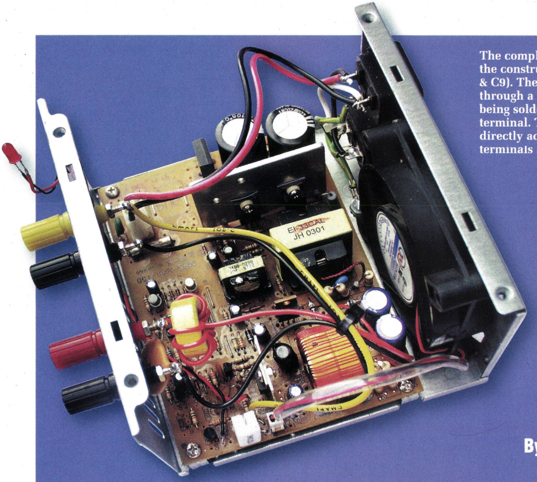

The completed prototype, highlighting the construction of the output filter (L4 & C9). The positive lead is threaded through a small toroid 5-6 times before being soldered to the rear of the output terminal. The capacitor is soldered directly across the positive and negative terminals as shown.

By LEONID LERNER

A regulated 125W HV supply for valve amplifiers

Looking for a low-cost high-voltage (HV) supply to run valve circuitry? Here's how to modify a PC power supply to produce a 700V or 400V HV rail.

VALVE CIRCUITS are not yet dead. While transistors are undoubtedly superior in most applications, the valve still offers several unique advantages. This applies first and foremost to its use in power circuits.

There exists a substantial body of opinion that valves outperform transistors in high-quality audio amplifiers, especially in the power output stages. The seriousness of these claims is reflected in the fact that some very reputable manufacturers offer valve amplifiers at the top end of their audio range. For the home constructor, reasonable-quality valve audio amplifiers can be made for a modest outlay using designs available freely on the Internet. These amplifiers are generally based on an EL34 or KT88 valve pair in the output stage, with both valves being readily available in Australia.

Another common application for valves is in the output stages of RF power amplifiers. They will operate satisfactorily at frequencies of up to about 30MHz, delivering up to 50W per valve. Their main advantage over RF power transistors, apart from being somewhat cheaper, is that they are much more tolerant of fault conditions.

When tuning a new power amplifier design, parasitic oscillations are often encountered which can easily destroy expensive RF power transistors. The valve, however, will live to see another day. Valves are therefore much more suitable for experimentation in new designs.

Although valves are readily obtainable, one of the main problems in their exploitation is the lack of suitable power supply transformers. Both the EL34 and KT88 are rated at a maximum plate voltage of 800V, with supply voltages in the order of 500-600V needed to extract maximum power and linearity. However, the only readily available high-voltage power transformers are isolating transformers, which deliver 240V, and magnetron transformers from microwave ovens, which deliver 1500V or more.

Clearly, both of these are unsuitable for our application.

The easiest way around this is to modify the switchmode power supply of a personal computer (PC), as explored in a previous issue of SILICON CHIP (October 2003). The older AT power supplies are readily available and have now become a surplus item. They are designed to produce about 200-300W, which is in the right ballpark for our application. For little cost, they include a ready-made PC board and almost all of the components we need for a HV switching power supply.

Moreover, due to its high operating frequency, the switchmode power supply offers much better regulation and far less ripple than can be obtained from a traditional valve power supply based on 50Hz AC rectification and smoothing.

Basic considerations

At first, it would appear that getting a PC power supply circuit to operate at high voltages involves just a few changes to the procedure outlined in the previous SILICON CHIP article. In particular, the number of power transformer secondary turns would have to be increased and all diodes, capacitors, and inductors would have to be replaced with high-voltage types.

