. The gate drive status LED only responds to the AC component of the signal, so if all is OK, it should blink when a burst of gate drive is fired.

The driver board connects to the HF power supply through a Cat5 computer network patch cable, and to the IGBT through a 10-way ribbon cable, with the cores connected to gate and emitter alternately to minimise inductance. This cable should be kept as short as possible. I decided to use a cable, rather than direct mounting, to allow flexibility in using the driver with many different kinds of IGBTs.

The power supply needs some explanation: it has a choke input filter rather than the usual capacitor input. The choke input filter is an old relic from the vacuum tube era, however it produces an output equal to the average voltage of the rectified sine wave, with much better regulation than a capacitor input filter. One traditional objection to this design was that it needed a big heavy inductor, but this really only applies to rectifiers running off 50/60Hz supply. This version operates at high frequency, so only a tiny inductor is needed. Fast diodes are essential too, but these are easy to buy nowadays.

The regulation is good enough that no regulator ICs are needed on board the driver, and the zener diodes hardly do much. They are still needed to make sure that the 36V DC supply splits up into +24 and -12, though. They also make sure that the current in the choke stays continuous: if it didn't, the output voltage would shoot up from the average value, to the peak, which is not too good. The supply does droop a little when you compare the idle condition to driving a large IGBT with long bursts.

Tesla coil application: H-Bridge and controlled rectifier

Schematic sheet xxxxx shows an example system driving a dual resonant Tesla coil with a H-bridge of IGBTs, with the DC bus voltage controlled by SCRs to avoid the need for a variac. Note that this needs the phase angle SCR controller board, which is still in development.

Construction, snags, etc

Driver board

There are several silly mistakes that will be corrected in the next PCB batch. See the "Errata" section of this document.

Choose the timing capacitor C22 according to your system's resonant frequency. With the 220pF shown, you get a tuning range of around 130 to 230kHz. The frequency is proportional to 1/C.

Choose R7 according to the FSD sensitivity of the meter you are using, and the sensitivity of your current transformer. (see below)

The voltage regulators should not need heatsinks, except for IC1 if you are using the gate drive booster system. Space is a bit tight around IC1, so you may need to mount it off board on flying leads in this case.

Current transformer

This varies depending on the expected current level. The transformer I supply as standard uses 33 turns on a 1" diameter ferrite ring (Fair-rite type 78) with a 0.33 ohm non-inductive burden resistor. This gives ~10mV/A and a maximum reading of about 500A (corresponding to 5V output from the peak hold rectifier)

In a DRSSTC, the same current transformer can be used to derive the feedback signal for the PLL. This is done by connecting two back-to-back diodes in series with the burden. CTSIG2 is taken across the burden alone, and CTSIG1 across the burden and the diodes. This gives a better signal at low primary currents.

A larger coil with IGBT bricks would probably use the same 0.33 ohm burden but a larger ferrite ring with 100 turns, giving 1500A FSD (for 5V output). A Rogowski coil with passive integrator is also a possibility at these high currents.

Interrupter

There is no PCB for the interrupter currently. If there is demand I could have one made, but everyone seems to have their own favourite circuit. The provided interrupter schematic gives <1-200Hz PRF and 5-300us pulsewidth, which is suitable for most sizes of DRSSTC. When the PRF control is turned to minimum, pressing the fire button produces a single pulse.

The interrupter schematic also shows how to wire the fine tune and DC link voltage controls. The DC link control is not functional unless you fit the optional phase angle controller board, and add thyristors (or a triac) plus their gate drive transformer(s), to the power module.

The interrupter input on the main board is an optoisolator, so you can use it with either active-high or active-low interrupter signals by wiring appropriately. Finn Hammer designed an add-on fibre optic link that converts the interrupter input to fibre.

Phase angle controller

This is still in the prototype stage. There are some issues with RFI on it, causing it to lose sync and increase the DC link voltage to full. This has only been seen when driving twin coils with the secondary bases flashing over to ground. (ie a short arc with a lot of capacitance behind it) No PCB is available.

A 2.2uF capacitor was required in the mains feed to the driver board transformer. This advanced the sync signal enough to cure problems with misfiring at startup, when the driver board and thyristor controlled rectifier were running off the same mains supply.

Gate drive booster: power supply

This is still in the prototype stage and has no PCB. Build according to the schematic supplied. The oscillator transformer must be a ferrite core with an airgap. Adjust the airgap to get a frequency of around 60-70kHz. The DC link inductor should be iron powder, or a ferrite bobbin type with plenty of airgap. I used a toroidal iron powder suppressor choke from an old lamp dimmer that was in my junk box. The resonating capacitor should be a low-loss type such as polypropylene film.

