SCRs, Triacs & SSRs – A Primer

Silicon controlled rectifiers (SCRs), Triacs and one main type of solid state relay (SSR) are all members of a family of very handy semiconductor devices known collectively as **thyristors**. Other members of the same family of devices are the Shockley diode, the programmable unijunction or 'PUT' and the gate-turnoff switch or 'GTO'. In this datasheet we'll give you an introduction to these devices, the way they work and how they're used. We'll also talk a bit about the Diac, which although not really a thyristor is often used with them as a triggering device.

The main characteristics of thyristor devices are that:

- They're all based on the semiconductor material *silicon*;
- They're all essentially designed for switching, rather than linear current control or amplification;
- They all use a more complicated physical construction than bipolar or MOS transistors, with at least three P-N junctions; and
- The main differences between most types of thyristor are in the way they're triggered into conduction — i.e., 'turned on'.

The simpler thyristor devices are all based on the three-junction PNPN structure shown in Fig.1, which we can use to understand the fundamentals of thyristor operation.

As you can see, the four layers of the PNPN structure can be visualised as if they were made by connecting together the base and collector regions of two complementary bipolar transistors: a PNP and an NPN. In other words, the inner N layer behaves as if it's both the base of the PNP transistor formed by the top three layers, and the collector of the NPN transistor formed by the three bottom layers. Similarly the inner P layer behaves as if it's both the collector of the PNP transistor and the base of the NPN transistor.

So the centre P-N junction is effectively the collector-base junction of *both* transistors, and the transistors are connected together so each one's collector current becomes the base current for the other. This means that there's an internal 'positive feedback loop', whereby they can each amplify the other's collector current — when there is any to amplify, of course.

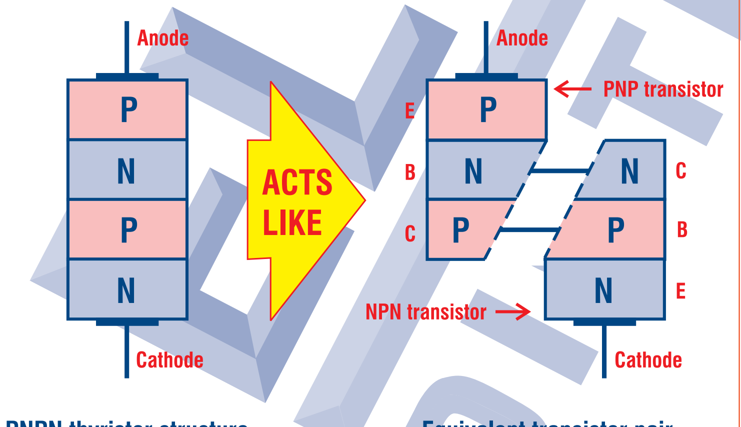

Fig.1: The simplest thyristor is the PNPN or Shockley diode, which behaves rather like two complementary transistors.

Because of this internal feedback action, the basic PNPN structure has an interesting property: it can operate in either of two very different, but equally stable states. In one state it passes only a very tiny leakage current, rather like a reverse-biased PN diode. This is known as the *blocking* or 'off' state.

In the other state it can conduct quite heavily, rather like a forward-biased diode or a transistor that is driven into saturation. Not surprisingly this state is known as the *conducting* or 'on' state.

Now let's see how these two states are possible, and how the PNPN structure can be made to flip from one state to the other, so it can be used as an electronic switch.

If it's connected in series with a load (say an incandescent lamp) and a variable source of voltage as shown in Fig.2, what do you think will happen as the voltage is slowly turned up from zero?

Right — *very little* happens for quite a while, because the centre P-N junction will be reverse biased, and reverse biased silicon junctions normally conduct only a very tiny leakage current. Even though those internal transistors will both be trying to amplify the current, they won't be very successful for quite a while because the current gain (beta) of silicon transistors is very low (less than 1) at low current levels.

So as the voltage is raised, the PNPN structure will only allow an extremely tiny current to flow through the load, like a switch that's in the 'off' state. And this will continue, unless or until the voltage is raised high enough for the leakage current to reach a level **Ith** where the current gain of those two internal transistors just sneaks up to more than 1.

If that point is reached, though, things do happen — and suddenly. Once the transistors have a gain of more than 1, they *do* begin to amplify each other's collector current, and the internal feedback loop begins working. Suddenly they both turn each other hard 'on', and the PNPN structure suddenly conducts heavily. Now it's like a switch in the 'on' state, allowing current to flow through the load.

It will tend to remain 'latched' in this conducting state too, because as long as a reasonable current is flowing this keeps the gain of the internal transistors high, and they keep each other turned on.

