$5.95*
NZ $6.95

eti

CIRCUIT techniques

volume four

The how, what, which, where, why, and how much anthology of electronic components, circuits and techniques.

An ELECTRONICS TODAY INTERNATIONAL publication

ETI Circuit Techniques

volume four

Editor: Roger Harrison
Assistant Editor: Jennifer Whyte
Advertising: Sydney 663 9999
Melbourne 662 1222
Editor-in-chief: Jane Mackenzie

CONTENTS

Circuit Techniques Volume 4 is published by the Electronics Division of the Federal Publishing Company, 140 Joynton Avenue, Waterloo, NSW 2017. Managing Editor: Bob Izzard.
It was printed by ESN - The Litho Centre, 140 Joynton Avenue, Waterloo, in February 1984, and distributed by Gordon and Gotch Ltd. Production by Vernon, Rivers & Associates Pty Ltd.

ISBN 0 86405-015-1
*$5.95 is the recommended and maximum price only.

Analogue delay lines

Analogue delay lines can be used to produce special audio effects such as echo, reverb, phasing, flanging, room expansion and predictive switching etc.

Ray Marston

SOLID STATE DELAY LINES are widely used in modern music and audio systems. They can be used to produce popular effects such as echo, reverb, chorus, phasing, flanging etc. in music systems, 'rare' effects such as ambience synthesis or 'room expansion' in expensive hi-fi systems and 'predictive' effects such as click/scratch elimination in record players or auto-switching in tape recorders etc.

Two basic types of solid state delay systems are available, analogue and digital. Digital delay systems tend to be more expensive and complex than analogue types, except where delay times are in excess of 250 ms, so we'll confine our present discussion to analogue systems only.

Delay line basics

Modern solid state analogue delay lines come in integrated circuit form and are almost universally known as CCD (Charge Coupled Device) or 'Bucket-Brigade' delay lines. In essence, these devices contain a stack of analogue memory (sample-and-hold) cells or 'buckets' (usually 512, 1024 or 4096), all wired in series. Analogue input signals are applied at the front of the bucket 'chain' and the delayed output is taken from the main's end.

Figure 1 illustrates the basic operating principle of an analogue delay line. Each bucket consists of a small capacitor and a tetrode MOSFET and acts like a sample-and-hold stage. An electronic switch is placed at the front of the chain which is externally biased to a preset voltage. Charges can be shifted down the chain, one step at a time, via an external two-phase clock signal; one phase of the clock is also used to activate the input sampling switch. The operating sequence is as follows.

On the first clock half-cycle, each existing bucket charge is shifted backwards one step to the next bucket in the chain and a sample of the instantaneous input signal is fed to the first bucket via SW1, where it is 'stored' as an analogue charge. On the second half-cycle, each existing charge (including the input one) is transferred backwards another step to the next bucket in the chain, but the input is NOT sampled via SW1. There is thus always an 'empty' bucket between each charged bucket in the chain. This double shifting process repeats on each clock cycle, with input samples repeatedly being taken and then clocked towards the back of the chain.

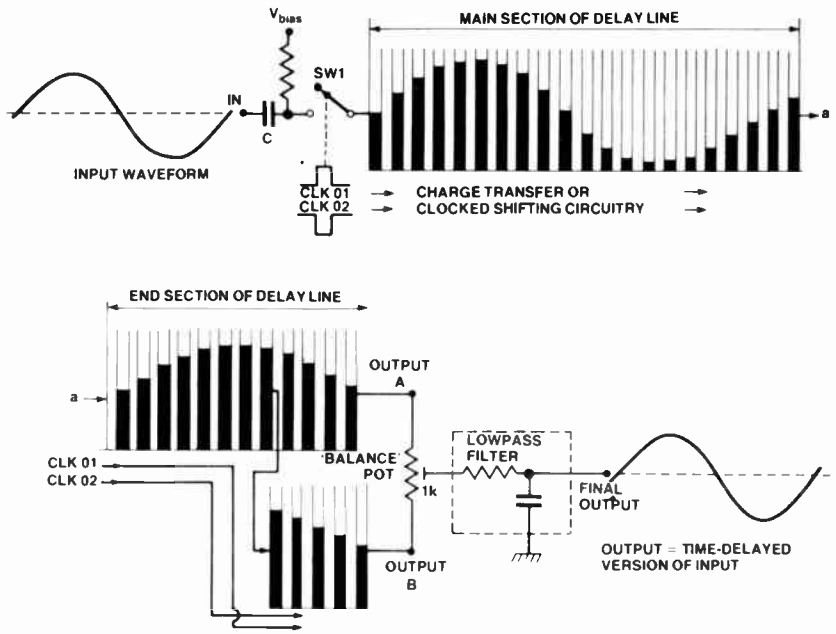

Figure 1. Basic operating principle of the 'bucket brigade' delay line.

In the final section of the delay line, a short section of buckets is wired in parallel with, and fed from, the main delay line, but has one bucket more than the corresponding section of the main line and is clocked in anti-phase. The IC thus has two outputs which, when added together, effectively fill in the 'gaps' in the main delay line bucket chain. The outputs can be 'added' either by shorting them directly together or preferably, by connecting them to a balance pot as shown in the diagram. The final output of the delay line is thus a quantised but time-delayed replica of the original input signal.

Figure 2 shows the essential 'usage' elements of an analogue delay line chip. The delay line MOSFETs use a tetrode structure, so the IC needs two supply lines (V_{DD} and V_{BB}), plus a ground or common connection. The input terminal must be biased into the linear mode by voltage V_{bias}. The two outputs of the device must be added together, as already described; in Figure 2 we've shown addition by direct-shorting. Finally, the IC must be provided with a two-phase clock

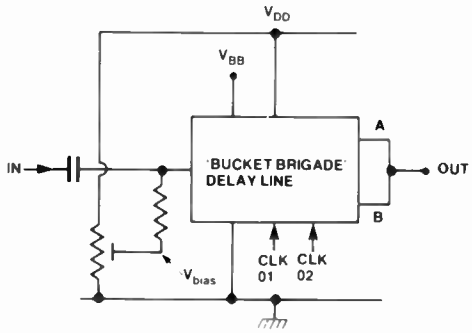

Figure 2. Essential 'usage' elements of an analogue delay line chip.

signal, normally consisting of a pair of anti-phase square waves that switch fully between the V_{DD} and GND (or 0 V) potentials.

How much delay?

We've already seen that the buckets of the analogue delay line are alternately 'empty' and 'charged' and that each complete clock cycle shifts a charge two stages along the bucket chain. Thus, the maximum number of samples taken by a line is equal to half the number of bucket stages (a 1024 stage line can take only 512 samples) and the actual time-delay available from a line is given by:

$$\text{Time Delay} = \frac{S.p}{2} \text{ or } \frac{S}{2.f} ,$$

where S = number of bucket stages, p = clock period, and f = clock frequency.

Thus, a 1024-stage line using a 10 kHz (100 us) clock gives a delay of 51.2 ms. A 4096-stage line gives a 204.8 ms delay at the same clock frequency. This seems pretty good, but there are two major snags. The first is that the maximum useful signal frequency of the delay line is equal to one third of the clock frequency, so a delay line clocked at 10 kHz has a useful bandwidth of only 3.3 kHz. The second snag concerns costs. Analogue delay lines are rather expensive. A 512-stage device will set you back around $30!

Figure 3 shows the block diagram of a basic, real-life analogue delay line system. The input signal is applied to the input of the delay line via a low-pass filter which has a cut-off frequency that is one third (or less) of the operating frequency of the clock generator, and is used to overcome 'aliasing' or intermodulation problems. The output of the delay line is passed through a second low-pass filter which also has a cut-off frequency one third (or less) of that of the clock. This serves the multiple purposes of rejecting clock break-through signals and integrating the delay line output pulses so that the final analogue output signal is a faithful (but time-delayed) copy of the original input signal.

We'll take a closer look at some of the elements of the Figure 3 circuit and at some practical delay line chips later in this article. In the meantime, let's digress slightly and look at the subject known as psycho-acoustics.

Psycho-acoustics

Many of the special effects that are obtainable with delay lines depend heavily on the human brain's idiosyncratic behaviour when interpreting sounds. Basically, the brain does not always perceive sounds as they truly are, but actually 'interprets' them so that they conform to a pre-conceived pattern. The brain can sometimes be tricked into misinterpreting the sounds. The study of this particular subject is known as psycho-acoustics. Here are some psycho-acoustic 'laws' that are worth knowing:-

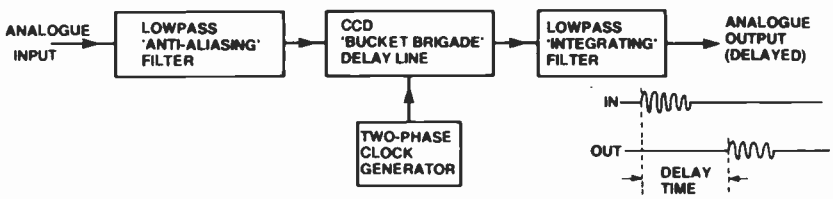

Figure 3. Block diagram of a basic analogue delay line system.

(1) If the ears receive two sounds that are identical in form but time-displaced by less than 10 ms, the brain integrates them and perceives them as a single (undisplaced) sound.

(2) If the ears receive two sounds that are identical in form but time-displaced by 10-50 ms, the brain perceives them as two independent sounds but integrates their information content into a single easily recognisable pattern, with no loss of information fidelity.

(3) If the ears receive two signals that are identical in form but time-displaced by greater than 50 ms, the brain perceives them as two independent sounds but may be unable to integrate them into a recognisable pattern.

(4) If the ears receive two sounds that are identical in basic form but not in magnitude, and which are time-displaced by more than 10 ms, the brain interprets them as two sound sources (primary and secondary) and draws conclusions concerning (a) the location of the primary sound source and (b) the relative distances apart of the two sources.

Regarding 'location' identification, the brain identifies the first perceived signal as the prime sound source, even if its magnitude is substantially lower than that of the second perceived signal (the Hass effect). Delay lines can thus be used to trick the brain into wrongly identifying the location of a sound source.

Regarding 'distance' identification, the brain correlates distance and time-delay in terms of roughly 0.3 metres per millisecond of delay. Delay lines can thus be used to trick the brain about distance information.

(5) The brain uses echo and reverberation (repeating echoes of diminishing amplitude) information to construct an image of environmental conditions, e.g: if echo times are 50 ms but reverb time is two seconds, the brain may interpret its environment as being a 15 m cave or similar hard-faced structure, but if the reverberation time is only 150 ms it may interpret its environment as being a 15 m wide softly-furnished room. Delay lines can thus be used to trick the brain into drawing false conclusions concerning its environment, as with ambience synthesisers or 'room expanders'.

(6) The brain is highly sensitive to sudden increases in sound intensity (transients of millisecond duration), such as clicks and scratches on discs, but is insensitive to transient decreases in intensity. Delay lines can be used to take advantage of this effect in record players where they can be used (in conjunction with other circuitry) to effectively predict the arrival of a click/scratch and replace it with a neutral or negative transient.

APPLICATIONS

Simple musical effects

Figures 4 to 15 illustrate a variety of analogue delay line applications. In these diagrams we have, for the sake of simplicity, ignored the presence of the usual input/output low-pass filters. Let's start by looking at some simple musical effects circuits.

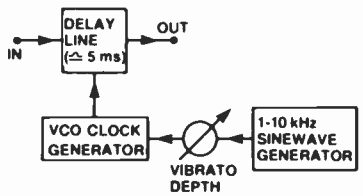

Figure 4. True vibrato circuit which applies slow frequency-modulation to all input signals.

Figure 4 shows how the delay line can be used to apply vibrato (frequency modulation) to any input signal. The low-frequency sinewave generator modulates the clock generator frequency and thus causes the output signals to be similarly time-delay modulated. Simple.

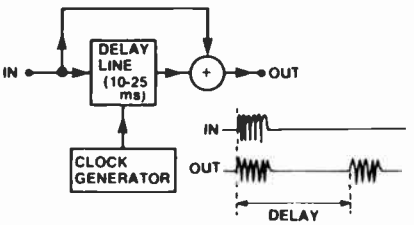

Figure 5. Double-tracking circuit.

Figure 5 shows the delay line used to give a double-tracking effect. The delay time is in the 'perceptible' range 10-25 ms and the delayed and direct signals are added in an audio mixer to give the composite 'two signals' output shown in the diagram. If a solo singer's voice is played through the unit it sounds like a pair of singers in very close

harmony. Alternative names for this circuit are 'mini-echo' and 'micro-chorus'.

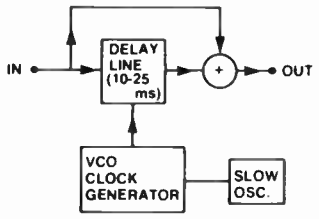

Figure 6. Auto-Double-Tracking (ADT) or mini-chorus circuit.

Figure 6 shows how the above circuit can be modified to act as an Auto-Double-Tracking (ADT) or mini-chorus unit. Clock signals are derived from a VCO that is modulated by a slow oscillator so that the delay times slowly vary. The effect is that when a solo singer's voice is played through the unit it sounds like a pair of singers in loose or natural harmony.

Comb filter circuits

Figure 8 shows a delay line used to make a comb filter. The direct and delayed signals are added together; signal components that are in-phase when added give an increased output signal amplitude and those that are in anti-phase tend to self-cancel and give a reduced output level. Consequently, the frequency response shows a series of notches, the notch spacing being the reciprocal of the line delay time (1 kHz spacing at 1 ms delay, 250 Hz spacing at 4 ms delay).

These phase-induced notches are typically only 20-30 dB

The two most popular musical applications of the comb filter are in 'phasers' and 'flangers'. In the phaser (Figure 9) the notches are simply swept slowly up and down the audio band via a slow-scan oscillator, introducing a pleasant acoustic effect on music signals.

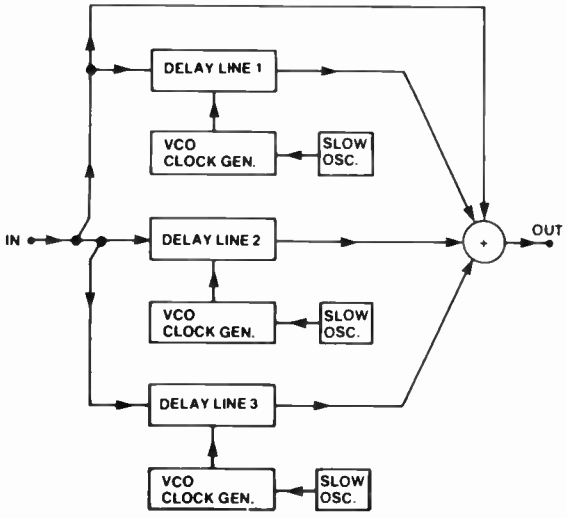

Figure 7. Chorus generator.

Figure 7 shows how three ADT circuits can be wired together to make a 'chorus' machine. All three lines have slightly different delay times. The original input and the three delay signals are all added together, the net effect being that a solo singer sounds like a quartet, or a duet sounds like an octet, etc.

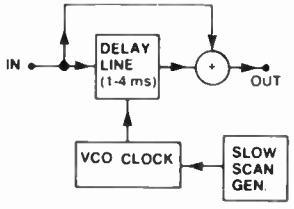

Figure 9. A phaser is a variable comb filter in which the notches are slowly swept up and down the audio band.

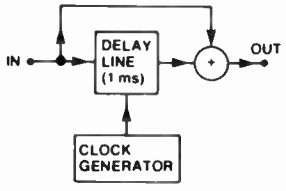

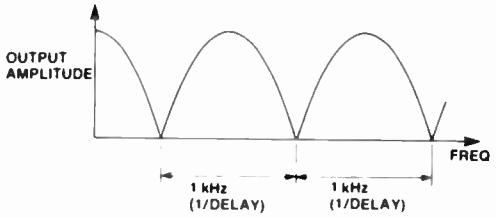

Figure 8. CCD comb filter. Notches are about 20-30 dB deep, 1 kHz apart.

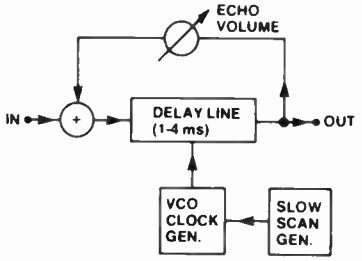

Figure 10. A flanger is a phaser with accentuated and variable notch depth.

The flanger circuit (Figure 10) differs from the phaser in that the mixer is placed ahead of the delay line and part of the delayed signal is fed back to one input of the mixer so that in-phase signals add together regeneratively. Amplitudes of the peaks depends on the degree of feedback and can be made very steep. These phase-induced peaks introduce very powerful acoustic effects as they are swept up and down through music signals via the slow-scan oscillator.

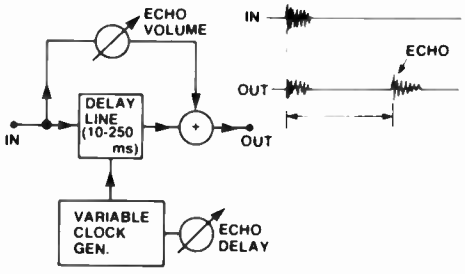

Figure 11. An echo unit.

Echo/reverb circuits

Figure 11 shows the basic circuit of an echo unit. The delay (echo) may vary from 10 ms to 250 ms and is usually adjustable, as is the echo amplitude. Note that this circuit produces only a single echo.

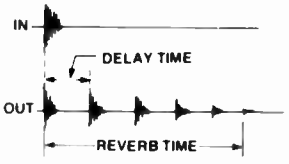

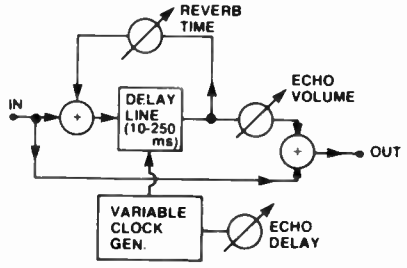

Figure 12. Echo/reverb unit.

The echo/reverb circuit of Figure 12 produces multiple or repeating echoes (reverberation). It uses two mixers, one ahead of the delay line and the other at the output. Part of the delay output is fed back to the input mixer so that the circuit gives echoes of echoes of echoes, etc. The reverb time is defined as the time taken for the repeating echo to fall by 60 dB relative to the original input signal and depends on the delay time and the overall attenuation of the feedback signals. Each delay time, echo volume and reverb time are all independently variable.

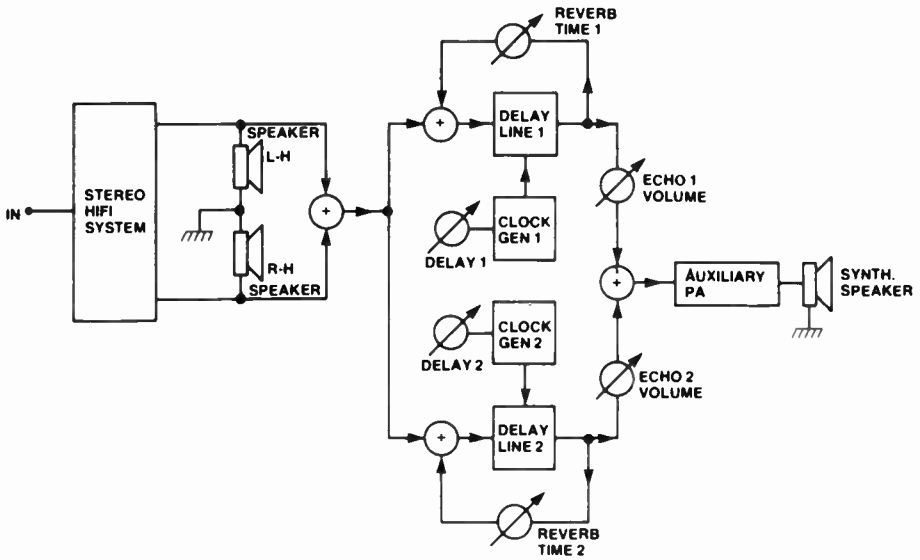

Figure 13. Ambience synthesiser or 'room expander'.

Figure 13 shows the basic circuit of an ambience synthesiser or room expander. Here, the outputs of a conventional stereo hi-fi system are summed to give a mono aural image. The resulting signal is then passed to a pair of semi-independent reverb units which produce repeating echoes but not the original signal. The reverb outputs are then summed and passed to a mono PA system and speaker which is usually placed behind the listener. The system effectively synthesises the echo and reverb characteristics of a chamber of any desired size so that the listener can be given the impression of sitting in a cathedral, concert hall or small club house etc, while in fact sitting in his own living room. Such units produce very impressive results.

There are lots of possible variations on the basic Figure 13 circuit. In some cases the mono signal is derived by differencing (rather than summing) the stereo signals, thereby cancelling centre-stage signals and overcoming a rather disconcerting 'announcer-in-a-cave' effect that occurs in 'summing' systems. The number of delay (reverb) stages may vary from one in the cheapest units to four in the more expensive units.

Predictive switching circuits

Delay lines are particularly useful in helping to solve 'predictive' or 'anticipatory' switching problems in which a switching action is required to occur slightly *before* some random event occurs.

Suppose, for example, that you need to make recordings of random or intermittent sounds (thunder, speech, etc). To have the recorder running continuously would be inefficient and expensive. It would not be practical to try activating the recorder automatically via a sound switch since part of the sound will already have occured by the time the recorder turns on.

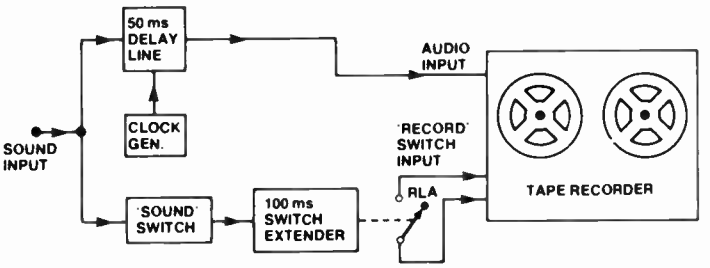

Figure 14. Automatic tape recorder with 'predictive' switching.

Figure 14 shows the solution to this problem. The sound input activates a sound switch which, because of mechanical inertia, turns the recorder's motor on within 20 ms or so. In the meantime, the sound travels through the 50 ms delay line towards the recorder's audio input terminal, so that the recorder has already been turned on for 30 ms by the time the first part of the sound reaches it. When the original sound ceases the sound switch turns off, but the switch extender maintains the motor drive for another 100 ms or so, enabling the entire 'delayed' signal to be recorded.

Finally, to conclude this 'applications' section, Figure 15 shows how 'predictive' switching can be used to help eliminate the sounds of clicks and scratches from a record player. Such sounds can easily be detected by using stereo phase-comparison techniques.

In Figure 15 the disc signals are fed to the audio amplifier via a 3 ms delay line, a bilateral switch and a track-and-hold circuit. Normally, the bilateral switch is closed and the signal reaching the audio amplifier is a delayed but otherwise unmodified replica of the disc signal. When a click or scratch occurs on the disc the detector/extender circuit opens the bilateral switch for a minimum of 3 ms, momentarily blanking the audio signal to the amplifier. Because of the presence of the delay line, the blanking period effectively straddles the 'click' period, enabling its sound effects to be completely eliminated from the system (see 'Psycho Acoustics').

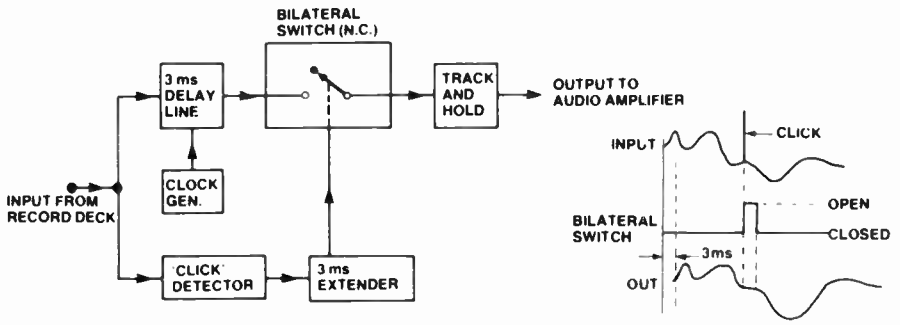

Figure 15. Record 'click' eliminator.

DEVICE NO.	STAGES	SAMPLES	DELAY TIME, ms, VS. CLOCK FREQ.	DELAY AT 7 kHz BANDWIDTH	NOTES
TDA1022	512	256	256/f	12.8 ms	Very popular low-cost delay line
SAD512	512	256	256/f	12.8 ms	512-stage delay line (obsolescent)
SAD512D	512	256	256/2xf	12.8 ms	Built-in clock divider uses single-phase clock
SAD1024A	1024	512	2 x 256/f	25.6 ms	Dual SAD512 delay line
SAD4096	4096	2048	8 x 256/f	102.4 ms	4096-stage delay line. Clock-terminal input Capacitance = 1000 pF

Figure 16. Basic details of five popular CCD delay lines.

PRACTICAL CIRCUITS

Delay lines

Figure 16 shows basic details of five popular CCD delay lines. The TDA1022 and the SAD512 are general-purpose 512-stage delay lines requiring two-phase clock inputs. They give 12.8 ms delay at 7 kHz bandwidth when driven at 20 kHz clock frequency.

The SAD512D is an 'updated' version of the SAD512 and incorporates built-in output drivers and a clock input divider. It requires a single-phase clock input.

The SAD1024A is a dual version of the SAD512. The two halves can be used independently or can be wired in series to give a delay of 25.6 ms at 7 kHz bandwidth.

The SAD4096 gives a performance equal to eight SAD512s in series. It provides a delay of 102.4 ms at 7 kHz bandwidth or 250 ms at 3 kHz bandwidth. The device requires a low-impedance two-phase clock drive, since its clock terminal input capacitance is about 1000 pF.

Figures 17 and 18 show a couple of practical delay line circuits using TDA1022 and SAD512D devices. Both circuits use a preset to adjust the input dc bias so that symmetrical clipping occurs under overdrive conditions and another preset to balance the two outputs for minimum clock breakthrough.

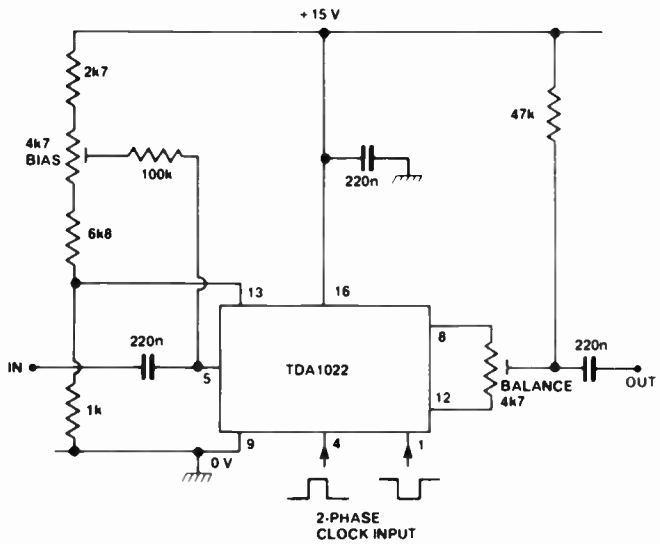

Figure 17. Delay line using the TDA1022.

Clock generators

The clock signals to a CCD delay line should be reasonably symmetrical, should have fairly fast rise and fall times and should switch fully between the supply rail voltages. CMOS devices make ideal clock generators and Figures 19 to 21 show three practical circuits.

The general-purpose two-phase generator of Figure 19 is inexpensive and can be used in most applications where a fixed or manually-variable frequency is needed. The frequency can be swept over a 100:1 range via RV1 and the centre frequency can be altered by changing the C1 value.

The high-performance two-phase generator of Figure 20 is based on the VCO section of a

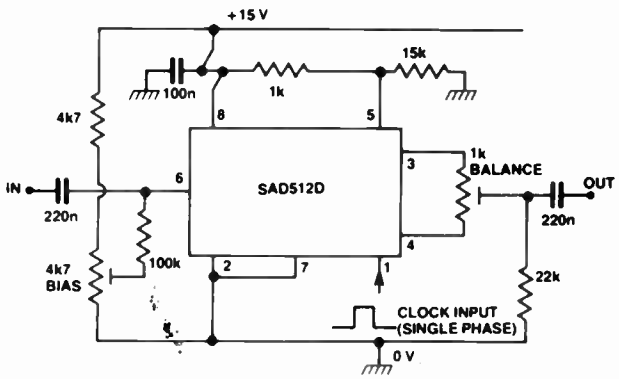

Figure 18. Delay line using the SAD512D.

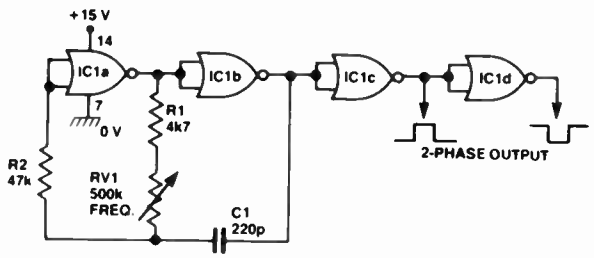

Figure 19. Variable-frequency general-purpose two-phase CMOS clock generator.

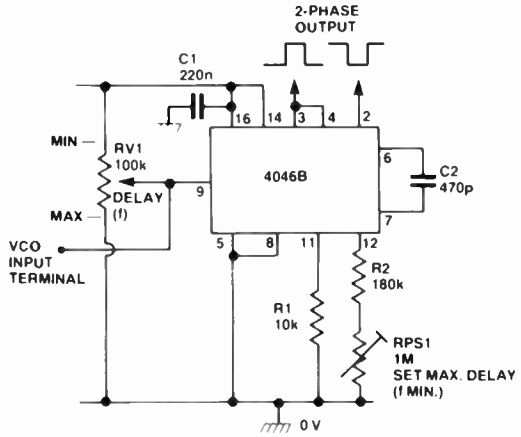

Figure 20. High-performance voltage-controlled two-phase CMOS clock generator.

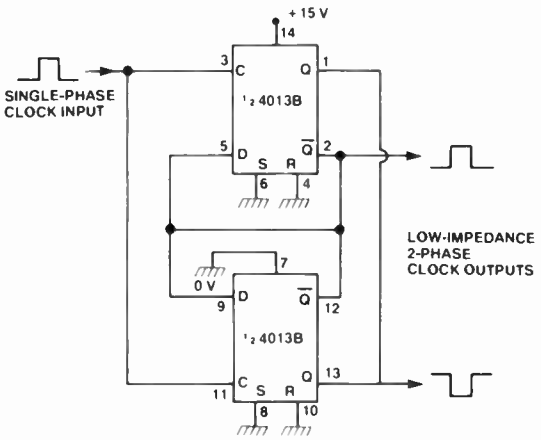

Figure 21. Single-phase to two-phase converter, with low impedance output.

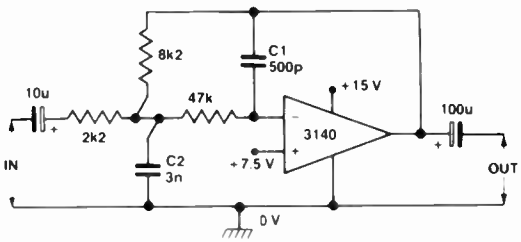

Figure 22. 25 kHz second-order maximally-flat low-pass filter.

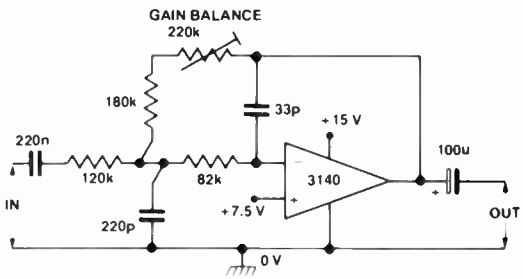

Figure 23. Adjustable-gain second-order low-pass output filter.

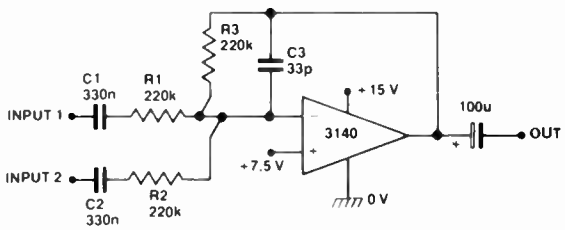

Figure 24. Combined two-input mixer/first-order low-pass filter.

4046B phase-locked loop chip and is useful in applications where the frequency needs to be swept over a very wide range, or needs to be voltage controlled. The frequency is controlled by the voltage on pin 9, being at maximum (minimum delay) when pin 9 is high and minimum (maximum delay) when pin 9 is low. Maximum frequency is determined by the C2-R1 value and minimum frequency by the value of C2 and the series values of R2-RPS1.

The Figures 19 and 20 circuits can be used to directly clock all CCD delay lines except the SAD4096, which has a clock terminal capacitance of 1000 pF (1n) and needs low-impedance clock drive. The SAD4096 is best driven by the circuit shown in Figure 21 which uses the two halves of a 4013 divider wired in parallel to give the required low-impedance two-phase output; the circuit is driven by a single-phase clock signal which can be obtained from either of the Figure 19 or 20 circuits.

Filter circuits

In most applications a low-pass filter must be inserted between the actual input signal and the input of the delay line, to prevent aliasing problems. Another must be inserted in series with the output of the line to provide clock-signal rejection and integration of the 'sample' signals. For maximum bandwidth both filters usually have a cut-off frequency that is one third (or less) of the maximum clock frequency used; the input filter usually has a first-order or better response and the output filter has a second-order or better response.

Figure 22 shows the practical circuit of a 25 kHz second-order low-pass filter with ac-coupled input and output. The non-inverting terminal of the op-amp is biased at half-supply volts, usually by a simple potential divider network. The cut-off frequency can be varied by giving C1 and C2 alternative values, but in the same ratio as shown in the diagram, e.g: cut-off can be reduced to 12.5 kHz by giving C1 and C2 values of 1n and 6n respectively.

All delay lines suffer from a certain amount of 'insertion' loss. Typically, if 100 mV is put in at the front of the delay line, only 70 mV or so appears at the output. Often the output low-pass filter is given a degree of compensatory gain to give zero overall signal loss. Figure 23 shows such a circuit. This circuit has a nominal cut-off frequency of about 12 kHz, depending on the setting of the GAIN BALANCE control.

Finally, to complete this look at CCD delay line circuits, Figure 24 shows how a two-input unity-gain 'mixer' (adder) can also be made to act as a first-order low-pass filter by simply wiring a roll-off capacitor (C3) between the output and the terminal of the op-amp. This type of circuit is often used at the front end of CCD flanger and reverberation designs. ●

A new type of dc motor controller

Ray Marston introduces a unique new circuit 'building block' that can provide very sophisticated power control of low-voltage dc loads such as heaters, blower motors and model locomotives.

Ray Marston

THE AUTHOR HAS for some years been experimenting with precision dc motor control circuitry and has, in the course of this work, evolved what is believed to be a brand new type of circuit element. In essence, this new circuit converts a dc input voltage signal into a switched-mode output signal of almost identical mean dc value and maintains that value independent of wide variations in load characteristics. It enables high-power dc loads to be precisely and variably controlled, with negligible power losses, via low-power local or remote input voltages.

When used to drive dc electronic motors, the new circuit continuously monitors the motor's speed via its 'dynamo effect'-generated voltages, and automatically adjusts the power feed to maintain the speed at a constant level, *irrespective of load variations*. This new circuit has no official name so I will, for the sake of simplicity, refer to it here as the SMVF *(Switched-Mode Voltage Follower)* circuit.

The SMVF circuit has lots of practical applications. It can be used to efficiently control the brilliance of lamps or the speeds of fan or blower motors in the car, or to give high-precision 'feedback sensing' speed control of servomotors, mini-drills and model locomotives, etc., in the home or workshop.

Figure 1. Basic Switched-Mode Voltage Follower (SMVF) circuit.

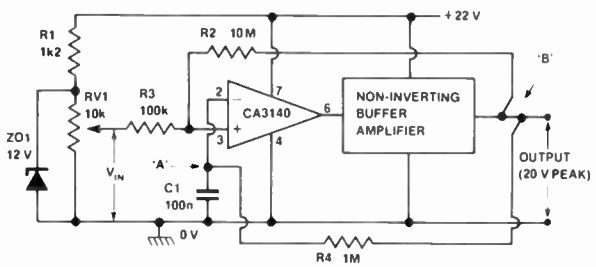

Figure 2. Waveform shapes, terminology and formulae of the Figure 1 circuit when driving a resistive load.

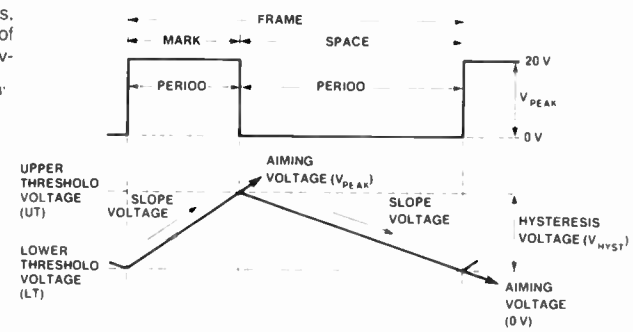

$$\text{UPPER THRESHOLD VOLTAGE (UT)} = V_{IN} + \left(\frac{V_{PEAK} - V_{IN}}{R2/R3}\right)$$

$$\text{LOWER THRESHOLD VOLTAGE (LT)} = V_{IN} - \left(\frac{V_{IN}}{R2/R3}\right)$$

$$\text{HYSTERESIS VOLTAGE } (V_{HYST}) = UT - LT = \left(\frac{V_{PEAK}}{R2/R3}\right)$$

$$\text{MARK SLOPE} = V_{PEAK} - LT\,(M_{SLOPE})$$

$$\text{SPACE SLOPE} = UT\,(S_{SLOPE})$$

$$\text{MARK OR SLOPE PERIOD} = \frac{V_{HYST} \times C(uF) \times R(MOHMS) \times 6500}{V_{SLOPE}}\ ms$$

$$V_{OUT}\,(MEAN) = \frac{V_{PEAK}}{\left(\frac{M_{SLOPE} + S_{SLOPE}}{S_{SLOPE}}\right)} = \frac{V_{PEAK}}{\left(\frac{M_{PERIOD} + S_{PERIOD}}{M_{PERIOD}}\right)}$$

BASIC PRINCIPLES

The basic SMVF circuit

Figure 1 shows the basic SMVF circuit, and Figure 2 shows its two generated waveforms, together with their special terminology and formulae. The circuit is powered from a single-ended supply and uses a 3140 op-amp as its active element. Unique and important features of this op-amp here are that its input and output signals can both swing all the way down to zero volts.

In Figure 1, a non-inverting power booster is interposed between the output of the op-amp and the final output of the circuit, to boost the available output current to a useful level.

In Figure 1, the op-amp is basically used as a voltage comparator, with a reference or control voltage applied to its non-inverting terminal from RV1, and a feedback voltage applied to the inverting terminal via R4-C1. When the *non-inverting* terminal voltage is *above* that of the inverting terminal, the *output* of the circuit switches *high* (to +20 volts). When the *non-inverting* terminal voltage is *below* that of the inverting terminal the output switches *low* (to zero volts if a resistive load is used).

The basic circuit operates as follows: Suppose that voltage V_{in} (in the range 1 to 12 volts) is set on the slider of RV1, that the C1 (inverting terminal) voltage is initially below this value, and that the circuit's output has just switched high. Resistors R2-R3 act as a potential divider between the output (+20 volts) and the RV1 slider potential, and apply a voltage slightly greater than V_{in} (the upper threshold voltage) to the non-inverting terminal of the op-amp.

Simultaneously, C1 starts to charge towards the 20 V 'aiming volts' via R4 until, eventually, it reaches the upper threshold value and the output of the op-amp comparator starts to switch low.

Because of the 'hysteresis' feedback action of R2-R3, a regenerative switching action is initiated at this point and the output of the circuit switches abruptly low (to zero volts if a resistive load is used). R2-R3 then pull the non-inverting terminal voltage to some value below that set on RV1 slider (to the 'lower threshold' value).

Simultaneously, C1 starts to discharge towards the 'zero volts' aiming value via R4

V_{IN}	UPPER THRES-HOLD	LOWER THRES-HOLD	MARK SLOPE	SPACE SLOPE	MARK PERIOD	SPACE PERIOD	V_{OUT} (MEAN)
PERFORMANCE WITH OUTPUT LIGHTLY LOADED (V_{PEAK} = 20 V)							
10.0 V	10.1 V	9.9 V	10.1 V	10.1 V	12.9 ms	12.9 ms	10.0 V
5.0 V	5.15 V	4.95 V	15.05 V	5.15 V	8.64 ms	25.24 ms	5.10 V
1.0 V	1.19 V	0.99 V	19.1 V	1.19 V	5.24 ms	109 ms	1.17 V
PERFORMANCE WITH OUTPUT HEAVILY LOADED (V_{PEAK} = 15 V)							
10.0 V	10.05 V	9.9 V	5.1 V	10.05 V	19.1 ms	9.70 ms	9.95 V
5.0 V	5.10 V	4.95 V	10.05 V	5.1 V	9.70 ms	19.12 ms	5.05 V
1.0 V	1.14 V	0.99 V	14.1 V	1.14 V	6.92 ms	85.5 ms	1.12 V

Figure 3. Actual performance of the basic SMVF circuit of Figure 1, at three values of input voltage, with light and heavy resistive loads, illustrating the good tracking and regulation characteristics of the design.

until, eventually, it reaches the lower threshold value, at which point the output of the op-amp comparator starts to switch high again, initiating another regenerative switching action in which the output abruptly reverts to high (+20 volts) state again. The whole process then repeats ad-infinitum.

Thus, the circuit acts as an oscillator and generates a rectangular or pulsed (switched-mode) output waveform, and maintains the MEAN values of the op-amp inverting and non-inverting terminal voltages at identical values.

Because the 'hysteresis' voltage generated by R2-R3 is fairly low however, the mean voltage on the non-inverting terminal is almost the same as that on the RV1 slider. Note, however, that R4-C1 actually integrates the switched-mode output waveform of the circuit, so that the *mean* value of the output waveform is identical to that on the non-inverting terminal of the op-amp, and almost identical to that on the RV1 slider. The circuit thus lives up to its title of a 'Switched-Mode Voltage Follower'.

Tracking and regulation

The Table of Figure 3 shows the actual performance figures of the Figure 1 SMVF circuit, at three values of input voltage and with very light and very heavy resistive loads, and illustrates the very good tracking and regulating characteristics of the design.

When the output is lightly loaded it has a peak value of 20 volts, giving a hysteresis value of 200 mV, and when the output is heavily loaded it is assumed to have a peak value of only 15 volts, giving a hysteresis value of 150 mV.

Note that the tracking of the circuit is very good, with the mean output voltage differing from the input by no more than 170 mV, and that the regulation is excellent, with a 25% drop in peak output voltage (caused by heavy loading) resulting in negligible drop in mean output voltage.

Also note that the *mark* period of the circuit increases only moderately as the input voltage is increased, but that the *space* period (and hence the *frame* period) decreases by very large amounts under the same condition. The SMVF circuit can thus be regarded as variable-frame type of pulse generator.

Control of dc motors

The basic SMVF circuit gives excellent self-regulating speed-control of dc electric motors; far better, in fact, than that obtainable from either variable dc-voltage or pulse-width control systems, the two best known alternative types of power control system. To understand why, we must digress slightly and look at the basic principles of motor-control and at the two alternative control systems.

Permanent magnet, dc electric motors of the types used in car fans, mini-drills and model locomotives, etc., are configured in exactly the same way as a dynamo. Consequently, when running, they generate a 'dynamo voltage' that opposes the externally applied voltage. Figure 4 shows the effect that this generated dynamo voltage (GDV) has on a dc-voltage control system when a 12 volt motor is powered from a six volt source.

When lightly loaded, the motor runs at medium speed and produces a GDV of five, which opposes the externally applied 6 V and gives an effective applied voltage (EAV) of a mere 1 V, so the motor consumes a fairly low running current (equal to EAV divided by the motor resistance).

When motor loading is increased, speed and GDV inevitably decrease, causing EAV and running current to increase, thereby tending to return the motor speed to its original value.