The resistive ladder used to sense output voltage would also have to be changed. However, after a few trials, I found that the switching power transistors did not last long and it soon became clear that getting the circuit to

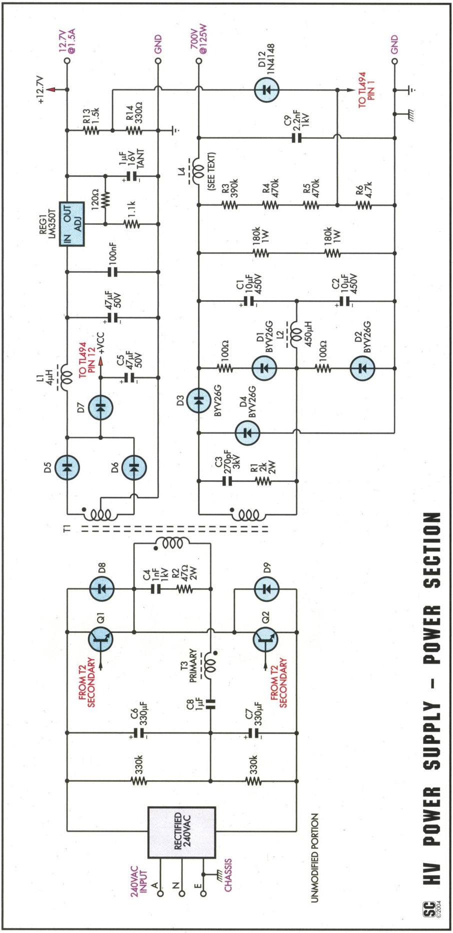

Fig.1: the power section of the modified high-voltage supply. Using the values shown, the output is a well-regulated 700V, suitable for driving two power valves. You can also build a 400V version by winding T1 & L2 accordingly and selecting alternate values for capacitors C1 & C2 and the R3-R5 divider string (see text).

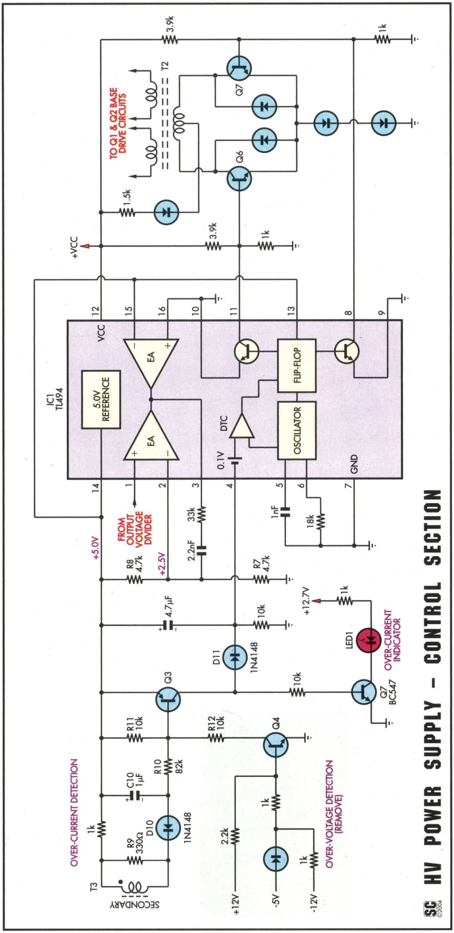

Fig.2: the schematic of a typical control section based on the TL494 PWM controller. The only changes needed here are the removal of the over-voltage detection circuitry and the addition of an over-current indicator, based on Q7 and an LED.

operate at 700V would entail a more substantial redesign.

The main problem is that the volts-per-turn ratio used in the secondary winding of a standard PC ferrite-cored transformer (operating in step-down mode) is about one turn per volt output. This means that 700 secondary turns would be required for an output of 700V.

And that's where we quickly run into problems. The power handling capacity of a coil, without considering insulation, is almost directly proportional to its volume. For example, if we wish to double the output voltage produced by a transformer, we have to double the number of secondary turns, and thus the coil length. The resistance of the coil will also approximately double.