The voltage-current characteristic of this basic PNPN thyristor therefore looks like the curve shown in Fig.3. As you can see this shows the forward current rising very slowly until the anode-cathode voltage is increased to reach the 'breakover' level (Vbo), where the current reaches the level Ith and the internal transistor gains rise to unity. Then the thyristor suddenly turns on, with its current rising to a much higher level and its voltage drop falling back to that of a forward-biased diode.

One way to turn on the PNPN structure, then, is to increase the anode-cathode voltage Vf to the breakover level Vbo. But there's also another way, which doesn't involve increasing Vf this far. Instead it's increased by a smaller amount, but *quite rapidly*, so that the capacitance of the reverse-biased centre P-N junction's depletion layer will itself draw enough

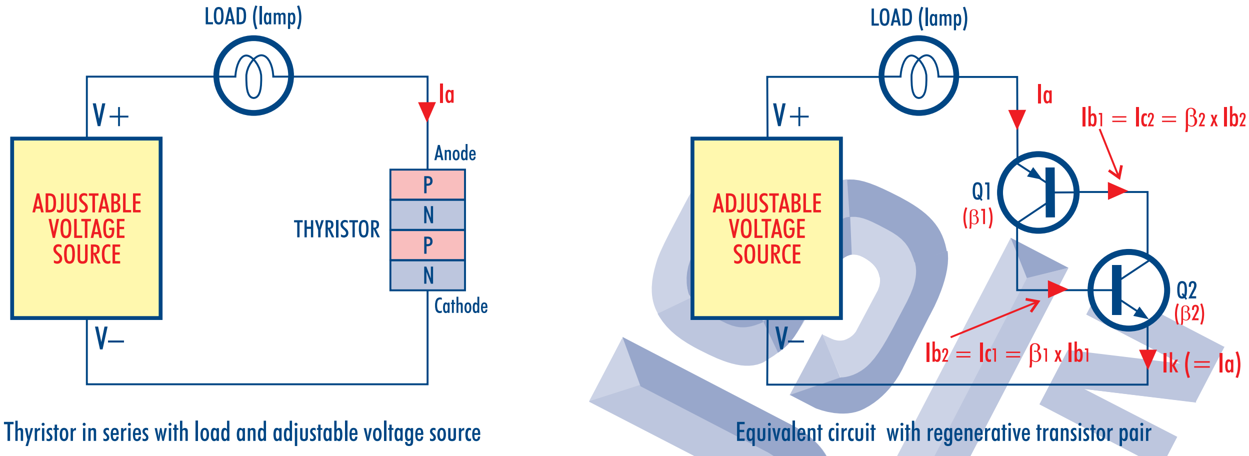

Fig.2: How to visualise the operation of the PNPN diode. The internal complementary pair of silicon transistors can each amplify the other's collector current. But this internal feedback loop or 'regeneration' effect only operates when their current levels are high enough to give enough gain...

current to raise the current gains of the internal transistors to just above unity.

Remember that the current drawn by a capacitor is directly proportional to both the capacitance C and the rate of change of voltage dV/dt:

$$I = C \times (dV/dt)$$

So for any given internal junction capacitance Cd, there will be a certain dV/dt (in volts per microsecond, say) where the capacitive current Ic will reach the turn-on level **Ith**. This is known as **dV/dt triggering**. Note that the actual dV/dt rate necessary to turn the thyristor on will vary with the anode-cathode voltage, because the depletion layer's width and capacitance also vary with voltage. In fact as Cd varies inversely with voltage, a larger dV/dt will be needed for triggering at higher voltages.

This then is the behaviour of the simplest thyristor device: the PNPN or Shockley diode.

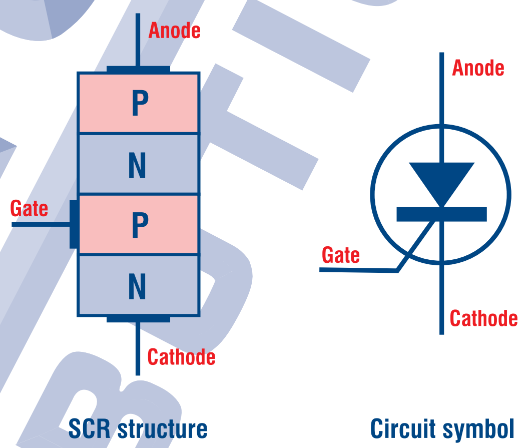

Fig.4: In an SCR, a gate electrode is added to the basic PNPN structure.