In Figure 4c, for example, the motor drive power (proportional to the square of EAV) is sixteen times greater than in Figure 4a; dc-voltage control systems are thus inherently 'feedback speed-sensing' and provide dc motors with excellent speed *regulation* characteristics. Unfortunately, however, their speed *control* characteristics are very bad at low and starting speeds.

An alternative way of controlling motor speed is to feed it with variable pulses of power. Most pulse-control systems feed fixed peak-amplitude, fixed-frame, variable-width pulses to the motor, and give good starting and speed control characteristics, but at the expense of poor regulation.

Figure 5 illustrates the reason for the poor regulation, assuming that the pulse has a peak amplitude of 12 volts and a frame width of 7 ms.

Figure 5a shows that, to give the same GDV of 5 V, mean EAV of 1 V and mean terminal voltage (MTV) of 6 V as in Figure 4a, the width of the Figure 5a pulse must be 1 ms. Consequently, when the motor is loaded so that its GDV falls to 4 V (Figure

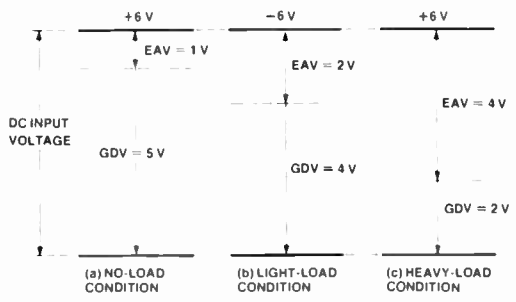

Figure 4. A voltage-controlled dc motor has excellent regulation because its speed-dependant 'Generated Dynamo Voltage' (GDV) decreases as the load increases and thus causes the Effective Applied Voltage (EAV) to increase. Its low-speed control is however, very poor.

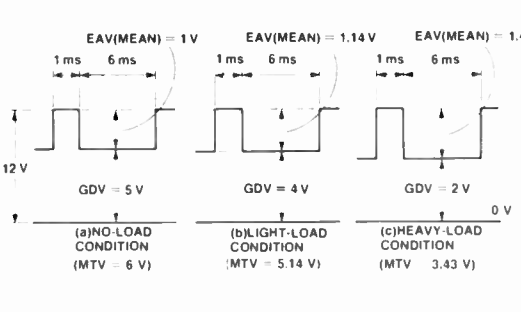

Figure 5. A conventional 'fixed-frame' pulse-controlled dc motor has good low-speed control, but poor regulation.

Figure 6. A SMVF-controlled dc motor has excellent regulation (equal to that of a dc-controlled system) and superb speed control (far better than that of a fixed-frame system).

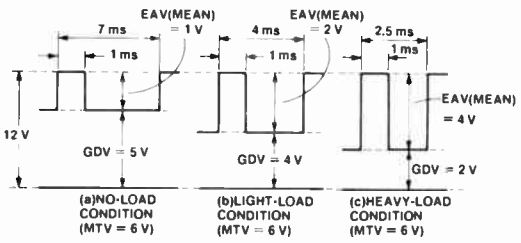

5b), this same pulse width gives an EAV of only 1.14 V and a MTV of only 5.14 V.

When the loading is increased so that GDV falls to 2 V, the MTV falls to a mere 3.43 V and the applied power is only twice as great as in the unloaded case.

Conventional fixed-frame pulse-width control systems thus have limited 'feedback speed-sensing' characteristics and give poor speed regulation, with the MTV falling as the motor loading increases.

The basic action of the new SMVF circuit, on the other hand, is such that it is fully 'feedback speed-sensing' and maintains its *mean* output voltage constant, irrespective of loading variations, as shown in Figure 6. It provides regulation that is as good as that of a dc-voltage control system, but with speed control that is greatly superior to that of a conventional pulse-control system.

In Figure 6a, to give the same GDV, EAV and MTV as the Figure 5a circuit, the widths of the mark pulse and frame are again 1 ms and7 ms, respectively. As the motor loading increases, however, the frame width reduces to maintain the MTV at a constant 6 V.

The SMVF circuit thus gives the same excellent speed regulation as a dc-voltage system. This regulation automatically occurs because, during the space part of each operating frame, the 'aiming voltage' of C1 (Figure 1) is equal to the motor's GDV, rather than zero volts.

The low speed control of the SMVF circuit is greatly superior to that of a conventional fixed-frame pulse-control system. This is because, at starting, the motor's applied pulse width must be greater than a certain minimum value, or the motor will not turn, but at medium to high speeds the frame width must be below a certain level to ensure smooth running.

Suppose that the minimum useful pulse width is 10 ms, and that at medium speed the maximum useful frame width is 100 ms. In a fixed-frame system, this means that the MTV of the circuit can be usefully varied over only a 10:1 range, giving (since speed is proportional to the square of MTV) a 100:1 speed-control range.

In the SMVF system, however, the frame width is variable, and may be one second at low speed but 100 ms at medium speed, in which case the MTV can usefully be varied over a 100:1 range, giving a 10 000:1 speed control range.

PRACTICALITIES

So far, I've dealt in depth with the basic theory of the switched-mode voltage follower circuit and seen that it gives superb dc-motor speed-control and excellent 'feedback sensing' speed regulation. Later in this article I'll introduce an 'amplified feedback' system that provides even better speed regulation.

In the meantime, let's look at some of the practical aspects of the basic SMVF circuit, dealing with matters such as component selection, power boosting and overload-protection techniques, and a few practical circuits.

Component selection

The pulse widths of the Figure 1 SMVF circuit are determined (apart from the input voltage) by the values of feedback components R4-C1 and hysteresis components R2-R3. Widths can be increased by increasing the values of R4, C1 or R3, or by reducing the value of R2. Note, however, that R2 and R3 also influence the 'tracking error' between input and output, the error increasing as the hysteresis voltage increases (the R2-R3 ratios reduce).

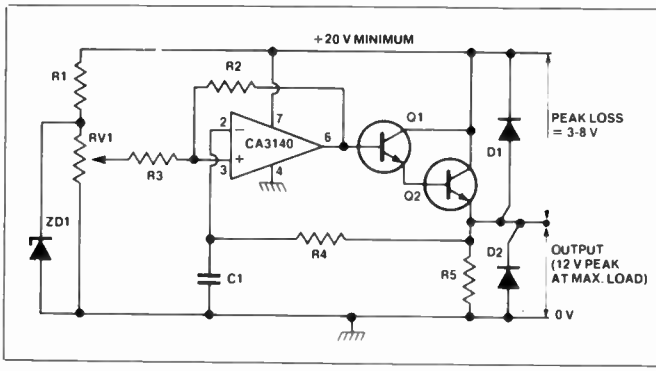

Figure 7. SMVF circuit with power boosting via a Darlington emitter follower.

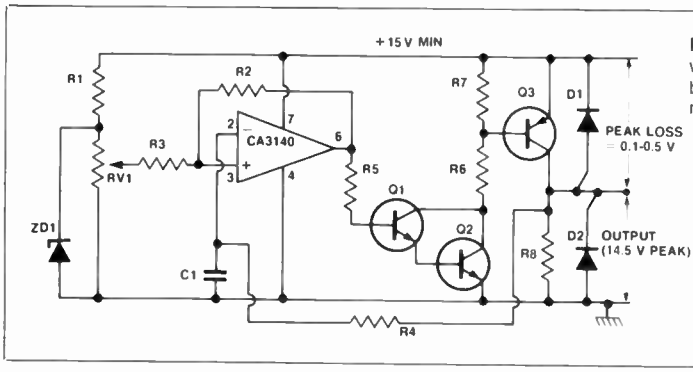

Figure 8. SMVF circuit with high-efficiency power-boosting via a pair of common emitter amplifiers.

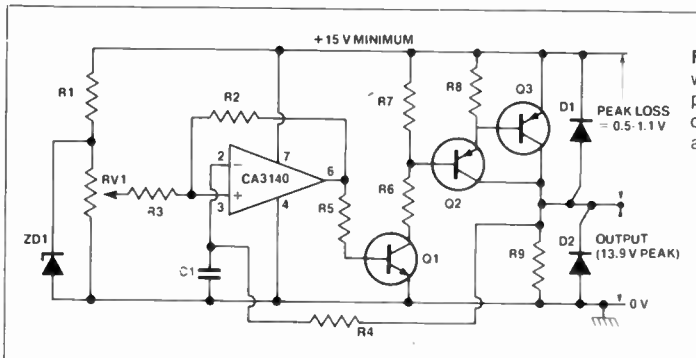

Figure 9. SMVF circuit with slightly less efficient power-boosting via a pair of common emitter amplifiers.

Thus, if (in Figure 1) R3 is reduced to 10 k, hysteresis falls to 20 mV and the tracking error at 1 V input falls to about 20 mV, but the C1 (or R4) value must be raised by a factor of ten to restore the pulse width.

The R3 value should ideally be large relative to the RV1 impedance, otherwise RV1 will affect the pulse width and hysteresis as its value is varied. If R2 is very large (greater than 1 M) it should be shunted by a 10 pF capacitor, to give sharp op-amp switching.

When driving dc motors, the minimum low-speed pulse width must be adjusted to suit the motor characteristics and adjustment is best effected by replacing R4 with a fixed and a variable resistor in series.

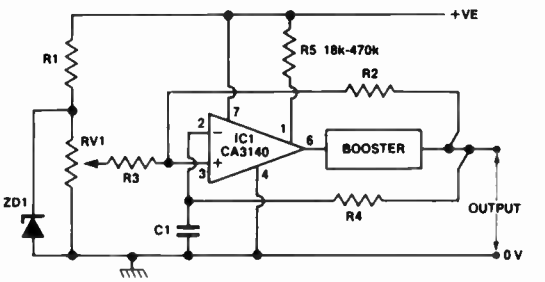

Figure 11. Offset biasing via pin 1 of the op-amp.

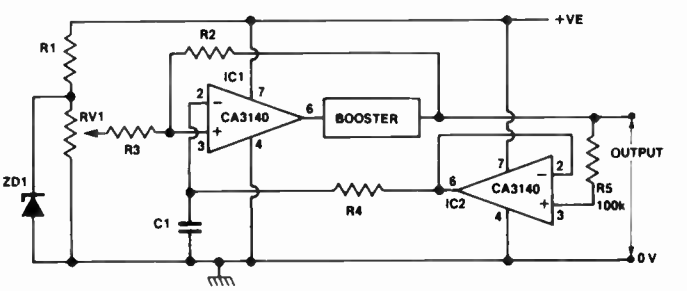

Figure 12. Offset biasing via an op-amp buffer.

Boosting the power

In Figure 1, a power booster is shown interposed between the op-amp output and the final output of the unit, to boost the available output current to a useful level. Since the SMVF circuit has a switched-mode output, this booster does not need to have linear characteristics, and its only essential requirements are that it should give zero overall phase inversion and should provide the required output current without excessive voltage loss.

Figures 7 to 9 show the basic circuits of three useful power boosters in use with SMVF circuits. Note that if these boosters are used to drive dc motors or other inductive loads, diodes D1 and D2 (with mean ratings at least equal to 30% of the peak motor current) must be wired to the circuits as shown, to protect the output transistors from the switch-off back-emfs of the motors, which may be as great as 100 volts.

In Figure 7, the booster takes the form of a Darlington emitter-follower. A disadvantage of this circuit is that a substantial amount of voltage (and thus power) is 'lost' across output transistor Q2 when the op-amp output is high (in the mark period). This loss is about 3 V off-load, rising to several volts under heavy loading. The circuit thus needs a fairly high supply voltage (at least 20 volts for 12 volts peak output) and good heat sinking of Q2.

Figure 8 shows an alternative booster, which provides very high efficiency, with minimal power loss across the output transistors. The op-amp output is inverted by the Q1-Q2 Darlington common-emitter amplifier, the output of which provides base drive to PNP common emitter output transistor Q3, giving zero overall phase inversion between the op-amp output and the load voltage.

The voltage 'loss' of the output is equal to the saturation voltage of Q3. If the Q3 base current is equal to at least one tenth of the peak load current, this loss may be as low as 100 mV off-load and 500 mV at maximum load.

The only disadvantage of the Figure 8 circuit is that Q3's base-current limiting resistor, R6, must pass a fairly high current and have a substantial power rating. The circuit's minimum supply voltage value is determined by the op-amp's requirements and by the output voltage requirement. With a 15 volt supply, the Figure 8 circuit can provide a maximum mean output of 12 volts and a peak output of 14.5 volts.

Figure 9 shows an alternative but slightly less efficient version of the Figure 8 configuration. In this case a Darlington pair are used as the output stage, and a single common-emitter amplifier is used as the driver.

The advantage of this circuit is that R6 has to pass only a fairly low current and can have a low power rating. The disadvantage is that the saturation loss of Q3 is greater than in Figure 8, being about 0.5 V off-load and 1.1 V at maximum load. With a 15 volt supply, the peak output is thus limited to 13.9 volts at full load.

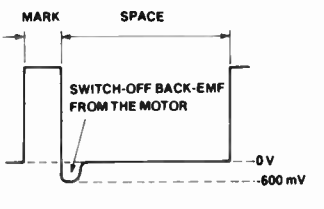

Figure 10. The negative back-emf of a motor prevents the SMVF circuit from turning fully off when its input is reduced to zero.

Offset biasing

A point to note about the basic SMVF circuit is that, when driving a motor or other inductive load, it does not turn fully off when the input is reduced to zero. To obtain switch-off, the voltage on the op-amp's non-inverting terminal must fall permanently below C1's lower space value, which is actually slightly negative under 'zero input voltage' motor driving conditions.

Figure 10 shows that this negative space voltage occurs because, at the end of each MARK pulse the motor produces a negative 'switch-off' back-emf that, even when clamped by an output diode, has a peak negative value of 600 mV which, when integrated by C1-R4 (Figure 1) equals a mean negative value of several millivolts.

To obtain complete switch-off from the SMVF circuit therefore, some form of 'offset biasing' must be applied to the design. Any one of three techniques can be used, as shown in Figures 11 to 13.

In Figure 11, offset biasing is applied to the internal circuitry of the op-amp via pin 1. The biasing resistor value must be in the range 18 k to 470 k (typically 100 k), its value being selected to suit the individual op-amp.

In Figure 12, offset biasing is obtained by wiring a 3140 unity-gain non-inverting amplifier (voltage follower) between the output of the circuit and the input of R4-C1. The op-amp output can not fall below a few millivolts positive, so (as far as R4-C1 are concerned) it effectively elimin-

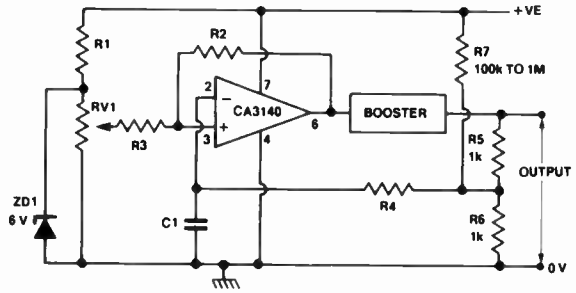

Figure 13. Offset biasing via potential divider.

ates the negative back-emf of the motor; R5 protects the op-amp input from excessive voltages.

Finally, Figure 13 shows a particularly useful offset biasing technique in which biasing is applied to the C1-R4 'integrator' via potential divider R7-R6, ensuring that the C1 voltage never falls to zero. Resistors R5-R6 form a 2:1 potential divider across the output of the circuit and feed C1-R4, so it is this point (rather than the output) that is *directly* regulated by this circuit. Since the output voltage is double that on R5-R6 junction, however, the output is *indirectly* regulated by this design. Note that the maximum input (RV1) voltage is 6 V, giving a maximum final regulated output of 12 volts.

Note, in Figure 13, that hysteresis components R2-R3 are fed from the op-amp's output, rather than from R5-R6 junction. Consequently, the peak voltage feeding R2 is double that on R5-R6 junction, so this design gives better pulse-width linearity than the basic Figure 1 circuit, and can be given almost any desired degree of offset via R7.

Practical SMVF circuits

Figures 14 and 15 show practical SMVF circuits that use techniques already described. These circuits are designed to provide fully variable power to 'fixed' loads such as lamps, heaters, blower motors and mini-drills, etc. They are not provided with short-circuit protection and are thus not suitable for driving loads such as model locomotives, etc., in which output shorts are likely to occur.

The Figure 14 circuit is designed for use in cars and is battery powered at 12 volts. It can supply maximum output currents of about five amps. Power boosting is achieved via a PNP output transistor and a Darlington driver, to give minimal power losses in the output stage. The output stage base drive is 250 mA, giving an output saturation 'loss' of only 500 mV at 3 A load, under which condition 34.5 watts are developed in the load and 1.5 watts are lost across Q2. Capacitor C3 is wired across Q1 base to enhance circuit stability.

The circuit uses potential divider offset biasing (via R9 and R11), with a 2:1 divider (R10-R11) across the output and with the maximum input limited to 6.2 V via ZD1. The maximum mean output is thus limited to 12.4 volts. At 12.6 battery volts, this is beyond the circuit's control range when RV1 is set to maximum, so the op-amp output locks high and turns Q2 fully on, giving maximum power drive to the load. If the battery is close to 15 V (under full charge), the output limits to 12.4 volts.

Preset RV2 is wired in series with R4, enabling the circuit's pulse widths to be preset to suit specific applications. When driving resistive loads such as lamps, RV2 can be preset to zero. When driving dc motors, RV2 should be set so that the motor turns with slight 'judder' (known as 'cogging') at minimum speed.

Figure 15 shows a mains-powered version of the SMVF circuit. It can power dc motors (such as mini-drills) with stall currents of up to 3 A; under the latter condition, the output impedance of the 30 VA mains transformer causes its output to fall to about 12 volts, giving about 30 watts of dissipation in the motor.

This circuit is similar to that of Figure 14, except that it uses a Darlington output stage with a single driver. This configuration results in a Q3 voltage loss of about 1.1 at 3 A load, but eliminates the need to give R7 a high power rating. The supply line has a value of about 20 volts off-load, and if the Figure 14 configuration were used here, R7 would need a power rating of 10 watts.

The supply to the op-amp is ripple-reduced via D1-C2, and the dc supply to the entire circuit is derived from a centre-tapped mains transformer and full-wave rectifier, rather than a single-ended transformer and bridge rectifier, since the latter option would result in an additional supply-volts drop of about 600 mV.

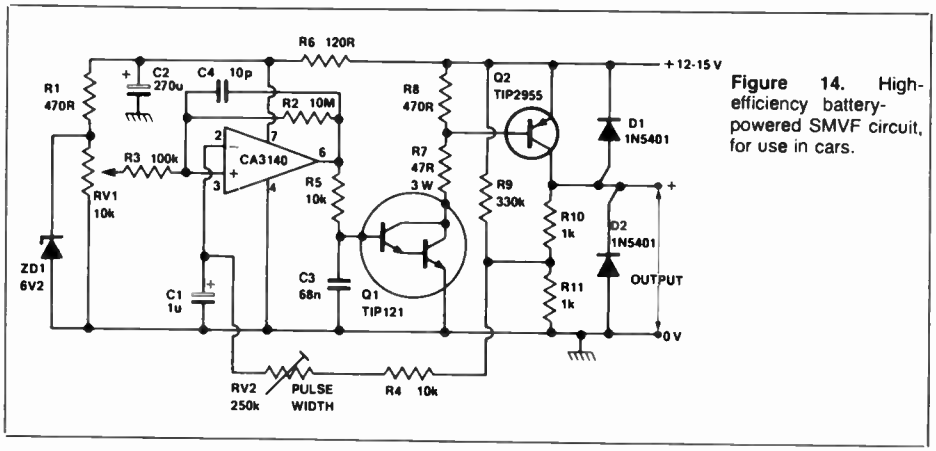

Figure 14. High-efficiency battery-powered SMVF circuit, for use in cars.

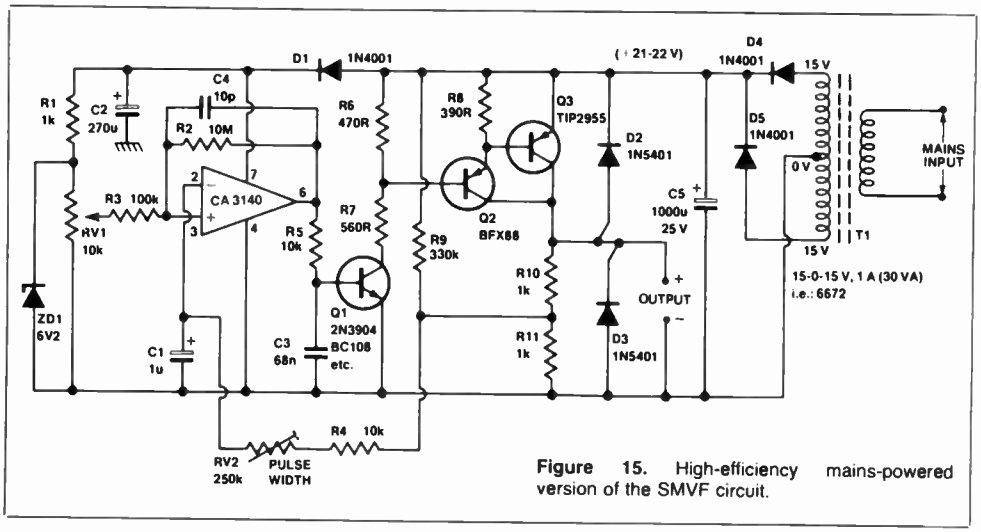

Figure 15. High-efficiency mains-powered version of the SMVF circuit.

Overload protection

The Figures 14 and 15 circuits are intended for use in fixed-load applications, and are provided with no form of overload or short-circuit protection. Such protection can be given by the various 'add-on' circuits shown in Figures 16 to 19.

Figure 16 shows a simple load-current limiter. The load current flows through monitor resistor R1, developing a current-proportional voltage. When this monitor voltage exceeds the base-emitter voltage (about 600 mV) of Q4, Q4 is biased on and starts to rob base-drive current from the Darlington output stage. The Darlington stage and Q4 form a negative feedback loop, causing the output current to self-limit at a value determined by R1 (about 3 A when R1 is 0.22 ohms). With a 12 volt supply and 3 A current limit, Q3 dissipates about 35 watts under short-circuit conditions, and must have a suitable heatsink.

Figure 17 shows a useful modification of the above circuit, in which two transistors (Q4 and Q5) turn on and under the 'overload' condition. Transistor Q4 limits the output current as already described, and Q5 activates an LED (or, better still, an audible alarm), to give a warning of the overload condition. Resistors R4 and R5 ensure that both transistors receive equal base-current drive.

At 'starting' speeds, the peak current of a pulse-driven dc motor is equal to its stall current, so a minor defect of Figure 17 is that Q4 and Q5 both pulse on at 'start' if the R1 value is such that it causes limiting at or below the stall current value of the motor. This defect is overcome in the circuit of Figure 18.

Here, the current is limited by either Q4 or Q5 turning on, but Q4 turns on at a peak current of about 9 A (via potential divider R4-R5), while Q5 turns on at a mean current of 3 A (via the R6-C1 integrator and R7).

If a short occurs at the output, Q4 instantly limits the current to 9 A peak, and a few tens of milliseconds later Q5 turns on and reduces the current to 3 A. The optional transistor Q6 can be used to activate a LED, indicating the short/overload condition.

Finally, Figure 19 shows an even more sophisticated circuit, in which the peak value is limited to 3 A (or whatever value is desired) via R1 and Q4, but the mean output current falls to only 3 mA or so under short-circuit conditions, thus eliminating the need for heavy heatsinking of Q3.

The circuit operates as follows: Both Q4 and Q5 turn on when an overload occurs. Transistor Q4 limits the peak output current to about 3 A, as already described, but the output of Q5 pulls R6 high and triggers the CMOS monostable multivibrator IC1, which applies a 500 ms positive pulse to the bases of Q6 and Q7 via limiting resistors R9 and R10. As Q6 turns

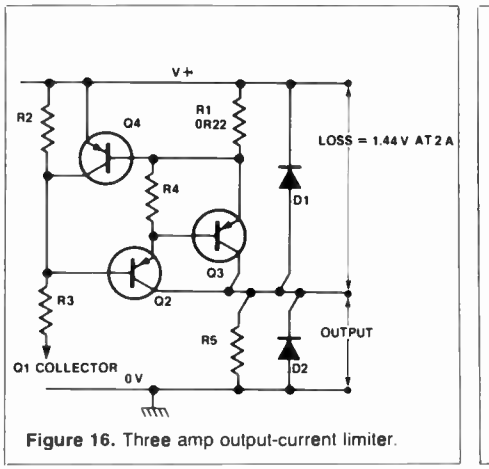

Figure 16. Three amp output-current limiter.

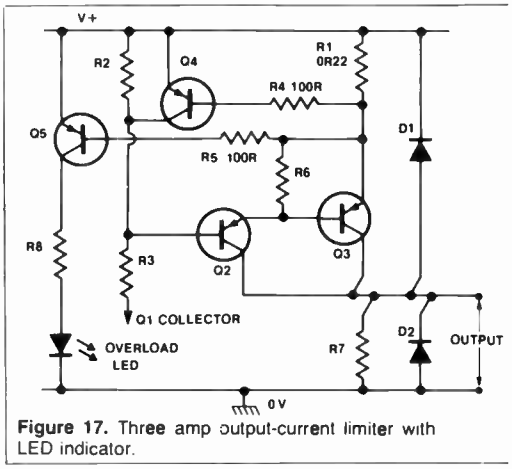

Figure 17. Three amp output-current limiter with LED indicator.

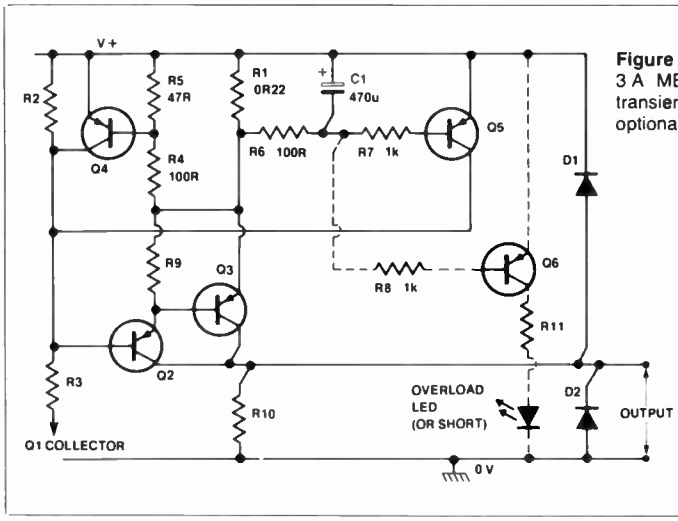

Figure 18. Circuit limits current to 3 A MEAN but can supply brief transients of 9 A. LED indicator is optional.

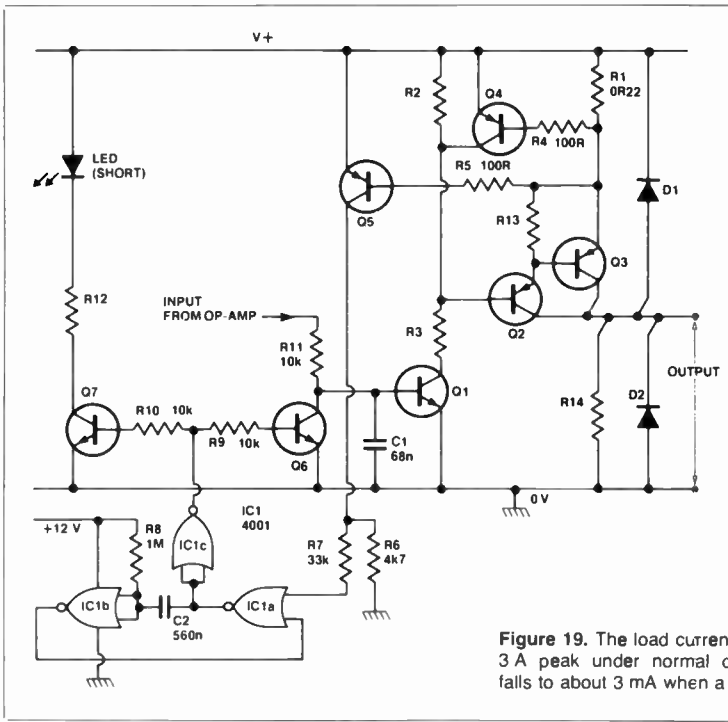

Figure 19. The load current of this circuit limits to 3 A peak under normal driving conditions, but falls to about 3 mA when a short occurs.

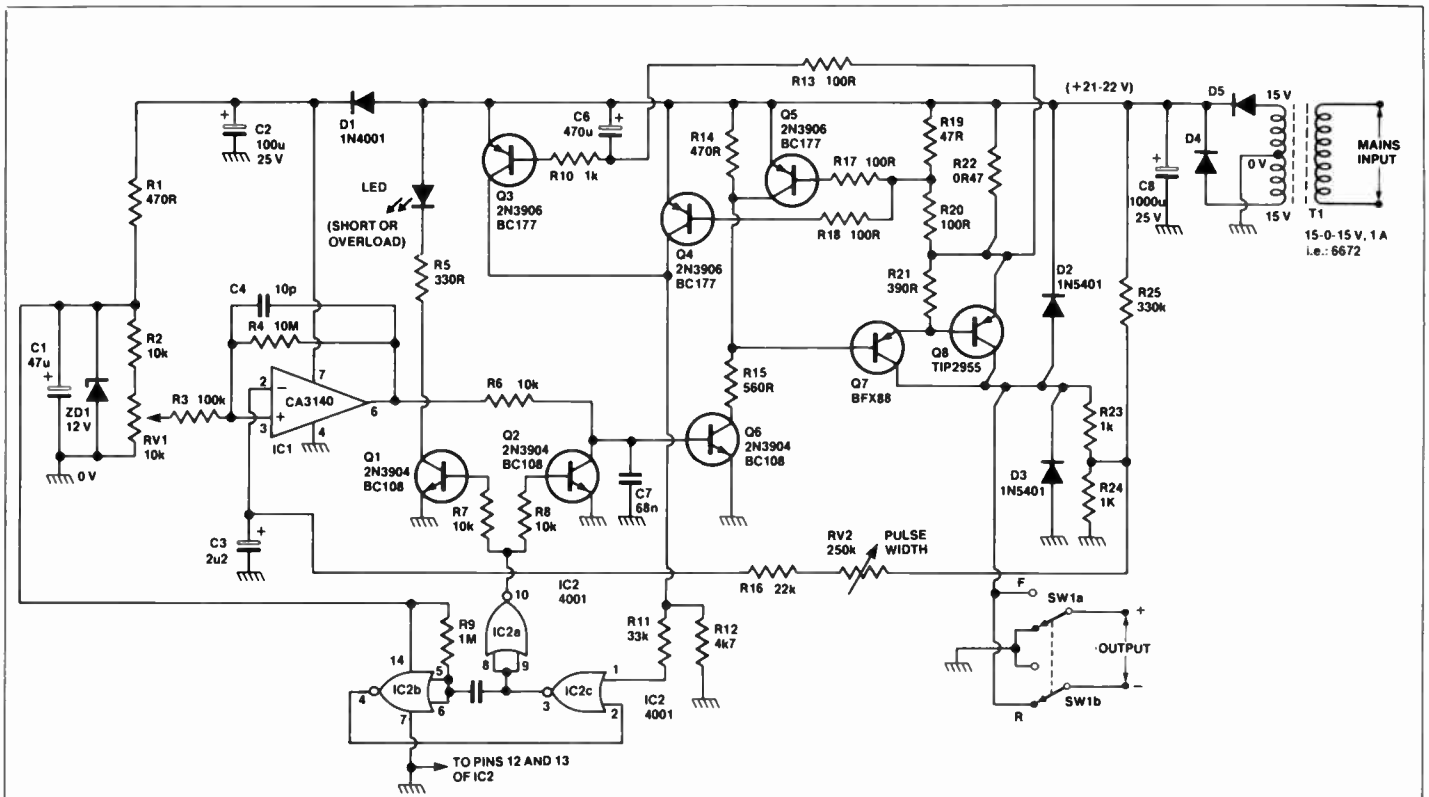

Figure 20. High-performance SMVF model-locomotive controller.

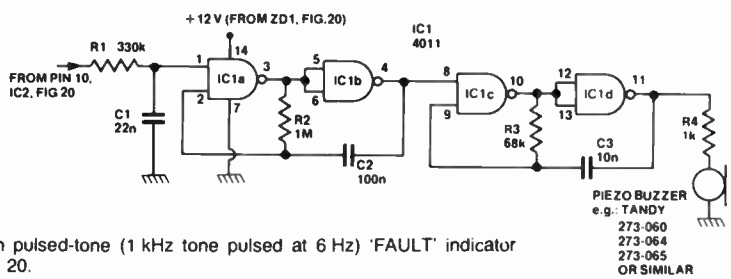

Figure 21. Add-on pulsed-tone (1 kHz tone pulsed at 6 Hz) 'FAULT' indicator for use with Figure 20.

on it robs Q1 of all base drive, causing the Darlington output stage Q2-Q3 to turn fully off, and as Q7 turns on it activates the short-circuit LED.

At the end of the 500 ms period, Q6 and Q7 turn off and the Darlington stage is re-enabled. If a short or overload still exist, however, the monostable fires again and turns the Darlington off for another 500 ms.

Suppose, then, that the 'delays' of the circuit are such that they cause an effective 5 us delay before the monostable activites. In this case, in each 'overload' cycle, the output is 3 A for 5 us and zero for 500 ms, giving a mean output current of only 3 mA.

Note that the Figures 16 to 19 circuits are all shown in use with a Darlington output stage in which the Darlington's base drive current is only a few tens of milliamps. These circuits can all be adapted for use with single-transistor output stages, but in such cases they may have to cope with base-drive currents of a few hundred milliamps. Also note, in all cases, that the presence of R1 causes a slight reduction in the maximum full-load output voltage that is available from the circuits.

Model locomotive controller

Figure 20 shows how the basic SMVF technique can be combined with overload-protection techniques to make a high performance model locomotive speed controller. The circuit has extensive overload and short-circuit protection, with visual fault indication. The maximum output current is limited to 1.5 A mean, 4.5 A peak, enabling locos to be run double- or triple-headed with ease.

If a moderate overload occurs, LED1 flashes on and off once a second, and regulated power is similarly pulsed to the track, so the loco s can still be controlled.

If an output short occurs, LED1 switches fully on and the mean output current reduces to about 3 mA; power restores automatically within 500 ms of the short being removed.

Switch SW1 enables the direction of the locos to be varied with ease, RV1 is the speed control, and RV2 enables the minimum pulse widths to be preset to suit the individual locomotive.

The Figure 20 circuit has one unique feature that has not yet been mentioned. If two identical units are made, and their outputs are shorted together in the same polarity, the unit with the highest mean output voltage automatically causes the other unit to shut down.

This feature greatly simplifies model railway control, since to change from one track to another (each provided with its own controller) it is simply necessary to throw the points and let the controller with the highest setting take over, rather than to go through a complicated procedure of sequentially operating the controllers and points, as in a normal system. The two SMVF controllers must, of course, be provided with independent power supplies.

The Figure 20 circuit is provided with a purely visual form of alarm indication (a LED). The control operation can be greatly enhanced, however, by also providing the circuit with an audible form of fault indicator, and Figure 21 shows a suitable circuit that can be powered from the 12 volt ZD1 supply and triggered from the output of IC2c (pin 10). The alarm generates a powerful pulsed-tone in a high-efficiency acoustic transducer.

AMPLIFIED-FEEDBACK TECHNIQUES

The basic SMVF circuit provides exceptionally good speed control of dc electric motors that is greatly superior to that of a conventional pulse-width controller and speed regulation that is as good as that of the very best dc-voltage control units.

In regulation, the SMVF system actually

regulates the output terminal voltage, holding its mean value constant and automatically varying the motor's *effective* applied voltage to suit varying load conditions, which are sensed via the motor's speed-dependent 'dynamo' voltage.

In the final analysis, the effectiveness of this form of motor-speed regulation is determined by the efficiency of the motor when it is functioning in the 'dynamo' mode. If dynamo efficiency is 100%, regulation efficiency will be infinitely good, and if dynamo efficiency is zero the regulation efficiency will be infinitely bad.

In practice, dynamo efficiency typically varies from about 40% in model locomotives to about 80% in good-quality minidrills.

It follows from the above that, if we could somehow amplify the *apparent* dynamo efficiency of a motor to (say) 95%, and then feed that signal to the SMVF circuit in place of the true dynamo voltage, the SMVF circuit would automatically provide that motor with near-perfect speed/load regulation, even if the motor were of the very poorest quality.

In practice, this action can actually be achieved by using an 'amplified-feedback' technique. Figure 22 shows the basic circuit, and Figure 23 shows the circuit waveforms.

Figure 22 is similar to the basic SMVF circuit, except that peak-voltage limiter D1-R5 is wired between the output and RV1 slider, and the R4-C1 integrator is fed from the D1-R5 junction. If we assume that the forward volt drop of D1 is 500 mV, that R5 is very large relative to RV1, that R4 is very large relative to R5, and that RV1 slider potential is at 2 V, it can be seen that during the mark part of the cycle the voltage at point 'A' limits to 2.5 V peak, while in the space part of the cycle the waveform at 'A' is equal to the motor's speed-dependent dynamo voltage.

Figure 23a shows the waveforms and formulae of the circuit under this condition. Note that the mark slope at point 'B' is fixed (at any given input voltage), but that the space slope depends on the speed-dependent dynamo voltage of the motor; the slope increases (and the space period decreases) as V_D decreases, so the mean output voltage at point 'C' increases as the motor voltage falls, thereby raising the motor drive and tending to return the speed to its original value.

The circuit actually regulates the mean voltage at point 'A', holding this value very close to the input voltage. The voltage at point 'A', however, is predominantly determined by the dynamo voltage, so the action of this 'amplified feedback' version of the SMVF circuit is such that it holds the generated dynamo voltage (and thus the motor speed) constant at the same value as the input voltage (from RV1 slider), auto-

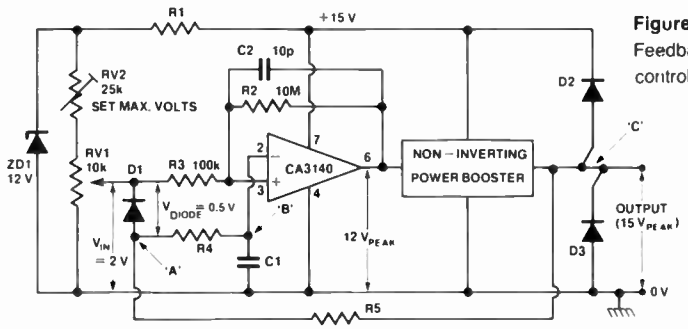

Figure 22. Basic 'Amplified-Feedback' SMVF motor-controller circuit.

$$\text{MEAN VOLTAGE AT 'A'} = V_D + \left(\frac{[V_{IN} + V_{DIODE}] - V_D}{\left[\frac{M_{SLOPE} + S_{SLOPE}}{S_{SLOPE}} \right]} \right)$$

$$\text{MEAN VOLTAGE AT 'C'} = V_D + \left(\frac{15\,V - V_D}{\left[\frac{M_{SLOPE} + S_{SLOPE}}{S_{SLOPE}} \right]} \right)$$

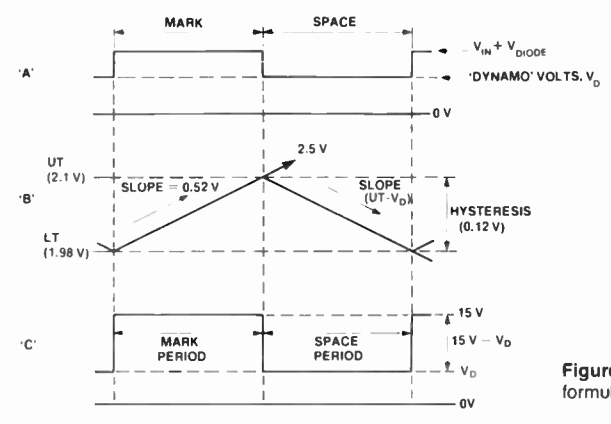

Figure 23a. Basic waveforms and formulae of the Figure 22 circuit.

SPEED-DEPENDANT 'DYNAMO' VOLTAGE	MARK SLOPE (VOLTS)	SPACE SLOPE (VOLTS)	MARK PERIOD (ms)	SPACE PERIOD (ms)	V_{MEAN} AT 'A'	V_{MEAN} AT 'C'	APPLIED-POWER FACTOR	NORM-ALISED POWER FACTOR
1.95 V	0.52	0.097	1.92	10.3	2.036 V	4.00 V	4.2	1.0
1.90 V	0.52	0.2	1.92	5.0	2.07 V	5.54 V	13.25	3.15
1.50 V	0.52	0.6	1.92	1.67	2.035 V	8.72 V	52.1	12.4
1.0 V	0.52	1.1	1.92	0.91	2.02 V	10.52 V	90.6	21.6
0 V	0.52	2.1	1.92	0.48	2.00 V	12.0 V	144	34.3

Figure 23b. Waveform values of the Figure 22 circuit, at V_{IN} = 2 V, at various 'dynamo' values.

SPEED-DEPENDANT 'DYNAMO' VOLTAGE	MEAN OUTPUT VOLTAGE USING: V_{DIODE} = 0.25 V	V_{DIODE} = 0.5 V	V_{DIODE} = 1.0 V	STANDARD SMVF CIRCUIT
1.95 V	4 V	4 V	4 V	4 V
1.90 V	7.47 V	5.54 V	4.05 V	4 V
1.50 V	10.8 V	8.72 V	6.5 V	4 V
1.0 V	12.2 V	10.52 V	8.25 V	4 V
0 V	13.3 V	12.0 V	10.07 V	4 V
REGULATION FACTOR	42.1	34.3	24.1	3.8

Figure 24. Performance of the Figure 22 circuit (at V_{IN} = 2 V) at three values of V_{DIODE}, and of the 'standard' SMVF circuit.

V_{IN}	MARK SLOPE	SPACE SLOPE	MARK PERIOD ms	SPACE PERIOD ms	V_{OUT}
0.5 V	0.505 V	0.615 V	1.98	1.62	8.24 V
1.0 V	0.51 V	1.1 V	1.96	0.91	10.25 V
2.0 V	0.52 V	2.1 V	1.92	0.47	12.02 V
4.0 V	0.54 V	4.08 V	1.85	0.24	13.25 V
6.0 V	0.56 V	6.06 V	1.78	0.16	13.74 V

Figure 25. Performance of the Figure 22 circuit, at various input voltages, when feeding a resistive load (roughly equivalent to a 'stalled' dc motor).

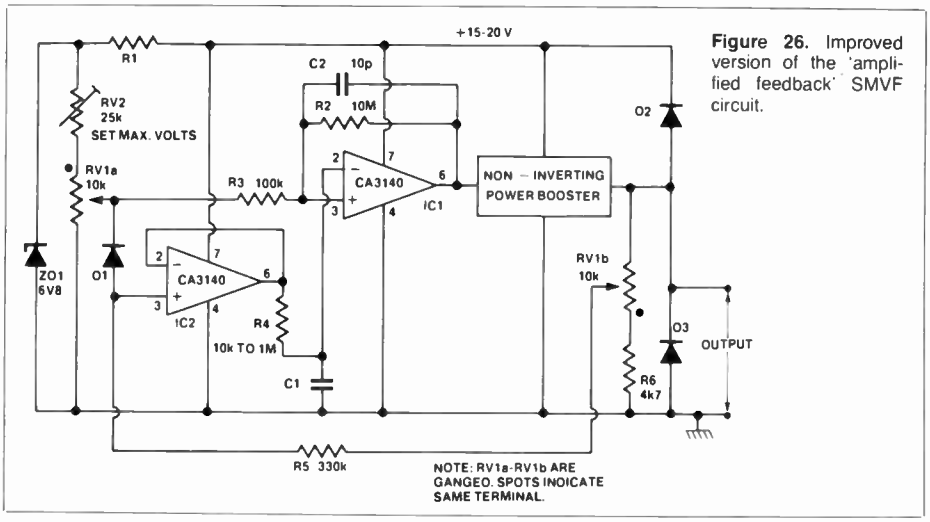

Figure 26. Improved version of the 'amplified feedback' SMVF circuit.

matically adjusting the mean output voltage to obtain this action.

Since V_D is inevitably less than the applied motor voltage, the maximum RV1 value must be less than 12 volts (when driving a 12 volt motor); the preset RV2 allows the maximum voltage to be set.