However, if the coil is to deliver the same power, the output current is halved so that the coil's "ohmic" (I^2R) losses are halved. To compensate for this, we can halve the wire's cross-sectional area so that the overall volume occupied by the coil is unchanged. Unfortunately, a multi-layered coil operating at high voltages and frequencies requires insulation whose thickness increases roughly proportionally to the voltage. As a result, our coil does not follow the volume law.

In fact, it is almost impossible to fit a 700-turn winding with adequate insulation into the space available around the core of a standard transformer.

The reason for the large number of secondary turns is that the original PC supply uses a full-wave centre-tapped rectifier configuration, which requires twice the number of turns of a full-wave non-centre-tapped configuration. However, even a non-centre-tapped configuration causes problems.

For a start, it is difficult to fit even 350 turns in the space available around the core. Also, the bridge configuration has no "cool" end of the secondary winding, with both ends alternatively switched between ground and maximum voltage. This means that heavy-duty insulation needs to be used between the primary and secondary windings.

Another problem is related to the mode in which the PC power supply operates. It relies on varying the duty cycle of the rectified mains pulses applied to the transformer to control the output voltage. This means that the secondary rectifier and filter network must be designed to supply an output voltage dependent on that duty cycle. A simple capacitive filtering network is unsuitable, as it would charge to the peak secondary voltage regardless of duty cycle.

The way this dependence is normally introduced is to place an inductor of appropriate value between the rectifying diode and the capacitor, forming an LC filter. However, combining an LC filter with a bridge rectifier does not clamp the secondary voltage, allowing large spikes to appear across the primary during transient conditions.

WARNING!

This is NOT a project for the inexperienced. Do not even think of opening the case of a PC switchmode power supply (SMPS) unless you have experience with the design or servicing of such devices or related high-voltage equipment.

Some of the SMPS circuitry is at full mains potential. In addition, the high-voltage DC output from this supply could easily kill you. Beware of any residual charge on the mains and output capacitors, even if turned off for some time.

The metal case and ground (0V) outputs of all PC power supplies are connected to mains earth. You should verify that these connections are in place after completing any modifications; under no circumstances should the output be floated!

DO NOT attempt to modify a SMPS unless you are fully competent and confident to do so.

Voltage doubler solution

The schematic diagram in Fig.1 shows a solution to these problems. It's based on a voltage doubler circuit fed by a relatively low secondary voltage, making the secondary winding easy to fit around the core. A filter inductor (L2) introduces the duty cycle dependence necessary for pulse-width modulation (PWM), while diodes D1 & D2 clamp the secondary voltage, thereby limiting voltage spikes across the primary.

There is sufficient space left around the former for a second 12.6V/2A secondary to feed the filaments of two power valves. This winding is also used to power the switchmode controller circuitry and the cooling fan.

The price we pay for going to the voltage doubler configuration is reduced power handling. The load current of the centre-tapped configuration has a large DC component and only about 20% ripple, whereas in the voltage doubler configuration current must drop to zero at some point in the cycle. This means that the average current is at best half the maximum current. And since the latter is limited by the saturation current rating of the transformer, the HV circuit can deliver just over 60% of the power of the original supply.

This does not apply to the 12.6V centre-tapped secondary, however. So, from an original power rating of 200W, 125W is now available for the HV supply. Alternatively, the unit can supply 20W for the filament supply and about 105W for the HV supply. This is more than sufficient to operate two power valves.

The circuit is capable of excellent performance. It maintains full regulation at up to 125W, with ripple at 2V peak-to-peak, or 0.3% at full power. This is quite acceptable, as most of the ripple is at twice the switching frequency (60kHz) and so is inaudible.

The 100Hz hum component is only 0.08%, which shows the excellent regulation of the TL494, since the rectified mains source contains 13% of 100Hz ripple at full power. Over-current protection is retained, with a LED added to indicate when it is active.

Circuit operation

The schematic of the power section of the HV supply is shown in Fig.1. The mains input and associated switching transistor circuitry remain unchanged, as indicated by the shaded portion of the circuit.