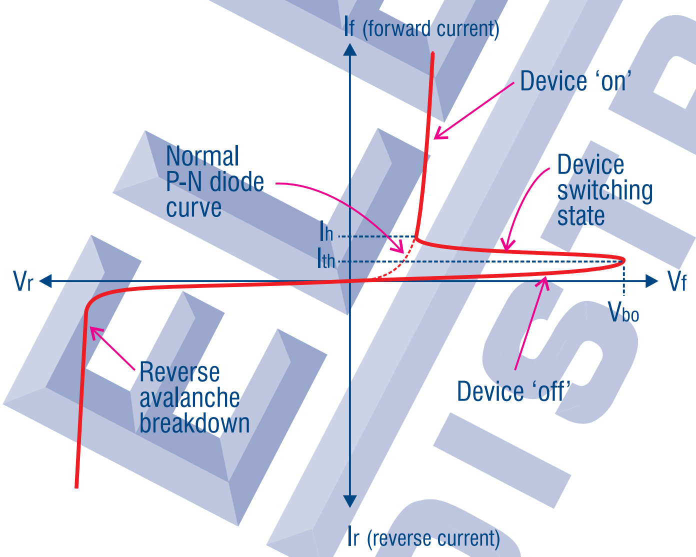

Fig.3: The voltage-current characteristic curve for a PNPN or 'Shockley' diode. The device is normally 'off' until the forward voltage Vf is increased to Vbo.

The SCR

Now that you know how the basic PNPN diode works, it's easy to explain SCR operation. That's because an SCR is essentially just a PNPN diode with an additional connection — to the lower P layer, which is also the base of the internal NPN transistor (Fig.4).

Can you guess what advantage this additional connection provides? That's right, it gives us yet another way to trigger the SCR into switching on. Instead of having to raise the overall anode-cathode voltage to the Vbo level, or applying it rapidly enough to get dV/dt triggering, we can make the thyristor's internal current levels rise to the switch-on level by injecting a small amount of base current into the NPN transistor. It will then amplify this current, the PNP transistor will amplify its current and the internal feedback loop will be kicked into operation.

The SCR's additional connection is known as the 'cathode gate' or **gate**, and with most SCRs only a very small gate current is needed in order to switch the device on.

Fig.5 shows the voltage-current characteristic of a typical SCR, and as you can see when the gate current Ig is zero, the SCR behaves exactly like a Shockley diode — turning on only when the anode-cathode voltage is increased to the Vbo level. But when we do inject gate current, increasing levels of Ig make the SCR switch on at lower and lower anode-cathode voltages.

So the SCR can easily be 'turned on' over a wide range of anode-cathode voltages, by feeding its lower P-N junction with a suitable value of gate current Ig. This value of Ig is known as the *gate triggering current*, and it's usually quite small. Even SCRs switching many amps of anode-cathode current typically only need a gate triggering current of less than 40mA.

The triggering current need only be applied for a short time, too, so it can be in the form of a short pulse. Once the SCR has switched into the conducting state it's no longer needed.

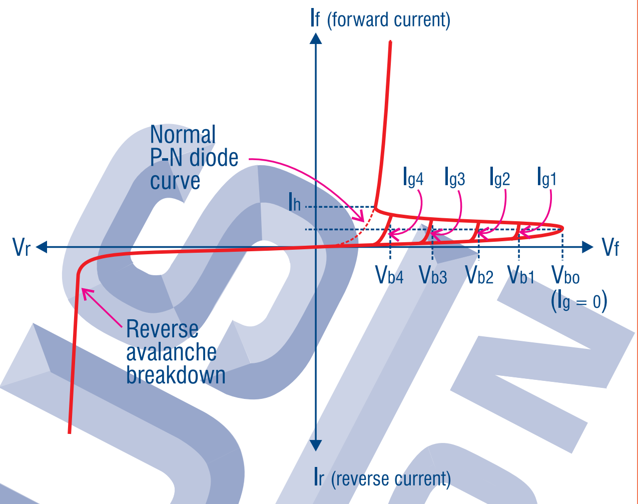

Fig.5: The voltage-current characteristic of an SCR. As you can see the forward breakover voltage varies with different values of gate current Ig.

Turning it off

So far, we've made no mention of how either the Shockley diode or the SCR can be turned off again, once they have been switched into the conducting state. That's because compared with turning them on, they're much harder to turn off.

In fact the only way to make the PNPN structure 'switch off' and swing back to the non-conducting state is to reduce the anode-cathode current level, so it falls back *below* the level where the current gain of those internal transistors is unity. Only then does the internal feedback loop stop holding it in the conducting state.

In short, the basic PNPN thyristor structure will remain in its conducting state unless the anode-cathode current is allowed to fall below Ih, the 'holding current' level. With both the Shockley diode and the SCR, the only way to do this is by external means, like turning down the voltage source or switching it off altogether.

So in general, neither Shockley diodes nor SCRs are very suitable for switching DC. They're much more suitable for controlling AC, where they are automatically 'turned off' at the end of each half-cycle when the applied anode-cathode voltage drops to zero.