Figure 23b shows actual waveform values of the Figure 22 circuit at various 'dynamo' values. Note that when V_D is below 1.95 V the space slope is large relative to the circuit's hysteresis voltage, so the slope is virtually linear. At 1.95 V however, the true 'linear' slope would actually be only 0.15 V (compared to a hysteresis value of 0.12 V), so the slope becomes exponential and takes up an effective value of 97 mV.

This exponential characteristic gives very sharp regulation near the 'lower threshold' value, which determines the maximum attainable dynamo voltage.

In Figure 23b, we've introduced an 'applied-power factor', which is the square of the effective applied voltage (mean output less V_D), together with a 'normalised' value in which all values are divided by 4.2. Note that, when the motor is stalled, the effective applied power is 34.3 times greater than when the motor is lightly loaded; the basic SMVF circuit would increase the applied power by a factor of only 3.8 under the same circumstances.

The regulation efficiency of the Figure 22 circuit is, in fact, heavily dependent on the forward voltage (V_{diode}) of the limiting diode or device. The Table of Figure 24 shows the output voltages obtained from the Figure 22 circuit, at various values of V_D, with three different values of V_{diode}. The table also shows the output of the basic SMVF circuit under the same conditions, and the 'regulation factors' (ratio of unloaded-to-stalled effective applied voltages).

Note from the table of Figure 24 that the 'stalled' output voltage depends on the value of V_{diode}. In practice, the stalled voltage is almost the same as the output voltage that is obtained with a purely resistive load.

The table of Figure 25 shows the performance of the circuit (with a V_{diode} of 0.5 V) at various values of input voltage when driving a resistive load.

The basic Figure 22 circuit suffers from a couple of practical defects, and Figure 26 shows how these can be overcome. The first defect is that, since R5 must be very large relative to RV1, and R4 must be very large relative to R5, unrealistically large R4 values are called for.

In Figure 26, this problem is overcome by interposing a unity-gain 3140 non-inverting buffer amplifier between the D1-R5 junction and R4, enabling R4 to be given any value in the range 10k to 1M. A useful benefit of this buffer is that it eliminates the need for additional offset biasing circuitry, as described earlier.

The second defect concerns the 'ceiling ratio' of the circuit. That is, the value of V_{diode} relative to V_D (as seen by R4 input), expressed as a percentage of V_D. Thus, in Figure 22, when V_D is 1 volt, V_{diode} is 0.5 V, giving a ceiling ratio of 50%, but when V_D is 5 V, V_{diode} is still 0.5 volts, giving a ceiling ratio of only 10%.

In practice, because of the 'dirty' nature of dynamo voltage, motor regulation becomes highly erratic if ceiling ratios fall significantly below about 15%.

In model railways, because contact is made with the motor via rotating wheels on dirty track, ceiling ratios of at least 30% are needed for stable operation.

Thus, for stable operation over the entire speed-control range, without loss of regulation sensitivity at low speeds, the ceiling ratio needs to be stabilised over the entire V_{in} range.

In Figure 26 this is achieved via RV1b, which is ganged to RV1a but connected in anti-phase. Thus, when RV1a is set low to give only 1 V input, RV1b is high and feeds almost the full 1 V V_D to R5, giving a ceiling ratio of 50%. When RV1a is set high to give (say) a V_D of 5 V, RV1b is low and feeds only one third of the 5 V V_D to R5, giving a ceiling ratio of 28%, and hence very stable operation and regulation.

Note that, because of the 'divider' action of RV1b, the circuit needs an input voltage (at maximum RV1a setting) of only 1.67V to produce a V_D of 5 V.

The Figure 26 circuit can be implemented using any of the power booster and overload-protection circuits described earlier, and gives superb motor speed-control and regulation. Appendix A shows the practical implementation of the system as a high-performance mini-drill controller, and Appendix B shows it implemented as an ultra-sophisticated model locomotive controller. ●

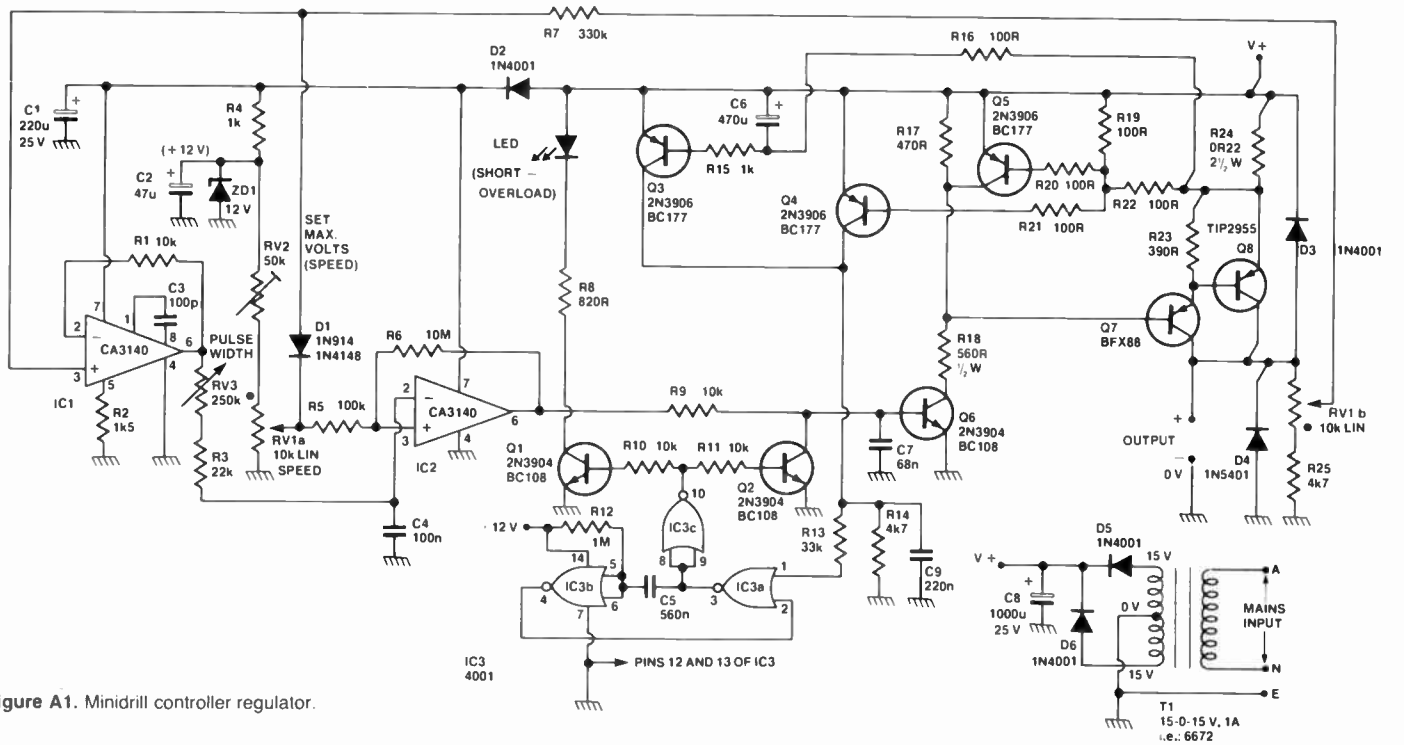

Figure A1. Minidrill controller regulator.

APPENDIX A

A Mini-drill controller-regulator

MINI-DRILLS are widely used in pc board production and in model engineering. They are normally designed for use with nominal 12 volt power supplies and draw maximum 'stall' currents of up to two or three amps. The speed of such drills can easily be varied by adjusting the mean voltage value that is fed to the drill motor, and many such 'mini-drill controller' circuits have been published in recent years.

Most previously published designs have, however, been of either the 'variable dc voltage' or the 'fixed-frame variable pulse-width' types, and have given less-than-best speed regulation, for reasons that have already been fully explained.

This new mini-drill controller, by contrast, works on the 'amplified-feedback switched-mode voltage follower' principle, and inherently gives fully-variable speed control with superb self-adjusting speed regulation.

The basic action of this circuit is such that it continuously compares the speed-dependent 'dynamo' voltage of the drill motor with the dc voltage that is set on an input control pot, and automatically adjusts the mean output voltage (to the motor) to maintain the dynamo voltage and speed at the desired value, irrespective of loading conditions.

Thus, the circuit may supply only 2 V (mean) to a mini-drill when it is running lightly-loaded at low speed, but may increase the mean drive to (say) 10 volts to maintain that speed when the loading becomes very heavy.

The new circuit feeds a variable-frame switched-mode output signal to the mini-drill. Power is actually fed to the motor (as a pulse of energy) during the mark part of each cycle, and under this condition the circuit's output transistor is fully saturated and dissipates very little power; the system is thus highly efficient.

During the space part of the cycle the circuit monitors the motor's dynamo voltage and thus measures the true speed of the motor; the results of this measurement determine the length of the space period and hence the mean power that is fed to the motor in that cycle.

The new controller/regulator circuit is shown in Figure A1. This particular design is specifically intended for use with '12 volt' mini-drills that draw maximum mean 'stalled' currents of up to 3 A. The circuit can be modified to accept motors of even heavier rating by simply reducing the value of one resistor (R24) and increasing the power rating of the supply transformer.

The Figure A1 circuit incorporates comprehensive 'overload' and 'short-circuit' protection, and has a visual fault indicator (a LED). The overload protection circuit comes into operation if mean load currents exceed 3 A, and under this condition power to the mini-drill and the LED is pulsed on and off once per second.

The short-circuit protection comes into operation if a resistance of less than two ohms nominal appears across the output terminals, and in this event causes automatic shut down that gives a mean output current of only 3 mA. In this mode, which is indicated by the LED turning fully on, the circuit samples the output once every 500 ms to see if the short still exists, and finally restores full power within 500 ms of the short being removed.

Construction and use

Construction should present very few problems. Note, however, that RV1 is a dual-gang linear pot, and that the two halves are wired to the circuit in anti-phase (as indicated by the 'spots' which indicate the fully anti-clockwise terminal). Also note that monitor resistor, R24, needs a power rating of 2.5 W. Output transistor Q8 dissipates about 3 W under maximum load conditions, and should be fitted with a small heatsink.

The power supply connections to the main circuit should be made at the specific points indicated, with the positive connection going to the 'hot' side of R24 and the negative to the negative output terminal.

The circuit has two control pots (RV1 and RV3) and one preset pot (RV2); RV2 and RV3 must be adjusted to suit the characteristics of the individual mini-drill motor that is in use, as follows:

(1) Set RV1 to zero, RV2 to mid value, RV3 to maximum. Connect the mini-drill and an analogue voltmeter (to read 12 volts nominal) across the output terminals.

Switch on and wind RV1 up to maximum value, noting that the drill speed increases. With RV1 at maximum, and with the mini-drill unloaded, adjust RV2 to give a reading of 12 volts on the meter.

(2) With RV1 still at maximum, reduce the value of RV3 (pulse width) until 'skip cycling' or erratic operation begins to occur and the output voltage begins to rise.

Increase the RV3 value slightly, so that 'skip cycling' stops, and note the RV3 setting. Now repeatedly and slowly vary the value of the speed control, RV1, from maximum to minimum to maximum, etc., to find the 'ideal' setting of RV3 (somewhere greater than the above-noted value) at which skip-cycling is minimal and drill rotation is acceptably smooth.

Note that, when the drill is unloaded, rotation will inevitably feel slightly rough at very low speed settings, but this roughness diminishes when the drill is loaded.

(3) Once the 'ideal' setting of RV3 has been found, adjust RV1 to give an output voltage reading of about 2 V with the drill unloaded. Now increase the loading and check that the output voltage increases and the drill speed stays almost constant.

Finally, stall the motor and check that the output rises to between six and 10 volts, indicating that the circuit is functioning correctly. All adjustments are then complete and the unit is ready for use.

Note that, once RV2 and RV3 have been initially set to suit a particular mini-drill, they are unlikely to ever need re-adjustment while that drill is in use.

If you wish to make a more powerful version of the Figure A1 circuit, simply increase the rating of the mains transformer and reduce the value of R24 to give the desired 'limit' current.

Circuit theory

The circuit of Figure A1 is simply a practical version of the 'amplified-feedback SMVF' circuit given in Figure 26 of the main article. IC1-IC2-RV1, etc., form the basic circuit, but with slight additional offset biasing given by R2 to compensate for poor quality RV1 pots that do not quite give zero resistance in the 'low' position.

The IC1-IC2 supply is decoupled from the main supply via D2-C1, and the switched-mode output of IC2 is fed to Darlington output transistors Q7-Q8 via the Q6 driver stage.

The circuit has comprehensive output protection, via monitor resistor R24. Mean output currents are sensed via Q3 and integrating components C6-R16; Q3 turns on when mean current exceeds 3 A and thence activates monostable IC3. Peak currents are sensed via the R19 and R22 divider, which feeds the bases of Q4 and Q5. When peak currents reach 6 A, Q5 starts robbing base drive from Q7-Q8, and Q4 turns on and fires the IC3 monostable.

When the IC3 monostable fires, it activates the fault indicator LED via Q1, and simultaneously removes all Q6 base drive via Q2, thereby disabling the output transistor for the duration of the 500 ms monostable pulse.

In the event of a shorted output, this circuitry limits the mean output current to about 3 mA.

Note that C9 reduces the rise times of the monostable trigger signals, enabling the circuit to power slightly capacitive loads without triggering the short circuit protection.

APPENDIX B

A model-locomotive controller-regulator

THE VERY SOPHISTICATED model-locomotive controller/regulator described here uses the amplified-feedback switched-mode voltage follower principle and gives a degree of speed control and regulation that is vastly superior to those available from any commercial unit or from any previously published designs, to my knowledge.

The controller can smoothly vary loco speeds all the way from an 'imperceptible' 10 mm/minute up to the maximum '12 volts' value, irrespective of whether the loco is unloaded or hauling a heavy load.

The unit can drive up to three locos simultaneously and is ideal for double- or triple-headed operations. Most important of all, the unit incorporates fully automatic speed regulation circuitry which continuously monitors the speed of the loco's motor via its 'dynamo effect' voltage and thence adjusts the power feed to hold the speed constant at the desired value, irrespective of loading conditions or the state of the track.

Thus, if the loco is running around a track hauling a load at a particular speed, and the loco then starts to run down an incline the circuit automatically reduces the power drive to hold the motor (and thus the drive-wheel) speed constant.

If the loco starts to climb an incline, the unit automatically increases the power drive to maintain the speed. If the incline is so steep that the loco cannot climb it, the power drive simply increases to such a level that the drive wheels spin at the same speed as they would if the loco were running on the level or down-hill.

The 'regulation' circuitry is fully effective, all the way from low speeds up to maximum.

The controller/regulator has lots of other attractive features. It incorporates fully automatic output protection circuitry, with audio/visual fault indication. If output loading is excessive (greater than 1.2 A mean), the feed to the track and the fault indicator pulses on and off once per second.

If a short occurs across the track, shut-down circuitry automatically reduces the mean output current to 1.5 mA and the fault indicator turns fully on, with a LED illuminating and the audio circuit generating a pulsed-tone alarm signal. Full power is automatically restored within 500 ms of the short being removed.

Another unique feature of this model locomotive controller-regulator is that if two (or more) such circuits are built and are individually powered, and have their outputs shorted together in the same polarity, the unit with the higher mean output voltage automatically causes the other unit to shut down.

This feature greatly simplifies model railway control, since to change a loco from one track to another it is simply necessary to throw the track points and let the controller with the higher setting take over, rather than to go through a complicated procedure of sequentially operating the controllers and points, as in a normal system.

If the two outputs are shorted together in opposing polarities, both units simply register short circuits and automatically shut down.

Construction and use

The circuit of the main unit is shown in Figure B1. Construction of this part of the unit should present few problems. Note, however, that RV1 is a dual-gang linear pot, and that the two halves are wired to the circuit in anti-phase (as indicated by the spots which show the fully anti-clockwise terminal). Resistor R24 is the current-monitor resistor, and needs a rating of 1 W. Output transistor Q8 dissipates negligible power and does not need a heatsink.

The power supply circuit is shown in Figure B2, and should be connected to the main circuit at the points indicated, with the positive connection to the 'hot' side of R24 and the negative to the negative output terminal.

When construction of the entire unit is complete, double-check all wiring and then test and adjust the unit as follows.

(1) Connect the unit's output to the track, and fit a loco in place. Turn RV1 slightly above zero until pulses of power can be heard reaching the loco. Adjust RV3 so that each pulse produces a small but finite movement of the loco wheels.

(2) Increase the RV1 setting, checking that the loco speed increases smoothly. Set RV1 to maximum and, with an analogue voltmeter connected across D5 (the output of the basic unit) adjust RV2 for a reading of 12 volts with the loco unloaded.

(3) Reduce RV1 to give an output of four to five volts. Run the loco on to a finger or buffer and check that the drive wheels keep turning at normal running speed, even though the actual loco is stopped, and that the output voltage increases. Stall the loco and check that the output increases further.

(4) Place a short across the track and check that the audio-visual fault alarm activates. Remove the short and check that the alarm condition ceases. All

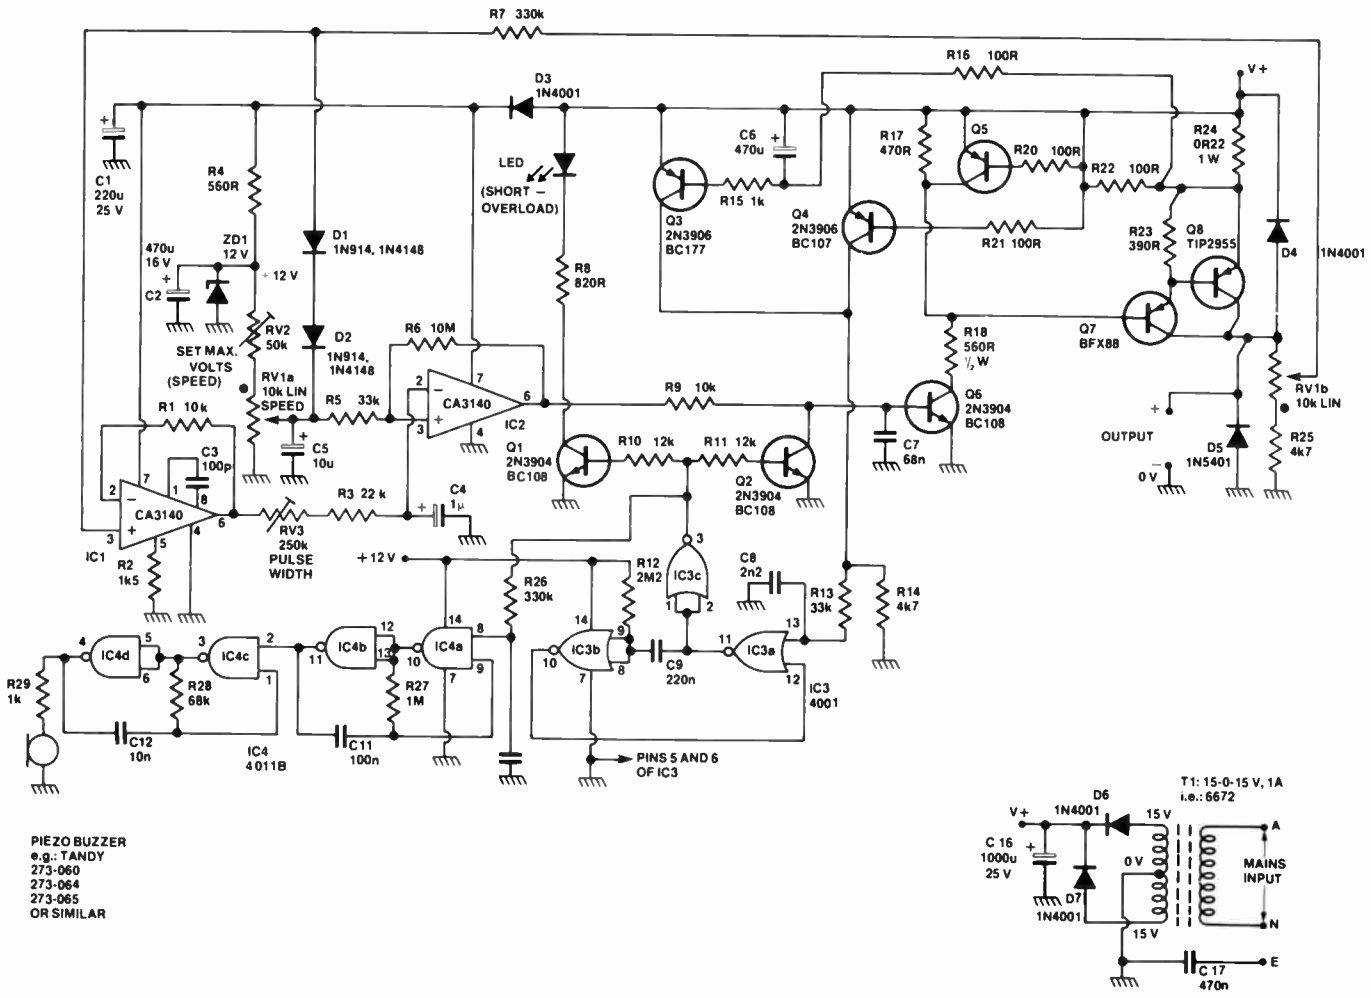

Figure B1. Model locomotive controller/regulator.

adjustments/tests are then complete, and the unit is ready for use.

Note that, once the above tests are complete, RV2 needs no further adjustment. RV3 may require re-adjustment to suit individual locos, and is simply set to give a definite but minimal motor movement with each pulse at minimum speed.

Circuit theory

The major part of the circuit is shown in Figure B1 and is simply a practical implementation of the 'amplified-feedback SMVF' circuit given in Figure 26 of the main article. IC1-IC2-RV1, etc., form the basic circuit, but with slight additional offset biasing given by R2, to compensate for poor-quality RV1 pots that do not quite give zero resistance in the 'low' position.

Note, however, that resistor R5 has a lower value than in Figure 26, to give a low value of hysteresis voltage and a wider-than-normal range of motor-speed control, and that diodes D1-D2 are wired in series to give a ceiling value of 1.2 volts, to accomodate the very 'dirty' track-derived loco motor signals.

The IC1-IC2 supply is decoupled from the main supply via C1-D3, and the switched-mode output of IC2 is fed to Darlington output transistors Q7-Q8 via the Q6 driver stage.

The unit has comprehensive output protection, via monitor resistor R24. Mean output currents are sensed via Q3 and integrating components C6-R16; Q3 turns on when mean currents exceed 1.2 A and thence activates monostable IC3.

Peak currents are sensed via the R19 and R22 divider, which feeds the bases of Q4 and Q5. When peak currents reach 3.6 A, Q5 starts robbing base drive from Q7-Q8, and Q4 turns on and fires the IC3 monostable.

When IC3 fires it activates the fault indicator LED via Q1, and also audio alarm via Q2, thereby disabling the Q7-Q8 output circuitry for the duration of the 500 ms monostable pulse.

In the event of a shorted output, this circuitry limits the mean output current to about 1.5 mA. Note that C8 reduces the rise times of the monostable trigger signals, enabling the circuit to power slightly capacitive loads without triggering the short circuit protection circuitry.

How to use digital voltmeter modules

Part 1

The popular DPM-05 digital voltmeter module with 3½-digit liquid crystal display is a very convenient 'workhorse' for a myriad of applications. This two-part feature tells you how to put it to use. The ETI-161 Digital Panel Meter is very similar and can be used in many of the circuits given.

Ray Marston

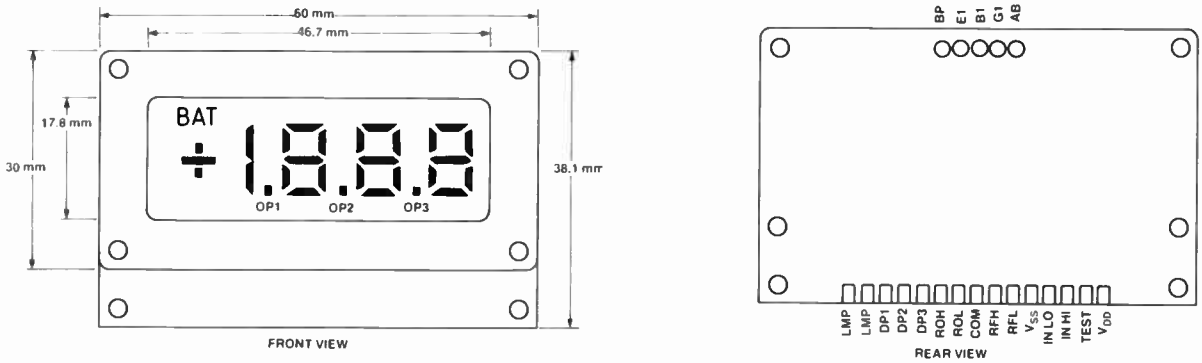

Figure 1. Physical details (left) and terminal notations (right) of the DPM-05 module.

MODERN DIGITAL VOLTMETER (DVM) modules can be used to replace moving coil meters in virtually all important 'analogue' measuring applications. Most of these modules combine an Intersil ICL7026, 7126 or 7136 analogue-to-digital (A-D) converter chip and a 3½-digit liquid crystal display plus a band-gap voltage reference and a few other components, into a compact module that consumes less than 1 mA from a 9 V supply and costs little more than a good quality moving coil meter.

Usually, these modules have a basic full-scale measurement sensitivity of ±199.9 mV, with 100 uV (2000-count) resolution and a typical calibrated accuracy of 0.1% ±1 digit, but can be used to read any desired current or voltage range by connecting suitable shunts or potential dividers to the input terminals. When connected to suitable external circuitry, the modules can be made to indicate ac voltage or current, resistance, capacitance, frequency, temperature, or any other parameter than can be converted into a linear analogue voltage or resistance. I'll show you how later in this two-part feature.

Several companies manufacture 3½-digit LCD DVM modules. Generally, these modules differ only in details of their internal circuitry and displays and in the number and notations of their user-available terminals. The DPM-05 module manufactured by Printed Circuits International Ltd, imported and distributed here by Jaycar, is probably

CHARACTERISTIC	DATA (at 25 C)
Display	3½ digit LCD
Full scale sensitivity	±199.9 mV
Power supply voltage	9 Vdc nominal (range 7 to 10 V)
Supply current	1 mA typ, 2 mA max
Initial calibration accuracy	better than ±0.15% of reading ±1 count
Zero-input reading	±000.0 typ
Display resolution	1 count = 100 μV
Input leakage current (at V_{in} = 0)	1 pA typ, 10 pA max
Operating temperature	0 C to +50 C
Clock frequency	40 kHz typ
Sample rate	2.5 readings/s
'Low battery' indication voltage	7.2 V typical

Table 1. Main parameters and features of the DPM-05 module.

Figure 2. Block diagram 'user view' of the DPM-05 module.

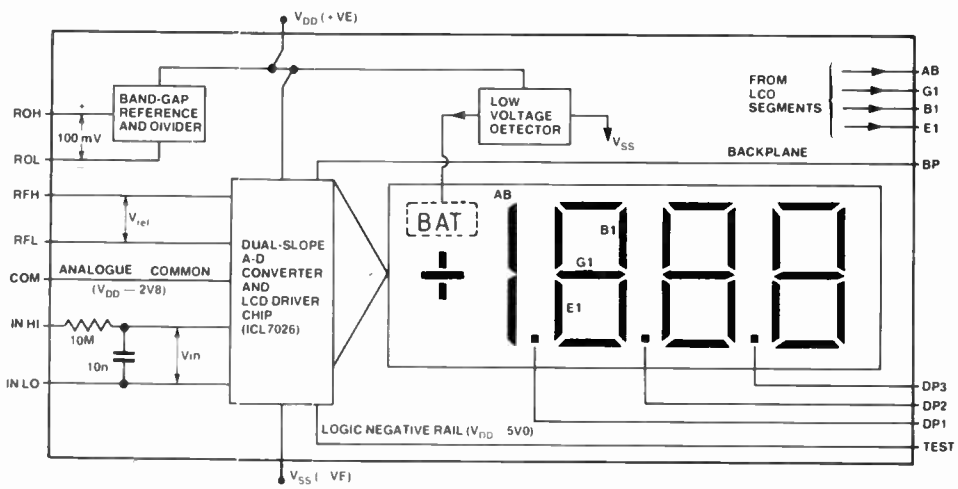

the best known and most widely available model, and is very typical of the genre, so we'll refer to this specific device throughout the rest of this article. Figure 1 shows the physical details and terminal notations of the DPM-05 and Table 1 lists its main parameters and features.

The ETI-161 Digital Panel Meter module is very similar. This was published in the August 1982 issue and kits are widely available. The accompanying panel shows the circuit and a rear view of the pc board with equivalent connections to those of the DPM-05 annotated. Note that the ETI-161 does not include the band-gap reference. Lab Notes in the November 1980 issue gives circuit details of a band-gap reference that could be adapted to the circuits in this feature, if necessary.

DPM-05 basics

Figure 2 shows the block diagram 'user view' of the DPM-05, which is normally powered from a 9 V battery connected between the V_{DD} and V_{SS} terminals. The heart of this particular unit is an ICL7026 chip which is a complete dual-slope analogue-to-digital converter and LCD driver. In essence, this chip automatically compares the relative values (ratios) of V_{ref} and V_{in} and produces an LCD display of 1000 x V_{in}/V_{ref}, updating the display about 2½ times per second.

Thus, if V_{ref} is 100 mV and decimal point DP3 is activated, the display reads 10.0 with an input of 10.0 mV, or 199.9 with an input of 199.9 mV. The module automatically displays the polarity of the input signal, gives automatic zero adjustment, and gives over-range indication by blanking the three least significant digits of the display. The three decimal points of the LCD are externally available at the DP1 to DP3 terminals, and can be turned on by pulling the appropriate terminal to V_{DD}. The module also houses a 'low battery' detector, which turns on an 'annunciator' in the display when the battery voltage falls below 7.2 volts.

It is important to note that the DVM module actually displays the relative *ratios* of the input and reference voltages. To give maximum versatility, each of the voltages is applied to the module via a *pair* of terminals (RFH and RFL for the reference, IN HI and IN LO for the input), and the integrator chip responds to the differential values of these inputs. In use, these terminals must be tied (either directly or indirectly) to within 500 mV of the COM terminal. When correctly used, the terminals have typical input impedances of about 5000 megohms, and pass typical leakage currents of only a few picoamps. The IN HI terminal incorporates an integrating ripple-reduction filter.

The module has two built-in reference-voltage sources. The voltage between the COM and V_{DD} terminals is zener-regulated at 2V8 and has a typical temperature coefficient of 80 ppm/°C, so any reference voltage below this value can be obtained by wiring a simple potential divider between

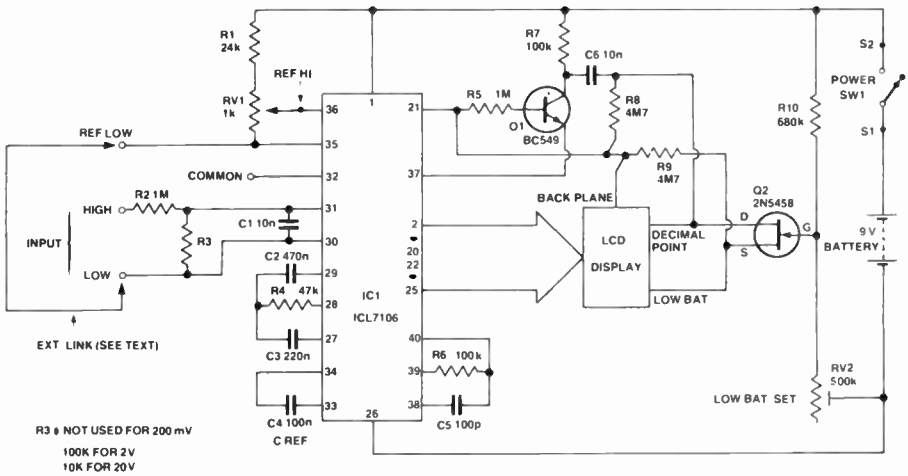

ETI-161 Panel Meter. Circuit of the ETI-161 panel meter project which can be used in almost all the applications circuits given in this two-part series. Many circuits require direct access to 'REF HI' (RFH), in which case delete R1 and RV1.

SPECIFICATIONS

Full scale readout	depends on setup. Full scale sensitivity is 199.9 mV
Resolution	100 uV
Accuracy	< 1 digit when correctly calibrated
Display	3½-digit LCD
Input impedance	> 10^{12} ohms
Input bias current	approx. 2 pA
Polarity indication	automatic
Conversion method	dual slope
Reference	internally generated
Power supply	±100 ppm 9 V @ approx. 1 mA

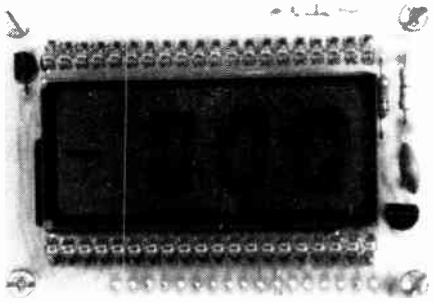

ETI-161. View of the Panel Meter project published in the August 1982 issue.

Connections. Rear view of the ETI-161 pc board(s) showing connections equivalent to the DPM-05 module.

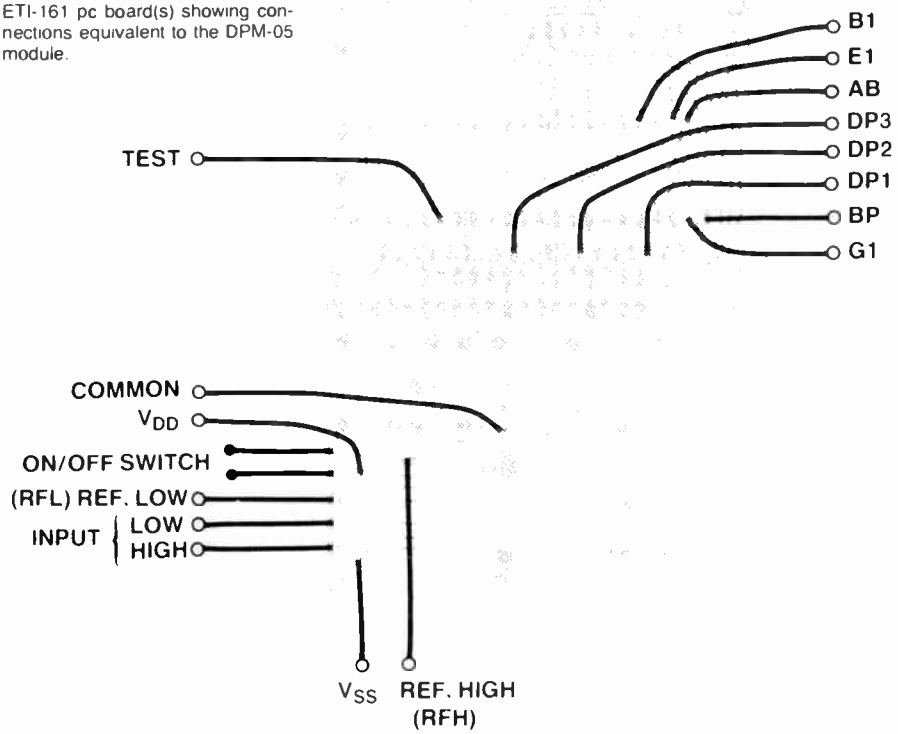

these terminals. The module also houses a precision band-gap reference. When ROL is tied to COM a stable 100 mV is generated between ROH and ROL and has a typical temperature coefficient of 50 ppm/°C.

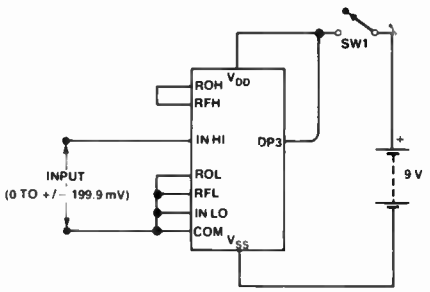

Figure 3. Standard '199.9 mV full-scale' connection of the DVM module.

Basic configurations

Figures 3 to 6 show four different ways of connecting the terminals of a DVM module to give different types of measurement action. Figure 3 shows the standard '199.9 mV full scale' DVM configuration. Here, the COM, IN LO, RFL and ROL terminals are all joined together, ROH is shorted to RFH so that the 100 mV band-gap reference is applied across the reference terminals, and decimal point DP3 is tied to V_{DD} so that the unit gives a reading of '100.0' when 100.0 mV is applied between IN HI and IN LO.

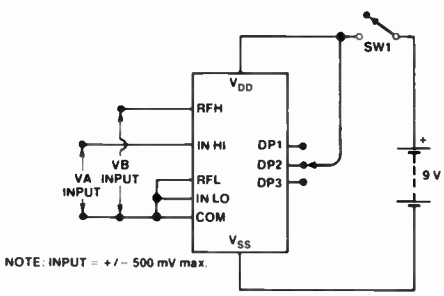

Figure 4. Basic ratiometric voltmeter connection. Display = 1000 x VA/VB.

Figure 4 shows the connections for making the module act as a ratiometric voltmeter which (ideally) gives a reading of '1000' when two input voltages have identical values, irrespective of the actual magnitudes of those values (up to a limit of 500 mV).

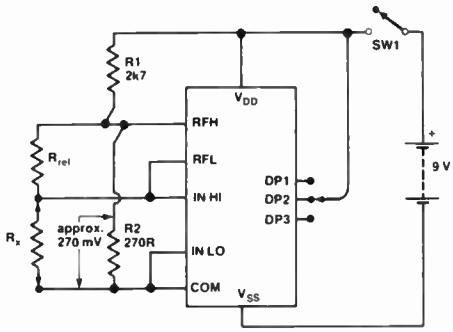

Figure 5. Precision resistance meter using ratiometric technique. Display = 1000 x R_x/R_{ref}.

Figure 5 shows the module connected as a precision ohmmeter. Here, potential divider R1-R2 generates roughly 270 mV between the R1-R2 junction and the COM terminals, and this voltage is used to energise potential divider R_{ref}-R_x. Identical currents flow through these two resistors, and the generated voltage of R_{ref} is applied across the RFH and RFL reference terminals, and the generated voltage of R_x is applied across the IN HI and IN LO input terminals. The display reading thus equals 1000 x R_x/R_{ref}. If R_x has a decade value (1k0, 10k etc), the display gives a direct readout of the R_x value, the reading being independent of the actual value of energising voltage developed across R2.

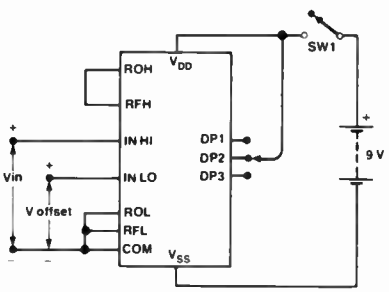

Figure 6. Method of applying zero-offset to the basic '199.9 mV' DVM circuit. Display = $V_{in} - V_{offset}$.

Finally, Figure 6 shows how an offset voltage can be applied to the basic 'DVM' circuit so that the display reads zero when the input voltage is at a value other than zero. This circuit is useful in temperature-reading applications for example, in which a special IC is used to give an output of 1 mV/°K, thus giving an output of 273.2 mV at 0 C and 373.2 mV at 100 C.

PERCENTAGE OF FULL SCALE	NOMINAL READING	TRUE READING ACCURACY A	B
100%	199.9 mV	±0.15%	±0.05%
50%	100.0 mV	±0.2%	±0.1%
25%	50.0 mV	±0.2%	±0.2%
10%	20.0 mV	±0.5%	±0.5%
5%	10.0 mV	±1.0%	±1.0%
1%	2.0 mV	±5.0%	±5.0%

Table 2. True reading accuracies of 3½-digit DVMs with calibrated accuracies of (A) ±0.1% and (B) ±0.01% of reading ±1 count.

By feeding the output of the IC between the COM and IN HI terminals and applying a 273.2 mV offset voltage between COM and IN LO, the module (which reads the *differential* value of the input) can be made to give a direct reading of temperature in degrees Centigrade.

Some finer points

If you intend to use a DPM-05 or similar module in a project, there are some fine 'usage' points that you will need to know. Let's deal with these points under various sub-headings.

Calibration accuracy. As supplied, a DVM module is pre-calibrated to read 199.9 mV full scale, with a typical accuracy of ±0.1% of reading ±1 count, at 25°C, this calibration being valid *only* when the module is used in the precise configuration shown in Figure 3. It should be noted that the best attainable accuracy of a 3½-digit (2000-count) meter is ±1 digit, and this corresponds to an actual reading accuracy of 0.05% at full scale, to 0.5% at 10% of full scale, and to 5% at 1% of full scale. Table 2 shows the reading accuracies of two meters, having different calibration accuracies, at various percentages of full scale.

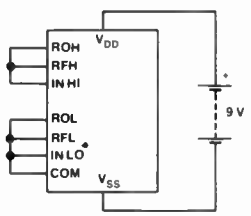

Figure 7. Ratiometric-accuracy test circuit. Ideally, the meter should read '1000'. Typically, the reading may be '998' (= 0.2% low).

Ratiometric accuracy. The DVM is a ratiometric reading unit. If connected as shown in Figure 7, with identical voltages applied to the RFH and IN HI terminals, it should ideally read '1000' ±1 count.

In practice, modules typically give a reading that is about 0.2% below this figure. This discrepancy is caused by the potential divider action of the internal 10M filter resistor and the input impedance on the internal IN HI line.

When the meter is supplied for use in the 'voltmeter' mode, it is calibrated to allow for ratiometric errors.

Reference accuracy. The built-in '100 mV' reference (between ROH and ROL) of the module is factory-calibrated so that the meter reads '100.0 mV' with 100.0 mV input applied. The precise value of the reference voltage depends on the ratiometric accuracy of the meter. Thus, if the ratiometric accuracy is 0.2% low (reading 998), the reference is also set 0.2% low (at 998 mV) to give the correct 'voltmeter' accuracy.

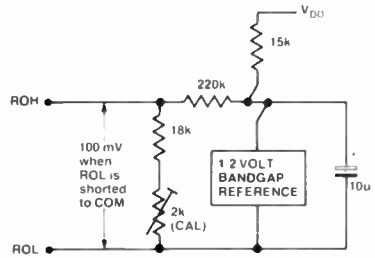

Figure 8. Typical bandgap reference circuit has an output impedance of about 20k.

input impedances and draw leakage currents of only a few picoamps. If the terminals are biased at voltages significantly different from COM, the input leakage currents may rise to several hundred picoamps, invalidating the auto-zero action of the chip. The chip may be damaged if the terminals rise above $V_{DD} - 0.5$ V or below $V_{SS} + 1$ V.

'COM' terminal. The COM terminal of the module is connected to the circuit of Figure 9 within the A-D chip, and this circuit enables the COM terminal to be used as either a precision voltage reference, as a current sink for external circuitry, or as an externally-biased analogue-reference point.

When used as a voltage reference, only very low external sink currents (below 100 uA) must be allowed to flow between V_{DD} and COM. Under this condition the basic calibration of the module is valid, and the COM terminal is held about 2V8 below V_{DD}, with a temperature coefficient typically less than 80 ppm/°C.

When used as a current sink, external currents of up to 30 mA can be allowed to flow between V_{DD} and the COM terminal (which has an impedance of about 15 ohms in this mode). In this mode, however, the basic calibration of the module may be invalid, and the RFH and RFL terminals may have to be driven from an external reference.

The COM terminal can source currents up to a maximum value of only 10 uA. Consequently, the common line of the A-D chip can be tied to a value that is more than 2V8 below V_{DD} by simply connecting the COM terminal to an external bias voltage of the required value. In this mode, the basic calibration of the module is invalid, and the RFH and RFL terminals must be driven from an external reference; the INPUT and REFERENCE terminals must be tied within 500 mV of COM (see Figure 12). Note that COM should not be allowed to fall to a value more than 4V7 below V_{DD}.

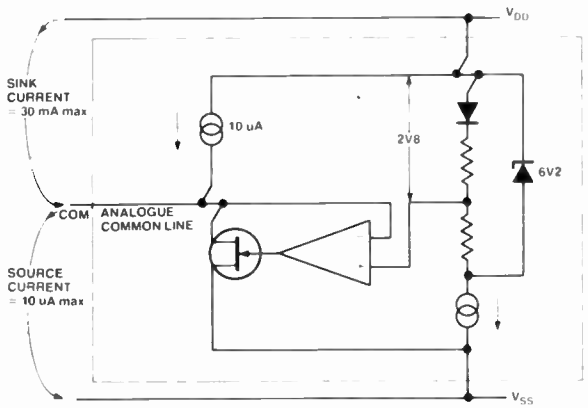

Figure 9. Analogue COMmon line biasing circuit within the A-to-D converter chip.