You will need to make up four Cat5 patch leads to connect the power supply to the gate driver boards. I recommend cutting two patch leads in half to get four “Cat5 to bare ends” leads. Be sure to get the phasing right: two of the Cat5 leads should have their gate drive cores swapped over compared to the other two. These should be labelled so you know which driver boards to plug them into.

Gate drive booster: driver board

I won't give you technical support on this unless you use the same types of ferrite cores, filter inductor, and MOSFETs that I designed it with. You can buy a pack of these parts from me. Other kinds may well work fine, but you're on your own with the troubleshooting. If you build it according to the schematic with the parts I recommend, it should just work with no tweaking required.

Be sure to get the phasing right on the gate drive transformer, and make sure the isolation clearances on both trannies are adequate. In particular you should mount the GDT so that the secondaries (that drive the MOSFETs) are on the same side as the Faraday shield on the PCB prong, otherwise the insulation may not be good enough. For the power transformer, it's the other way around: the primary (connected to the oscillator in the HF PS unit) goes on the same side as the faraday shield.

Testing: The +24 OK and -12 OK LEDs should light and the DRIVE OK LED should blink when the interrupter is fired. The gate voltage should sit at -12V normally and pulse to +24 to switch the IGBT on. Rise and fall times should be <500ns.

I highly recommend setting up all of the drivers and checking the phasing of each arm using a dual trace scope, before applying any DC link voltage. If you get the phasing wrong, it's possible to turn both IGBTs in an arm on at once and destroy them. To change phasing, swap over the gate drive wires in the Cat5 cable where they enter the HF power supply unit.

Setup (driver board)

You will need a dual trace oscilloscope, 10:1 probe, and ferrite cored current transformer. A 100:1 high voltage probe is useful.

Wire up the tank circuit including the feedback CT.

*HOW TO PHASE THE FEEDBACK
Identify the leg of your H bridge whose output is in phase with the gate drive signal on pins 1 and 2 of J2. (This gate drive signal is also on pins 6 and 7 of IC3 if it's more convenient to probe here.) Take the wire coming from this H-bridge output, and pass it through the current transformer from the side with the phase dot. (this is also the side with the screw-on lid). See IMGP0554 for example wiring. In this picture the in-phase output of my H-bridge is marked by a red dot.

If you made the current transformer yourself instead of buying a premade unit from me, you'll have to figure out which side carries the phase dot by yourself. :P

Start by turning LOOP GAIN to minimum and TARGET PHASE to centre. (see errata- all trimmers are backwards- anticlockwise to increase- except CURRENT LIMIT. This will be fixed in the next board run)

Place the extra CT on a tank circuit wire and connect it to the scope. If you have the high voltage probe, connect it to one leg of the H-bridge and display the voltage waveform on the other scope channel. If you have no HV probe, use a normal 10:1 probe and be careful

If you are working on a Tesla coil, set the primary coil to the intended tapping point, and remove the secondary coil, or short it.

Set the remote FINE TUNE control, if fitted, to centre.

Set the interrupter for the longest bursts possible and enable it.

Power up the H bridge at reduced voltage. (about 50V DC.)

Looking at the scope, adjust TUNING for maximum tank circuit current.

Now increase the LOOP GAIN. One or other of the tuning LEDs may light as the PLL tries to take control. Adjust TARGET PHASE and TUNING until the tank current reaches a maximum and both tuning LEDs are off. Keep increasing the LOOP GAIN to maximum.

Phase check: If everything is phased correctly, you should be able to get both tuning LEDs extinguished with the loop gain turned up full. If the phasing is wrong, then it will be impossible to get both lights out. Either one will light, or the other will, or both will flicker.

Now fine tune TARGET PHASE for the cleanest zero current switching towards the middle/end of the burst. This adapts the PLL to cancel out the delays in your power switches.

Move the primary tap a little (after powering down) then power up again and check that the PLL tracks the new frequency and the switching remains clean. If you move it far enough, one of the tuning LEDs should light to tell you that the TUNING control needs readjusting.

The initial setup and checkout is now complete. From now on you should only need to adjust the TUNING, unless you make radical changes to the resonator system.

In a dual resonant Tesla coil, the next step is to fit the secondary. When you power the system (and be careful as it will produce some high voltage output even at low DC link voltage) you should now see two settings of the TUNING control that extinguish both tuning LEDs. These are the two resonant modes or "poles". Both are valid operating points, however they react to streamer loading in different ways. It is a matter of debate which pole is "better".