For example a typical use for SCRs is in a **phase-controlled rectifier** ciruit, like that shown in Fig.6. Here a pair of SCRs are triggered into conducting on alternate half-cycles, but the triggering pulses are varied in timing/phasing to control how long each SCR conducts within their half cycles. This varies the output level as you can see, because when the SCRs are triggered on *early* in the half cycles (upper curve) they conduct for much more of the time than when they're not triggered until later in the half cycles (lower curve). So the average DC output level across reservoir capacitor C can be varied simply by varying the timing of the triggering pulses fed to

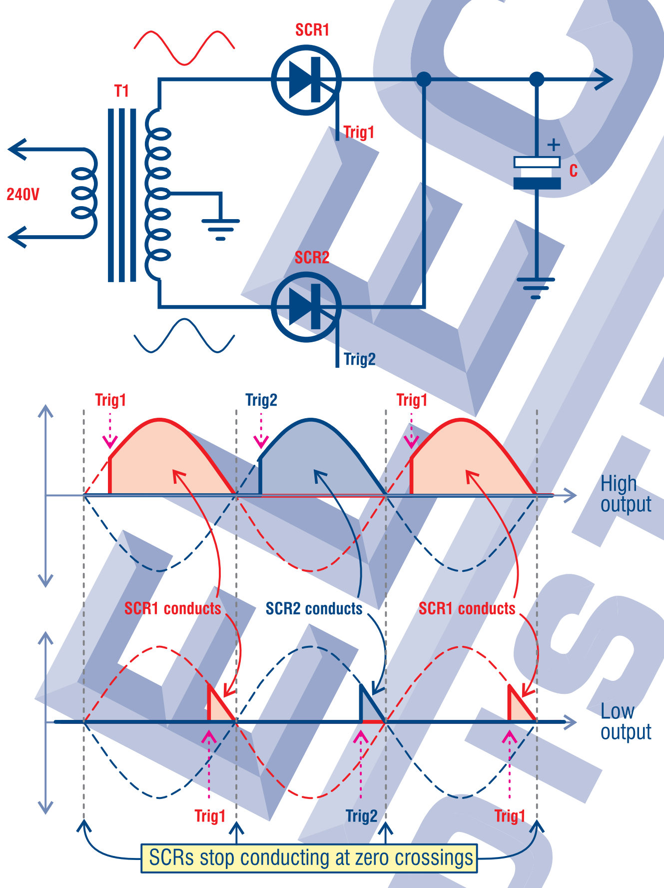

Fig.6: A simple phase controlled rectifier using a pair of SCRs. Note that each SCR turns off at the end of every half cycle, when the AC input drops to zero.

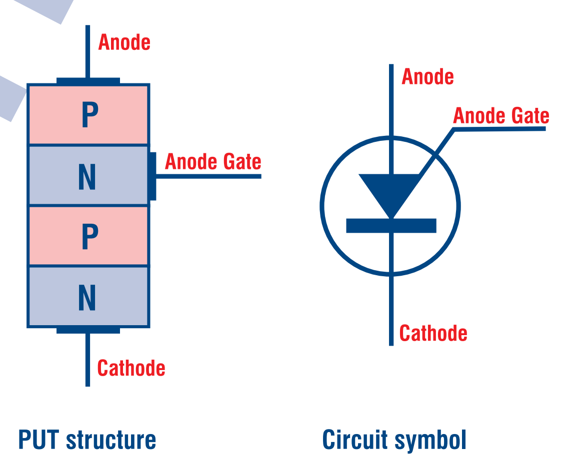

Fig.7: The programmable unijunction or 'PUT' is really just an SCR with an anode gate...

SCR1 and SCR2, in each half cycle.

Note, though, that this circuit only works because both SCRs always turn off again at the end of their conducting half-cycles, when the AC input drops to zero. This is known as 'automatic commutation'.

The PUT

As we've seen, the SCR is nothing more than our basic PNPN thyristor structure with an additional connection to the base of the lower NPN transistor, to allow cathode gate triggering. But wouldn't it be possible to have the additional connection to the PNP transistor's base instead, so the thyristor could be turned on via the upper transistor rather than the lower one? Absolutely. In fact that exactly describes another member of the thyristor family: the **programmable unijunction transistor** or 'PUT' (sometimes also called the 'complementary SCR', because it's a kind of 'upside down' SCR). Instead of a cathode gate it has an anode gate, so it's turned on by applying a triggering current or pulse of the opposite polarity.

The basic structure and symbol for a PUT are shown in Fig.7. Most PUTs are quite small devices, made for switching very low currents.

The SCS

So an SCR is a PNPN thyristor with an cathode gate, and a PUT is one with an anode gate instead. But isn't it possible to have a thyristor with *both* gate connections, so you can use either to trigger them into conduction?

Yes, it is. Such a device is usually called a **silicon controlled switch** or 'SCS', and its structure and symbol are shown in Fig.8. Again most SCS devices are quite small and made for switching very low currents.