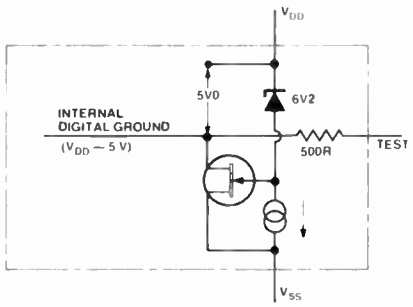

Figure 10. Internal digital ground biasing circuit of the A-to-D converter chip.

The reference output is accurate only when ROL is tied directly to COM (which is normally 2V8 below V_{DD}) and when ROH is loaded by an impedance greater than 50 megohms or so. Figure 8 shows the typical circuit of a band-gap reference. The output impedance of the circuit is about 20k so an external loading of 2M would introduce an error of 1%, and a loading of 20M an error of 0.1%. The high input impedance of the RFH terminal causes negligible loading.

Input connections. The A-D converter chip houses analogue and digital circuitry. All analogue action is internally referenced to the COM (common) line of the chip. Normally, the INPUT and REFERENCE inputs should be tied (directly or indirectly) to within 500 mV of the COM line, and under these conditions the terminals have very high

TEST & BP. The negative or ground rail of the digital circuitry of the A-D chip is internally biased at about 5 V below V_{DD} by the circuit of Figure 10 and is coupled to the TEST terminal via a 500 ohm resistor. This terminal can be used as the negative rail of external digital circuitry that is powered from V_{DD}, provided that the TEST currents do not exceed 1 mA.

If TEST is shorted directly to V_{DD} the LCD should read '−1888'; under this condition 10 mA flows into the TEST terminal and a steady dc voltage is applied to the LCD; this voltage may burn the display if sustained for several minutes.

The back-plane (BP) drive signal to the display switches fully between TEST and V_{DD} at the clock frequency divided by 800. With a 40 kHz clock, BP has a frequency of 50 Hz (giving a period of 20 mS). Note that the calibration accuracy of the module is independent of the clock frequency, which is thus not designed to be particularly stable.

Auxiliary terminals. The DPM-05 has a number of auxiliary terminals that are used only in special applications. The two LMP terminals give access to a backlight bulb fitted to the LCD in some special modules.

The AB terminal connects to the '1000' digit of the LCD, and the E1, B1 and G1 terminals connect to the E, B and G segments respectively of the '100s' digit of the LCD. These terminals can be decoded with the BP signal to detect the over-range (O/R) and under-range (U/R) states of the module and thence activate auto-ranging circuitry, etc. Figure 11 shows the external decoder circuit that must be used; the two IC's are powered from the V_{DD} and TEST terminals.

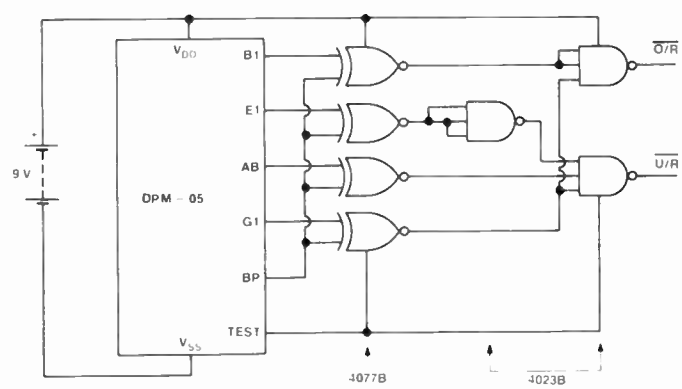

Figure 11. Circuit for developing under-range and over-range signals from the DPM-05.

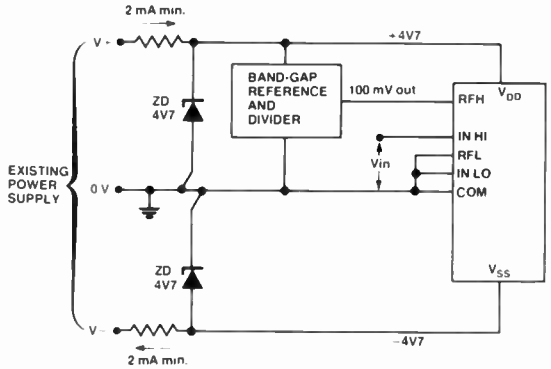

Figure 12. Method of building the module into existing equipment that is powered from split supply rails.

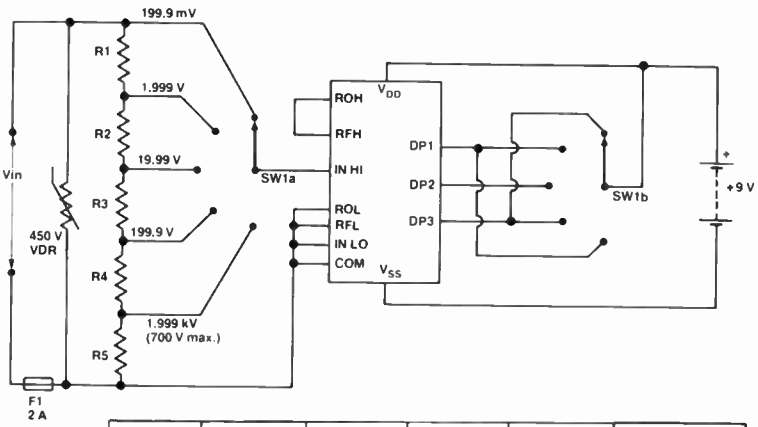

R1	R2	R3	R4	R5	INPUT Z
9M 10M	900k 1M	90k 100k	9k 10k	1k 1110R	10M 11.11M

Figure 15. Five-range dc voltmeter.

Power supplies

The most popular application of the DVM module is as a self-contained multimeter which is used as a general purpose test instrument. In this type of application the module is simply powered from a 9 V battery connected between V_{DD} and V_{SS}.

The module can, however, be built into existing equipment and used in dedicated measuring/indicating applications.

If the equipment is powered from a single-ended supply, the module must be powered from its own 'floating' supply, derived from either a battery or from a separate winding of a mains transformer.

In the case of a battery-powered instrument, the supply to the meter can be switched by a spare pair of contacts on the main switch.

If the equipment is powered from split supplies, the module can be powered from the existing power rails by using the connections shown in Figure 12, in which COM is tied to the common rail, V_{SS} is fed from −4V7, V_{DD} from +4V7, and the REFERENCE and INPUT terminals are referenced to the COM terminal. The RFH terminal must be driven from an external reference, as shown.

PRACTICAL APPLICATIONS

DC volt & current meters

The DVM module is supplied ready-calibrated to give a full scale reading of ±199.9 mV dc. The module can be made to give alternative full scale dc voltage readings by connecting the input voltage to the module via a decade potential divider, as shown in Figure 13, or can be made to act as a dc current meter by wiring a suitable shunt resistor across the input terminals, as shown in Figure 14. Note in both diagrams that the appropriate decimal point of the display must be tied high on each range, as indicated.

The module can be used as a five-range dc voltmeter by using the connections shown in Figure 15; the table shows alternative potential-divider component values to give input impedances of 10M or 11.11M.

Precision '9'-decade (9M, 900k, etc) resistors are used in most multimeters and are available from several component suppliers. Note that in multi-range applications the circuit should be provided with some form of overload protection, and in the diagram this is given by fuse F1 and by a voltage-dependent resistor (VDR) or 'transient suppressor' across the divider. Also note that on the '1.999 kV' range the maximum input is actually limited to 700 volts by the VDR.

The module can be used as a five-range dc current meter by using the connections shown in Figure 16. Note here that the generated voltages of the shunts are directly monitored by the DVM module, and that variations in the switch resistance of SW1a have no effect on the accuracy of measurement; a separate

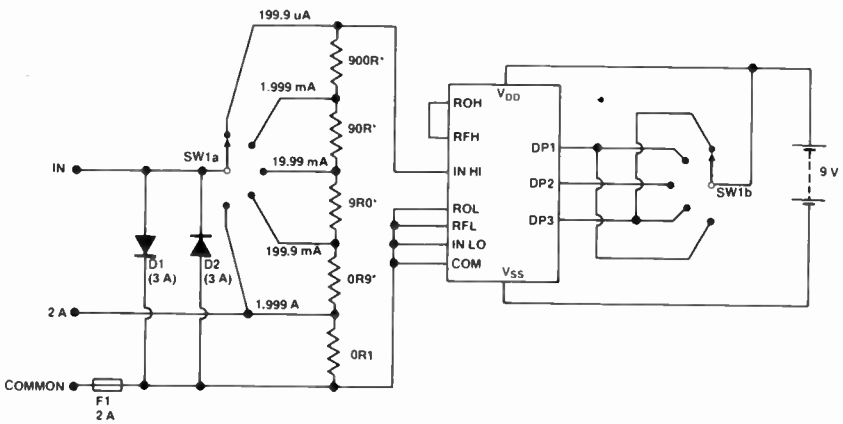

Figure 16. Five-range dc current meter.

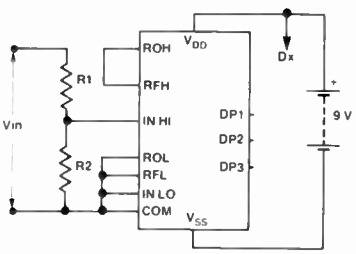

Figure 13. The DVM module can read alternative dc voltage ranges by connecting the input via a potential divider

R1	R2	V FULL SCALE	DECIMAL POINT HIGH
0	10M	199.9 mV	DP3
9M	1M	1.999 V	DP1
9M9	100k	19.99 V	DP2
10M	10k	199.9 V	DP3
10M	1k	1.999 kV	DP1

R1	I FULL SCALE	DECIMAL POINT HIGH
10k	19.99 μA	DP2
1k	199.9 μA	DP3
100R	1.999 mA	DP1
10R	19.99 mA	DP2
1R	199.9 mA	DP3
0R1	1.999 A	DP1
0.01R	19.99 A	DP2

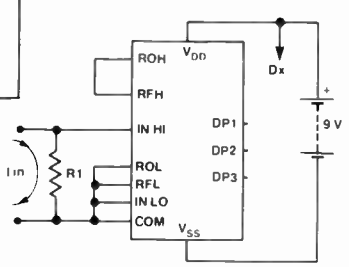

Figure 14. The DVM module can be made to read dc current by connecting a shunt resistor across its input.

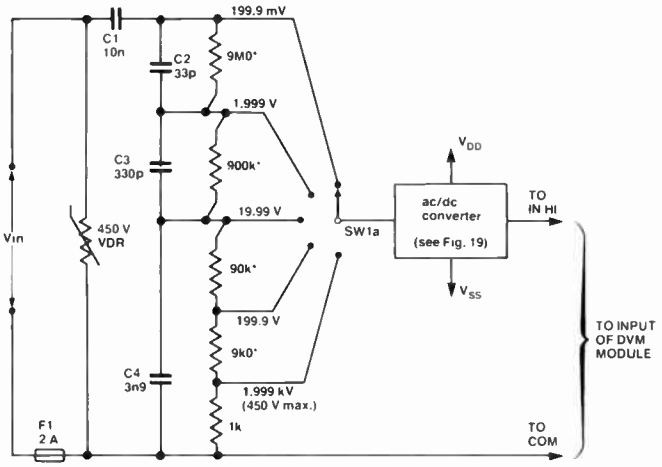

Figure 17. Modification of the Figure 15 circuit, to act as a five-range ac voltmeter.

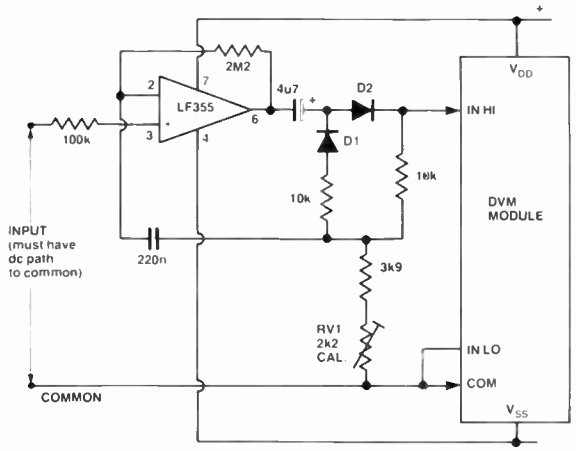

Figure 19. Precision ac/dc converter.

input terminal is used for the '2 Amp' measurement. The circuit is protected against positive and negative overloads by diodes D1-D2 and fuse F1.

AC volt & current meters

Figure 17 shows how the Figure 15 circuit can be modified to act as a five-range ac voltmeter that has a frequency response flat within 1 dB to about 120 kHz.

Input signals are fed to the attenuator via dc-blocking capacitor C1, and the attenuator is frequency compensated by C2 to C4. The attenuator output is fed to the input of the module via a precision ac/dc converter, which gives a dc voltage output equal to the RMS value of a sinewave input.

Figure 18 shows how the Figure 16 circuit can be similarly modified to act as a five-range ac current meter. In this case it is not feasible to prevent dc currents feeding into the shunts: instead, dc-blocking is done at the output of the shunts via C1-R1, and the resulting ac signals are fed to the input of the DVM module via a precision ac/dc converter.

Note that the input protection network of this circuit differs from that of Figure 16 in that pairs of diodes are wired in series.

Figure 19 shows the circuit of the precision ac/dc converter for use with the above two circuits. The gain of the converter can be set to precisely 2.2 via RV1, to give a dc output voltage that is equal to the RMS value of a sinewave input.

The converter is powered from the supply rails of the module, and is designed around an LF355 op-amp, which can operate quite happily from the 2V8 between V_{DD} and COM.

Resistance meters

The easiest way to use a DVM module as a resistance (ohm) meter is to use it in the ratiometric configuration shown in Figure 5. This technique has two major advantages. First, it is very stable and inherently self-calibrating, the meter reading being equal to $R_x \times (RV/R_{ref})$, where RV is the ratiometric value of the meter when used in the Figure 7 test circuit. RV is typically only 0.2% low, so measurement accuracy is determined primarily by R_{ref}. The second advantage is that very low test voltages are generated across R_x, the maximum voltage being ⅔ of the energising voltage (typically 100 to 300 mV) at full scale. Figure 20 shows how the module can be connected as a practical five-range ohmmeter ●

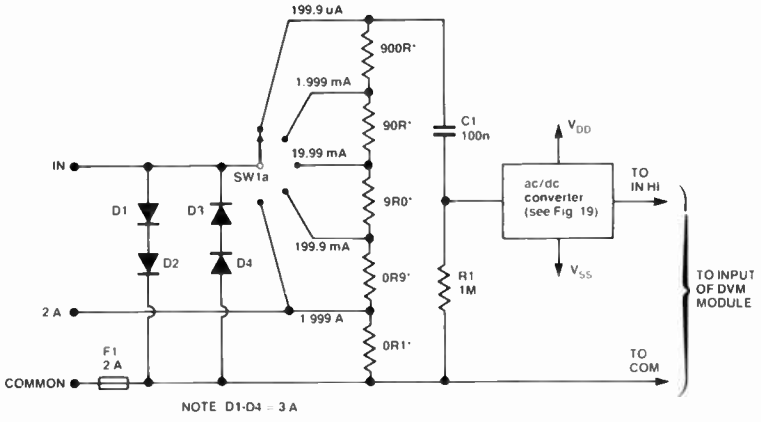

Figure 18. Modification of the Figure 16 circuit, to act as a five-range ac current meter.

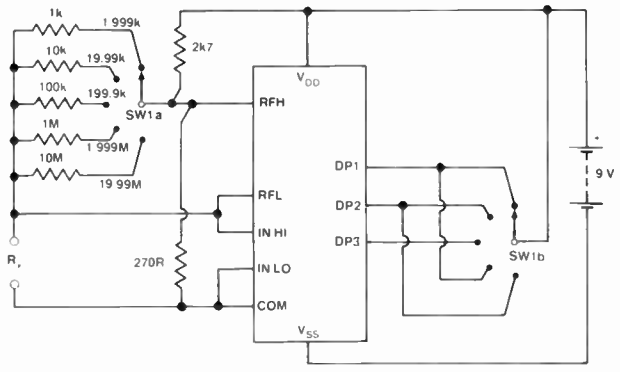

Figure 20. Five-range ohmmeter.

How to use digital voltmeter modules

Part 2

In this concluding part of the series, Ray Marston shows how to use 3½-digit LCD DVM modules to build a multimeter and measure temperature, capacitance, frequency and a whole lot more.

Ray Marston

IN PART 1 I explained the basic characteristics and usage rules of 3½-digit LCD digital voltmeter modules and showed how these units can be used to measure voltage, current and resistance.

This part kicks off with the complete circuit of a five-function, 25-range digital multimeter and continues by showing how these modules can be used to give accurate measurement of temperature, capacitance, frequency and various other parameters. It finishes up with some vital hints on actually constructing projects that are designed around the modules.

A 25-range multimeter

Figure 21 shows how the circuits of Figures 15 to 20, in Part 1, can be joined together to make a complete five-function 25-range multimeter. Table 3 details the ranges/functions of the meter.

The reader should have little difficulty in following the Figure 21 circuit. Functions are selected by SW1, ranges by SW2. SW1a connects the inputs to the voltage, current or resistance measuring networks, and SW1d activates the ac/dc converter or energises the 'ohms' circuitry when necessary.

Voltage ranges are selected by SW2a, current ranges by SW2b, and resistance ranges by SW2c.

SW2d and SW2e control the decimal point

Figure 21. Complete five-function 25-range multimeter.

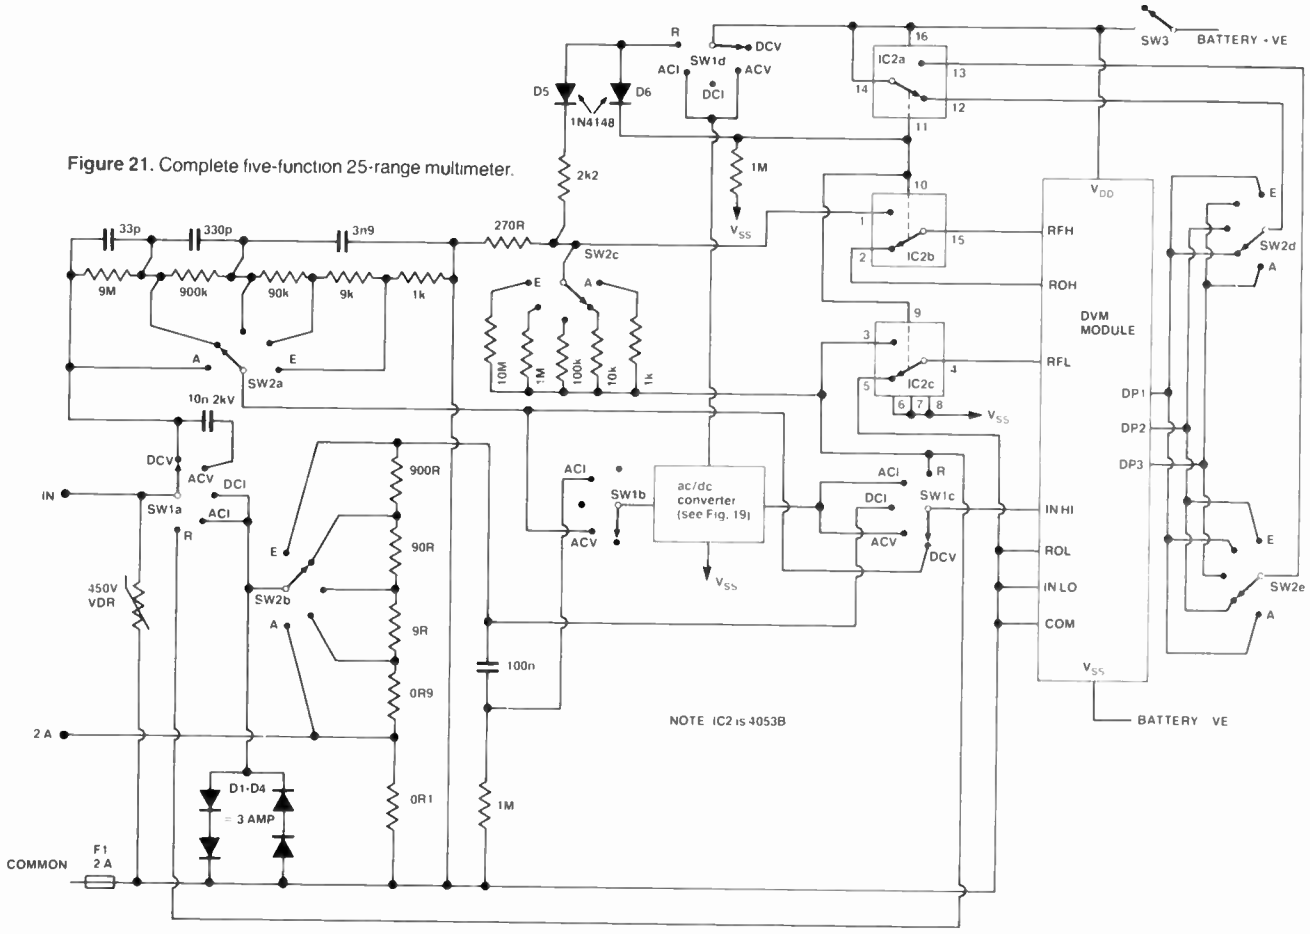

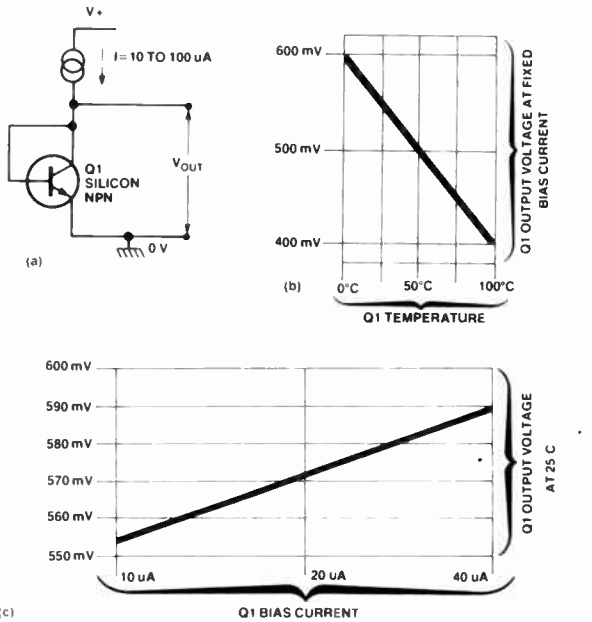

Figure 22. When a transistor is connected as in (a), its output voltage varies at a rate of about −2 mV/ C as shown in (b). The output voltage also varies with drive current as shown in (c).

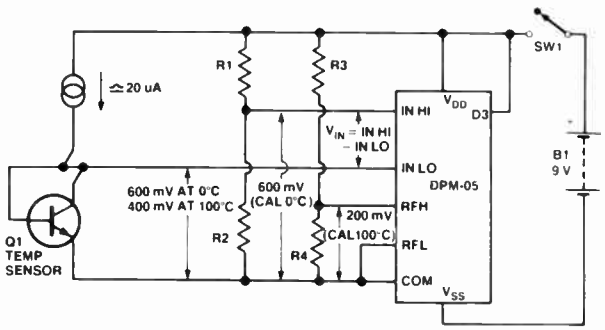

Figure 23. Basic digital thermometer circuit using an idealised (Figure 22b) transistor sensor.

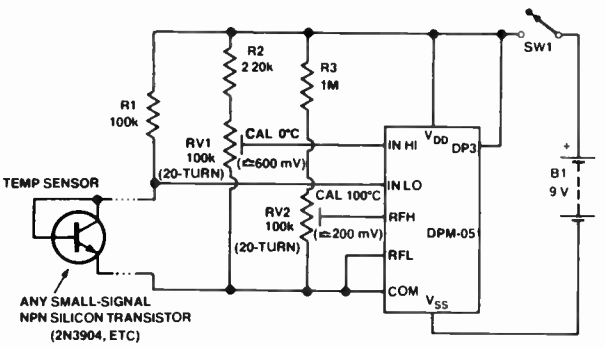

Figure 24. Simple digital thermometer using a transistor sensor. Linear accuracy is about 1.5 C.

positions on each range, the appropriate switch being selected automatically by IC2a. IC2b and IC2c control the basic configurations of the DVM module. IC2 (a triple two-way analogue switch) is activated via SW1d.

Digital thermometers

A DVM module can be made to act as a wide-range (−50°C to +150°C) digital thermometer by feeding the output of a linear voltage-generating temperature sensor to its inputs. Two suitable types of sensor are readily available; the first of these is the ordinary bipolar silicon transistor and the second is a dedicated IC. In either case, the resulting digital thermometer has a temperature discrimination of 0.1°C. Linear accuracy varies from 0.5°C to 1.5°C, depending on the sensor and circuitry used.

Because of the low mass of a transistor sensor, the device has a thermal response time some 10 to 100 times faster than a normal mercury thermometer. When used to measure a sharp change in the temperature of free air, a transistor-sensor circuit typically settles to within 0.1°C of the new temperature in less than one minute; a mercury thermometer takes some 20 minutes to attain the same accuracy.

Transistor-sensor circuits

When an ordinary NPN silicon transistor is connected as shown in Figure 22a and driven from a constant-current source, it generates an output voltage that varies in direct proportion to the transistor temperature. This voltage has a negative temperature coefficient of about −2 mV/°C and typically varies from about 600 mV at 0°C to 400 mV at 100°C, as shown in the idealised graph of Figure 22b.

In practice, the 'straight line' of the Figure 22b graph is linear within 1 mV or so over the 200 mV '0°C to 100°C' temperature variation range, but the precise voltage generated at any given temperature depends on the individual transistor and its operating current. If operating currents are kept below 100 uA, errors due to self-heating are negligible. Figure 22c shows the measured variation in voltage at 25°C of a small sample of transistors at currents ranging from 10 to 40 uA.

MODE (SW1)	RANGE (SW2) A	B	C	D	E
DCV	199.9 mV	1.999 V	19.99 V	199.9 V	1.999 kV (700V max)
ACV	199.9 mV	1.999 V	19.99 V	199.9 V	1.999 kV (450 V max)
DCI	199.9 μA	1.999 mA	19.99 mA	199.9 mA	1.999 A
ACI	199.9 μA	1.999 mA	19.99 mA	199.9 mA	1.999 A
R	1.999 k Ω	19.99 k Ω	199.9 k Ω	1.999 M Ω	19.99 M Ω

Table 3. Ranges and functions of the Figure 21 multimeter circuit.

Figure 23 shows the basic method of connecting the idealised transistor sensor of Figure 22b to a DVM module so that the meter gives a direct readout of temperature in °C. The output of the sensor is fed directly to the module's IN LO terminal and a 600 mV offset voltage (equal to the sensor voltage at 0°C) is fed to IN HI. The module actually responds to the *differential* value (IN HI minus IN LO) of the input, so under this condition it sees an input of 600 mV − 600 mV = 0 mV, and gives a reading of '00.0'.

At 100°C the module sees an input of 600 mV − 400 mV = 200 mV. Since a reference voltage of 200 mV (equal to the difference in voltage between 0°C and 100°C) is fed to RFH, the meter gives a reading of '100.0' under this condition.

Figures 24 and 25 show two practical examples of digital thermometers. The Figure 24 circuit is virtually the 'standard' one published in many magazine articles and application sheets and has a typical linear accuracy of 1.5°C over the 0°C to 100°C temperature range. A stable 2V8 is generated between V_{DD} and COM of the DVM module, so R1 drives the sensor with a current of about 22 uA at 0°C, rising to about 24 uA at 100 C. This current variation, combined with the basic linear error of the transistor, causes the 1.5°C linear error of the circuit.

The 'CAL 0°C' voltage feeding IN HI is variable from zero to 875 mV via RV1, and the 'CAL 100°C' voltage feeding RFH is variable from zero to 255 mV via RV2. These two controls are for the calibration of the meter, using the technique to be described shortly.

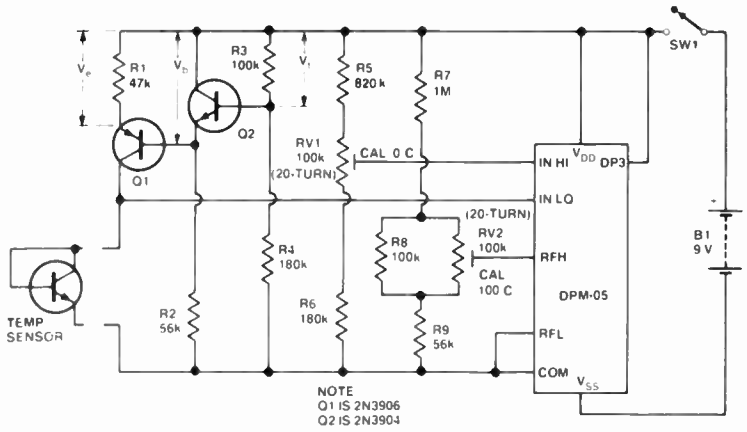

Figure 25. Precision digital thermometer using a transistor sensor. Linear accuracy is 0.5 C

Figure 25 shows a precision version of the digital thermometer, giving a linear accuracy of about 0.5 C. In this case the transistor sensor is energised at about 20 uA via constant-current generator Q1 which is temperature compensated by Q2. This section of the circuit works as follows.

Potential divider R3-R4 is wired between V_{DD} and COM and generates voltage V_t (about one volt) across R3. This voltage is 'followed' by NPN transistor Q2, causing $V_t + V_{be}2$ to appear on Q1 base; let's call this voltage V_b. The voltage V_e appearing on the emitter of PNP transistor Q1 is equal to $V_e - V_{be}1$ and it is the product of V_e and R1 that determines the magnitude of the constant-current output of Q1. Note, however, that V_e is, in fact, equal to $V_t + V_{be}2 - V_{be}1$ and that since Q1 and Q2 operate at virtually identical temperatures and at similar current levels, the $V_{be}1$ and $V_{be}2$ values automatically cancel out at all temperatures and V_e thus equals V_t. The output current of Q1 is thus independent of ambient temperature.

Other points to note about the precision circuit of Figure 25 are that the 'CAL 0 C' control, RV1, is adjustable over the limited range of 460 mV to 710 mV. Also, the 'CAL 100 C' control, RV2, is adjustable over the limited range 140 mV to 260 mV, thus giving very fine adjustment of each calibration point.

Calibration procedure

The procedure for calibrating the Figures 24 and 25 circuits is as follows. First, solder the base and collector leads of the sensor transistor together. Then solder the sensor to a pair of flexible leads and connect it to the meter circuit. Paint all visible transistor leads and solder joints with insulating varnish (Humbrol clear varnish No 35 is excellent). Next, set RV1 and RV2 at mid value, mix a quantity of crushed ice and cold water in a tumbler (to act as a '0°C' standard) and immerse the sensor in the tumbler. Now adjust RV1 to give a reading of '00.0' on the meter. Finally, remove the sensor from the tumbler and immerse it in gently boiling water (to act as a '100°C' standard), then adjust RV2 to give a meter reading of '100.0'. Basic calibration is then complete.

If the meter is to be used mainly around some mid-scale value, such as 25°C etc, RV1 can (after initial calibration) be used to set the meter 'spot on' at that value by immersing the probe and a standard thermometer in a liquid that is raised to the desired temperature.

An IC-sensor circuit

Intersil make a special two-terminal IC for use as a temperature sensor in digital thermometers. The device is the AD590 and gives an output current of 1 uA/°K which, when fed through a 1k resistor, gives a voltage of 1 mV/°K. Uncalibrated accuracy of the device varies from 0.5°C to 10°C. Linearity error varies from 0.3°C to 1.5°C, depending on the grade of the device (indicated by a suffix number). Figure 26 shows how an AD590 can be used with a DVM module.

The AD590 needs a supply voltage of at least four volts and this is obtained by wiring the IC between V_{DD} and TEST (which is internally biased at about five volts below V_{DD}) via D1. The COM terminal is biased about 600 mV above TEST via D1. R1 is wired in series with the AD590 and generates approximately 1 mV/°K (= 273.2 mV at 0°C, 373.2 mV at 100°C). This voltage is fed to the IN HI terminal. Bandgap reference IC2 generates a temperature-stable 1.2 V via R2 and this voltage is divided down via R3-RV1-R4 to give a 'SET 0°C' offset voltage of 273.2 mV nominal at IN LO. The bandgap reference voltage is also divided down by R5-RV2-R6 to provide a 'SET 100°C' scaling voltage of 100 mV nominal at RFH. The circuit must be calibrated in the way already described for the Figures 24 and 25 circuits.

Digital capacitance meter basics

A DVM module can be made to read capacitance values by connecting the unknown capacitance to the module via a linear capacitance-to-voltage converter. The easiest way to make such a converter is to use the technique shown in Figure 27. Here, the unknown capacitor and a standard resistor are used as the timing elements in a precision monostable which produces an output pulse with a width, W, that is directly proportional to the C-R product. The monostable is triggered at a fixed frequency via a clock generator and the output of the monostable is converted to a mean dc value by a simple C-R integrator.

The mean dc value of the monostable output equals the peak pulse amplitude multiplied by W/P, where W and P are the width and the period of the pulse respectively. Thus, since R_x and P are fixed, the mean dc voltage out-

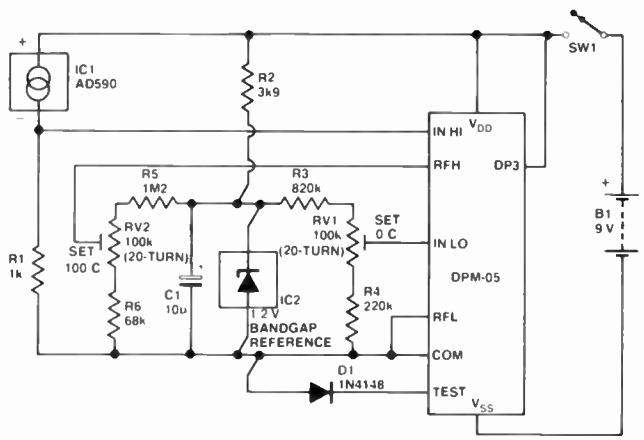

Figure 26. Digital thermometer based on the AD590 temperature-sensor IC

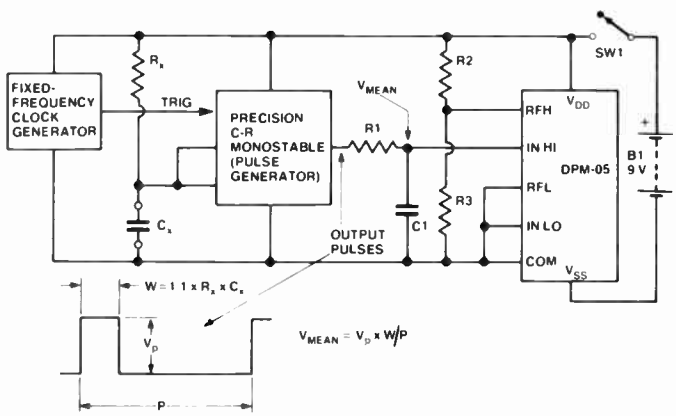

Figure 27. Basic operating principle and circuit of a digital capacitance meter

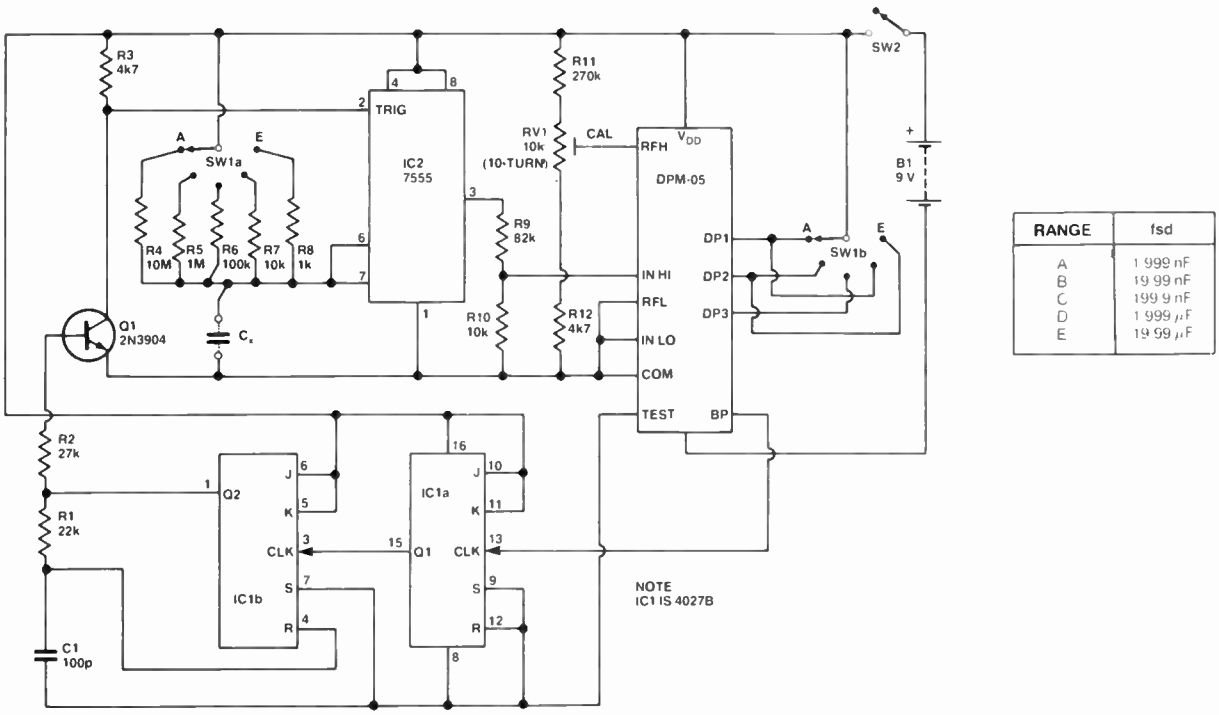

RANGE	fsd
A	1.999 nF
B	19.99 nF
C	199.9 nF
D	1.999 μF
E	19.99 μF

Figure 28. Digital capacitance meter.

put is directly proportional to C_x and when this voltage is fed to the DVM module, the module acts as a digital capacitance meter.

In Figure 27, the reference (RFH) voltage of the module is derived from the mono's supply rail via potential divider R2-R3. Since the meter reads the *ratio* of the input and reference voltages, the calibration of the unit is independent of variations in supply rail voltage, but can be varied by altering the R2-R3 ratio. The circuit can be made to read different capacitance ranges by switching R_x in decade multiples.

Practical capacitance meters

The basic Figure 27 circuit is quite easy to implement and gives very accurate results. Figures 28 and 29 show two practical versions of the circuit. Both of these designs use a 7555 timer IC (a CMOS version of the 555 timer) as the precision monostable element, and use decade values (1k to 10M) of R_x for range selection.

The 7555 monostable generates a pulse with a width of 1.1 x C x R, giving a full-scale pulse width (at '1999' on the DVM module) of 22 ms with C and R values of 1999 pF and 10M, or 19.99 uF and 1k etc. To give the 7555 adequate recovery time between pulses, the clock period must be at least 50% longer than the maximum pulse width and must have a period of at least 33 ms.

The 7555 mono is triggered by pulling pin 2 of the IC low. If the pin is not returned high again by the time the output pulse ends naturally, the trigger pulse extends the output pulse artifically. The trigger pulse must thus be shorter than the minimum output pulse. In our application, the shortest pulse width that can be indicated by the DVM module is 22 ms/1000 = 11 us. So in the circuits of Figures 28 and 29 it is a design requirement that the 7555 must be triggered by negative-going pulses with widths less than 11 us and periods greater than 33 ms. In Figure 28 these requirements are met as follows.

In the DVM module, the TEST terminal is internally biased at about 5 V below V_{DD} and the BP (backplane) terminal switches between TEST and V_{DD} at about 50 Hz (= clock frequency divided by 800), giving a period of 20 ms. In Figure 28 IC1 is powered via the TEST terminal and the BP signal is divided-by-2 by flip-flop IC1a. The resulting 25 Hz (40 ms) signal is used to clock IC1b, which is configured as a monostable and generates positive-going output pulses with widths of 2 us via R1 and C1. These pulses are level-shifted and inverted via R2-Q1-R3 to produce negative-going 2 us trigger pulses with periods of 40 ms on the pin-2 TRIG terminal of IC2, the 7555 monostable generator.

The pulse width of the 7555 is controlled by C_x and precision range resistors R4 to R8. The 7555's output is attenuated by R9-R10 to give a mean value of about 100 mV at the midscale ('1000') setting of the DVM module. The resulting signal is fed to the module's IN HI terminal where it is integrated by the internal 10M — 10 nF filter. Divider R11-RV1-R12 feeds 100 mV nominal to the RFH terminal of the module and RV1 is used to adjust the precise calibration of the capacitance meter.

Accuracy of the Figure 28 meter is determined mainly by the precision of the R4 to R8 range resistors, which should be 1% or better hi-stab types. To calibrate the meter, simply connect a precision capacitor (say 100 nF) in place of C_x, switch to the appropriate range and adjust RV1 to give the appropriate meter reading. Calibration is then valid on all ranges.

The Figure 28 circuit has two minor defects. First, the clock signals of the 7555 are derived (via BP) from the clock signals of the DVM module. These signals are not highly frequency-stable and the calibration of the circuit may thus shift by up to 0.5% or so over the normal range of operating temperatures and supply voltages. If precise accuracy is needed, calibration should be checked before use.

The second snag is that the circuit reads *all* capacitance, including residuals, appearing between the C_x terminals. These residuals include stray capacitance and the internal capacitance of IC2 between pins 6/7 and 1, and typically total 32 pF. With no external capacitance connected, the meter thus gives a typical reading of '.032' on range A and '0.03' on range B. These residuals are too small to give readings on the remaining ranges of the meter, but must be subtracted from all readings obtained on ranges A and B.

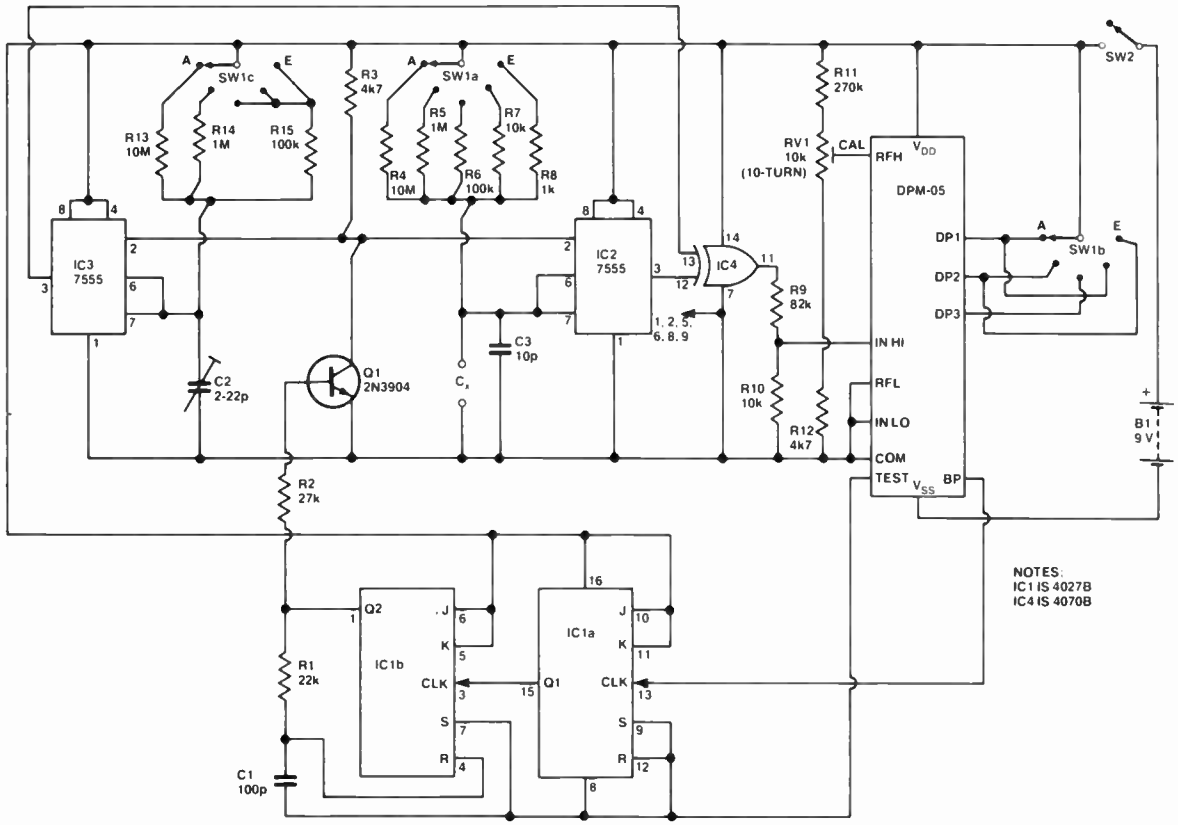

Figure 29. Precision capacitance meter with zero residual reading.