Typical control circuitry based on a TL494 PWM controller is shown in Fig.2. There is quite a bit of variation in the control circuitry between different manufacturers, so your circuit might differ somewhat. This is especially true if the over-voltage and over-current protection in your supply is based on the LM339 comparator rather than on discrete transistors, as shown. Fortunately, there are few modifications to this part of the circuit.

Operation is quite straightforward, with the design based on a conventional half-bridge "forward converter" topology. The 240VAC mains is first rectified and then filtered by the capacitive divider C6 & C7 to provide two supplies at ±170V DC. This is switched alternately through the ferrite transformer by power transistors Q1 and Q2.

A 1μF capacitor connected in series with the transformer primary limits the current by forming an 8Ω load with inductor L2. This provides some protection in case of a shorted secondary, which effectively occurs at startup before C1 & C2 are charged as well as during fault conditions. Transformer T3 is used to sense the magnitude of the primary current for over-current protection.

The secondary winding develops a voltage of 502V using the specified turns ratio. For 400V designs, the secondary voltage reduces to 319V. This is rectified in the voltage doubler (D3 & D4) and smoothed by an LC filter (L2, C1 & C2).

During the "on" period, energy coupled to the secondary winding finds a current path through L2 and into the load and output filter capacitors. During the "off" period, the energy stored in L2 is discharged into the load.

The inductance of L2 is chosen so that current continues to flow for most of the "off" period at full load. You can see this effect in the SPICE simulation (Fig.3). As previously described, the use of an LC filter ensures that the output voltage depends on the duty cycle, as required for PWM control.

Diodes D3 & D4 have to withstand a reverse voltage of about 900V during the transistor "on" period, as well as

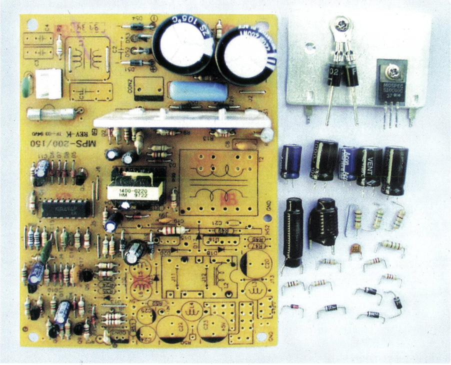

The first step is to remove all of the low-voltage components on the secondary side in preparation for the HV rebuild.

some voltage spikes passed from the primary to the secondary by interwinding capacitance. Note that these spikes are generated during the "off" period by primary leakage inductance - they do not transform to the secondary inductively. Hence, the BYV26G fast avalanche diode with a peak reverse voltage of 1400V was chosen for the job. These are available locally from RS Components (Cat. 216-9397).

Diodes D1 & D2 provide a low impedance return path for inductor (L2) current during the switch-off period. They also combine in the D2-C2-C1-D3 and D4-C2-C1-D1 circuits to clamp the secondary voltage to $\pm V_{OUT}$.

One of the advantages of this clamping method is that it passes much of the energy stored in the core of T1 to the load. This energy would otherwise recirculate through the primary side protection diodes (D8 & D9), as well as dissipate in a more aggressive clamp or snubber network with higher losses.

At power up, the clamp forms a short circuit across the secondary until C1 & C2 are charged, so 100Ω resistors have been inserted to limit the maximum current. The clamp is important in reducing the inductive kick of the primary winding (as opposed to the primary leakage inductance whose kick can not be avoided). The effect of the secondary clamping can be seen as a plateau during the "off" period in the SPICE simulation and the measured primary voltages of Fig.3 and Fig.4(a), respectively. This waveform resembles a square wave at any duty cycle.

An important parameter in the design of the power sections of the circuit is the choice of the secondary voltage to output voltage differential. This is needed to provide headroom to compensate for a drop in the secondary voltage with increased power output, the difference being made up by the by duty cycle variation controlled by the TL494.