The LASCR

Yet another type of thyristor is the **light-activated SCR** or 'LASCR'. As the name suggests, this is basically an SCR which is triggered into conduction not by supplying it with gate current, but by shining light on the PNPN structure — so that additional electron-hole pairs are generated in the semiconductor crystal. This increases the 'leakage' currents to a level where the gains of the internal transistors reach unity, and the thyristor conducts as before.

The semiconductor chip in an LASCR is made flatter and thinner than for other types of thyristor, so light photons can reach the internal junctions more easily. The LASCR chip is also fitted in a package with a quartz 'window', to allow the light to reach it.

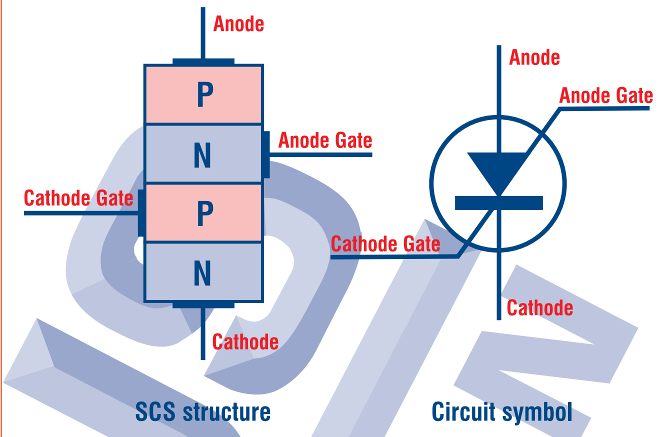

Fig.8: The silicon controlled switch or 'SCS' goes one better, with connections to both gate layers...

The GTO

So far, all of the thyristor devices we've looked at are a lot easier to turn on than to turn off. They can really only be turned off by externally reducing the anode-cathode current below the holding current level **Ih**, as we've seen. However there is one type of thyristor which *can* itself be turned off, by applying a large reverse bias to the gate electrode.

This is the **gate-turnoff switch** or 'GTO', which is essentially a special kind of SCR. Here the doping levels of the various layers inside the device and their exact physical structure are arranged so that with a large reverse bias on the gate, the current through the internal NPN transistor is reduced and its current gain pulled back to less than unity.

The Triac

All of the thyristor devices we've looked at this far have been *unidirectional* in terms of their conduction behaviour — i.e., they only perform as thyristors when anode-cathode voltage is connected to them with the

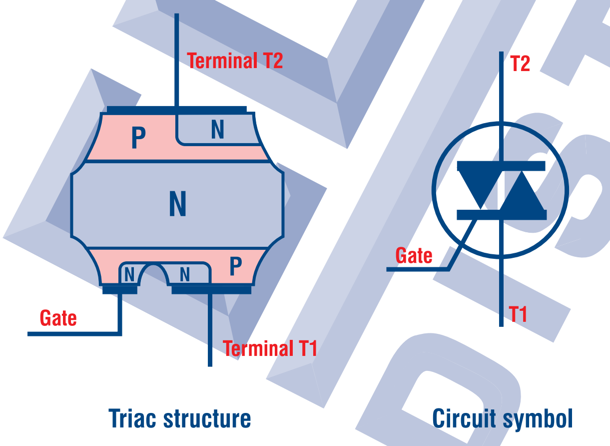

Fig.9: Triacs are more complex inside, but behave like two SCRs connected in inverse parallel.

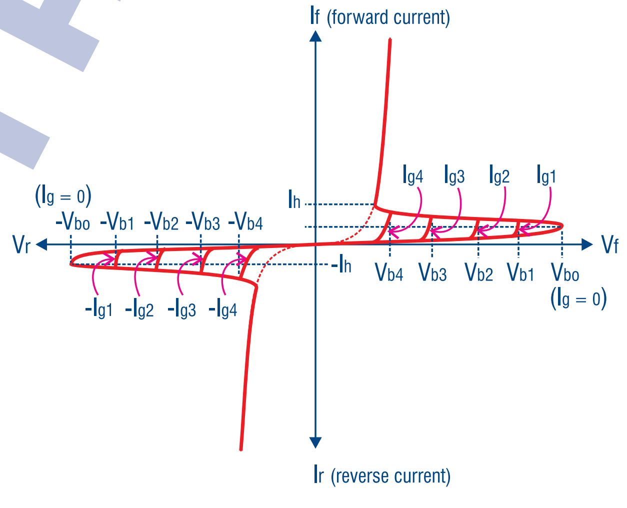

Fig.10: This allows them to be used for controlled switching of currents in both directions: i.e., AC.

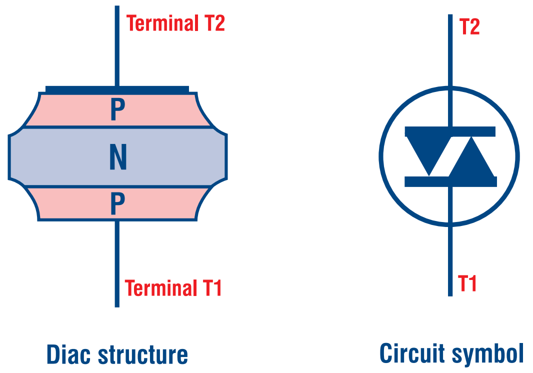

Fig.11: The Diac has a structure like a PNP transistor, but is used as a breakdown diode.