Figure 29 shows how the circuit can be modified so that residual capacitance is effectively cancelled and the meter gives a zero reading on all ranges when no external capacitance is connected to the C_x terminals. In this case the BP-derived signal is used to synchronously trigger *two* 7555 monostables. Their outputs are EX-ORed via IC4 to give a pulse with a width equal to the *difference* between the pulse widths (and thus the residual capacitances) of the two monostables. This pulse is fed to the IN HI terminal of the DVM module via R9-R10. Thus, if the monostables have identical residuals, the EX-OR pulse width is zero and the meter gives a zero reading with zero external C_x applied.

In Figure 29 monostable IC2 is connected to the C_x terminals and functions in the same way as in Figure 28, except that an additional 10 pF is permanently wired across the terminals. The IC3 monostable, however, has C2 wired across its input terminals and the value of C2 can be adjusted to equal (and thus cancel) the residual of the input of IC2. IC3 is range-switched in parallel with IC2 via SW1c. Precise ganging is provided on ranges A, B and C only and on all other ranges the residuals are too small to influence the meter readings.

Frequency measurement

A DVM module can be made to read frequency by connecting the unknown frequency to the module's input via a f-to-V converter. A suitable converter can easily be made by using a 7555 monostable; Figure 30 illustrates the principle. The input signal is first fed to an input conditioner and trigger-pulse generator which triggers a fixed-period 7555 monostable on the arrival of each new input cycle. The output pulses of the mono are converted to mean dc values by integrator R2-C2 and fed to the input of the DVM module which is scaled via R3-R4.

The mean dc value of the 7555 output pulses equals V_p (the peak amplitude of the pulses) multiplied by W/P, where W and P are the width and period of the pulses respectively. V_p and W are, however, fixed. Only the pulse period is variable and this is inversely proportional to the input frequency, f, so the mean output voltage is equal to V_p x W x f and is thus directly proportional to f. Therefore, when the DVM module is suitably scaled via R3-R4 it gives a direct reading of input frequency.

In practice, the lowest convenient full scale frequency range of a DVM-based $3\frac{1}{2}$-digit frequency meter is 1.999 kHz. In this case, the 7555 pulse has a period of 500 us at full scale. For maximum accuracy the pulse width must be as large as possible but must not be greater than two thirds of P. A pulse width of about 300 us is necessary and this can be obtained from the 7555 by choosing values for R1 and C1 of 27k and 10 nF respectively.

Figure 31 shows how the basic Figure 30 circuit can be modified to act as a multi-range frequency meter. In this case the input signal is fed to an input conditioner and Schmitt trigger and the Schmitt output is used to ripple-clock four decade dividers. The 7555 300 us monostable is provided with a trigger generator than can be fed from the output of the Schmitt or from any of the dividers. Thus, when the 7555 is triggered directly from the Schmitt the meter reads 1.999 kHz full scale and when fed from the

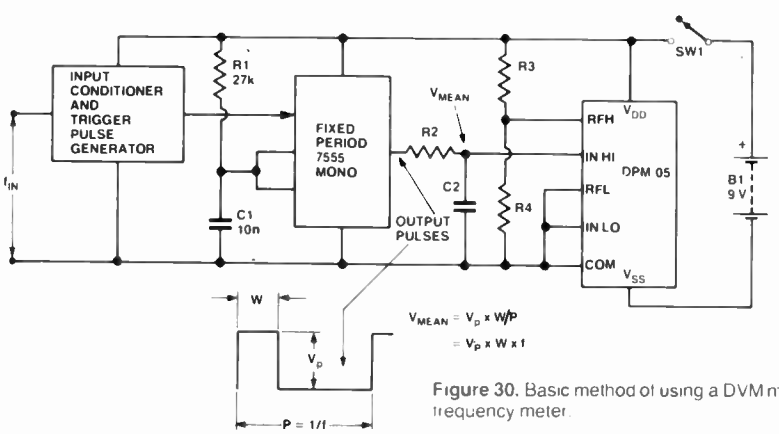

Figure 30. Basic method of using a DVM module as a frequency meter.

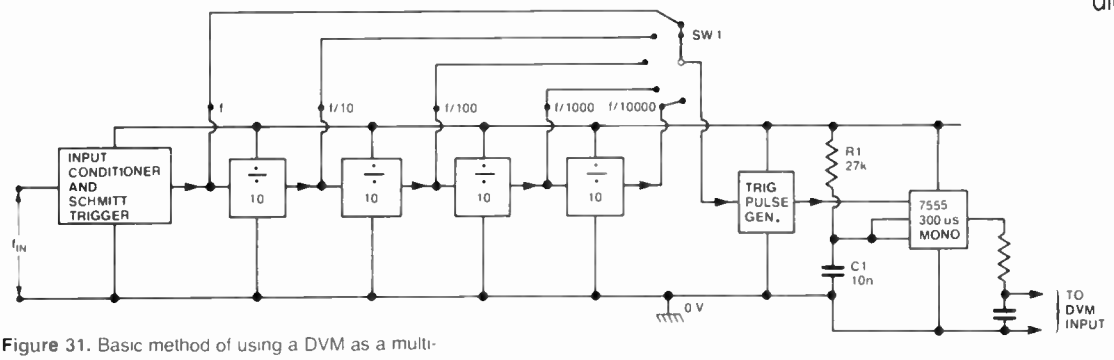

Figure 31. Basic method of using a DVM as a multi-range frequency meter.

output of the last divider stage reads 19.99 MHz full scale.

Practical frequency meters

The 7555 IC can operate from supplies as low as 2 V. Standard CMOS counter ICs, however, need supplies of at least 3 V. Consequently, if a DVM module is to be used as a frequency meter sharing supplies that are common with those of the CMOS divider stages, the DVM module must be used in the 'split-supply' mode with its COM terminal pulled below the normal 'V_{DD} — 2V8' value by external circuitry.

In other words, COM must operate at '0 volts' and V_{DD} and V_{SS} at nominal values of +4V5 and −4V5 respectively. Figure 12 in Part 1 showed how these supplies can be obtained if the module is built into existing equipment that has split supplies. Alternatively, Figure 32a shows how the supplies can be obtained from a stack of six 1V5 cells. Figure 32b shows how the supplies can be obtained from a single 9 V battery via an op-amp supply-splitter. The supply-splitter of Figure 32b adds a quiescent current consumption of about 2 mA to the DVM circuit, but can supply additional supply currents of tens of milliamps to circuitry connected between +4V5 and 0 V.

Figure 33 shows the practical circuit of a DVM-based digital frequency meter that reads up to 19.99 MHz full scale in five

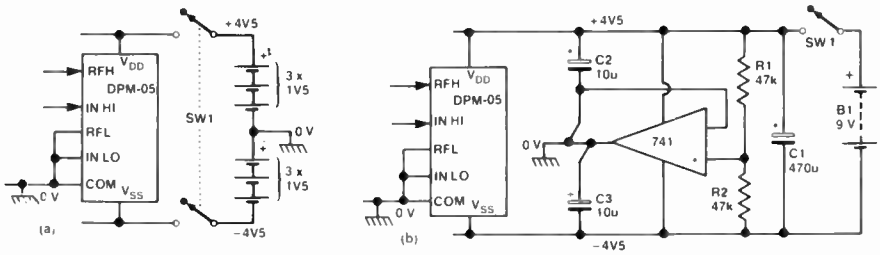

Figure 32. A DVM frequency meter needs split supplies. These can be obtained from either a stack of 1V5 cells (a) or from a single 9 V battery and an op-amp supply-splitter (b)

decade ranges. When used with the Figure 32b power supply, the circuit consumes about 3 mA quiescent from the 9 V battery, rising to 4 mA at 1 MHz, and (when calibrated) has a reading accuracy of +/− one digit. The circuit accepts input signals in the range 200 mV to 5 V RMS and operates as follows.

Input signals are fed, via C1-R1, directly to the input of IC1a, a very fast Schmitt trigger, which is biased at half-supply volts via R2-R3. The Schmitt output is used to ripple-clock four decade-divider stages. Ordinary CMOS dividers typically operate at maximum speeds of only 800 kHz or so when powered from 4V5 supplies. To give the required fast operating speeds the very latest 'HC' types of silicon-gate CMOS counters are used in the first two (IC2 and IC3) counter positions. On the prototype unit they clock at frequencies up to about 18 MHz.

The output of the IC1a Schmitt and of the four divider stages are fed to range-selector switch SW1a. The output of SW1a is fed to 4 us trigger-pulse generator C4-R4-IC1b-IC1c which triggers the 7555 monostable via Q1. The output of the 7555 is fed to IN HI of the module via R8-R9, and a calibration 'reference' voltage is fed to RFH via RV1. The circuit is calibrated by feeding in a signal of known frequency, switching to the approriate range and trimming RV1 for the appropriate reading on the DVM module.

Once RV1 has been initially calibrated, calibration is influenced only by variations in the pulse width of the 7555 and these may be caused by thermal variations in the values of R7 and C5. For optimum calibration stability R7 should be a metal-glaze resistor and C5 should be a polycarbonate capacitor.

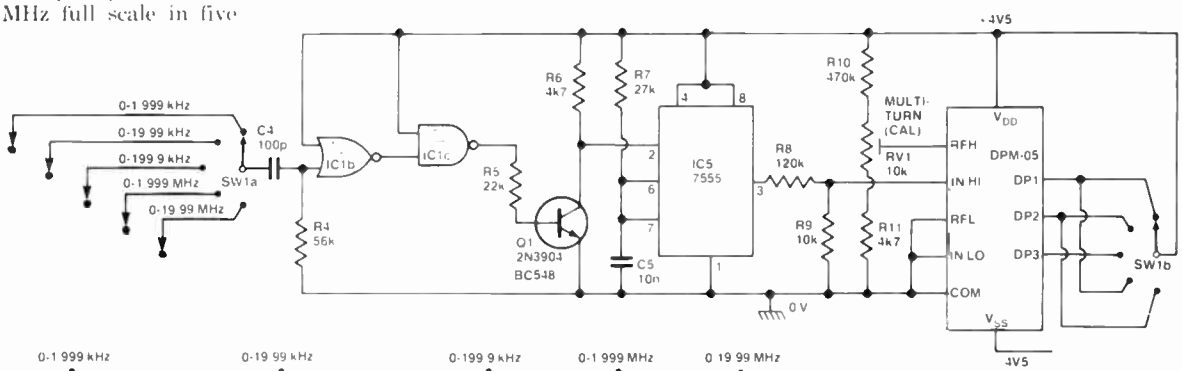

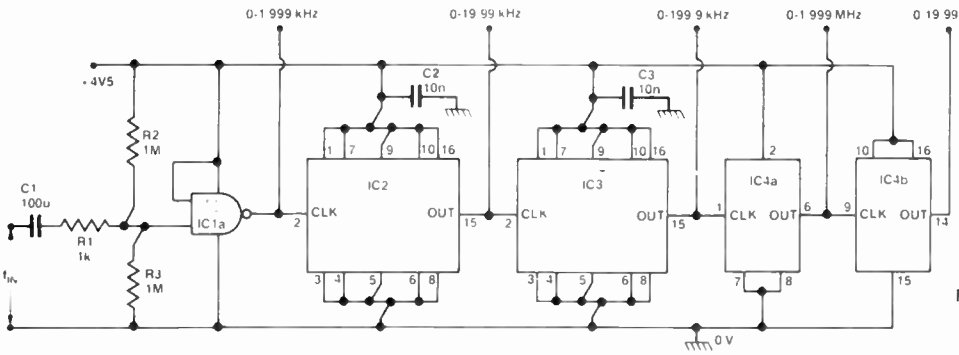

NOTES
IC1 IS 74HC132
IC2,3 ARE 74HC160
IC4 IS 4518

Figure 33. Digital frequency meter reading 0 to 9.99 MHz in five decade ranges

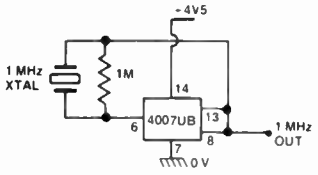

Figure 34. This 1 MHz crystal calibration oscillator can easily be added to Figure 33 circuit.

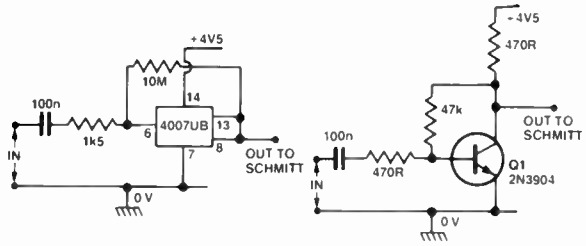

Figure 35. Two simple preamplifiers that can be used with the frequency meter

The Figure 33 circuit can be modified in a variety of ways to satisfy individual requirements. Figure 34 shows a 1 MHz crystal calibration oscillator, designed around one section of a 4007UB CMOS IC, which can be easily added to the frequency meter and consumes a mere 300 uA when active. Figure 35 shows two simple preamplifiers which can be used to improve the basic sensitivity of the meter. The Figure 35a design, based on one section of a 4007UB, has an input impedance of about 1M and improves sensitivity by about 20 dB (to 20 mV RMS) at audio frequencies, but is useful to only a few hundred kHz. The simple Figure 35b design also gives a gain of about 20 dB at low frequencies, but has a low input impedance (about 2k2) and is useful to several MHz. Both circuits consume a couple of milliamps.

Figure 36 shows, in basic form, how the DVM module can be used to read both frequency and ac volts (or any other desired parameter). With SW1 switched to 'f', the input is switched to the input of the f-meter circuit and IN HI and RFH of the module are switched to the outputs of the circuit. When SW1 is switched to 'Vac', the input is switched to the input of the frequency-compensated attenuator, which has its output fed to IN HI via SW2 and a precision ac/dc converter (see Figure 19 in Part 1). RFH is switched to a 100 mV standard voltage derived from a bandgap reference.

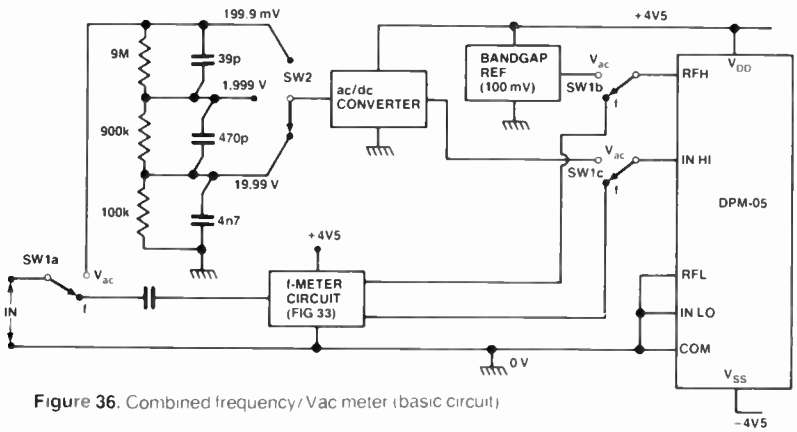

Figure 36. Combined frequency/Vac meter (basic circuit)

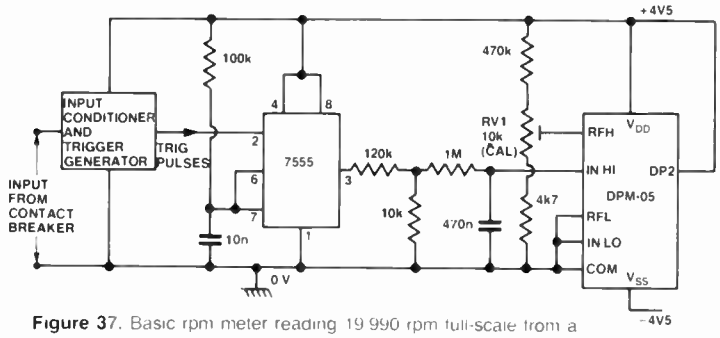

Figure 37. Basic rpm meter reading 19 990 rpm full-scale from a four-cylinder four-stroke petrol engine

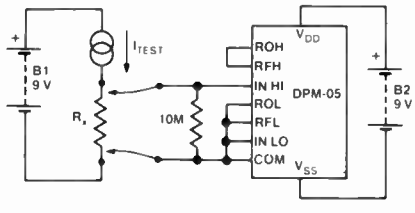

Itest	Rx fsd
1 mA	19.99 Ω
100 uA	1999 Ω
10 uA	19.99 kΩ

Figure 38. Basic milli-ohmmeter using four-terminal measurement technique

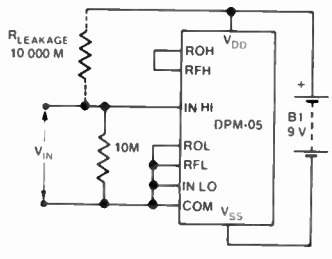

Figure 39. In this 199.9 mV dc voltmeter circuit pc board or module leakage resistance causes the meter to indicate 28.0 mV with no external input applied

Miscellaneous applications

DVM modules can be used to indicate the value of any parameter than can be converted into a predictable (linear or log) voltage, current or resistance. Linear transducers are readily available for measuring values of pH, light intensity, and radiation etc.

Cyclic parameters such as rpm and heartbeat rate etc, can be measured by adapting the frequency meter technique already described. The rpm of a petrol engine, for example, is directly proportional to contact-breaker (CB) frequency, f. On a four-stroke engine, f = N x rpm/120, where N is the number of cylinders. Thus, on a single-cylinder engine 10 000 rpm gives a CB frequency of 83.3 Hz, and on a four-cylinder engine a frequency of 333.3 Hz. Figure 37 shows the basic circuit of a digital rpm meter designed to read 19 990 rpm full scale (10 000 rpm at mid scale) on a four-cylinder four-stroke engine. The 7555 monostable gives an output pulse width of about 1 ms.

When measuring low values of resistance care must be taken in circuit design to ensure that the resistive effects of range switches, fuses and terminals etc, are excluded from the measurement results. The only way of achieving this is to use the four-terminal measurement technique shown in Figure 38, in which two independent circuits are used. Here, the unknown resistor is connected between the R terminals and fed with a constant current from B1. The volt drop *directly across* R_x is measured via a 199.9 mV full scale dc voltmeter powered from B2. Thus, when 10 mA is passed through R_x, the voltmeter indicates 19.99 ohms at full scale.

Constructional notes

When using DVM modules two vital usage points must be noted. The first of these arises from the high sensitivity of the module and is illustrated in Figure 39, where the module is wired as a 199.9 mV full scale dc voltmeter with a 10M input resistance. Thus, if a leakage resistance of 10 000 megohms appears between V_{DD} and IN HI, the meter will read 28.0 mV with no external input applied. Leakage resistances of this magnitude (and lower) can be caused by minute amounts of moisture or dirt appearing between the terminals of the module or the tracks of a pc board to which it is connected. To eliminate the possibility of this effect, the entire module and pc board must be cleaned and dried after project construction is complete. Then both must be thoroughly coated with insulation varnish. Humbrol clear varnish No 35 (available from model and art shops) is excellent for this purpose.

The final usage point concerns external components. General-purpose resistors and capacitors have very poor thermal stability. Consequently, in all practical DVM-based designs great care must be taken to ensure that all critical resistors are metal-glaze or similar hi-stab types, and all critical capacitors are polycarbonate types. ●

Low dropout voltage regulator

Three-terminal regulators require at least 2.5 V drop from input to output to operate correctly. This circuit only requires 0.75 V which makes it suitable for those critical applications where you don't have input voltage to spare.

E. Smeda

I NEEDED a voltage regulator with very low dropout voltage, capable of supplying 5 V at about 3 A. Since many three-terminal regulators have a dropout of 2.5 V or more, the circuit described here was devised using discrete components.

Although very simple, its performance was found to be comparable to that of IC regulators. The circuit of my prototype is shown in Figure 1, and it gave the following results:

Dropout voltage (@ 3 A)	0.75 V
Load Regulation (0-3 A)	less than 10 mV
Line regulation (Vin 6 — 15 V)	less than 10 mV
Ripple rejection (@ 3 A)	−63 dB
Output (no load)	4.96 V

Changes in the output voltage due to ambient temperature variations will be entirely dependent on the characteristics of the zener diode, ZD1 and transistor Q3. Thus, these should be kept clear of heat producing sources, e.g.: the heatsink for Q1 and the power transformer.

The circuit, as it stands, has no well-defined current limit, but this feature can be included with the addition of the components shown in Figure 2. Doing this, however, causes the dropout voltage to increase by 0.5 V. An alternative method of providing current limit without increasing the dropout voltage is to load the output to the required maximum output current and gradually increase the value of the resistor R1 until the output voltage just starts to drop. The disadvantage with this method is that R1 must be selected on test and will need to be re-adjusted if Q1 is ever replaced. Thus, if current limit is desired and a slight increase in dropout voltage can be tolerated, the method shown in Figure 2 is the preferred one.

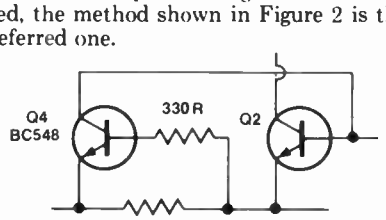

Figure 2. Adding current limit.

Another characteristic of the circuit is that, should a heavy load cause the output voltage to drop below approximately 1.2 V, the regulator will automatically shut itself off and can be restarted by removing the input voltage (or switching off at the mains), waiting several seconds, and then re-applying power. Merely removing the load will not allow a restart.

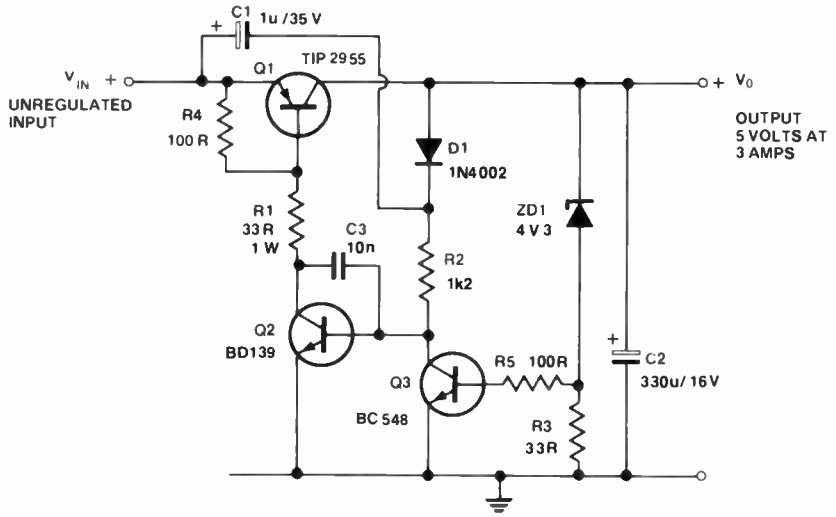

Figure 1. Circuit of the low dropout regulator. Design information is given so that the circuit can be arranged for other voltages.

Components C1 and D1 are the startup components and allow reliable starts even with heavily capacitive loads. In order that the circuit may be adapted for any voltage and current (up to about 5 A), the following simplified design procedure is given.

(1) Select the output voltage, Vo (5 V)
(2) Select the maximum current, Io (3 A)
(3) Select a suitable transistor for Q1 (TIP2955 70 V, 10 A)
(4) R1 max. = $Vo/(Io/h_{FE\ min}\ Q1)$
= 5/(3/20) = 33.3
use 33R
(5) Dissipation of R1 = $Vo^2/R1$ = 25/33
= 0.75 W (use 1 W)
(6) Select a suitable device for Q2 (BD139 80 V, 1 A)
(7) R2 max. = $Vo/(Ioh_{FE\ min}\ Q1 \times h_{FE\ min}\ Q2)$
= 5/(3/20×40) = 1300
use 1k2
(8) Dissipation of R2 = $Vo^2/R2$
= 25/1200
= 21 mW
(use ¼ W)
(9) Select a suitable device for Q3 (BC548 25 V, 100 mA)
(10) Select a suitable zener diode, ZD1.
Voltage = Vo − V_{be} (Q3) = 5−0.65
= 4.35 V (use 4V3 zener)

As the current flowing in Q3's base will usually be very small, it may be ignored. Thus, a low wattage, 400 mW or 1 W, zener may be used.

R3 should be chosen to bias the zener well into its operating region. A good rule of thumb is to select the current through the zener (I_z) to be a fifth of its maximum.

$R3 = V_{be}(Q3)/I_z$
$I_z = 0.2 \times (P_z/V_z)$ (P_z = zener power, V_z = zener voltage)
= 0.2 × (0.4/4.3) assuming 400 mW zener
= 18.6 mA.
Thus, R3 = 0.65 V/18.6 mA
= 35 ohms (use 33R, ¼ W)

(11) R4 ensures Q1 is not turned on by leakage. Its value is not critical. 100 ohm, ¼ W is usually suitable here.
(12) R5 is used as a precaution, preventing excessive current through Q3's base; 100 ohms is suitable.
(13) C1 is the startup capacitor. Any value between 1u and 4u7 should be satisfactory. Its voltage rating should exceed Vin.
(14) C2 should always be used for stability. A value of 100u per amp of load current is suitable. 330u was used in the prototype. Its voltage rating should exceed the output voltage rating.
(15) C3 is required for stability. 10n should suit.
(16) Mount Q1 on a suitable heatsink, according to the power it dissipates. (This should be low if Vin-Vout is low). ●

Gain control

Tim Orr continues his occasional series of circuits, methods and explanations with a detailed look at how gain can be controlled by another electronic signal, be it squarewave, sinewave or voice signal. This leads to some interesting circuits – from ducks to filters!

THERE ARE MANY cases in signal processing where the control of the gain is necessary. Some common examples are automatic volume controls in cassette recorders and in the IF sections of radio receivers. Also in professional audio equipment there is a whole range of compressor, expander, limiter and noise gate devices which find great use in recording and broadcast studios. Maybe you have wondered how the volume of the music drops when the DJ starts to talk and then fades up again when he stops. This process known as voice over or "ducking", uses *voltage control of gain.*

Noise reduction systems such as dolby and dbx employ voltage controlled amplifiers. Synthesisers and sound processors obtain effects such as ring modulation, automatic panning, frequency shifting, dynamic filtering, tremolo and envelope shaping also by the use of this technique.

Gaining gain

There is a wide variety of methods which can be used to obtain the gain control. This can be anything from constructing the variable gain element yourself from basic parts, to buying ICs or modules designed specifically to solve your particular problem. Generally the solution is some sort of compromise, because unfortunately the problem of making high performance controlled gain cells (multipliers), is rather difficult and therefore the ICs tend to be rather expensive.

However with a bit of care a cost effective solution can usually be produced.

A good example is the AGC in a transistor radio. The transistors in the IF section have an h_{fe} that varies widely with collector current. Thus, by sticking three transistors in series it is possible to vary their overall gain by about 40 dB, (x 100), merely by controlling their collector currents. The AGC stops the audio output of the radio from varying as the radio reception conditions alter.

Electronic multipliers

When it is required to control the level of one signal with that of another, an electronic multiplier is used. This process is analogous to arithmetic multiplication. If input A is positive, fig. 1, and input B is positive, then the product (the output), will also be positive. If A goes negative then

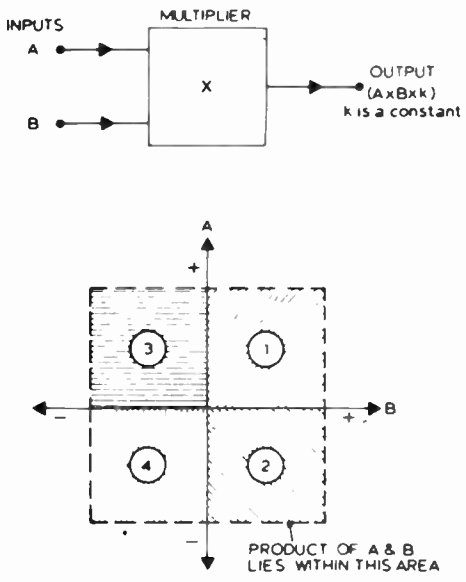

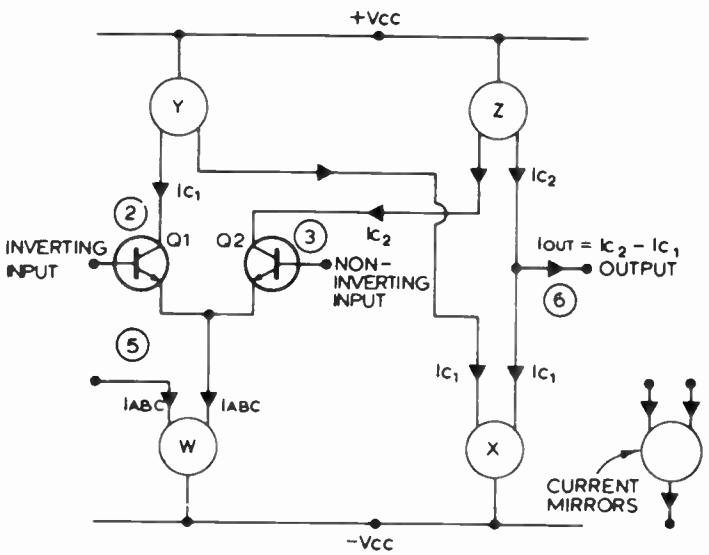

Fig. 1, left: the principle behind electronic multipliers. The graph shows the possible outputs for a variety of combinations of input polarities.
Fig. 2, above: internal workings of a CA3080, an Operational Tranconductance Amplifier.

the product will be negative. If both A and B are negative then the product will be positive thus preserving the arithmetic rules.

If A and B are limited to be only one sign each then the multiplier is known as a one quadrant multiplier. This is the product can only be in one quadrant. If A can be both +ve and -ve, and B only of one sign then the multiplier is known as a two quadrant multiplier. This is what is called an amplitude modulator. The audio signal which is bipolar is A and the control voltage is B.

If A and B can be both +vé and -ve, the product can lie anywhere in the four quadrants and hence the multiplier is known as a four quadrant multiplier. This type of device is found in frequency shifters and ring modulators.

CA3080 — An OTA!

The CA3080 is a two quadrant multiplier, or to give it its full title, it is an Operational Transconductance Amplifier. It has a differential input and a single quadrant current input known as I_{ABC}, (amplifier bias current), Fig. 2. The differential transistor pair is used to steer the I_{ABC} current between the two transistors Q2. There is a region where the input differential voltage is linearly proportional to the percentage of current steered between the two transistors. This voltage region is fairly small, being about 20 mV, but using the CA3080 in this area then a reasonably linear 2 quadrant multiplier can be obtained.

What has happened is the the I_{ABC} current has been multiplied by the input voltage. The product is the difference between the two collector currents. This difference is extracted by the use of mirrors, current mirrors that is. The current mirrors can be attached to either the +ve or the —ve supply rail.

They have two terminals, and whatever current flows into one terminal, then the same flows into the other, which is why they are called mirrors.

What we want to do is take the difference between the collector currents of Q1 and Q2, I_{C1} is reflected from mirror Y and then from mirror X and then appears at the output. I_{C2} is reflected from mirror Z and then appears at the output. The two currents are substracted from each other and the output current is thus $(I_{C2}-I_{C1})$, which is the product of $I_{ABC} \times V_{in} \times K$, where K is a constant. Note that the I_{ABC} current is also reflected from a current mirror on the negative rail.

The CA3080 is a low cost two quadrant multiplier and can be used to perform a wide variety of multiplication functions. The linearity of the device holds true for I_{ABC} variations of over three decades. When using this device keep I_{ABC} below 0.5 mA.

VCA Using CA3046 Array

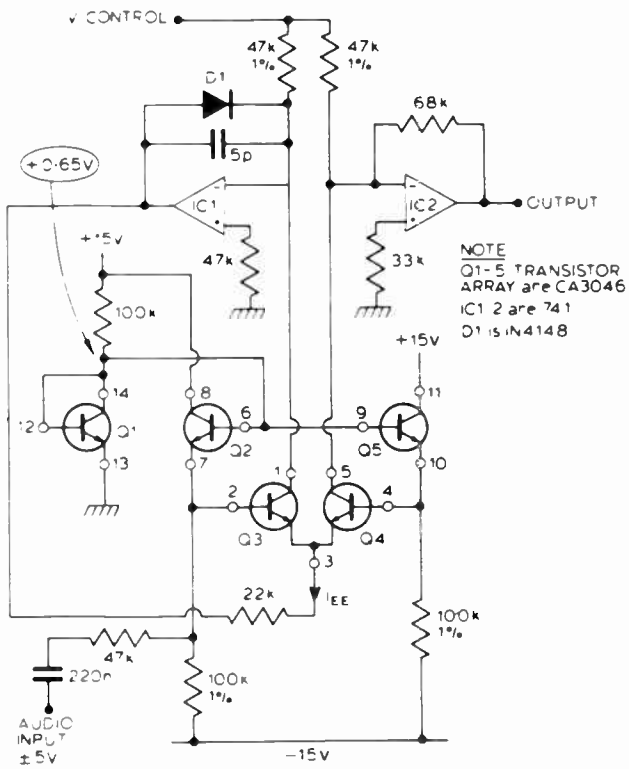

The CA3046 is an array of 5 transistors which are all well matched and relatively cheap. Q3, 4 forms the differential transistor pair, IC1 controls the current and IC2 extracts the differential output current and turns it into an output voltage. The audio input is inserted into the base of Q3 but also connected to this node is the emitter of Q2. Q2 and Q5 serve to predistort the input signal, but they distort the signal the opposite way to which the multiplier distorts it. This is known as distortion cancelling, and it allows a larger signal level to be applied to the multiplier for the same percentage of distortion at the output. The larger input signal allows a higher signal to noise ratio to be obtained. Transistor Q1 is used to bias the bases of Q2, 5 to a suitable operating region.

Stereo Voice Over (Ducking) Circuit for Disco Unit

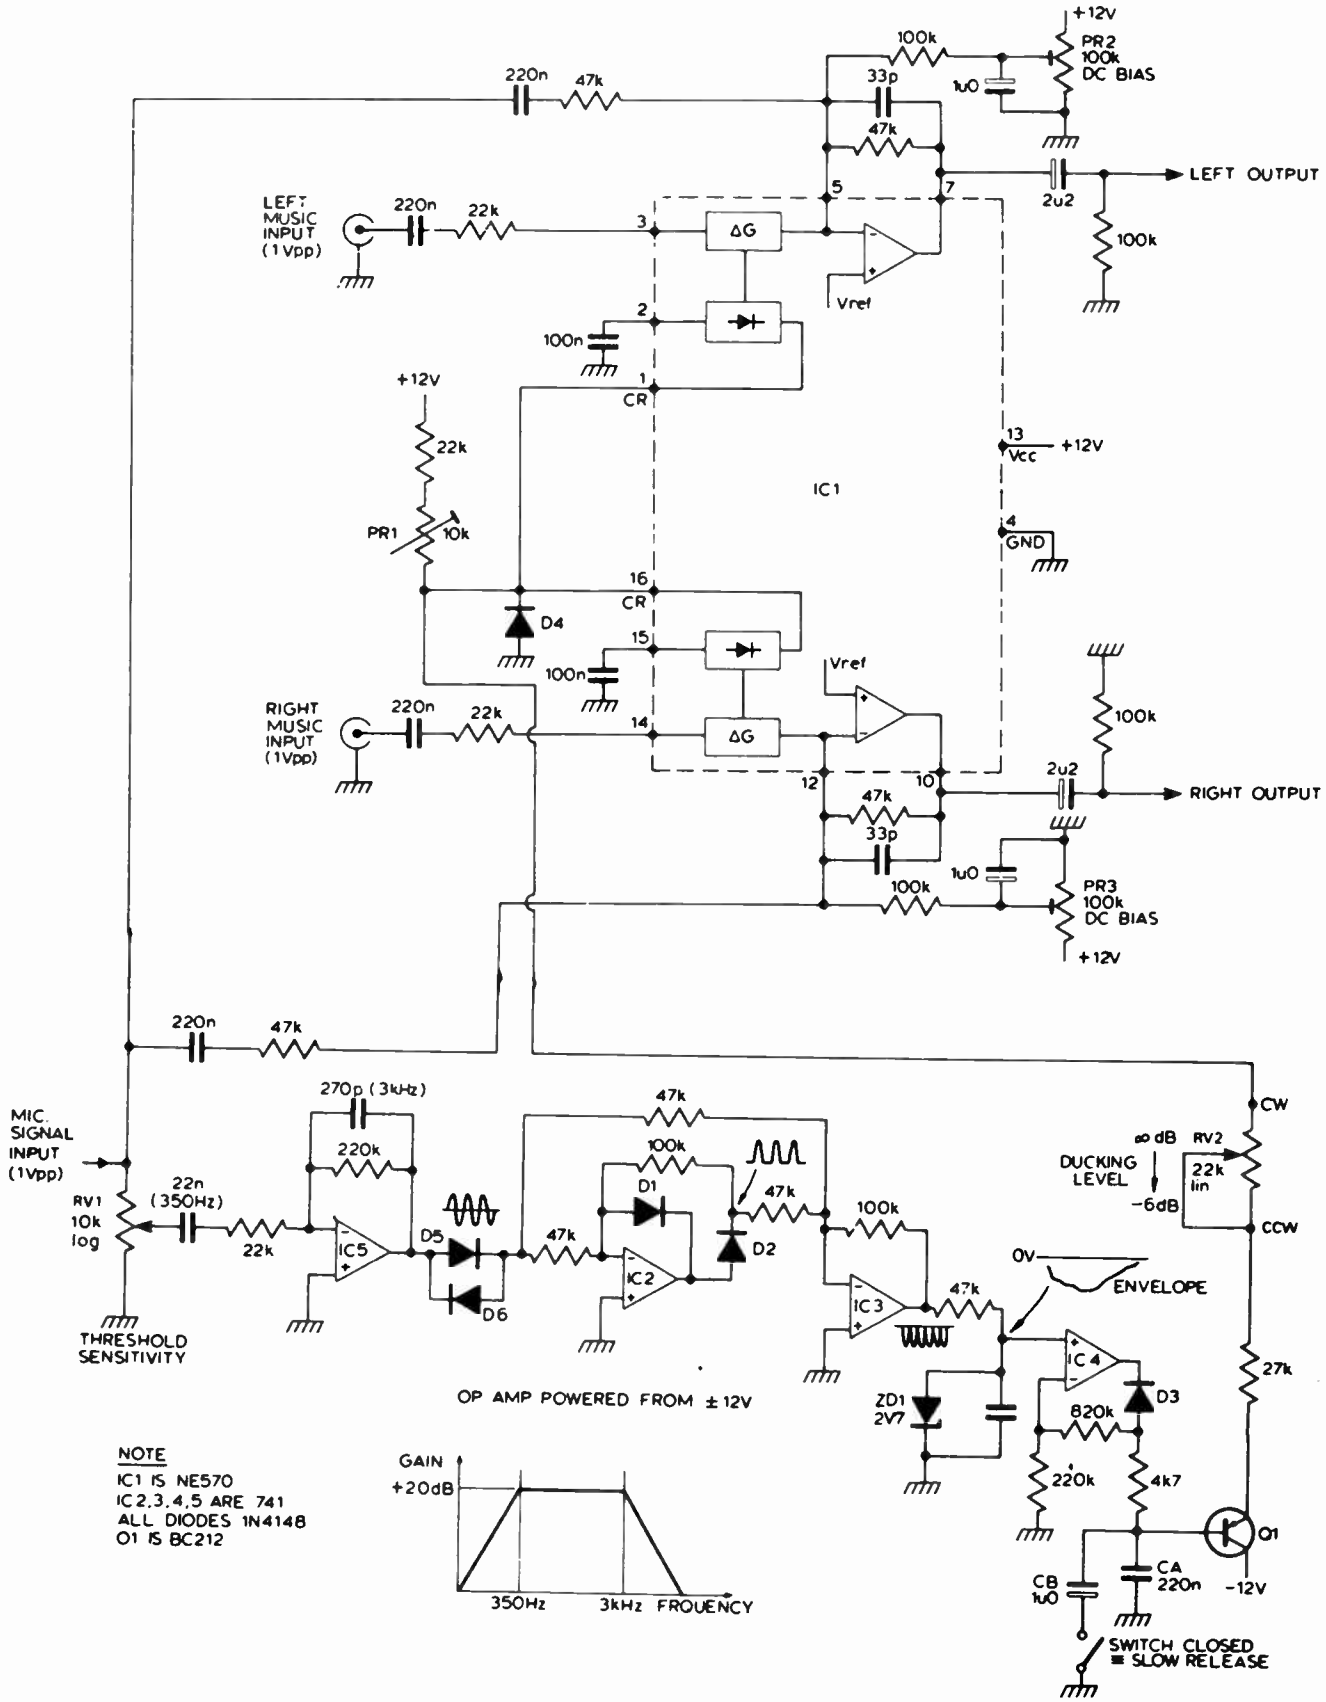

The circuit operation is as follows. The microphone signal comes via VR1. This pot sets the sensitivity of the circuit to the microphone signal. If it is too sensitive the unit will be 'ducking' every time the DJ breathes. IC5 is an amplifier and filter. The filter has been specifically tailored to fit the characteristics of speech, thus making the ducking unit less sensitive to spurious noise. IC2, 3 forms a precision full wave rectifier, the output of which is low pass filtered and then fed to IC4. This wave form is the envelope of the microphone input signal.

IC4 is a peak, negative going, voltage detector with a gain of x 5. When the DJ begins to speak, IC4 goes negative and in doing so pulls the base of G 1 negative. When the DJ stops speaking the base of G 1 rises back towards O V with a time constant determined by CA or CA + CB.

This is the release time and it controls the speed with which the faded down music comes back to full volume. G 1 is an emitter follower and is job is to rob current from the gain cells in the NE570.

This current sets the volume of the two music channels. When the base of G 1 is pulled down to the negative rail, the amount of robbed current is maximum, and when no current flows into pins 1 and 16 of the NE570 and all of it flows into g 1, then both nusic channels are turned off.

To set up PR1, put a large signal into the microphone channel, set RV2 so that it is a short circuit and then adjust PR1 so that the two music channels just close off. PR2 and PR3 should be adjusted so that pins 7 and 10 Of the NE570 are both + 6 V.