Secondary voltage drop has several sources: ohmic losses in inductor coils, non-linearity of the cores, 100Hz ripple due to discharge of mains storage capacitors C6 & C7 and voltage drop across C8 which is charged and discharged every switching cycle. The latter two effects contribute to a primary voltage ripple of 22V and 9V peak-to-peak respectively at full output power, which manifests as a 63V ripple across the secondary. Choosing a 100V differential allows output voltages of up 800V to be delivered by this supply at full power and regulation.

The output voltage is smoothed by a capacitive divider consisting of two 10μF capacitors (C1 & C2) rated at 450V. Alternatively, the 400V version has a higher output current and so 47μF capacitors rated at 350V (or higher) should be used.

At this rating, they are each available in a small package which is easily accommodated in the space provided on the PC board for the original 5V supply components. Their capacitance contributes only about 700mV of 60kHz ripple (0.1%) at full load.

Two 180kΩ resistors across the output set a minimum load current, ensuring that the PWM controller switches Q1 & Q2 on for at least a small portion of each period. Resistor chain R3-R6 divides down the high-voltage output, developing a lower voltage feedback signal that is applied to the non-inverting error amplifier input of the TL494.

Output voltage regulation is achieved by varying the duty cycle so that the voltage applied to the TL494's non-inverting amplifier input (pin 1) equals the voltage on the inverting input (pin 2). In this case, 2.5V is applied to pin 2 via resistive divider R7 & R8. Hence, if R6=4.7kΩ, then R3 + R4 + R5 should equal 1311kΩ for 700V output, or 747kΩ for 400V.

The filament supply is provided by a simple modified version of the original 12V secondary. Unfortunately, we can't use the TL494 to regulate the 12V supply because the original circuit used a coupled inductor shared between the secondaries for this purpose. Our two secondaries now have a high voltage between them. Hence, an LM350T adjustable 3A regulator is used to derive the 12.7V supply. It also powers the 12V cooling fan and must be fitted with a heatsink.

This secondary also supplies power to the TL494 via D7 & C5, as in the original circuit. If a 24V filament supply is required, the more common 7824 1A regulator can be used, as less current is required. The cooling fan can be run from 24V using a 47Ω 5W series dropping resistor.

When the HV supply is only lightly loaded, the duty cycle is so small that the filament supply is not able to deliver its rated current. This can occur at power-on because no plate current flows when the filaments are cold. However, without HV current the filaments can not warm up. To avoid this stalemate, an auxiliary voltage control circuit consisting of resistors R13 & R14 and diode D12 is employed.

During normal operation, D12 is

reverse-biased and the voltage at pin 1 of the TL494 is derived from the HV supply alone. However, when the filament voltage drops, the cathode of D12 becomes less positive until, at about 1.9V, the diode conducts and prevents the filament voltage dropping any further. With the resistor values shown, this threshold is set at about 10.5V.

Circuit protection

Care is required to ensure that the deadtime control circuit connected to pin 4 of the TL494 operates correctly in the modified circuit. The function of the deadtime control is to provide over-voltage and over-current protection if the transformer core saturates.

Primary side current is sensed using T3, a small current transformer. Its primary winding is connected in series with the primary of the main switching transformer. T3 employs a very large transformation ratio (n of about 180), combined with a relatively small resistance across its secondary winding.

This resistance swamps the effects of primary inductance, such that the voltage drop across the transformer is due only to the resistance. The secondary voltage is then proportional to the primary current at about 2V per amp with n = 180 and R9 = 350Ω.

The resultant signal is rectified by D10, smoothed by C10 and applied to potential divider R10 & R11. When the voltage at the midpoint of this divider exceeds about 0.6V, Q3 conducts and a positive voltage is applied to pin 4 of the TL494 through diode D11. A voltage of 0V on this pin sets a minimum deadtime of 4% while at 3.3V, Q1 and Q2 do not turn on at all.