Operation (Tesla coil application)

Some guidelines for operation.

General/safety:

Always wear eye protection when working near a live H-bridge. When the supply voltage is turned up, the capacitors store large amounts of energy and can cause devices to explode violently without warning.

When tuning a DRSSTC, it can help to run two long wires from the bridge to the tank circuit, to keep the sparks away from you. Provided the wires are thick (eg 2.5mm2 or more) and tightly twisted together to minimise stray inductance, it doesn't make much difference. Or use the remote fine tune pot. But whatever you do, don't peer into the H-bridge when it's running!

Using the tuning LEDs:

The tuning LEDs indicate if the PLL has failed to lock on. In normal operation, both should be extinguished.

DO NOT run the coil at high power if either tuning LED is lit. Zero current switching is not being achieved, and the power devices may be damaged.

If you change the primary tapping point or topload capacitance, one of the LEDs will generally light and the tuning will need readjusted.

If either tuning LED lights during a run, stop and investigate. The PLL lock range was chosen to accommodate heavy streamer loading and ground arcs, so the tuning LEDs should not light unless there is mistuning or a fault, such as shorted secondary or primary turns, a faulty tank capacitor, or the like.

One of the LEDs may light if the DC link voltage is zero or very low. This just means that the primary current is so small the PLL can't get enough signal to lock. It should go out as the DC link voltage is increased: in tests, 30V DC usually gave enough signal.

One of the LEDs may light slightly under heavy streamer loading. If the PLL was near the edge of its lock range, streamer loading can pull it out. Adjust the FINE TUNE or TUNING control to extinguish it.

When tuning, use as long a burstlength as possible from the interrupter. The longer the burst, the more sensitive the LED indication.

Note: The TUNING and FINE TUNE controls have more or less the same effect except the TUNING control has a much wider range. The FINE TUNE is preferred for adjustments on a live coil but only because it's mounted remotely on the handset and can be used at a safe distance.

Using the overcurrent detector:

If the primary current rises too high, the FAULT LED will flicker. The current is being limited and the coil will continue to operate. Recommended procedure is to choose an operating current you feel comfortable with, and use the peak current meter- or a CT and oscilloscope- to set the current limit such that this current is not exceeded. Worst-case current draw can be tested by shorting out the secondary coil or removing it altogether. It is a good idea to ramp up the DC link voltage gradually when you do this test! You don't want an explosion if the current limiter is set wrongly, or not working for some reason.

If you know the sensitivity of your current transformer, the current limiter can be preset by measuring the voltage across VR2. The limiter will activate when the peak-to-peak output of the CT exceeds the voltage across VR2. The CT I supply has a sensitivity of 0.01V/A, so 8V will give 400A peak (=800A peak-peak). DO NOT adjust it to more than 9V or the limiter may never kick in at all! (I replace D6, D7 with zeners to stop the comparator latching up and this clips the input to 9V p-p)

Errata (Driver board)

1. Four of the five trimpot locations at the board front have the pin-out backwards. CURRENT LIMIT is correct. Depending on the type of trimpot, they may operate the wrong way round (anticlockwise to increase)
2. The protection diodes were left out from the gate drive chip outputs. These are not always needed, but it is easy to solder some 1N5819s to the board underside.
3. A silly design error rendered the UVLO non-functional. To get it fully working, you need to cut a track, and add a resistor (R51) and 7.5V zener diode (ZD2: see schematics)
4. Under very severe RFI/EMP (such as strikes to the primary on a Tesla coil) something in the driver circuit has been observed to latch up once in a while. The gate drivers stick ON, causing R2 to smoke and possibly activating the UVLO due to the heavy current drain. The cause is not yet known, but thought to be the 4046 latching up and stopping its oscillator, due to spikes coming in through CTSIG1.
5. D6, D7, D8, D9 were changed for 12V zeners to suppress interference that was causing the overcurrent detector to latch up on primary strikes. This had the side effect of limiting the meter reading to 450A, and making the current limiter totally non-functional if it is set higher than 450A. I should have used a higher voltage of zener but at the time of writing it was too late to fix.

The PCB errors (1, 2, 3) were mostly fixed in the second board run. Now all the pots are the right way round, maybe.

The new boards are marked with the code number “dr060207_revised” and also have a smiley face near IC2 and IC3.

Error 4 was fixed by changing D6, D7, D8, D9 for zeners. Error 5 can be fixed by using 13V zeners instead of 12V.