Clever Fuzz Box

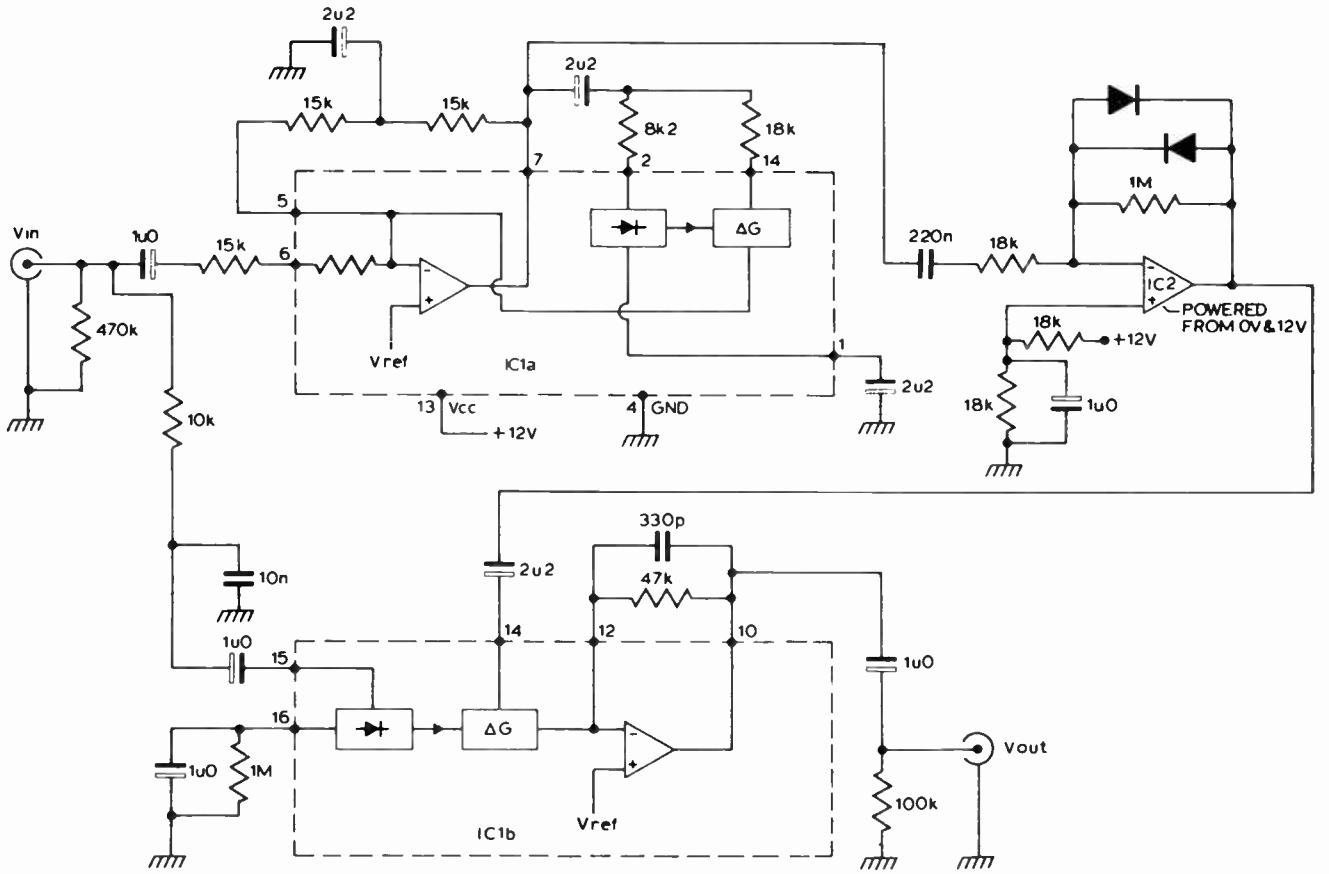

Fuzz boxes are used by guitarists to produce harmonic distortion and sustain. If you want to produce only the distortion, but to retain the original envelope of the signal then this is the circuit for you.

IC1 is a 2:1 compressor as described previosuly. This produces a relatively high level signal which then drives IC2, which is a x 50 amplifier with diode clamping. IC2 produces the distorted (fuzz) found. This is then fed into the IC3 gain cell, the output of which drives the op amp. This gain cell is driven by the rectified original signal (low pass filtered at 1k5 Hz), so that the distorted sound is given the envelope characteristics of the original sound.

If a fuzz sustain sound is required rather than a dynamic fuzz then IC3 could be modified (by the inclusion of a clamped high gain amplifier driving pin 15) so that it acts as a low level expander. This will squelch the noise at the end of the fuzz period.

Track and Hold

In this example the CA3080 is used as a current controlled switch. When the control voltage is high, I_{ABC} is maximum, (0.44 mA) and the OTA gain is maximum. The voltage at pin 2 of IC1 adjusts itself so that it is the same as that on pin 3, this being due to the 100 per cent feedback via the high input impedance voltage follower IC2. When the control voltage is OV, I_{ABC} is zero and hence the gain of the OTA is zero. Therefore no current comes out of its output and so the voltage at the output of IC2 remains frozen (Hold mode). The maximum differential input voltage is 5 V and this must not be exceeded. The capacitor C should be selected to suit the speed of the operation.

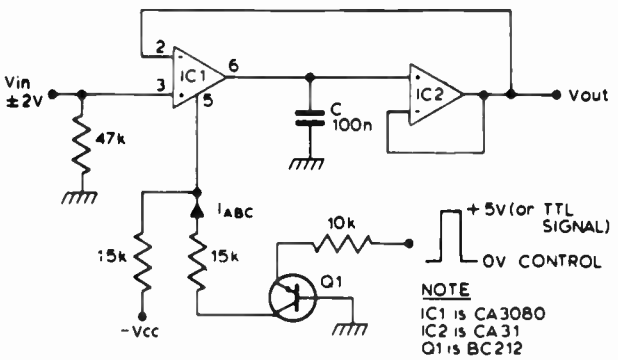

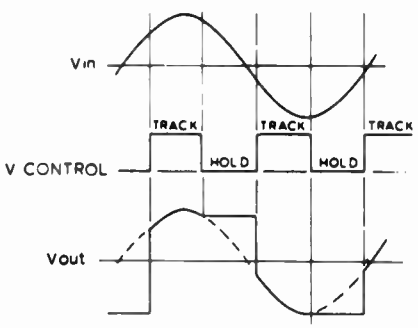

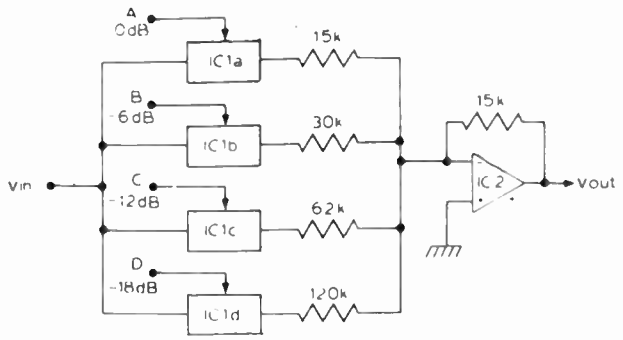

0 ≡ -6V
1 ≡ +6V

A	B	C	D	GAIN
1	0	0	0	0dB
0	1	0	0	-6dB
0	0	1	0	-12dB
0	0	0	1	-18dB

POWERED BY ±6V

NOTE
IC1 is 4016
IC2 is 741

Voltage Controlled (Switched) Attenuator

The CD4016 is a quad analogue transmission gate. That is, it is a quad voltage controlled switch. When the control is high the switch is ON, having an effective resistance of about 400R. When the control is low the switch is off and ot looks like a 100M resistor. Thus by using 4016 switches it is possible to 'Switch' the voltage gain of an amplifier. The resistors in this example are selected to give 6 dB changes in gain.

Filter

A state variable filter produces three outputs: highpass, bandpass, and lowpass. It is thus a very versatile filter structure, even more so if the resonant frequency can be varied. This frequency is linearly proportional to the gain of the two integrators in the filter. Two CA3080's, (IC2, 4) have been used to provide the variable gain, the resonant frequency being proportional to the current I_{ABC}. Using 741 op amps for IC3 a control range of 100 to 1, (resonant frequency) can be obtained. If CA3140's are used instead of 741's then this range can be extended to nearly 10,000 to 1.

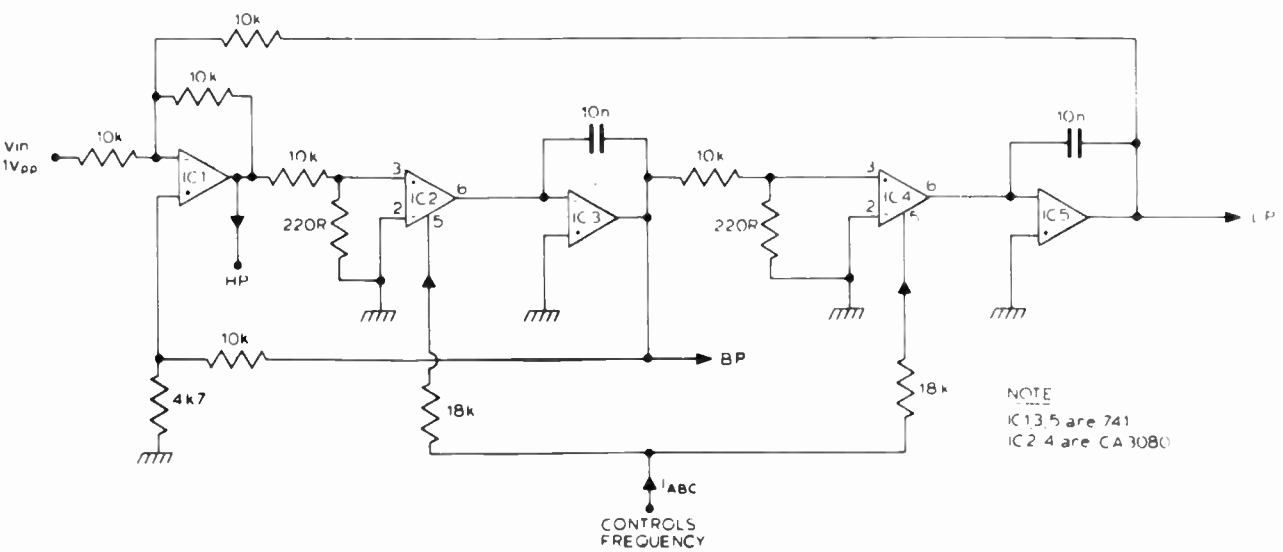

Basic Limiter Circuit

Most professional limiter circuits use a FET as the variable gain element. Relatively low distortion with a reasonable signal to noise ratio can be obtained. A basic limiter circuit is shown this being no different to previous circuits except for the variable gain element.

When a relatively small voltage (20 mV) is applied to the drain source of a FET, it acts like a fairly linear resistor. As the gate source voltage is varied, this resistor (RDS) also varies.

In fact the channel resistance RDS is inversely proportional to gate source voltage V_{GS}. When V_{GS} is OV, then RDS is at its generally minimum resistance (R_{ON}) which can be as low as 5R, but it is generally more like 100R. When V_{GS} exceeds the pinch off voltage (Vp or V_{GS} off) the channel resistance goes up to several hundred megohms. So a junction FET can be used as a voltage controlled resistor, except that R_{ON} and V_{GS} (OFF) tend to vary widely from device to device. However with a bit of perseverance suitable devices can be selected and made to work.

One circuit trick that greatly reduces distortion is shown here. Half of the audio signal at the drain of the FET is presented to the gate. This is superimposed on top of the control voltage and produces a distortion cancelling effect. Distortion levels below 0.1% can be achieved using this technique.

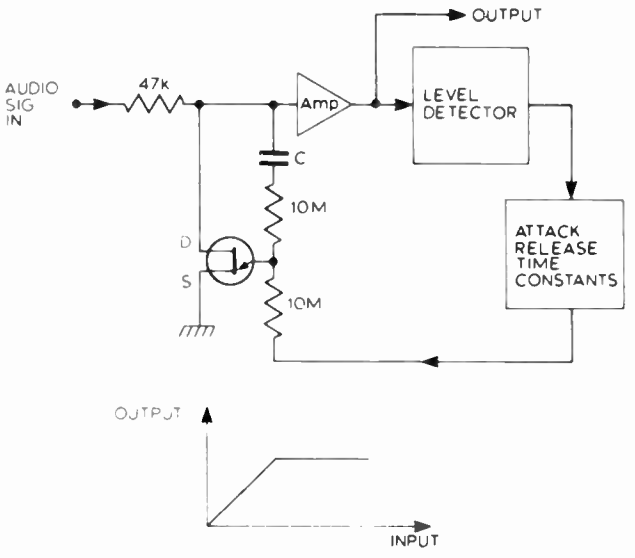

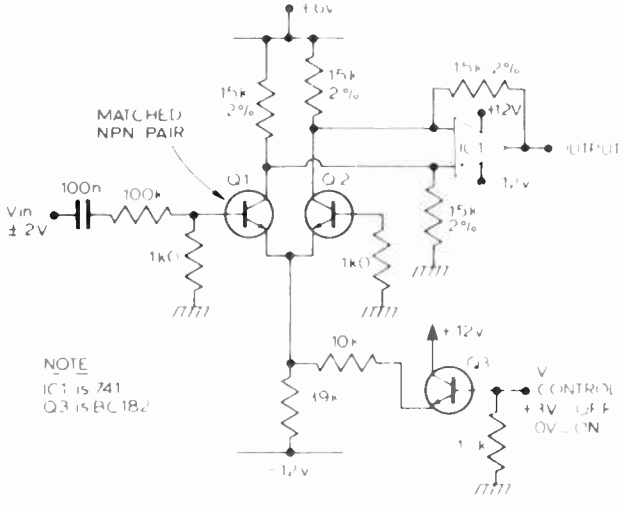

Transistor VCA

A circuit similar in operation to a CA3080 can be constructed with a matched pair of transistors and an op amp. Transistors Q1, 2 form a differential transistor pair which is used to steer whatever current is available between the two collectors, just as in the CA3080. The difference between the collector currents is equal to the product of the input voltage times the current I_{EE} times a constant. This difference is extracted by the differential amplifier IC1. The current I_{EE} is controlled by Q3. As the control voltage goes positive, Q3 robs most of the current flowing down the 39k resistor, and hence I_{EE} and the output of IC1 decrease.

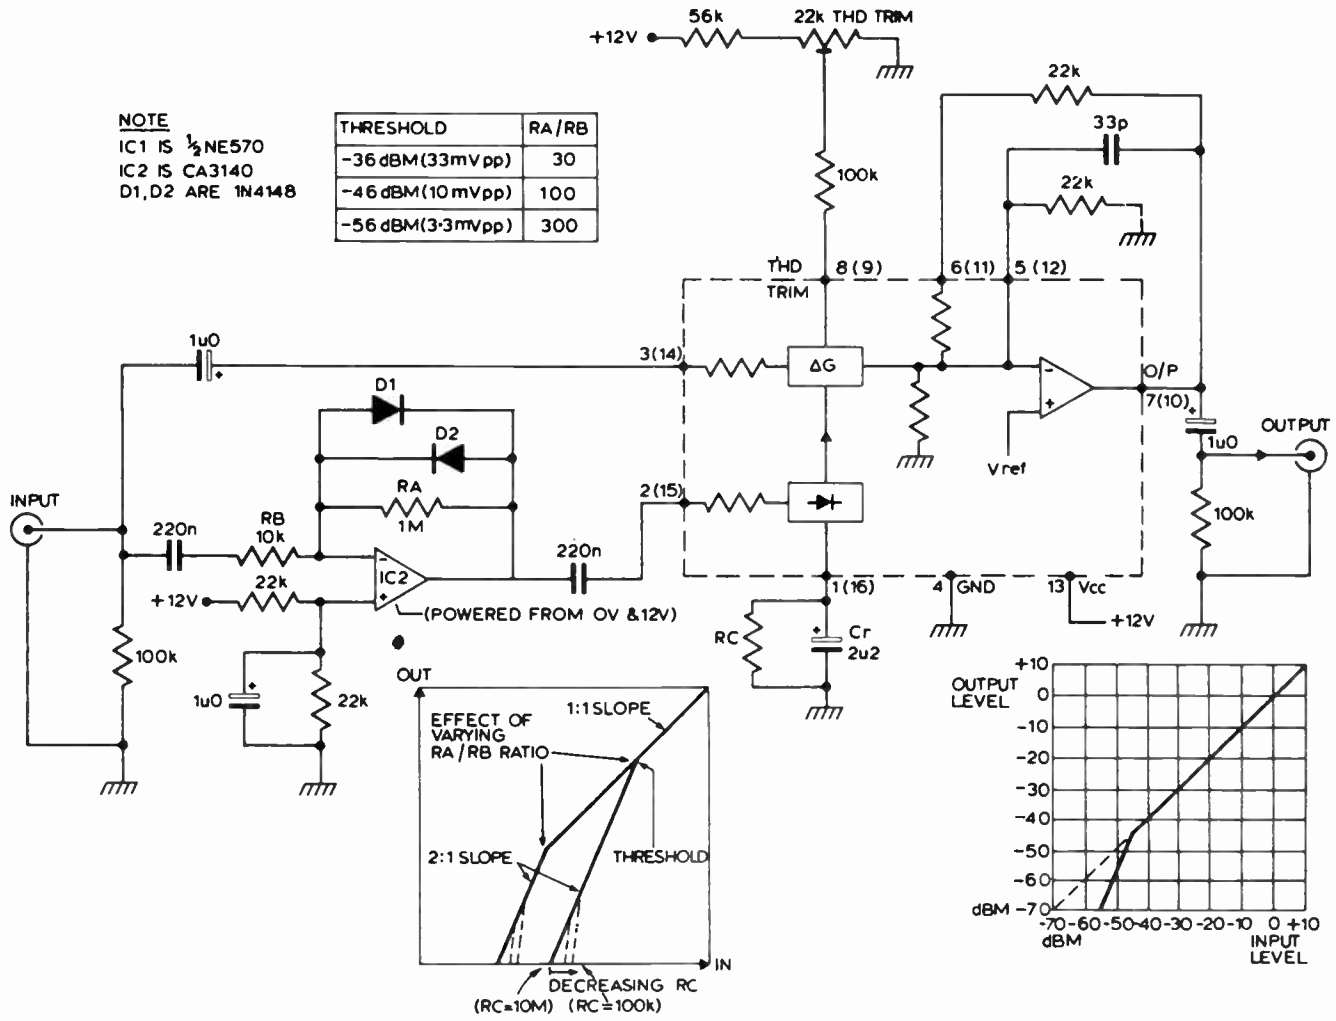

Two Channel Low Level Expander/Noise Gate

It is often required that a rather noisy signal be cleaned up a bit. This is not possible to do continuously, but it is possible to clean up noise in what was initially the gaps. The results of this cleaning up process can quite often be heard when telephone conversations from "foreign correspondents" are broadcast.

By turning down the signal level in the gaps, (by performing a low level expansion) the perceived sound quality improves dramatically.

The circuit performs just such an expansion. The inputs signal passes through the variable gain cell and then appears at the op amp output. The gain of the gain cell is controlled by the signal coming from IC2. This is a high gain amplifier with diode clamping, so that the output swing is limited to about 1VO ptp. Therefore for input signals of 10 mV pp to 10 V pp, the output of IC2 remains at about 1 VO ptp to 1V2 ptp.

So, for this range of input voltages the gain of the gain cell remains roughly static. Now when the inout level drops below 10 mV, the output of IC2 will start to fall and so will the gain of the gain cell. This produces a 2:1 downwards expansion curve, which means that the output then gets quieter at a rate faster than the input. To accentuate this effect, a bleed resistor can be placed in parallel with Cr.

The resistor robs some of the current that would have otherwise gone to the gain cell and causes the input/output curve to roll off much more rapidly at low signal levels. Also, by varying the resistor ratio of RZ/RB, the expansion threshold level can be altered.

Incredibly Simple Compressor

Not all gain control systems need to be complicated or indeed active. One product which I saw advertised was a compressor to help prevent loudspeaker/overloads. All it was was a lightbulb in series with the loudspeaker. When the power exceeds a certain level, the lamp will turn on, glow, its resistance increase dramatically and hence a bigger percentage of the power output is dissipated in the lamp. A nice, simple solution, but I think it would require some experimentation to find the right sort of car headlamp bulb!

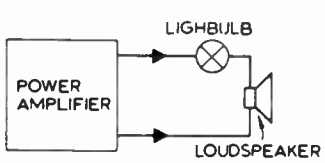

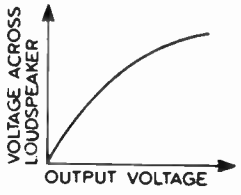

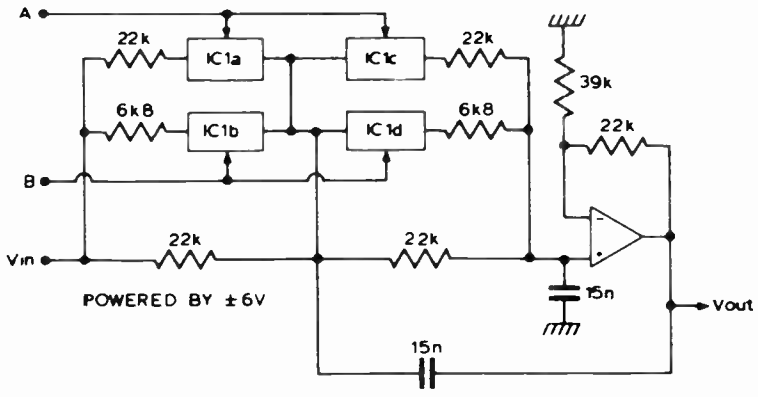

	①	②	③	④
A	OFF	ON	OFF	ON
B	OFF	OFF	ON	ON
Fc	500Hz	1kHz	2kHz	2·5kHz

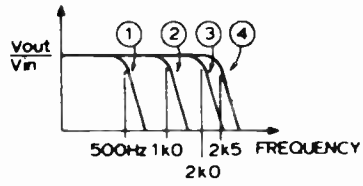

Switched Frequency Low Pass Filter

In this example the effective resistance is switched by using 4016 gates. The filter is a lowpass Butterworth and by turning gates A or B ON or OFF the cut off frequency can be altered. This allows the filter control to be physically remote or even to be computer controlled. Mark Space modulation of A and B would enable continuous control over the cut off frequency.

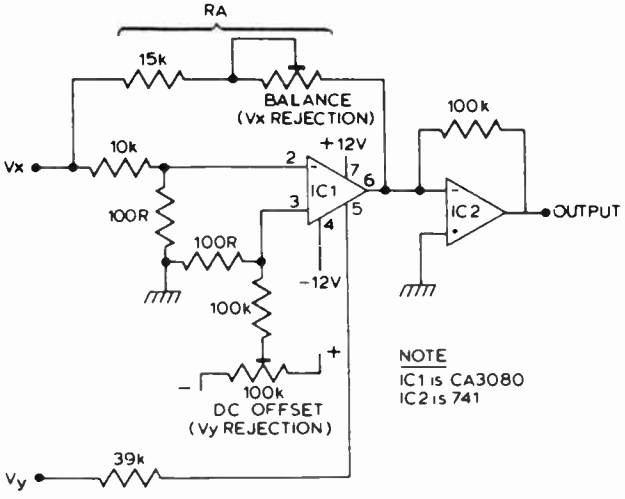

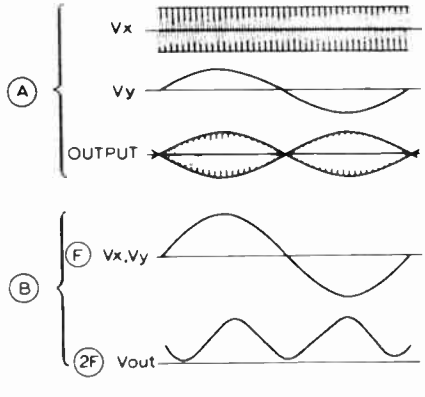

Four Quadrant Multiplication

By using a few circuit tricks, the CA3080 can be made to perform 4 quadrant multiplication. In fact the CA3080 performs 2 quadrant multiplication and the trick is to move the axis on the multiplying graph. If we ignore the RA resistor chain then we have a 2 quadrant multiplier circuit similar to that shown previously. Imagine that V_x is a 1kHz sine wave. 1 Vptp and V_y is a 0V. The output of IC2 is a sine wave of fixed amplitude. Now if we connect RA , and adjust the balance control, it will be possible to cancel out the output, because the signal coming from IC1 is out of phase with that from the RA resistor chain. So with V_y set at 0 V there is no output for IC2. If V_y goes + ve, the output of IC1 will become greater than the current via the RA chain and the output if IC2 will grow.

If V_y goes–ve the current through the RA chain will exceed that from IC1 and the output of IC2 will grow, the phase being opposite to that when V_y was a sinewave from an oscillator, then this circuit could be used to generate ring modulation effects.

When V_x is set at OV there may be some V_y breakthrough and this can be minimised by adjusting the V_y rejection preset.

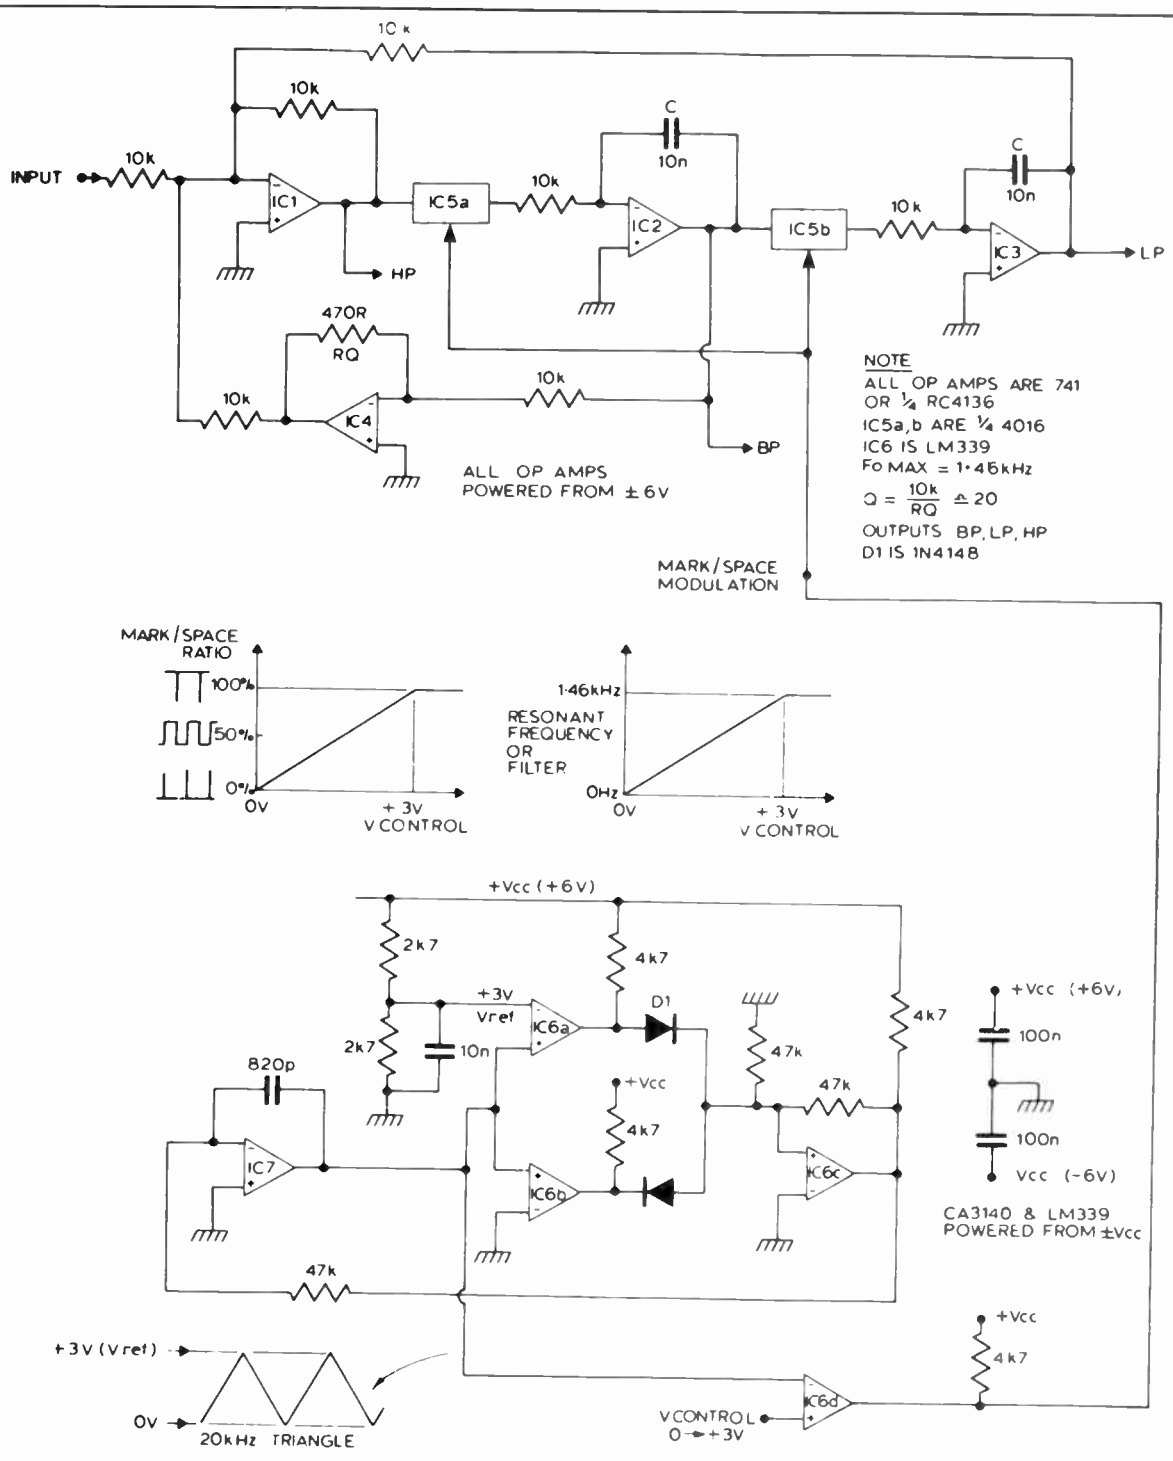

It is possible to change the gain of an amplifier by effectively altering the input resistor. This can be done by markspace modulating a voltage controlled switch in series with the resistor.

When the markspace ratio is low, the switch is OFF most of the time and the effective resistance is large. When the markspace ratio is high the switch is ON most of the time and the effective resistance approaches that of the series resistor.

Having generated a markspace control waveform, it is possible to gang up together literally hundreds of voltage controlled switches. This enables large numbers of variables to be simultaneously changed.

The circuit is a markspace modulated universal filter (IC-6) and the markspace generator itself (IC-11).

IC7-10 forms a triangle square wave oscillator. IC7 is an integrator whose outout ramps up and down between OV and a + 3 V reference. IC8-10 are all fast comparators. IC8 detects when the integrator outputs of IC8 & 9 are used to flip over a schmitt trigger IC10, which then drives the integrator. Thus the integrator output ramps up and down between OV and + 3 V at a rate of 20 kHz.

It is important that the frequency of the markspace oscillator be relatively high. As a rule of thumb it should be 2 ½ times the highest frequency components of the signals that you hope to process. The triangle output is fed into IC11's inverting input, the control voltage into the non inverting input. The output of IC11 is the markspace modulation which is used to drive the switches IC5,6. The filter resonant frequency is directly proportional to the mark space ratio that drives these switches

The number of IC's used is a quad package, and so is the 4016 and so can be the op amps (use RC4136). Thus the whole circuit can be realised with only 4 IC's. Also the mark space oscillator canbe used to drive other independent comparators

Why some CMOS circuits don't work as you expect

Stephen Dolding

A 4093 is a 4093, right? Well . . . yes, and no. There are quite a few pitfalls in the CMOS 'jungle' and it's handy to know about them before venturing forth.

CMOS '4000' SERIES integrated circuits are manufactured by at least six major manufacturers and the 74C series by at least two major manufacturers, but it must not be assumed that a 4XXX from one manufacturer is interchangeable with a 4XXX device from another manufacturer. This article explains some of these differences.

Schmitt gate oscillator

What could be simpler than the oscillator circuits shown in Figures 1(a) and 1(b)? There are so few components that you would expect these circuits to work first time.

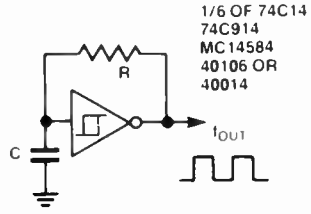

Figure 1a. Different manufacturers ICs will produce different results

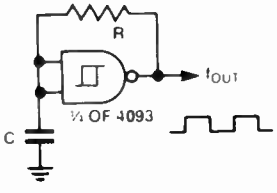

Figure 1b. Such a simple circuit but the value of f_{out} can vary from 52.9 Hz to 249 Hz, depending on the brand of 4093 used

You have selected stable, close tolerance components and calculated the frequency according to the formula:

$$f_{out} = \frac{1}{CR\log_e\left[\left(\frac{V_{cc}-V_{t-}}{V_{cc}-V_{t+}}\right)\left(\frac{V_{t+}}{V_{t-}}\right)\right]}$$

where C = capacitor value in uF
R = resistor value in kilohms
f_{out} = frequency out in kHz
V_{cc} = supply voltage
V_{t+} = upper trigger level of Schmitt trigger
V_{t-} = lower trigger level of Schmitt trigger

Now, the question is, "what are the values of V_{t+} and V_{t-}?" We need to refer to the manufacturer's data sheet for an answer. But which manufacturer? There are at least six different manufacturers to choose from and each one gives a different range of possible values for V_{t+} and V_{t-}.

Considering the 4093 IC (Figure 1b) it all the databooks are consulted it is found that for V_{cc} = +5 V the *highest* typical value of V_{t+} = 3.6 V (extreme = 4.3 V).

The *lowest* typical value of V_{t+} = 2.7 V (extreme = 1.7 V).

The *highest* typical value of V_{t-} = 2.2 V (extreme = 3.3 V).

The *lowest* typical value of V_{t-} = 1.4 V (extreme = 0.7 V).

Now if V_{t+} = 3.6 V, V_{t-} = 1.4 V, V_{cc} = +5 V, C = 100 nF and R = 10k, then from the formula above f_{out} = 52.9 Hz.

With the same type from another manufacturer, V_{t+} could be 2.7 V and V_{t-} could be 2.2 V, in which case recalculation gives a frequency of 249 Hz!

These figures are based on typical values of trigger level. Extremes of high and low trigger levels could give frequencies ranging from 27.5 Hz to 961 Hz with the same values of C = 100 nF and R = 10k. This gives a frequency range of almost 35:1 if the whole spectrum of possibilities is considered.

Now it will be clear why the circuit may oscillate at a frequency which is considerably different from what was expected or intended by the designer who only consulted one manufacturer's databook!

The monostable

Now let us look at another CMOS circuit often used by the hobbyist — the 4528 dual monostable, shown in Figure 2.

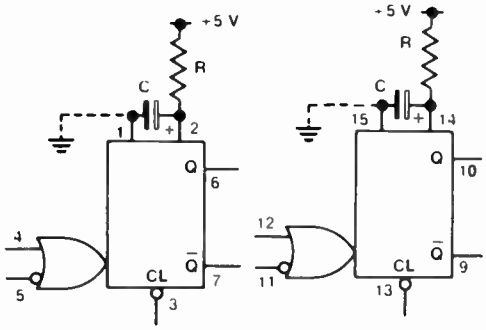

Figure 2. The 4528 dual monostable can have variations in its performance of up to ±50%

Here again, different manufacturers give different formulae for the monostable time constant. A typical formula (for +5 V supply) is

$$t = 0.37CR$$

However, it could vary from $t = 0.32CR$ to $t = 0.42CR$ typical, with variations up to ±50%.

The pulse width depends very much on the supply voltage. Some manufacturers' 4528s give increasing pulse width with increasing supply voltage, others give a reverse effect.

The formula given above depends on the value of C being greater than 10n. For smaller values of capacitance the manufacturers' data sheets need to be consulted.

It should also be noted that some manufacturers require pins 1 and 15 to be grounded externally for correct operation. To overcome the variations in timing formula, a CMOS 4538 integrated circuit can be used in place of the 4528. The 4538 is pin compatible with the 4528 IC and the formula is:

$$t = CR$$

with variations of only ±5%.

In all timing circuits using CMOS ICs it is wise to make provision for trimming the value of the timing resistor to allow for adjustment.

The counter/divider

Next we come to a well-known decade counter/divider IC — the 4017.

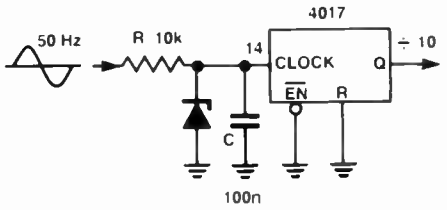

Figure 3. This circuit will only work correctly with particular brands of 4017 decade counter/dividers. The wrong choice can result in false counting.

Will the circuit in Figure 3 always work correctly? No, only with a Motorola 14017 or an RCA 4017, because these have an internal Schmitt trigger on the clock input. Other manufacturers' 4017s do not, so false counting may result.

BCD decoder

Another fairly common integrated circuit likely to give problems is the 4028 BCD-to-10-line decoder. We now come to see that actual logical differences can occur between one manufacturer's 4XXX and another manufacturer's 4XXX.

With the 4028, some manufacturers (Motorola and RCA) do not decode the six 'illegal' binary codes 1010 to 1111 (i.e: 10-15), while other manufacturers (including National, Fairchild and Philips) decode these outputs as if the input was 8 (1000) or 9 (1001).

The problem of logical differences between one manufacturer's device and another manufacturer's device (with supposedly the same type number) applies also to the 4585 four-bit comparator and even to the ubiquitous 555 and 556 timers. There may be other examples too. The problem fortunately does not occur with the range of quad gates.

The moral

The above-mentioned examples were all encountered during the design of one piece of industrial equipment which made use of these common CMOS parts.

You may well ask "If design engineers, who have ready access to all the data books, can run into such problems, what about the unsuspecting hobbyist, who has no data?"

The moral of this article is that "forewarned is forearmed". It is hoped that this article may at least prevent some construction projects from being abandoned because they do not appear to work correctly at first sight. Designers who publish projects should check that there are at least two manufacturers' ICs which will work in the circuit as intended and, if necessary, spell out the names of suitable manufacturers in the parts list. Best of all, only design circuits that will work with all manufacturers' devices of the same basic type number (though this may not always be possible).

If problems occur, all that may be required is to try an IC from a different manufacturer. ●

Surface acoustic wave devices — fundamentals and applications

Surface acoustic wave (SAW) devices have been used for the last decade in professional and military equipment and are now beginning to be found in consumer electronic products. In these articles the basic physics of SAW devices are explained and some important current applications of the technology are examined.

P. J. Hall
Physics Department, University of Tasmania

Part 1

IN 1887 Lord Rayleigh, the great pioneer of modern acoustics, showed that a unique mode of wave propagation can exist at the surface of an elastic solid. This Rayleigh, or surface, mode wave has four important characteristics:

(1) The wave has a longitudinal component (solid displacement parallel to the direction of propagation) and a transverse component (displacement perpendicular to the propagation direction) each 90° out of phase. This gives rise to a backward elliptical motion of solid elements at the surface (Figure 1a). The amplitude of both wave

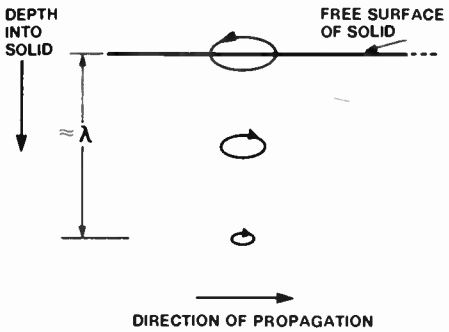

Figure 1(a). Motion of solid elements in an isotropic solid due to the propagation of a surface wave. Similar disturbances occur along axes of high symmetry in the anisotropic crystalline solids used for SAW device substrates.

components falls rapidly within the solid, confining the wave to within one or two wavelengths of the surface. The wave generates an oscillation of the surface as shown in Figure 1b.

(2) The wave velocity of propagation is of the order of 1500-4000 metres per second for most materials and is thus about 100 000 times as slow as electromagnetic wave velocities. The SAW wavelength is reduced by this factor, allowing a physically small length of solid to contain a great number of wavelengths.

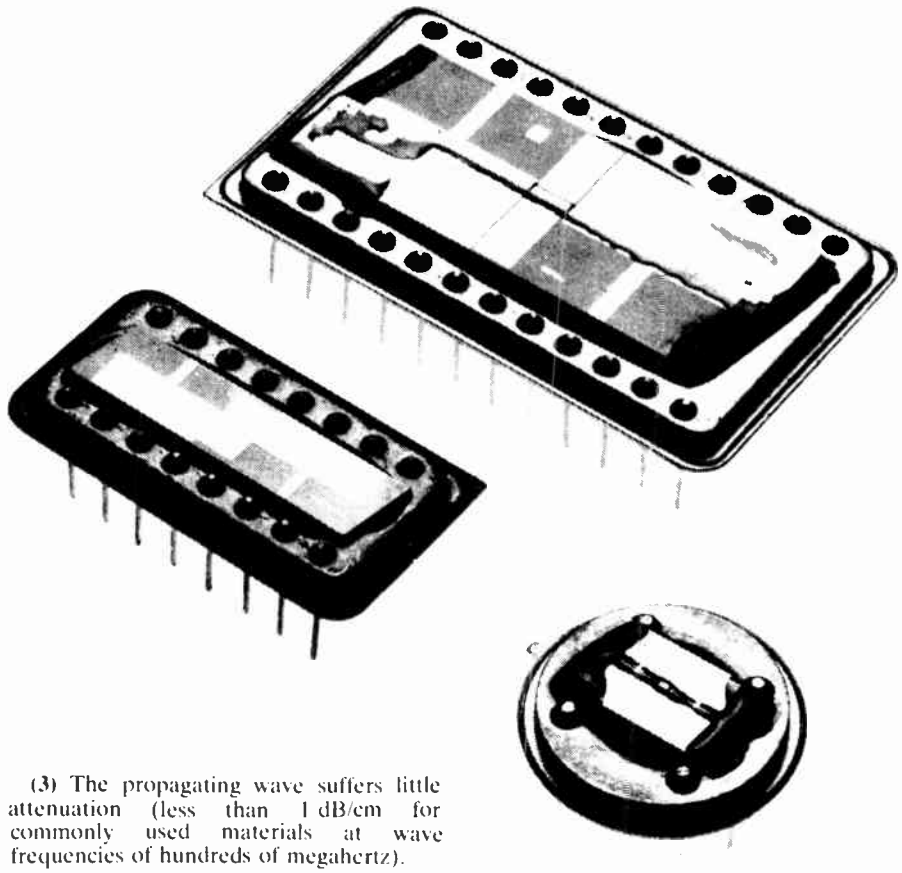

(3) The propagating wave suffers little attenuation (less than 1 dB/cm for commonly used materials at wave frequencies of hundreds of megahertz).

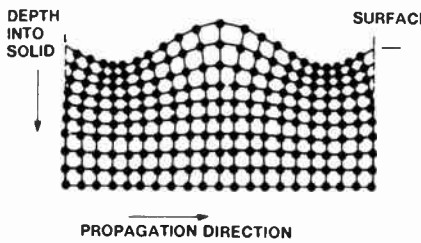

Figure 1(b). Oscillations in a vertical plane through a solid. Note the surface oscillations.

(4) The wave mode is essentially non-dispersive i.e: waves of different frequencies propagate with the same velocity.

The usefulness of SAW technology in electronics arises because it is possible to convert radio frequency electrical signals (over the range extending from a few megahertz to about one gigahertz) to surface acoustic waves and then back to electrical signals in a fairly efficient transduction process.

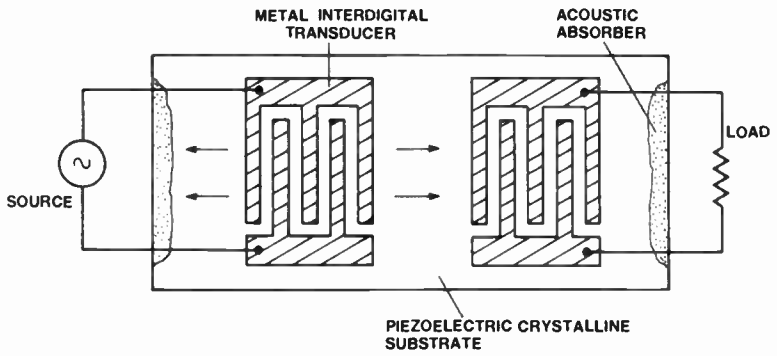

Figure 2. Basic SAW device configuration.

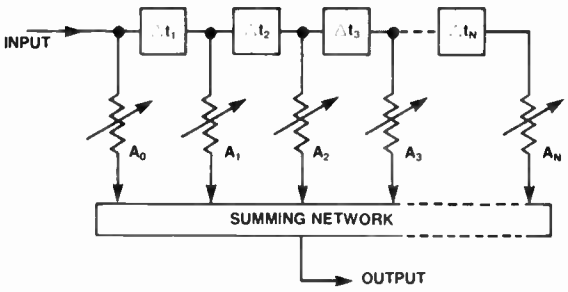

Figure 3. A transversal filter. In the SAW realisation the tap amplitudes A_0-A_N are set by the length of the IDT fingers and the time delay between taps is controlled by the relative positions of the fingers.