The values shown for R10 & R11 set a threshold current of about 2.8A but you could vary this by altering these resistors. Transistor Q7 and LED1 were added to the circuit to indicate activation of the over-current protection.

Transistor Q4 and the voltage divider connected to its base provides protection against output voltage imbalances by injecting current into the base of Q3 under fault conditions. The voltage divider in the original circuit was designed to produce about 0V at the base of Q4 under normal conditions. However, since the modified supply no longer generates the negative voltages of the original circuit but still has the +12V circuit, this would upset the current balance in the resistor network.

Catastrophic failures aside, output voltage regulation prevents over-voltage anyway, so the easiest solution is to disable this part of the circuit by removing the associated components (shown shaded in green on the circuit diagram).

Your circuit may use a different configuration to the one shown here. For example, the LM339 comparator is frequently used for over-voltage detection. If the voltage on pin 4 of the TL494 exceeds about 0.3V under no load, simply disconnect any resistor running from the 12V supply to the control circuitry connected to this pin.

Note that some power supplies do not use discrete components in the protection circuitry at all. Unfortunately, this article can not hope to cover all possible variations. If you do not feel confident in modifying the existing circuitry, then it is recommended that you construct the circuit shown in Fig.2 and use it to replace the protection circuits connected to pin 4.

Selecting component values

Of major importance in this design is the correct selection of filter inductor L2. If the inductance of L2 is too small, the circuit reduces to a standard capacitive voltage doubler configuration and the dependence of output voltage on duty cycle is lost. Alternatively, if it is too large, the voltage developed across it each half-cycle is insufficient to raise the current required by the load.

In practice, the optimum value is about 450μH for a 700V output (150μH for 400V).

Another challenge is choosing appropriate values for the primary and secondary damper networks – R2 & C4 and R1 & C3. The former is needed to damp the leakage inductance component of the primary, which exists in all coils due to a small amount of primary flux that's not coupled to the secondary. The energy stored in this flux during the "on" period ($1/2LI^2$) generates a current which charges the transistor output capacitance and transformer stray capacitance (C_0) when the transistors turn off. In the absence of resistive losses, this energy is fully transferred into capacitive energy ($1/2CV^2$).

If the transformer is rewound as described in the construction section, it

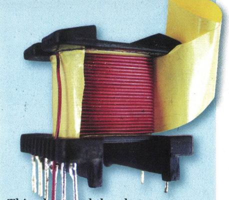

This photo and the photo immediately below show how to wind and insulate one layer of the HV secondary. The layer must start and finish on the secondary face of the former, adjacent to the PC board pins. Starting on the pin end of the former, close-wind one complete layer (no overlaps). After the layer is complete, apply about 1 and 1/4 turns of high-voltage insulation tape. Position the start of the tape approximately as shown.

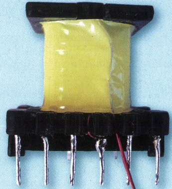

Return the free end of the wire back to the start side, and then bind over it with the end of the insulation tape. The aim is to insulate the return wire from the layer beneath and the one to follow. With the return wire sealed between the two layers of tape, continue winding the next layer, or terminate at the pins if it's the final layer.

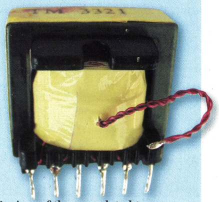

A view of the completed transformer. Note how the centre-tapped connection to the final (12V) winding exits through a small hole in the tape, rather than being terminated at the pins.