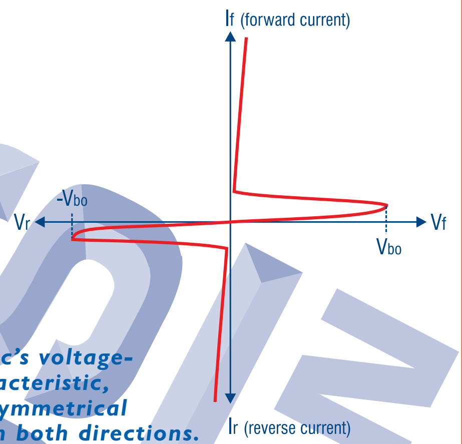

Fig.12: A Diac's voltage-current characteristic, showing its symmetrical breakdown in both directions.

anode positive with respect to the cathode as shown in Fig.2. They don't behave in the same way with reversed polarity, so they're really only suitable for switching rectified AC or as a controlled rectifier.

Luckily there's another type of thyristor again, which *can* perform equally well for voltages of either polarity and is therefore very suitable for switching and controlling AC. This is the **Triac**.

The internal structure of a Triac is shown in Fig.9, along with its usual circuit symbol. As you can see it's fairly complex inside, but there's no need to worry about this because in operation the Triac behaves almost exactly as if it were two SCRs connected in inverse parallel, sharing a common gate electrode.

The voltage-current characteristic for a typical Triac is shown in Fig.10, showing that its performance with reversed voltage (lower left) is virtually the mirror image of that for forward voltage (upper right).

The Diac

As mentioned earlier, the Diac is actually not a thyristor at all, but will be mentioned here because it's often used as a triggering device for thyristors — especially for Triacs.

The Diac is a three-layer PNP device, rather like a bipolar PNP transistor that is symmetrical — i.e., with the two P-N junctions identical in terms of size and doping levels. There's also no external connection to the centre 'base' layer, so the Diac has only two electrodes (Fig.11). In fact it's designed to operate as a *bidirectional breakdown diode*.

Because of its symmetrical structure, one of the two P-N junctions is always reverse biased. This means that the device draws very little current when small voltages are applied — just the reverse leakage current of that junction. However if the voltage is increased, a point is reached where the junction's depletion layer entends over the full thickness of the N layer, and the device enters avalanche breakdown. At this point it suddenly turns on, and conducts heavily (Fig.12).

This behaviour makes the Diac very suitable for triggering thyristors like the Triac, because its ability to conduct only when a particular voltage level is reached allows it to be used to produce narrow and carefully timed trigger pulses from simple R-C timing circuits. An example is the simple AC lamp dimming of Fig,13, where the Diac is used to trigger the Triac for phase control of the lamp's current (in virtually the same way as the two SCRs are used in Fig.6, although here we're not rectifying the AC — just controlling it).

This circuit works by again varying the exact point that the Triac is turned on during each half cycle of the AC waveform. This is achieved by adjusting VR1, which controls how rapidly capacitor C takes to charge up from the start of each new half cycle, to reach the voltage where the Diac conducts. When VR1 is set to the low end of its resistance range, C charges rapidly and its voltage reaches the Diac conduction point very early in each half cycle — so the Diac delivers a triggering pulse to the Triac gate, making it turn on early and conduct for most of the half cycle. The average current through the lamp is therefore high, and it glows brightly (Fig.14 upper curve).

On the other hand if VR1 is set to a higher resistance, C will charge more slowly during each half cycle, and will take longer to reach the Diac conduction point. So the Diac won't conduct until later in each half cycle, and in turn the Triac also won't be triggered until later. So the Triac won't conduct for as much of each half cycle, reducing the average current through the lamp and dimming its brightness (Fig.14 lower curve).

In fact if VR1 is set to its full resistance, C doesn't even reach the Diac conducting point before the end of each half cycle, so the Diac never conducts and the Triac is never triggered on either. So the average lamp

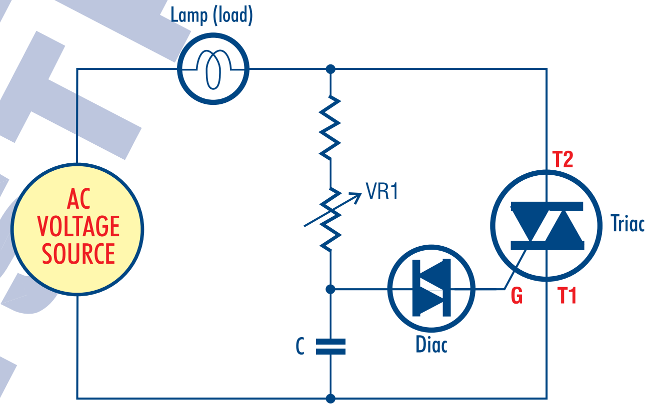

Basic 'phase control' AC lamp dimmer circuit

Fig.13: In the simplest kind of AC lamp dimmer, a Diac is used for adjustable Triac triggering.

current is zero, and it remains dark.