A mechanism for achieving this is shown in Figure 2. Using standard microelectronic fabrication techniques a metal interdigital transducer (IDT) is deposited on a piezo-electric crystalline substrate, usually quartz or lithium niobate. Although complex in detail the basic IDT operation is fairly easy to understand.

The electrical input signal is applied in such a way as to ensure that, at any instant, the charge on neighbouring IDT 'fingers' alternates in sign. The resultant electric field between the fingers leads to a mechanical stress in the piezo-electric solid and surface acoustic waves are launched in both directions. For maximum efficiency the physical spacing of the IDT fingers ought to be one half of an acoustic wavelength at the input frequency.

Wideband transducers can be designed by optimising the electrical and acoustic loads on the IDT. The propagating waves create electric fields in the solid which cause a voltage to be induced in the output transducer when the waves pass underneath.

The device illustrated in Figure 2 is an electrical delay line, the delay being simply the separation of the IDTs divided by the SAW velocity. Typical delays range from microseconds to hundreds of microseconds.

In many devices the unwanted energy propagated by the bi-directional IDT is either absorbed by an acoustic termination or scattered away from the functional surface. This leads to a minimum device insertion loss of 6 dB and also means that the portion of the signal reflected back from the receiving transducer, then back from the transmitter to finally arrive at the receiver IDT again, is suppressed by only 12 dB relative to the wanted output.

This 'triple transit' signal is the main device spurious response and can be further suppressed by deliberately mismatching the impedance of the IDTs, leading to an increased insertion loss. Fortunately the trade is reasonable: increasing the insertion loss from 6 dB to 12 dB improves the triple transit suppression from 12 dB to 33 dB.

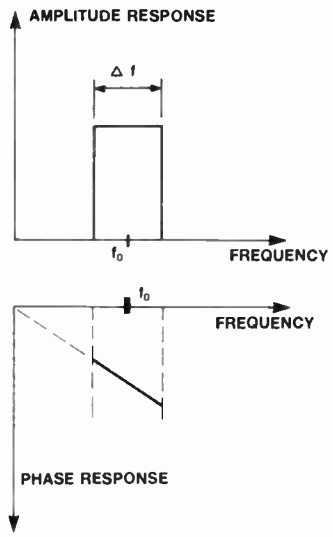

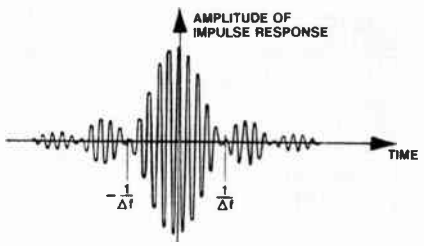

Figure 4. Responses of an ideal bandpass filter. The amplitude and phase responses are 'single-sided' (no negative frequencies) representations and the impulse response is derived by considering a psuedo-impulse (a very short pulse) modulated onto a carrier at the filter centre frequency. The impulse response is infinite in extent in time.

Polyphase IDTs are now available which launch a unidirectional surface wave and lead to devices with much lower insertion loss and better triple transit suppression. Nevertheless, the relatively high insertion loss of most currently available SAW devices is due to IDT mismatch.

To appreciate the signal processing potential of SAW devices it is necessary to look on the IDT as a sampled delay line. Surface waves travel from one pair of electrodes to the next at a finite velocity so the input signal is repetitively delayed and added to itself. The constructive or destructive signal addition creates the amplitude and phase response of the SAW device.

The amplitude of a particular delay line tap can be set by choosing the length of the IDT fingers at that point. The phase, or time delay, of the tap can be varied by changing the relative position of the fingers. In a given SAW device either the transmitter or receiver IDT, or a combination of both, may be used as the signal processing element.

It turns out that such a delay line is exactly what is required to form the basis of a transversal filter. In such a device a given frequency response is obtained by synthesising the time domain impulse response of the desired filter. The impulse response is the output from the filter when the input is a very short pulse and is, in fact, the Fourier transform of the desired frequency domain characteristic.

Figure 4 shows the impulse and frequency response of an ideal band-pass device, the type of filter to be discussed later.

The impulse response synthesized by a transversal filter is necessarily finite in length and sampled rather than continuous, but by following the usual principles applying to the design of sampled systems, the desired impulse response can be well approximated.

To use a SAW device as a filter it is only necessary to 'draw' (using metallization) the sampled impulse response on the substrate. The second IDT can be a broad-band type designed to pass all frequency components of interest (Figure 5).

One major constraint applies to all SAW filters. Since the polarity of the charge must reverse on alternate IDT electrodes in order to establish surface wave propagation, the sign of alternate samples of the impulse response reverses.

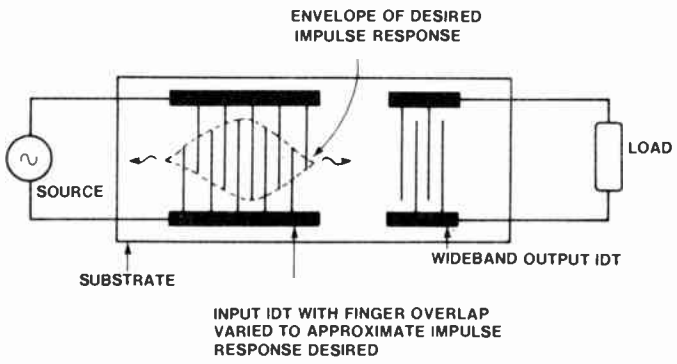

Figure 5. SAW transversal filter.

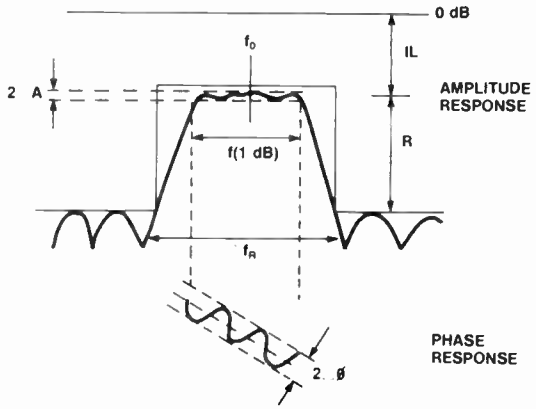

Figure 6. Major parameters of a bandpass filter.

Sampling in this way produces a pass band displaced from baseband i.e: a bandpass filter.

A major advantage of SAW filters is that they are not 'minimum phase' devices, therefore the amplitude and phase responses can be specified independently. Hence, a fairly sharp frequency response can be combined with a linear phase response, often important in video and data distribution networks.

Before examining some specific applications of SAW filters it is worth summarizing some of the advantages of the technology. These advantages can be listed as:

(1) Versatile response with essentially no adjustment or tuning.
(2) Straightforward fabrication using standard techniques.
(3) Excellent repeatability from device to device due to the photo-lithographic processes used in making the IDTs.
(4) Small size and weight. Mechanically rugged.
(5) Passive.
(6) Wide frequency range of operation (10 MHz to 1 GHz).
(7) Good temperature stability.
(8) Graceful degradation (small faults cause small performance changes).
(9) Radiation resistant.

SAW bandpass filters

The primary specifications for a bandpass filter are the centre frequency, f_0, and the bandwidth, $\triangle f$. These and other secondary specifications are explained in Figure 6. It is convenient to express the bandwidth in a normalized form, the percentage fractional bandwidth of the filter being simply $(\triangle f\ f_0)X100$.

Table 1 details the achievable performance with typical modern SAW filters but does not include data for the new devices with polyphase IDTs.

The most commonly used substrate material is lithium niobate, but in filters with a small fractional bandwidth (less than 5%) quartz may be used to ensure greater temperature stability. The centre frequency range is restricted at the low frequency end by large device die sizes and a poor spurious response performance by the filter. At the high frequency limit, photo-lithographic processes used to manufacture the IDTs are severely taxed.

Table 1. Typical Parameters of SAW transversal bandpass filters

Substrate material	Lithium niobate or quartz
Centre frequency (f_0)	10-1000 MHz
1 dB bandwidth ($\triangle f$)	60 kHz to 50% fractional bandwidth
Insertion loss (IL)	10-30 dB
Amplitude ripple ($\triangle A$)	± 0.2 dB
Phase ripple ($\triangle Ø$)	± 1°
Ultimate rejection (R)	50-70 dB
Shape factor ($\triangle f_R/\triangle f$)	1.1
Size	10-30 mm

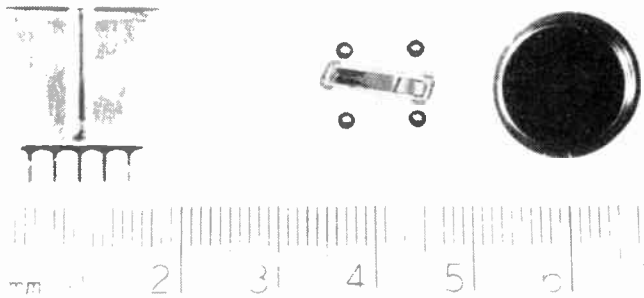

Figure 7(a). Two versions of a SAW TV IF filter. The filter on the right has the top removed to reveal the substrate and IDTs (devices courtesy of Signal Technology Ltd).

The filter bandwidth limitations are set essentially by the number of fingers in the IDT. For very narrow bandwidths (less than 60 kHz) the number of fingers needed is very large, causing an intolerable number of reflections and spurious responses.

In situations where such narrow bandwidths are needed, a second type of SAW filter is used. These 'resonator' filters are the SAW analogy of microwave cavities and can achieve fractional bandwidths of 0.01% or so.

At the other extreme, fractional bandwidths of more than 40% are difficult to synthesize with SAW transversal filters because of the very small number of IDT fingers employed to approximate the wanted response.

The insertion loss figures quoted in Table 1 are typical but the adoption of polyphase IDT techniques can reduce the insertion loss to as low as 1.5 dB. As a general rule, increasing the fractional bandwidth of a SAW filter increases the insertion loss since the resistive loading of the IDT must be increased.

Amplitude and phase ripple in a well-designed filter are due mainly to spurious device responses. These responses also govern the ultimate rejection capability of the filter. The 50-70 dB rejection quoted is achieved routinely in modern filters.

Figure 7a shows a typical SAW bandpass filter, in this case a 38.9 MHz TV IF filter. The same device is available in a choice of the two packages shown and both encapsulations sell for a few dollars.

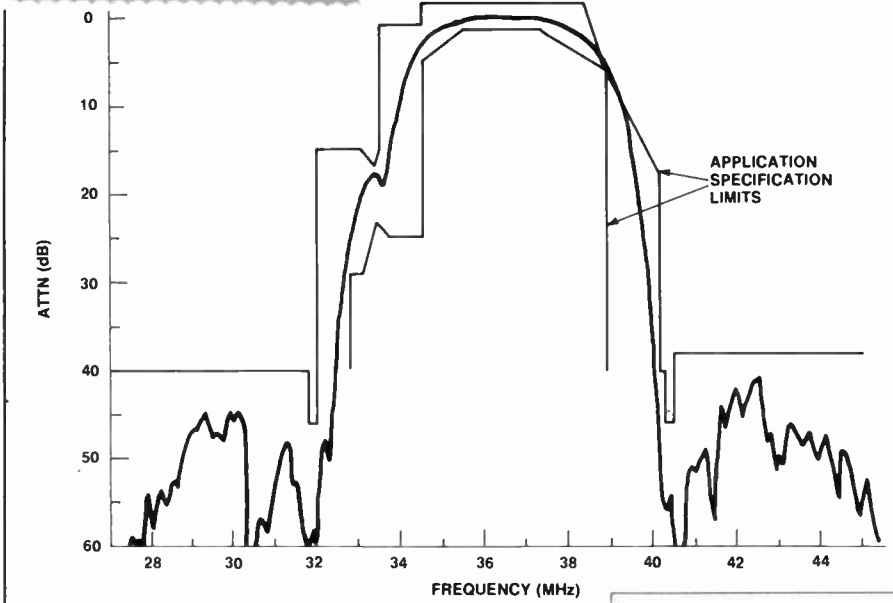

Figure 7(b). Response of TV IF filter (courtesy of Signal Technology Ltd).

The frequency response of the filter is shown in Figure 7b and is close to the ideal response for the TV application.

Most SAW devices are not produced in the same numbers as TV filters but nevertheless the price is attractive in many situations, especially when the total cost of alternative filters is considered.

In addition to the advantages of high performance, no alignment filtering, SAW devices offer exceptional device reproducibility. Figure 8 demonstrates the exceedingly small variation in response expected in production SAW filters. This advantage is present in all SAW devices and is often of great benefit in coherent (phase sensitive) communication and signal processing systems.

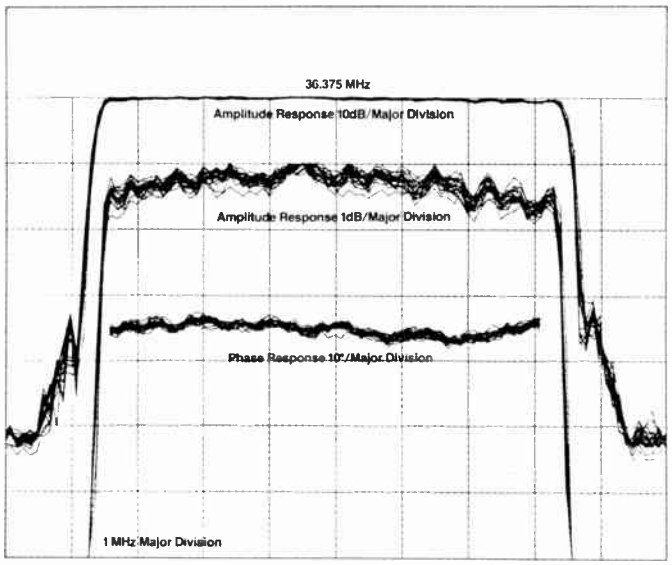

Figure 8. Overlay of the responses of 20 SAW filters (courtesy of Anderson Laboratories Inc).

Time coded filters

When examining SAW bandpass filters it is convenient to retain a conventional frequency domain specification of the filters even though the response is actually synthesised in the time domain. It is also stressed that SAW filters may have independent specifications for the phase and amplitude responses.

Bandpass filters are normally designed with a linear phase characteristic or, equivalently, with a constant group delay (i.e: all inband frequencies entering the filter emerge at the same time). A device that exhibits a flat group delay is said to be 'non-dispersive'. In fact, there are applications where dispersion (the emergence of different frequencies at different times) is required and where SAW dispersive filters are now used extensively. These devices are best characterized by their time domain rather than their frequency domain responses.

Chirp radar

One of the most common professional uses of SAW technology is in pulse compression or chirp radar systems. An impulse, which in theory contains an infinite number of frequency components, is applied to a SAW device known as a dispersive delay line or DDL (Figure 9). The DDL output is a dispersed version of the input and is a frequency modulated pulse or chirp centred at f_0 and extending over a range B in frequency and T in time.

Note that the dispersive nature of the device arises because of the non-uniform spacing of the IDT fingers rather than as a consequence of any dispersion during propagation. The chirp slope of the DDL is simply B/T and the device time-bandwidth product, N=BT, is an important parameter.

In the radar application (Figure 10) the DDL output is amplified, transmitted, reflected from the target and the returned echo applied to the receiver. The heart of the receiver is a second DDL, the dispersion characteristic of which is the exact inverse (i.e: chirp slope of −B/T) of the transmitter DDL.

The effect is to compress the returned signal back to a pulse resembling the transmitter impulse, except for 'side lobes' due to the finite bandwidth of the expansion-compression process. The time between the excitation impulse and the appearance of the compressed pulse is directly related to the range of the object in the radar beam.

This may appear to be an unnecessarily complex scheme but in practice a major advantage emerges. In all radar systems the detection sensitivity is proportional to the transmitted energy and the time resolution is proportional to the transmitted bandwidth.

In real radars the peak transmitter output cannot be increased indefinitely and once the limit is reached the only way to increase the energy output of a normal pulsed radar is to lengthen the pulse. Unfortunately, the bandwidth of the pulse is then decreased and the time, and therefore the range, resolution is reduced accordingly.

In a chirp radar the transmitted pulse bandwidth is artificially increased by the superimposed frequency modulation. A pulse long enough to ensure adequate detection sensitivity can be used while at the same time retaining sufficient bandwidth for good time resolution.

It is found that the time resolution of a chirp radar is the actual pulse length, T, divided by the time-bandwidth product, N=BT. Since N can be large (>1000 in many cases), the advantage is real. Note

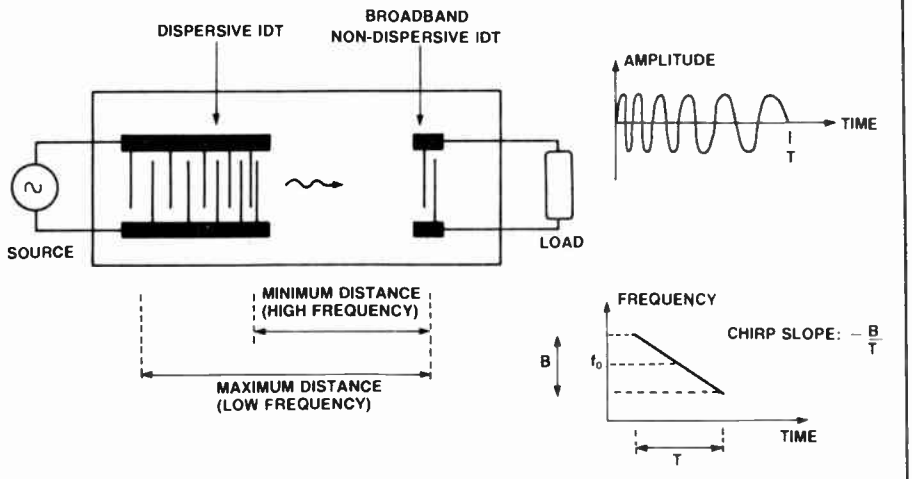

Figure 9. A dispersive delay line with two representations of its impulse response. When an impulse excites the dispersive IDT, the high frequency components arrive at the output before the low frequency components. The resultant output is a 'chirp' sweeping down in frequency. The complementary 'un-chirp' can be generated by reversing the high and low frequency ends of the dispersive IDT.

that since T/BT = 1/B, the compressed pulsewidth or time resolution is of the order of the reciprocal bandwidth.

The side lobes of the compressed pulse may be a problem in some radar applications since a weak echo can be obscured by the side-lobes of a nearby strong echo. It can be shown that if the frequency spectrum of the chirps is uniform, the first side-lobe is down only 13 dB with respect to the peak of the compressed pulse.

A partial solution is to taper the chirp spectrum of the receiver DDL impulse response, a process known as 'weighting' the chirp. The result is a broader compressed pulse exhibiting side-lobes down 40 dB or so.

SAW technology has been embraced enthusiastically by radar designers and most major SAW device manufacturers sell DDLs (weighted and unweighted) primarily intended for radar use. The ruggedness of the devices makes them ideal for use in environments such as aircraft, ships and other surroundings usually considered hostile to electronic components.

Before returning to SAW DDL's in their second major application, spectrum analysis, it is worth expanding on the concept of matched filtering, of which pulse compression is a particular example.

SAW devices as matched filters

If a signal is written as a function of time, denoted by s(t), then the impuse response of a filter 'matched' to the signal is s(−t) i.e: a time-reversed form of the signal. The output from such a filter is known as the auto-correlation function of the input signal and is the output exhibiting the best signal-to-noise ratio obtainable in the presence of white noise.

In general, a matched filter cannot be built because causal principles would need to be violated. The output would need to precede the input but, in practice, excellent approximations are possible.

Referring to the chirp radar system in Figure 9, it is clear that the receiver DDL is a matched filter for the transmitter chirp since the compressor impulse response is the same as the expander chirp reversed in time. The classic (sin x)/x compressed pulse can indeed be shown to be the auto-correlation function of the transmitted chirp.

SAW matched filters can of course be designed for any waveform and such filters will undoubtedly find application in both radio and line circuits used for the transmission of digital data. Already SAW devices are being used as code generators and correlators in modern spread spectrum communication systems. ●

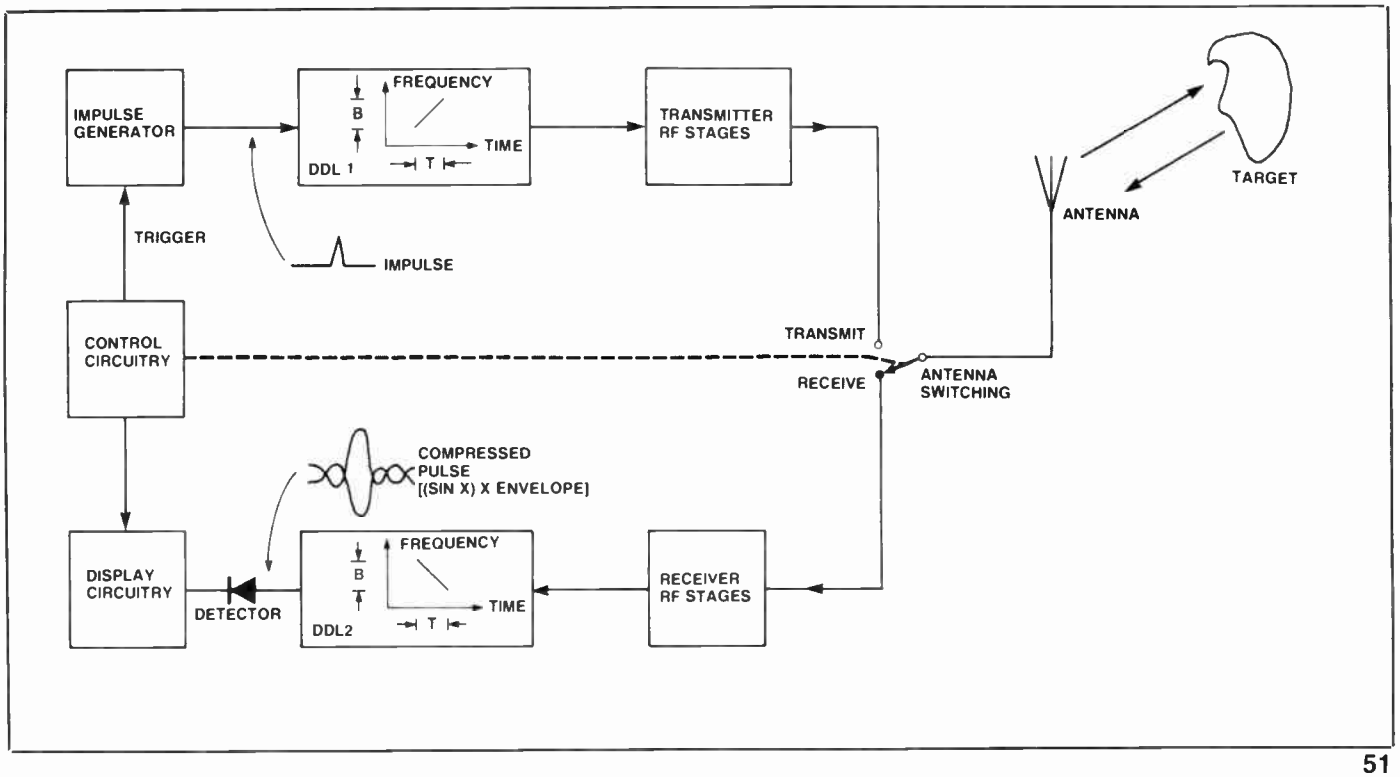

Surface acoustic wave devices — fundamentals and applications

The basic physics of SAW devices was examined in part 1. In this article a major application of time coded SAW filters, spectrum analysis, is explained.

P. J. Hall

Physics Department, University of Tasmania

Part 2

THE THIRD MAJOR use of time-coded SAW filters is in spectrum analysis and, as in chirp radars, the basic component is the dispersive delay line (Figure 10).

Before examining SAW spectrum analysis methods, it is useful to consider alternative techniques and their limitations. This section is mainly concerned with analysis systems designed for on-line signal processing applications. In such applications, spectrum analysis is used to obtain information about signals on a real-time 'as the action happens' basis, perhaps in order to control some other process e.g: the electrical steering of an antenna array to minimise interference to a radio link.

The most obvious way of analysing a given bandwidth is to assemble a bank of frequency contiguous filters covering the analysis bandwidth and to relate the output of each filter to the spectral energy falling within the passband of that filter. Filterbanks are commonly used in audio, radio astronomy and some military applications with refinements dependent upon the specific use.

Unfortunately, when a large bandwidth is to be analysed with good frequency resolution, the number of filters becomes prohibitively great and the analyser becomes large, expensive, difficult to align, and difficult to keep in alignment. Filterbanks consisting of SAW bandpass filters are now available and offer an excellent alternative to LC filters in some RF applications, although in most situations the more elegant SAW techniques to be described would normally be used.

Spectrum analysis

Conventional scanning spectrum analysers are acceptable for laboratory use but are too slow and too wasteful of signal energy to be used in most signal processing roles. A signal shorter than the time the analyser takes to scan the analysis bandwidth can come and go without ever registering on the instrument.

In military jargon, the analyser has a very low probability of intercept (POI). If the scan rate is increased to raise the POI, the available frequency resolution is decreased. The scanning analyser wastes energy because if 'N' frequency resolution cells are resolved across the analysis bandwidth, the time spent actually looking at each cell is only 1/Nth of the total observing time.

Digital spectrum analysers based on the Fast Fourier transform (FFT) algorithm are

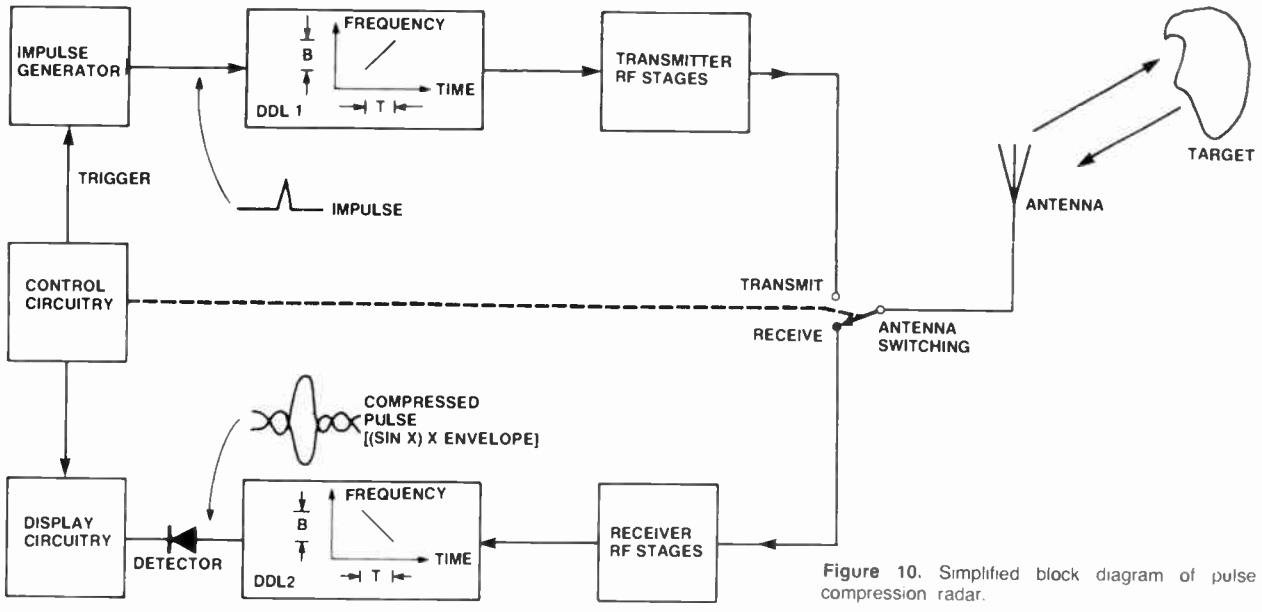

Figure 10. Simplified block diagram of pulse compression radar.

now familiar in audio test and development laboratories as well as in some signal processing applications. FFT methods produce accurate spectra and are capable of conveniently providing a true Fourier transform i.e: phase as well as amplitude, or power of spectral components. Unfortunately, the analysis bandwidth for real-time FFT processors is currently rather limited with most commercial units being essentially audio analysers.

A second digital method of power spectrum analysis, the auto-correlation technique, can be extended to wide bandwidths but is not suitable as a general approach, particularly in applications where high time resolution i.e: large number of spectra per second is needed.

Another instrument sometimes used for spectrum analysis in radio astronomy and military applications is known as the acousto-optical spectrum analyser. This device relies on the interaction of laser light with RF sound waves in a translucent piezoelectric crystal. Wide bandwidths can be analysed with good frequency and time resolution but current designs are plagued with a number of practical problems, the most serious being mechanical and thermal instability and limited dynamic range.

Coherent (phase sensitive) acousto-optical processors have been described but with existing technology, practical problems in such systems are acute. Integrated acousto-optical processors which utilise the interaction of light from a solid state laser with surface acoustic waves may substantially avoid the problems of existing optical analysers.

SAW dispersive delay line spectrum analysers overcome many of the disadvantages of the other types of instruments and offer a state-of-the-art solution to the problem of real-time, wideband Fourier analysis.

Compressive receiver

A block diagram of a SAW 'compressive receiver' spectrum analyser is shown in Figure 11. This configuration, which is only one of several possible arrangements, uses two SAW dispersive delay lines (DDLs) similar to the type used in the chirp radar system described earlier.

One DDL is used to expand an impulse into a chirp, and this chirp is then mixed with signals in the band to be analysed. Consider CW input signals for simplicity. Each signal produces a chirp from the mixer which is time co-incident with, but offset in frequency from the other chirps. This is illustrated in the frequency-time diagram for the process (Figure 12).

Signals spaced in frequency at the input produce chirps which enter the passband of the second DDL, the convolver, at different times. The chirp slope of the convolver is the exact inverse of the expander DDL chirp slope so chirps entering the convolver emerge as compressed pulses, just as in the pulse compression radar. In this case however, chirps due to different input frequencies enter the convolver at different times, so the compressed output pulses are separated in time.

Hence, separation in frequency at the compressive receiver input produces separation in time at the output, so the instrument functions as a spectrum analyser. The time axis on a display device e.g: an oscilloscope can be read as frequency by using the convolver chirp slope as a scaling factor, i.e: $f = (B/T)t$. This process is an example of the 'chirp transform' and, as mentioned earlier, other practical implementations are possible.

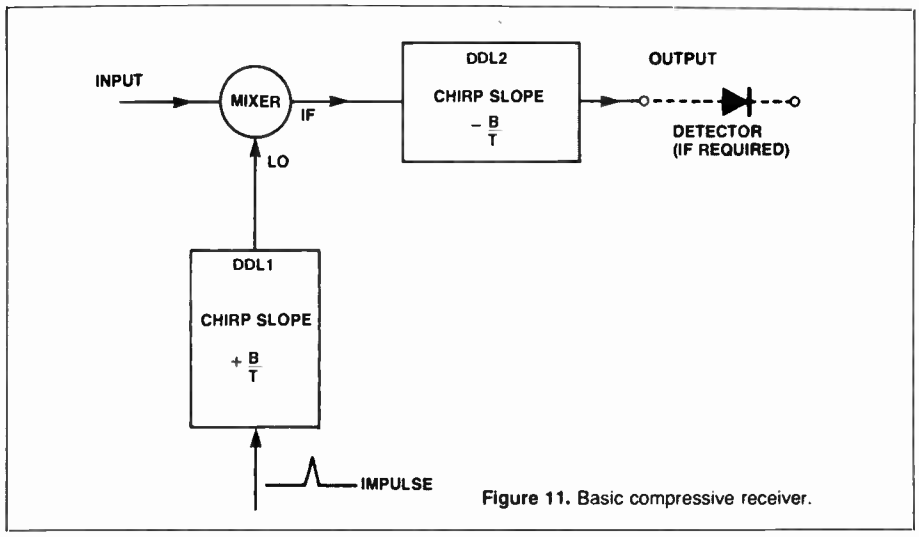

Figure 11. Basic compressive receiver.

Several key characteristics of the compressive receiver are related to the time-bandwidth product, $N = BT$, of the convolver. As expected, the width in time of the compressed pulses is of the order of the reciprocal bandwidth, $1/B$. The receiver takes T seconds to analyse the total bandwidth, B. The number of identifiable frequency resolution cells is therefore about $T/(1/B)$ or N, so the time-bandwidth product is analogous to the number of channels in a filterbank or the number of points in a digital FFT analyser. The equivalent width in frequency is simply B/N or $1/T$.

It can be shown that if a compressive receiver is designed to give a frequency resolution equivalent to a scanning spectrum analyser, the compressive receiver can analyse the input bandwidth a factor of N faster than the scanning analyser. Since N can be up to 20 000 with currently available SAW devices, the advantage can be very large. As shown in Figure 12, it is usual to use an expander chirp with twice the bandwidth and duration of the convolver impulse response.

This ensures that chirps due to signals at either edge of the analysis bandwidth emerge from the convolver fully compressed, but it also means that chirps due to some very short duration input signals can fall outside the convolver bandwidth for part of the active cycle of the receiver.

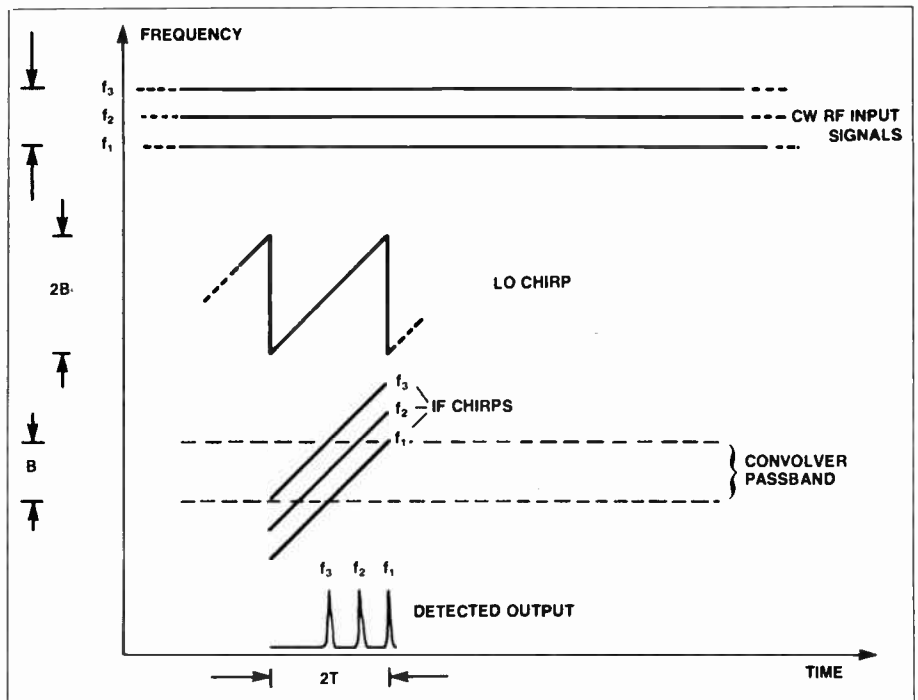

Figure 12. Frequency-Time diagram for compressive receiver. Note the duality between frequency at the input and time at the output.

The effect is to lower the POI to 50% for signals of duration less than T (typically microseconds). Nevertheless, this figure is much higher than the POI factor for almost all other analysers. A high POI for SAW compressive receivers makes the instruments popular with military users, who often use more advanced techniques to raise the POI to 100%.

SAW analysers

When compared with digital spectrum analysis techniques, the most obvious advantages of SAW analysers are high speed and wide analysis bandwidth. It has been estimated that SAW analysers exhibit a speed advantage of from 100 to 10 000 depending on the DDLs used and the type of digital instrument used as the basis for comparison. In certain applications the overall benefit of adopting SAW technology can be very great.

For example, the authors of a U.S. paper describing a sophisticated SAW signal processing module claim an advantage factor for SAW technology over digital methods of 12 000. This figure is based on a system figure of merit defined as 'the equivalent number of complex multiples per second per dollar of development cost'.

SAW processors are similar to digital instruments in that both are capable of providing amplitude and phase spectra and, as such, are useful in roles requiring coherent signal processing. Part of the recent work in radio astronomy at the University of Tasmania has been to demonstrate that coherent SAW spectrum analysers are viable now, using existing technology and techniques.

In contrast, coherent optical methods (the main rival to SAW technology for wideband, real-time applications) are some way from being fully developed. As a bonus, SAW technology offers a dramatic improvement over optical methods in the areas of mechanical and thermal stability, and dynamic range.

The limitations associated with SAW spectrum analysers are fairly minor. The analog transformation process exhibits an accuracy equivalent to a 6 or 7 bit digital transform. For many applications this is sufficient, particularly when the speed advantage of a SAW instrument is considered.

The second practical problem is that SAW analysers are simply too fast to be directly useful in applications where the instruments need to be interfaced to computers or other digital post-transformation devices. For example, two University of Tasmania spectrum analysers output data at a rate of around 40 ns per frequency resolution cell, and this is the rate at which a digital processor needs to sample the detected output.

We have solved the problem by producing 'digital video integrators', or DVIs, to act as an interface between the SAW analysers and slower logic. A DVI samples the analyser output at high speed and adds together many spectra before the integrated spectrum is output to a computer or other device. The DVI runs continously so there is no loss of new data while the integrated spectrum is output.

This process degrades the time resolution slightly but, even so, the integrated spectrum time resolution can approach 100 μs with existing data acquisition computers. Of course, the wide-band, coherent processing capability of SAW spectrum analysers is retained.

The DVI techniques developed in Tasmania have created interest overseas, and it is hoped that with the local VLSI program underway, further work can be undertaken on the design of an Australian produced parallel processor ship for use in the back-end of a DVI.

Conclusion

These articles have summarised some important current uses of SAW technology but the treatment has been necessarily brief. Discussion of important devices such as reflective array compressor (RAC) dispersive delay lines, non-linear acoustic convolvers, SAW oscillators and SAW memories has been omitted.

Interested readers are referred to the general references listed below. For the benefit of designers considering using SAW devices I have provided a list of suppliers known to me. Most companies can supply customised as well as standard stock devices.

GENERAL REFERENCES

1. Butler, M.B.N., 'Surface acoustic wave devices', *Electronic Engineering*, June 1980, July 1980, February 1981.
2. Dieulesaint, E. and Royer, D., '*Elastic waves in solids*', Wiley, 1980.
3. Hartmann, C.IS. et al., 'SAW devices for military, communications, radar and EW systems', *Microwave Journal*, July 1982.
4. Oliner, A.A. (Ed.), '*Acoustic surface waves*', Springer-Verlag, 1978.
5. Penimuri, D. and Havens, D.P., 'Surface acoustic wave filters prove versatile in VHF applications', *Electronics*, July 5, 1979.

SAW DEVICE SUPPLIERS

1. Andersen Laboratories Inc., 1280 Blue Hills Avenue, CT 06002, USA. TWX 710-425-2390.
2. Crystal Technology Inc., 1035M E. Meadow Cir., Palo Alto, CA 94303, USA.
3. GEC — Marconi Research Centre, Great Baddow, Essex, CM2 8HN, U.D. Telex 99201.
4. Hughes Aircraft Company, Gound Systems Group, P.O. Box 3310, Fullerton, CA 92634, USA.
5. Racal-MESL Ltd., Lochend Industrial Estate, Newbridge, Midlothian, EH28 8LP, U.K. Telex 72384.
6. RF Monolithics Inc., 4441 Sigma Road, Dallas, Texas 75234, USA.
7. Siemens Ltd., Australian Office, 544 Church Street, Richmond, Victoria, 3121. Telex AA30425.
8. Signal Technology Ltd., Cheney Manor, Swindon, Wiltshire, SN2 2PJ, U.K. Telex 444410.
9. Thomson-CSF, Chemin des Travails, B.P. 53, 06802 Cagnes-Sur-Mer, France. Telex 204.780F. ●

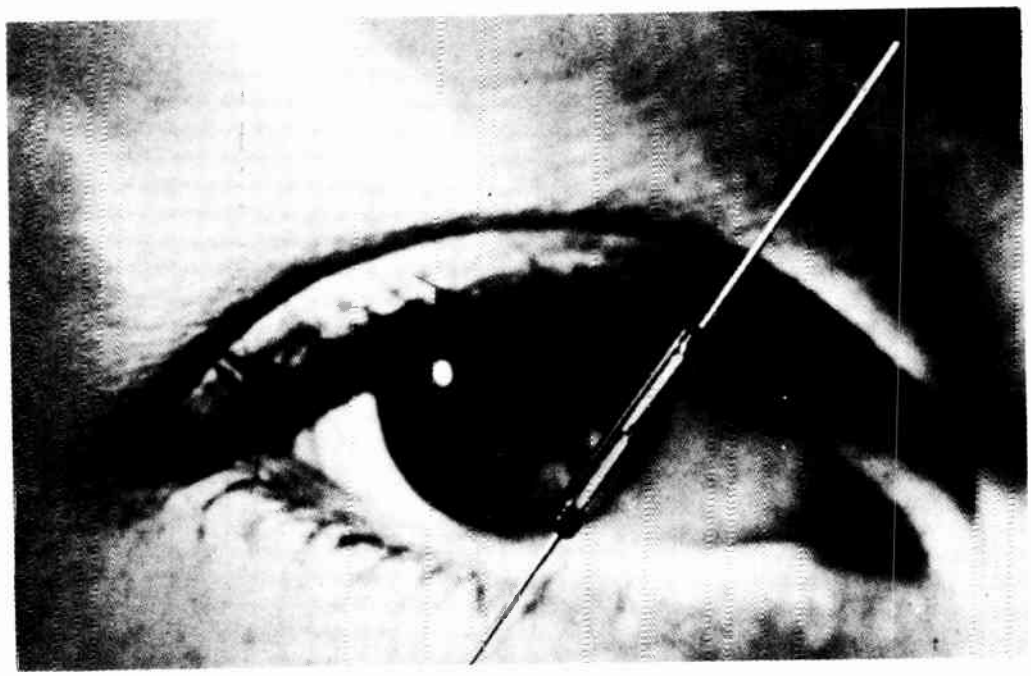

Practical guide to reed switches

THE DRY REED is an almost perfect low-current switch.

It is fast – operating times of less than one millisecond are typical. It is reliable – as many as one billion operations can be achieved. And it is cheap – quantity price is well under 50 cents.

The dry reed switch is not by any means a new device for it was invented back in 1945 by Dr. W. B. Ellwood of the USA's Western Electric Corporation.

But it was ahead of its time. It remained practically unnoticed by the engineering world until only a few years ago when it was 'rediscovered' by the telephone industry.

And since then reed switches are receiving interest and acceptance at an ever increasing rate.

In its basic form, a reed switch is a magneto-mechanical relay. In other words it relies upon a magnetic force to initiate a mechanical switching action.

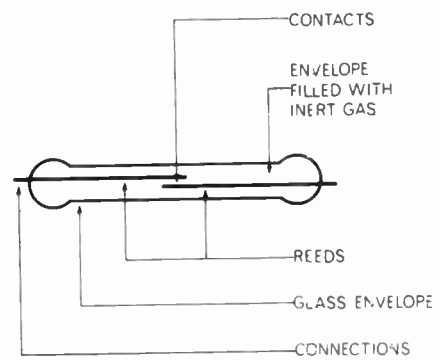

Fig. 1. The basic reed switch.

THE BASIC SWITCH

A typical reed switch is shown in Fig.1. It consists of two flattened ferromagnetic reeds sealed in a glass tube. The reeds are fixed, one at each end of the tube, so that their free ends overlap in the centre but with a 0.01" gap between them.

During the sealing operation the air inside the tube is pumped out and replaced by dry nitrogen so that the contacts operate in an inert atmosphere.

When the reed switch is brought within the influence of a magnetic field (either from a coil or a magnet) the reeds – being ferromagnetic – become a flux-carrying portion of the magnetic circuit. The extreme ends of the reeds will assume opposite magnetic polarity, and if sufficient flux is present, the attraction forces overcome the stiffness of the reeds and they flex towards each other and touch.

When the magnetic field is removed the reeds spring back to their original positions. There is however a difference between the value of field required to close the reeds, and the reduced value that will allow them to open again.