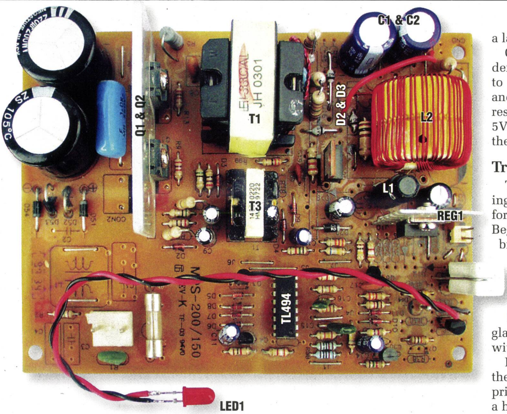

A view of the reassembled PC board showing the newly rewound transformer (T1), HV filter inductor (L2) and HV capacitors (C1 & C2). L2 can be secured to the board using non-acidic silicone sealant.

will have a leakage inductance of about 10μH. As this is less than 1% of the 3.5mH total primary inductance, it is quite acceptable. C_0 is about 270pF, while at full 150W load, the primary current can reach 1.7A.

Plugging in the values gives a voltage spike of about 350V. This adds to the voltage drop across the primary inductance during the "on" period and can destroy the output transistors if the ringing is not damped (protective diodes D8 & D9 offer limited protection due to their finite resistance and turn-on time).

The damper network has the side effect of dissipating energy not only when the transistors switch off but also when they turn on. Making the damper capacitor too large leads to the energy dissipation at turn-on far outstripping the parasitic energy.

The parasitic energy is just the energy stored in the leakage inductance and equates to a power of 0.9W at full load. We've selected a 2W resistor for R2, which leaves 1.1W to be dissipated during switch-on. A capacitor of 1nF will dissipate about 1W in R2 during the "on" period.

R2 should be 50Ω for critical damping. Making R2 smaller does not increase damping; rather, the damper capacitor effectively acts in parallel with the primary winding to change its ringing frequency.

A second source of ringing occurs when current through L2 drops to zero during the "off" period. Depending on the polarity of the half-period, either diode D1 or D2 stop conducting. However, the voltage across L2 can not change instantaneously, due to the energy stored in the diode and switching transistor output capacitance. The resultant ringing is dissipated by the damper network across the secondary and by hysteresis losses in L2.

Construction

You should read the earlier Silicon Chip article (October 2003) on modifying a PC power supply prior to commencing construction. Note that quite a few more components need to be removed here, since most of the secondary side is unsuitable for HV operation.

Begin by removing the large low-voltage secondary capacitors and inductors. The 5V rectifier and associated heatsink also need to be removed, as well as the secondary RC damping network. The multiple power supply leads for the various output voltages are best unsoldered and removed using a large (60W) soldering iron.

Care must be taken when desoldering the ferrite transformer (T1) to avoid melting the former plastic and loosening the pins. Remove all resistors and links leading from the 5V supply to the control circuitry of the TL494.

Transformer preparation

The next job is to remove the existing windings from the ferrite transformer in preparation for the rewind. Begin by carefully removing the tape binding the core sections together, as it can be reused later. Soak the transformer in methylene chloride paint stripper overnight to remove most of sealing varnish.

Note that gloves and protective eyeglasses should be worn when working with paint stripper.

If you don't wish to wait overnight, then the transformer can be warmed prior to dipping for a few minutes with a hair drier held at close range. After about 15 minutes, the transformer can be gently removed and light pressure applied by means of a screwdriver between the slab section of the core and the former, allowing the latter to be released.

It is advisable to remove the E-section out of the former immediately by pressing gently on the centre prong of the "E" (the outer prongs are fragile and easily broken). Care needs to be taken here since this is the only component that is not easily replaced.

If the E-section won't separate with light pressure, then wash the transformer thoroughly and use a razor blade and sidecutters to slice and remove sections of insulation and copper wire to free it up. Complete the transformer disassembly by washing all components and removing all the wire and insulation from the former.

Transformer rewind

Great care must be taken with the transformer rewind to ensure primary to secondary isolation. **In particular, make sure that each layer is completely covered with the tape, right up to the shoulders of the former, so that turns from different layers can not touch.**

Except where noted, there should be no gaps between the start and finish of a layer and the shoulders of the former. This ensures that wire from the next layer can not creep into the gap