Snubber networks

We saw earlier that one way of triggering the simplest thyristor (the Shockley diode) is to apply its anode-cathode voltage at a fast enough rate (dV/dt), so that the capacitance of its inner blocking P-N junction draws enough current to raise the current gains of its internal transistors to more than unity. As it happens, the Shockley diode isn't the only type of thyristor which can be triggered into conduction this way — the same applies to the SCR, the PUT, the SCS and the Triac as well.

With these other thyristors this 'alternative' triggering mechanism can often be a problem though, because when we're trying to control the conduction of an SCR or Triac via its gate electrode, the fact that they can *also* be triggered via the dV/dt 'rate effect' can easily mean that they're likely to be triggered into conduction at times when we don't want them to be. For example an SCR in a variable-output power supply may turn on as soon as the anode-cathode voltage is applied, rather than waiting for the pulse from our gate triggering circuit — giving higher output than intended. Similarly a Triac in a lamp dimmer circuit like that in Fig.13 may conduct too early in some half cycles due to sharp noise pulses ('spikes') on the AC supply, causing premature dV/dt triggering and thus making the lamp flash briefly or 'flicker' in brightness.

To prevent this type of unwanted dV/dt triggering, many thyristor circuits include what is called a **snubber network** across the thyristor itself, to 'slow down' the rate at which its anode-cathode voltage can be applied. Mostly these snubber networks consist of a series capacitor and resistor combination, as shown in Fig.15. The values of Cs and Rs are chosen so that they represent a fairly high impedance at the normal working frequency (i.e., 50Hz), but a much lower impedance for faster changes in voltage. So noise pulses with a relatively high dV/dt are shunted around the thyristor, preventing it from being triggered by them.

RFI suppression

We also saw earlier that when thyristors *do* switch into conduction, they do so quite quickly. So quickly, in fact, that the sudden rise in current can result in noise radiation, due to rapid changes in the magnetic fields around the wiring. This is particularly true in circuits where the thyristor can be turned on when the anode-cathode voltage is relatively high, like AC phase-control circuits (in contrast with **zero-voltage switching**, where the thyristor is turned on only when the anode-cathode voltage is very low — effectively zero).

Many thyristor circuits thus tend to generate radio frequency interference (RFI) or 'hash' when they're operating, unless steps are taken to suppress the noise radiation.

One common way of reducing the RFI radiation is to fit small inductors or 'RF chokes' in series with the main thyristor leads, as shown in Fig.15. The choke or chokes are usually fitted as close to the device leads as possible, to minimise the radiation.

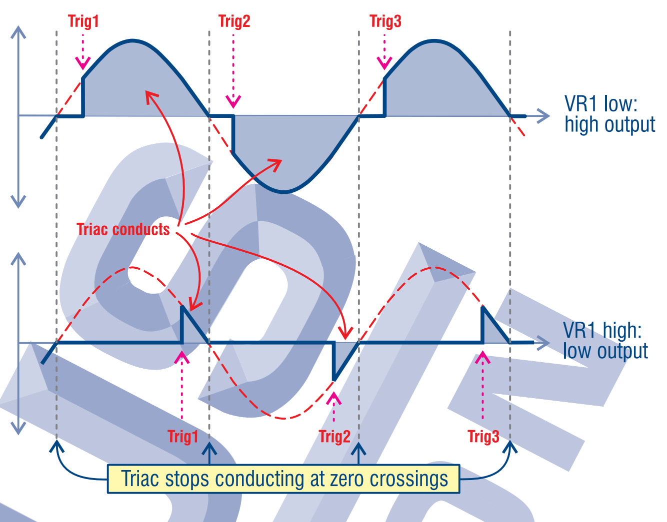

Fig.14: Conduction of the Triac in Fig.13 is phase controlled to vary the average output current, without needing to rectify the AC.

Often the device leads are bypassed to earth as well via low-value capacitors as well, to give additional protection. Where the thyristor is working at mains voltage the capacitors need to have the correct voltage rating of course, and also have a value small enough to prevent problems with residual-current or 'earth leakage' fault detectors. In 240V/50Hz circuits this usually limits the capacitor values to 10nF.

Solid State Relays

Now we've come to solid state relays or SSRs. As the name suggests these are electronic replacements for traditional electromechanical relays, with the advantage of 'no moving parts' and hence much greater reliability. Nowadays they're widely used in many appliances and other equipment where relatively high voltages and currents need to be controlled by low power circuitry.

For example in modern washing machines where everything is controlled by a microcomputer, SSRs are used to allow the micro to control the AC mains voltage fed to the machine's tub, rotor and pump motors, and also the solenoids for controlling water supply. Similarly in a microwave oven SSRs are used to allow the micro to control the power for the magnetron and the turntable motor.

In short, an SSR is a compact electronic relay which typically allows high voltage power to be controlled using low power DC — such as 5V from a microcomputer output line, and at a very low current level (often 15mA or less).

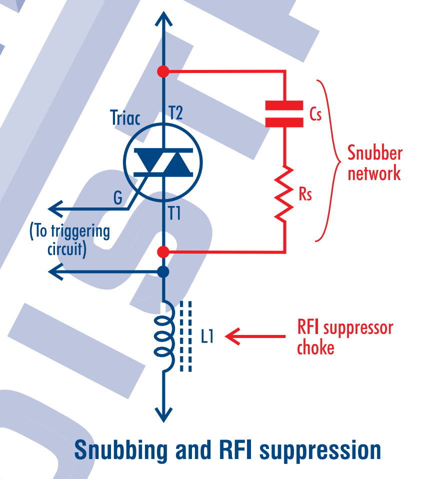

Fig.15: A snubber network across the thyristor stops false triggering.

Two different types

There are actually two main types of SSR, differing in terms of the type of semiconductor device used to perform the output circuit switching. Some modern SSRs designed for switching high current DC use power MOSFETs as the output switch, as these devices offer a very low 'on' resistance — often lower than 10mΩ (milliohms). This gives a low voltage drop, and minimises power loss.

On the other hand SSRs designed for switching AC power can't use a MOSFET as the switch (or not alone, anyway). Instead they generally use a Triac, because these are more suitable for AC.

Fig.16 shows what's inside a fairly basic SSR of the Triac output type, as used for general purpose switching of 240V AC power. You can see that the Triac is controlled by an opto-coupled Diac trigger device. As a result a small DC input current (or pulse) fed through the opto's input LED can turn the Triac on or off, to switch the mains voltage. And because of the high isolation provided by the opto-coupled Diac, this control is achieved very safely and reliably. There's very low risk of transients on the mains voltage arcing over and damaging the micro or other low voltage circuitry supplying the input control current.

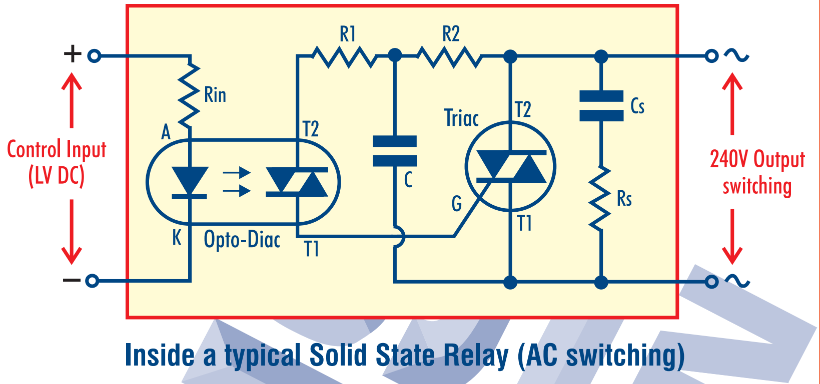

Fig.16: What's inside a typical solid state relay of the type using a Triac for output switching. The opto-coupled Diac prevents mains transients from leaking back into the low voltage control circuitry.

Note that this type of SSR typically contains a built-in snubber circuit (Cs and Rs) across the Triac, to prevent spurious triggering due to mains transients. It may also contain a suppressor choke or two, to minimise RFI.

Important: use the right type of SSR!

Finally, a quick note about when to use the two main types of SSR and when NOT to use them. Although they're much more reliable than conventional electromagnetic relays, they're also not quite as flexible. In fact for some applications you have to be very careful which type you use.

As we've already noted, the type of SSR which uses a power MOSFET as the output switch isn't really suitable for switching AC. For that type of application, you need to use the type of SSR which uses a Triac switch.

On the other hand SSRs which *do* use a Triac (or an SCR) as the switch are not really suitable for switching DC, because as we've seen earlier thyristors (except the GTO type) are much harder to turn off than to turn on. In fact with this type of SSR the only way to turn them off is by removing the source of output current externally, or at least reducing the current to a level below the Triac's 'holding' current Ih.

So if you're switching DC, the correct type of SSR to use is the type using a MOSFET switch, while if you're switching AC or 'raw' rectified AC use the type of SSR which has a Triac switch.

Thyristors & Solid State Relays from Electus and Jaycar