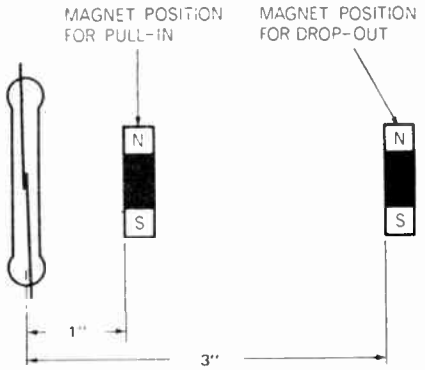

Fig. 2. The reeds close when the magnet is brought within one inch, and will remain closed until the magnet has been moved at least three inches away.

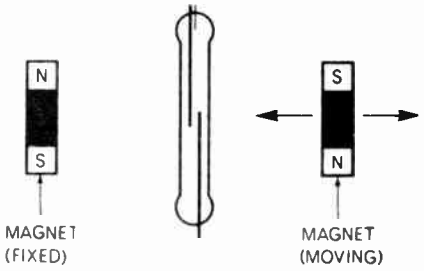

Fig. 3. A fixed magnet of opposite polerity to the moving magnet may be used to reduce pull-in, pull-out differential.

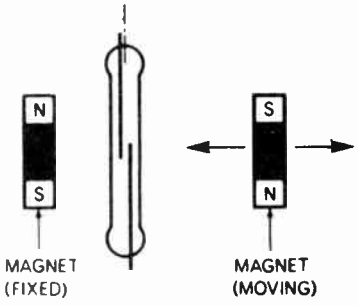

Fig. 4. 'Normally closed' operation can be obtained by biasing a 'normally open' reed switch with a fixed magnet. The moving magnet cancels out the fixed magnet and thus allows the switch to open.

Fig. 5. This type of reed switch may be used for either change-over, or normally closed operation.

A typical example of this is shown in Fig.2. In this example the reeds close when the magnet is brought within one inch, but they will remain closed until the magnet has been moved about three inches away.

This phenomena – which is caused by magnetic hysteresis in the reeds – can be considerably reduced by introducing a second magnet, of opposite polarity, on the further side of the switch. This is illustrated in Fig.3. The fixed magnet must not be mounted within the normal pull-in position for single magnet operation, otherwise the reed switch will be held in a closed position by the second magnet and will open when the moving magnet is brought close to the switch. By selecting the correct types and strengths of magnets the differential can be set to practically any required value.

OPERATING MODES

As can clearly be seen in Fig.1, the reed switch is 'normally open'. The reeds close when a magnet is brought close to the switch enclosure.

However there are many applications where the switch is required to be 'normally closed' and to open when the magnet is introduced. This can be done either by biasing the switch with a second magnet (as shown in Fig.4), or by using a reed switch with change-over contacts (Fig.5).

In most applications where a reed switch is opened or closed by a permanent magnet, the magnet is fitted to a moving part, and the reed is fitted to a stationary part.

There are, however, a number of applications in which both the magnet and the reed must be located on a stationary component. Operation may then be effected through distortion of the magnetic field by an external moving ferrous mass. If the magnet and the reed are sufficiently close, the reeds switch will be normally closed, but will be opened by the magnetic shunting effect of the external ferrous object. Alternatively, the magnet may be located so that the reeds are normally open and the external ferrous object used to 'reinforce' the field and thus close the reeds.

There are many different ways in which a moving magnet may be caused to operate a reed switch.

Linear planes of operation are shown in Fig. 6.; movement of the magnet in any of the planes a-a, b-b, and c-c will operate the switch. Magnet selection is fairly critical if the switch is operated in mode b-b, spurious operation may be caused by negative peaks on the magnet's field pattern curve. If these are large, the reeds will pull-in three times as the magnet is moved from one end of the switch to the other.

Rotary motion may also be used. Various ways of achieving this are shown in Fig.7. (A most versatile and simple impulse generator can be put together in a few minutes by placing one or more magnets on a gramophone

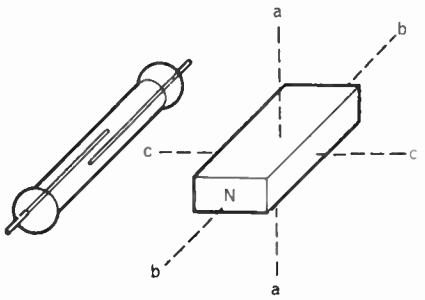

Fig. 6. Linear planes of operation; movement of the magnet in any of the planes indicated may be used to actuate the switch.

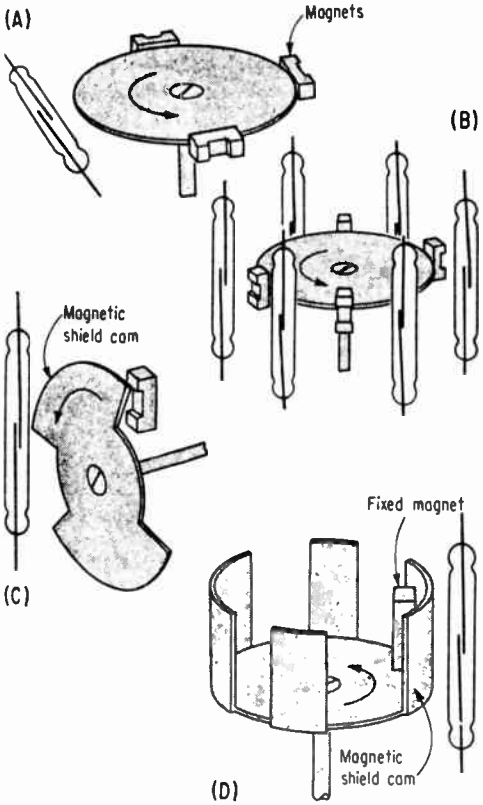

Fig. 7. Rotary motion may also be used to actuate a reed switch. In A and B the switches are stationary and the magnets rotate. In examples C and D both the switches and the magnets ere stationary and the switch operates whenever the cutout portion of the magnetic shield is between magnet and switch.

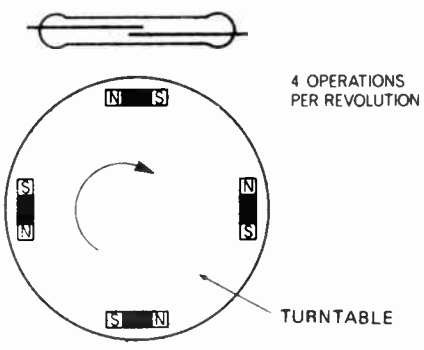

Fig. 8. Simple yet versatile impulse timer can be improvised by placing one or more magnets on a gramaphone turntable.

turntable and fastening a reed switch to the motor base board. (Fig.8). Switching rates from approx one every two seconds to well over 2000 a minute can be selected merely by changing the turntable speed and/or using more magnets!)

Since the reed switch is truly a sealed device, it can be used in applications where conventional switches are not permitted, or where they have very limited life. Reed switches are frequently used in simple on/off push buttons, and outdoors, in dusty areas such as cement plants, especially in areas where explosive gases may be present.

OPERATING LIFE

The operating life and load carrying characteristics of reed switches are interrelated. A switch may operate for 100 million or even 1000 million closures providing it is switching very low currents. But the same type of switch may fail after half a dozen switching cycles if the load greatly exceeds the designed rating. The majority of reed switches are manufactured with contact ratings between 0.1A and 3.0A.

The current handling capacity of reed switches varies from type to type. In general the rating will be determined by the size and surface plating of the reeds, for the reed is an electrical conductor, and current rating will be a function of contact area.

The maximum rated contact loading is only applicable for purely resistive loads. If the load is capacitive or inductive the switch must either be drastically derated, or the switch contacts protected in a suitable fashion.

Four suitable methods of contact protection are shown in Fig.9.

In dc circuits all that may be required is a resistor shunted across the load (Fig.9A). Where the load is a relay coil or operating solenoid a resistor of approximately eight times the coil resistance is adequate to absorb a major portion of the induced energy when the circuit is interrupted. The addition of the resistor will of course increase the steady-state current flow but this extra load is negligible.

Another cheap and simple way to protect the reed switch is to wire a capacitor across the contacts. The required value depends upon load current, but something between 0.1 uf and 1.0 uf will be sufficient. (Fig.9B).

The most generally used method of protection is the resistor-capacitor series network shown in Fig. 9C. This circuit must be used if the switched load current is ac. The resistor should be approximately 160 ohms and the capacitor somewhere between 0.1 uf and 1.0 uf. That this is an extremely effective method was proven by a recent trial during which a motor starter was switched 50 million times without failure.

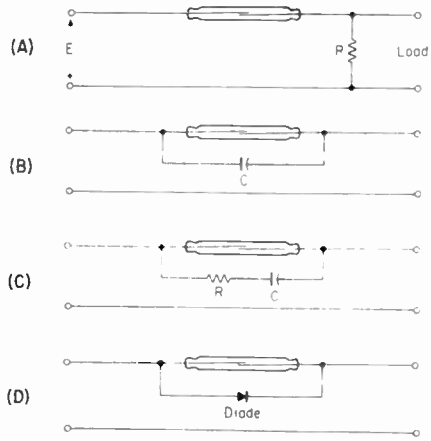

Fig. 9. Contact protection techniques: A – Resistor shunting load. B – Capacitor shunting contact. C – Resistor-capacitor series network for ac loads. D – Diode shunting.

The component values may either be determined empirically (as described below) or mathematically. In the latter case, the component values can be obtained from –

$$C = \frac{I^2}{10}\ \mu F \quad , \quad R = \frac{E}{10 \times I\left(1+\frac{50}{E}\right)}\ \Omega$$

Where I is the closed circuit current in amps and E is the open circuit voltage in volts.

A fourth method of protection is to connect a diode across the switch contacts. (Fig.9D). This method is effective only with dc; diode polarity must of course be preserved.

Suitable protection circuits are often best determined empirically. One way is to connect the switch to the normal operating voltage and load, and then to actually observe the arcing across the reeds whilst the switch is in use.

HEAVY CURRENT SWITCHING

There will be many applications in which a reed switch can usefully be used to switch very large currents. This

Switching in explosive atmospheres, obviating ignition risk; in dust filled atmospheres where conventional contacts would be unreliable; and in extremely cold conditions where ordinary switches would freeze up. In radioactive environments, magnetic operation can maintain integrity of shielding.

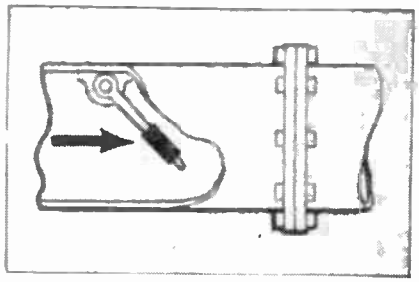

Flow control and indication, minimising restraint on the moving part and avoiding perforation of the container wall.

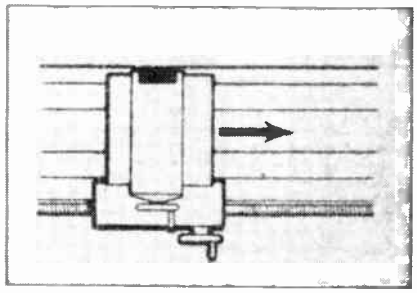

Position control and indication, obviating mechanical contact with its implications of wear, and simplifying mounting.

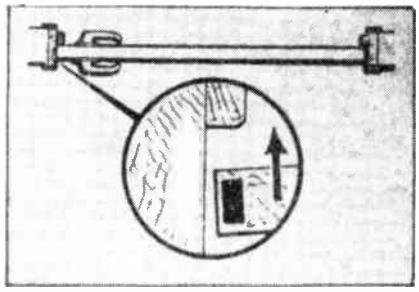

Door switches, obviating mounting and adjustment problems, and offering total concealment for security devices.

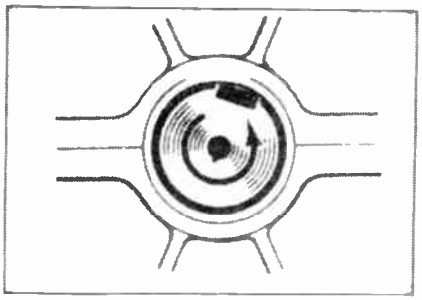

Tachometer applications, requiring the simplest addition to the moving part and offering ability to work in unfavourable conditions, plus high speed operation.

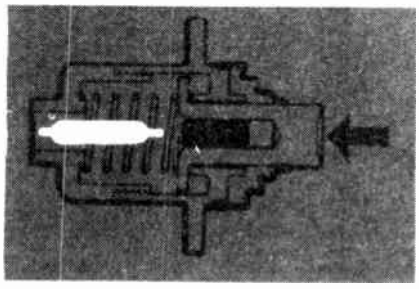

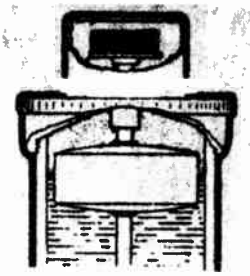

Hydraulic brake fluid level indicator, where feasibility depends on simplicity and ease of application.

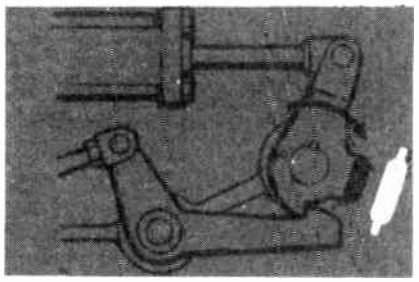

Safety interlock switching, giving extreme reliability and simplicity of application to complex mechanical layouts. Reed insert completes circuit to illuminate warning lamp or permit further stage of operation.

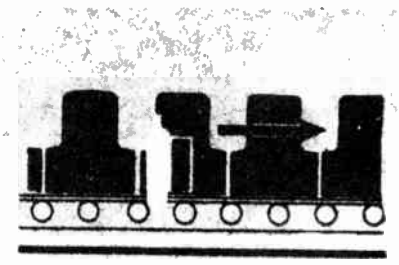

Proximity counting, providing a very easy method of recording the passage of ferrous items past a point.

can be done quite simply by combining a reed switch with a Triac. (Fig.10). Even miniature reed switches will safely carry the gate current required to trigger the largest Triacs, and by using this system it is possible to switch single phase loads of whatever Triac rating is used. Triacs can be readily obtained with ratings from 1 amp to 125 amps.

Three phase loads can also be switched by using the reed switch to energize a miniature three pole relay that in turn triggers a Triac in each of the phases of the supply.

SWITCHING AT LOW LEVELS

One great advantage of the reed switch is its ability to operate reliably when switching currents and voltages at very low levels. This is a major problem with standard switches because there is insufficient energy to break down non-conducting films on the switch contacts. But a reed switch – due largely to its gold-plated contact surfaces and inert atmosphere – will perform satisfactorily for at least a billion operations.

Some idea of the extraordinary reliability of reed switches was shown during a series of tests undertaken by the Bell Telephone Company in the USA. In one test four switches were operated at 120 closures a second carrying a load of 500 micro-volts, 100 microamps, dc. Each switch completed 50 million consecutive closures without a single instance of closed resistance exceeding 5 ohms.

FAILURES

A reed switch rarely fails completely. As load currents are increased the contacts suffer the same form of contact erosion experienced in conventional switches. The resultant particles are magnetic and collect in the air-gap. If these fragments become numerous enough they intermittently bridge the gap and cause a failure-to-open. It is also possible for these fragments to alter the closed contact resistance.

The most common cause of contact failure is the mechanical locking of a spike on one reed and a corresponding crater on the other. This type of failure is commonly called a 'weld' but it is not a weld in the true sense. The contacts are not joined by molten metal but are held by friction or interlocking, between the spike and the crater.

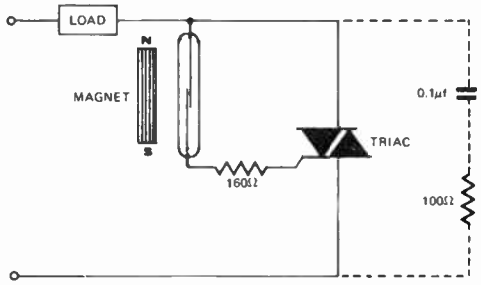

Fig. 10. Reed switch/Triac combination may be used to switch single phase loads as high as 125 Amps. Components shown in dotted lines must be included if the load is reactive.

SPECIFICATIONS	STANDARD	MINIATURE
Maximum voltage	150 Vdc 250 Vac	50 Vdc 150 Vac
Maximum current	2.0A	0.5A
Maximum power	25W	6W
Max. initial resistance	50 m.ohms	100 m.ohms
Max. end-of-life resistance	2 ohms	2 ohms
Peak breakdown voltage	500 V	300 V
Closure rate	400 Hz	2000 Hz
Insulation resistance	5000 M.ohms	1000 M.ohms
Temperature range	-55°C to +150°C	-55°C to +150°C
Contact capacitance	1.5 pF	0.5 pF
Vibration	10G at 10-55Hz	10G at 10-55 Hz
Shock	15G minimum	15G minimum
Life at rated load	5×10^6 operations	5×10^6 operations
Life at zero load	500×10^6 operations	500×10^6 operations

Table 1. Typical specifications for standard and miniature reed switches.

Reed switches are actuated by a magnetic field.

This field can be generated by a permanent magnet, or by an electrically energised coil. When coils are used, the reed switch is simply inserted within the coil former and it is then closed (or opened) when current is passed through the coil.

It operates, in fact, as a relay, and in this form reed switches are used by the million, in telephone systems around the world.

When a reed switch is to be electrically actuated, an indication of the magnetic field strength that is required is generally quoted by the manufacturer in terms of so many ampere-turns. This figure may range from 50 AT to 250 AT (but as explained later, this may be substantially reduced by the judicious positioning of a bias magnet).

Various combinations of turns, wire sizes and dimensions may be used to close any specific type of switch, and these parameters will in turn be determined not only by the required number of ampere-turns, but also by the circuit voltage and current that is available. For example a switch that requires 100 ampere-turns may be actuated by a 220 ohm winding drawing 13 mA at 3.0 V., or by a 25,000 ohm winding drawing 1.2 mA at 31 V.

Table 2 provides all the data required to design operating coils for a wide variety of standard sized reed switches, (i.e., 2.75" overall, 2.0" long, 0.217" diameter).

The operating coil may either be wound on a bobbin manufactured specifically for the purpose (Fig. 11) or made up from a length of paper, aluminium or plastic tubing that is a neat fit over the outside diameter of the glass reed.

Another method of making operating coils is to wind them, using a cement coated wire, onto an arbour that is shaped to create the desired final form. After removal from the arbour, the winding should be protected by a layer of insulating tape.

EXTERNAL MAGNETIC FIELDS

A reed switch is influenced by a magnetic field regardless of whether that field is produced by the operating coil or by some other magneto-motive force. The magnetic force generated by the field winding can be modified or even completely cancelled by the field from a nearby permanent magnet, or by the alternating flux from a nearby choke, transformer or other inductive device. Even the proximity of a sheet steel chassis may affect the energy at which a reed switch will just actuate.

But the effect of external magnetic influences may be usefully exploited to modify the characteristics of the basic reed switch assembly.

For example the ampere-turns required to close any given switch, can be halved by placing a magnet a short distance away from the coil – the magnet's polarity must be the same as that of the operating coil. The positioning of the magnet is fairly critical and is best determined by trial and error, (Fig. 12).

A similar method can be used to obtain a latching action. In this case the method exploits the magnetic hysterisis of the reed switch. The magnet is placed far enough away from the coil so that it does not close the reeds magnetically, but sufficiently close so as to hold the reeds closed once they have been actuated by an electrical signal through the coil.

In this example the reed relay can be unlatched only by physically removing the magnet, or by applying opposite polarity drive through the operating coil.

A further modification of the magnetic latching principle is shown in Fig. 13. Here, whilst magnetic latching is still used, the operating coil has two windings, one of which is used to actuate the relay, and the other, which is connected in opposite polarity, is used to unlatch the relay.

A magnet may also be used to convert a normally open reed relay to normally closed operation. This is done by locating the magnet sufficiently close to the reed so that the contacts are held closed. (Fig. 14). The coil is wound so as to produce a magnetic flux of opposite polarity to the magnet. When the operating coil is energised, the resultant magnetic flux will cancel out that from the permanent magnet, and the reed will open.

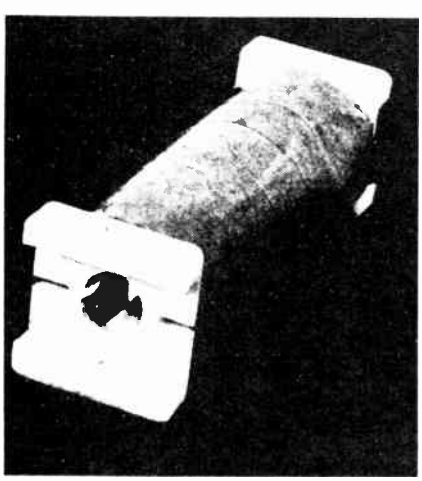

Fig. 11. Bobbin for standard sized reed switch.

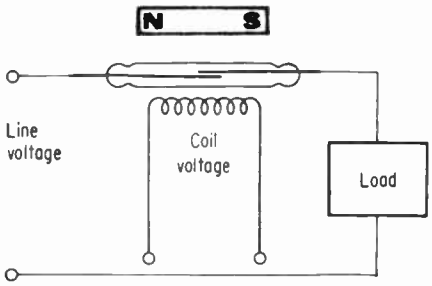

Fig. 12. Magnet assists reed switch to close, and thus reduces coil energy requirement.

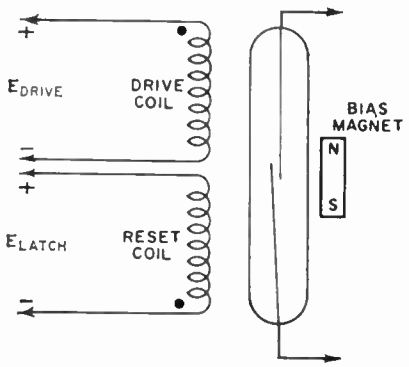

Fig. 13. In this circuit, reed switch is latched by aiding permanent magnet, and reset by magnetic opposition from reset coil field.

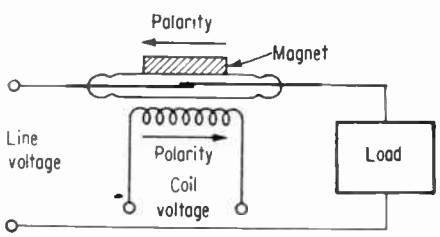

Fig. 14. Normally closed operation using magnetic bias.

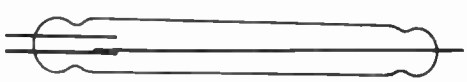

Fig. 15. This type of reed switch may be used for either change-over or normally closed operation.

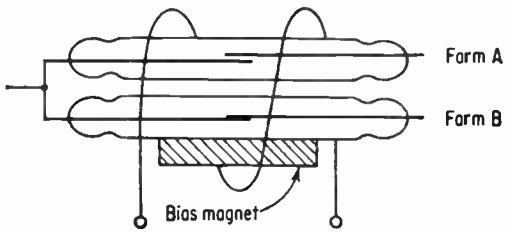

Fig. 16. Change-over action may also be obtained by combining a magnet and two normally open switches actuated by a common operating coil.

Turns		7,520		9,600		11,900		15,000		19,500		25,000		36,000		43,100		53,000		66,500		86,800	
Ohms		219		355		550		855		1,390		2,380		4,320		6,380		9,900		16,000		25,000	
Switch Sensitivity	50AT	1.5	6.5	1.8	5	2.3	4.2	2.8	3.3	3.5	2.5	4.8	2	6.1	1.4	7.3	1.2	9.4	0.8	12	0.75	15.7	0.6
	100AT	3.0	13	3.6	10	4.6	8.4	5.7	6.7	7.1	5.1	9.5	4	13	2.8	14.7	2.3	19	1.9	24	1.5	31	1.2
	150AT	4.5	20	5.3	15	6.9	12.6	8.5	10	11	7.6	14	6	18	4.2	22	3.5	28	2.7	36	2.3	47	1.8
	200AT	6.0	26	7.0	20	9.2	17	11	13	14	10.2	19	8	24	5.6	29	4.6	38	3.8	48	3.0	94	2.4
	250AT	7.5	33	8.8	25	11.5	21	15	16	18	12.7	24	10	30	7.0	37	5.8	47	4.6	60	3.8	110	3.0
		v	ma	v	ma	v	ma	v	ma	v	ma	v	ma	v	ma	v	ma	v	ma	v	ma	v	ma

Table 2. Data for standard size reed switch operating coils – bobbins to be 2" long x 0.220" inside diameter, winding build up will be approx. 0.2".

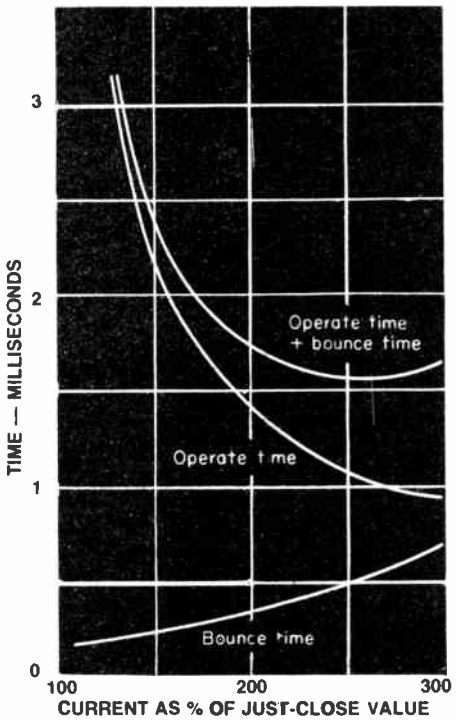

Fig. 17. Variation of operate and bounce time with energizing current.

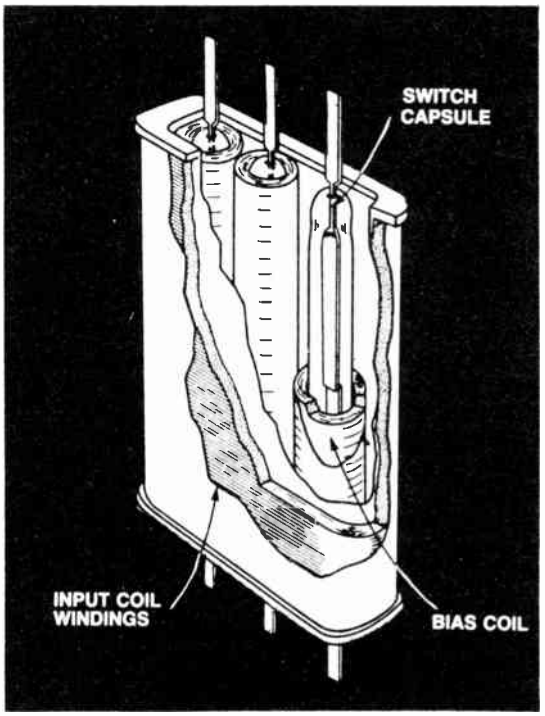

Fig. 18. Reed switch logic module.

Change-over action may be obtained either by using a reed switch specially made for the purpose (Fig. 15) or by using a magnet and two normally open reed switches actuated by a common operating coil (Fig. 16).

It is possible to actuate a number of separate reed switches located inside one large operating coil, but due to variations in the sensitivity between one reed and another, and the positioning of individual reed switches within the operating coil, it is not possible to predict the contact action sequence. All other things being equal the most sensitive reed will operate first. This will then act as a magnetic shunt, retarding the operation of the remaining reed switches. This is a major difference from conventional electro-mechanical relays where a single armature or card drives all of the movable contacts and any pair of contacts can be adjusted to ensure synchronous operation or a specific contacting sequence.

Nevertheless if a current at least 150% of the just-operate ampere-turns of the highest rated switch of the group is applied suddenly, this effect is less noticeable, and in most applications may be virtually neglected.

SWITCHING CHARACTERISTICS

The moving blades inside a reed switch have very low mass and move only a few thousands of an inch. The operating coil is not iron cored and so has little self-inductance thus allowing a magnetic field to build up very rapidly. These factors combine to ensure that a reed relay is an inherently quick-acting device, in fact, operating times of less than one millisecond are quite typical.

The speed at which any specific reed relay closes is primarily a function of the number of ampere-turns in the operating coil. But when the contacts close they normally bounce two or three times and the harder the relay is driven (i.e., the greater the number of ampere-turns) the greater the number of times that the contacts bounce. In general reed relay coils are designed so that nominal rated voltage produces approximately 50% more ampere-turns than the just-operate value. This gives optimum total operate time including the contact bounce time. (Fig. 17).

After the contacts have closed and have stopped bouncing, the reeds continue to vibrate for a short time. This vibration produces magnetostriction contact noise – a damped oscillatory voltage that decays to zero – and this may cause problems in low signal level circuits.

With no suppression devices across the operating coil, reed release time is very fast – it may be as short as 25 micro-seconds. Adding a suppression diode has little, if any effect on the operate or bounce times, but it does significantly lengthen the release time. For example, the release time of a standard type of reed without a suppression diode may be 50 micro-seconds, but with a diode the release time may be extended to a milli-second or so.

Due to the geometry of the reed switch construction, the capacitance between contacts is low, and with standard sized reed relays this will be about one pico-farad. The capacitance between the reeds and the operating coil will be about 2.5 pico-farads but this can be reduced to approximately 0.5 pico-farads by interposing a grounded electrostatic shield between the coil and the reed.

Some thermal EMF will be generated at the junction of dissimilar metals in reed switches due to the heating produced by energizing the coil. This thermally-generated EMF may be undesirable if the reeds are used to switch low level analogue signals – as for example in data-logging.

For applications where the thermal EMF must be held to the minimum a bi-stable latching reed relay should be used. A short pulse to the set coil operates the relay, no heat generating holding current is then required. Another short pulse to the reset coil releases the relay. Using this type of operation, the latching relay thermal EMF remains below five micro-volts, compared to as much as 100 micro-volts for continuously energized relays.

REED SWITCHES AS LOGIC ELEMENTS

The reed relay is almost an ideal buffer between solid state devices and higher power output elements. The winding impedance and current levels is well suited for the collector or emitter circuits of standard transistors, and it can also be driven by many IC elements.

With the recent introduction of the 'pico-reed' switch, a single pole relay is now available in a dual-in-line package, making it both physically and electrically compatible with integrated circuit components. And as reed relays become smaller their operating speeds increase, so that a pico-reed relay can be made to follow 1 kHz pulses.

Reed switches are finding increasing use as logic elements for use in adverse environments. They are capable of performing a large variety of logic functions including AND, OR, EXCLUSIVE OR, and NOT operations. The relays can be used to construct flip-flop circuits, and these can be used in binary, binary coded decimal and decimal counters, ring counters, up-down counters and shift registers.

Whilst operating speed is very considerably slower than with solid state logic elements, there is not the same necessity for precise voltage and frequency regulation, nor the susceptibility to voltage transients. And for these reasons reed relay logic circuits are becoming increasingly used in industrial equipment.

A wide range of reed relay logic elements are commercially available – generally in a configuration similar to that shown in Fig. 18. The individual reed switches are surrounded by bias coils, and these in turn share a common input winding coil. By adjusting the ampere turns level to both input winding and the individual bias windings, a multiplicity of functions can be obtained. A permanent magnet can also be used in this type of logic element to provide memory or latching functions.

When designing reed relays as logic elements, the amount of magnetic flux that is required to close the relay is regarded as one flux unit. Thus a two input AND gate consists of one reed switch surrounded by two windings each of which can generate one half a 'flux unit'. (Fig. 19).

It is possible to expand the concept to produce three, four or five input AND gates by providing a separate winding for each input, such that each winding provides $\frac{1}{n}$th of the total required flux.

An OR gate is produced by providing two windings either of which can be energised to the level of one flux unit, (Fig. 20). Thus a voltage in either winding can cause the relay to close. Again, as with the AND gate, a number of windings may be used provided each one can provide one full flux unit.

The basic OR gate can be used as an exclusive OR gate simply by reversing the direction of one winding. In this application if either one or the other winding is energised then the relay will close, but if both are energised then the resultant magnetic fields will cancel out and the relay will remain open.

Inverted operation is provided by using relays that are magnetically biased into normally closed operation.

Cross-bar matrix switching is readily achieved by using a double wound relay, in which each winding provides half a flux unit, at each selection point, i.e., A1, A2, A3, etc., (Fig. 21).

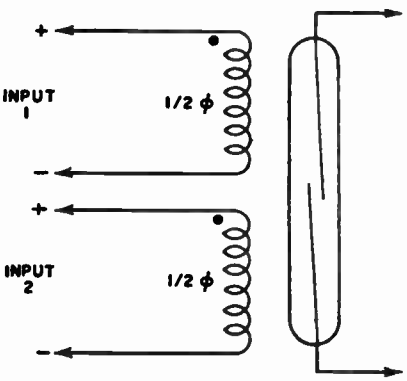

Fig. 19. The reed relay as an AND gate. Each coil can produce ½ a 'flux unit', and an output is obtained only when voltages are applied to inputs.

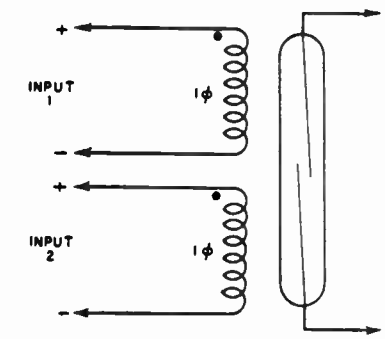

Fig. 20. Reed relay OR gate. Either coil can close reed switch.

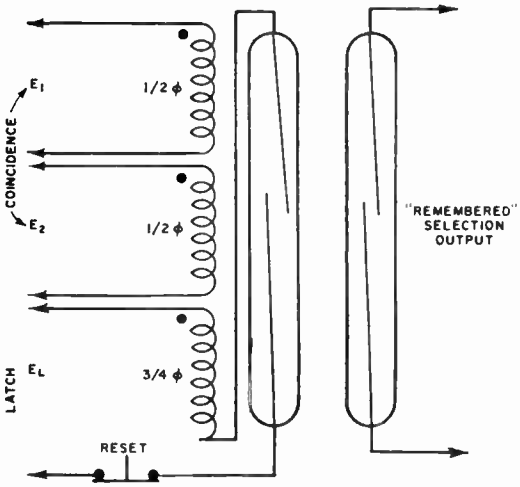

Fig. 22. How a 'memory' may be incorporated in a matrix element.

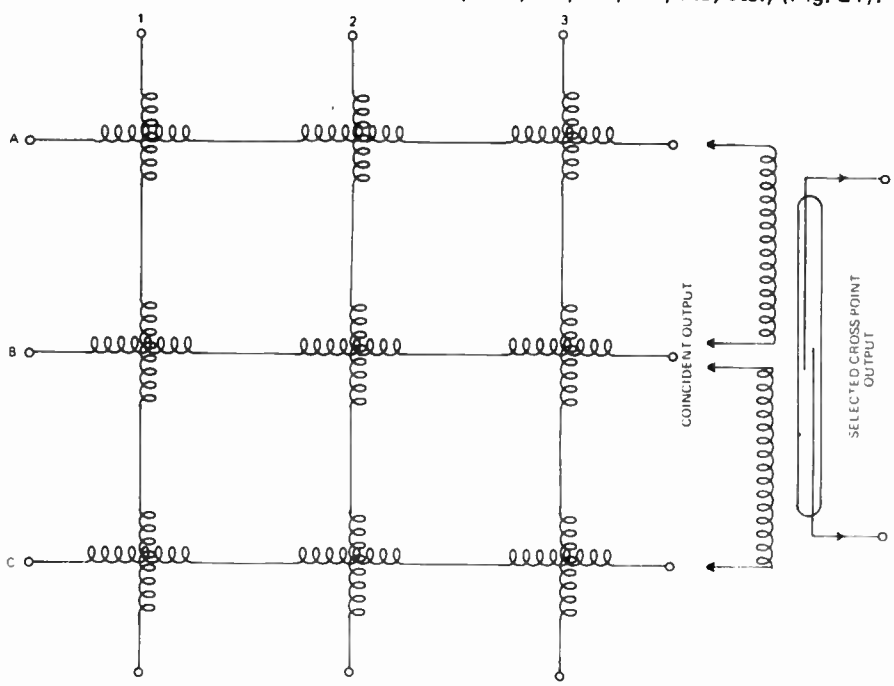

Fig. 21. Cross-bar matrix switching, the appropriate relay will close whenever both coils of any given relay are coincidentally energised.

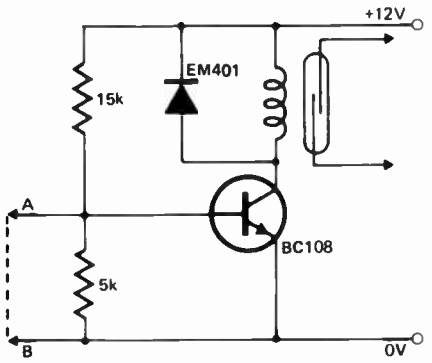

FIG. 23

The appropriate relay will close whenever both coils of any given relay are coincidentally energized.

The cross bar switching system may be used with a magnetic or electrical memory if required. Fig. 22 shows how two reed switches can be used, together with a latching winding, to provide an electrical memory in a reed relay cross bar switching system. In this form the matrix will remember the inputs after they have been removed, until the latching power supply is interrupted.

Reed switches can be combined with solid-state electronic components to provide extremely reliable and maintenance free circuitry.

The low operating current of the actuating coil is well within the collector current rating of practically any transistor (and most linear integrated circuits). Many simple practical circuits can be constructed using a single transistor and a reed switch.

The circuit shown in Fig. 23 is commonly used to open or close a relay when an external circuit is made or broken. It is commonly used in simple burglar alarm installations.

In operation, the transistor is cut off by a short circuit across points 'A' and 'B' (shown as dotted lines). Because the transistor is cut off, the reed relay operating coil in the transistor's collector circuit is not energized, and the relay contacts are open.

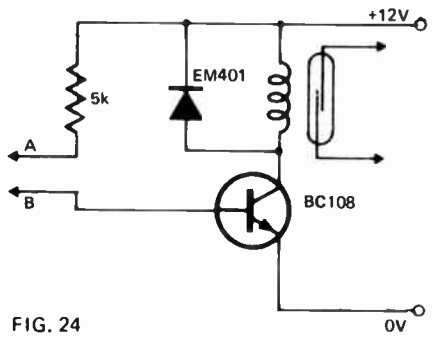

FIG. 24

If the short circuit is removed from points 'A' and 'B', the transistor is biased on via the 15k resistor, the relay coil is energized and the reed switch is closed. Current consumption of this circuit – whilst the relay is de-energized – is less than one milliamp.

The circuit shown in Fig. 24 has a similar function to that of Fig. 23, except that the relay will close when a short is placed across points 'A' and 'B'.

It is often necessary to arrange for the relay to remain closed even though the actuating signal is only momentary. This can be done by using an actuating coil containing two reed switches, and using one of the reed switches to short out the transistor the moment the coil is actuated – Fig. 25 refers.

A very sensitive circuit that can be used as a moisture sensing switch is shown in Fig. 26. This circuit has a gain of well over 2500.

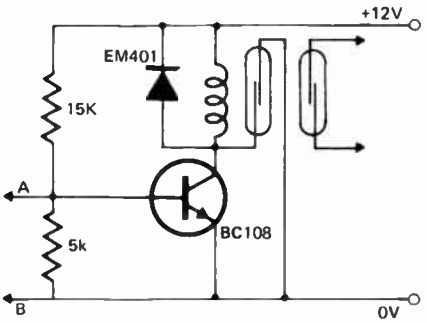

FIG. 25

The relay will close whenever the resistance between points 'A' and 'B' falls below a few hundred thousand ohms. The 100k potentiometer is not an essential part of the circuit, but may be included as a 'sensitivity' control. The current consumption of this circuit, when the relay is de-energized, is less than one micro-amp.

Any of these circuits (Figs. 23, 24, 25, 26) may be combined with the Triac actuating circuit shown in Fig. 27 and used to switch very high current loads.

For example the moisture sensing circuit shown in Fig. 26 can be combined with the Triac switching circuit to energize a large motor driven pump. If necessary, three reed switches may be combined in one energizing coil to switch three Triacs in a three phase circuit. Using this principle loads of several hundred Amps may be switched without using a single contactor.

An unusual application for a pair of reed switches is shown in Fig. 28.

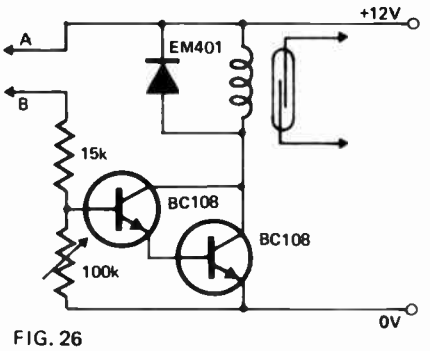

FIG. 26

This circuit can be used to switch a common antenna to either a transmitter or receiver. As the capacity between the contacts on the open reed is less than 0.2 pF, the system may be used at very high frequency.

Time delays of up to 10 seconds can be obtained using the simple circuit shown in Fig. 29. The delay is adjusted by the 50k potentiometer. It is not practicable to obtain longer delays than 10 seconds by increasing the size of the capacitor.

RESONANT REEDS

Resonant reed switches are basically similar in construction to normal reed switches, except that one reed is designed to resonate mechanically when its operating coil is energized at a specific frequency. At all other frequencies the reed will not move to any extent.

As the reed only makes contact for a portion of each cycle, it is usually necessary to arrange for latching action, or for some form of storage or pulse lengthening circuit.

Resonant reed switches are used for a variety of applications where response is required only to one specific frequency – these include communications, selective signalling, data transmission, telemetry, frequency monitoring etc.

Reed switches may also be used in very sophisticated logic circuits, usually in applications where their immunity to noise causes them to be chosen in preference to the generally cheaper solid-state components.

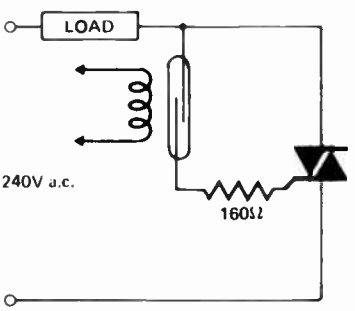

FIG. 27

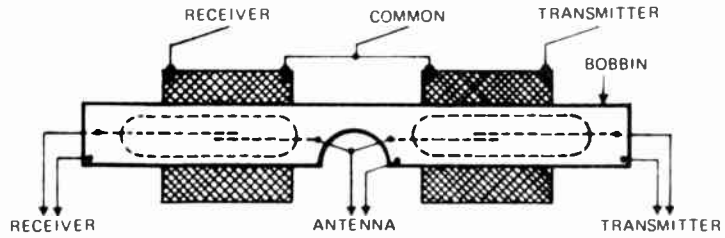

FIG. 28

Fig. 28. Here reed relays are used to switch a common antenna to either a receiver or transmitter. As the capacitance between the reeds is less then 0.2pF, the arrangement may be used at very high frequency.

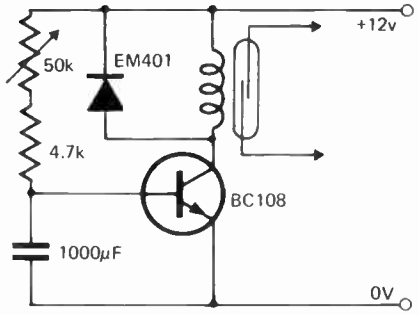

FIG. 29

Fig. 30 shows a four-stage shift register which uses reeds as magnetically latched devices in simple magnetic circuits. The information in each stage is stored as closed or open switches, and the condition of each stage is transferred to the next as the shift control is operated. Only a single contact set is used for control purposes in each stage, as the state of a stage is stored as a capacitor charge during the shifting interval. However each stage as auxiliary switching contacts for output purposes.

The basic principle of operation can be considered as a series of latching relays. Momentarily closing the 'set' contact energizes coil S1 and closes reed switches STG1. The associated bias magnet latches these switches closed, thus allowing the 'set' contact to be re-opened. The logic state of the first stage may now be shifted to the second stage by operating the shift contacts in this sequence:

a. Closing contact A, thus charging capacitor C2.
b. Closing B for a few milliseconds and thus unlatching STG1 switches.
c. Opening contacts A and B.
d. Immediately and momentarily closing contact C. Capacitor C2 now discharges through coil 2S via contact C. Switches STG2 now close and are latched by the associated bias magnet. The switches associated with the second stage are now closed and those of the first stage are open.

Sequential operation of the shift circuits in this manner moves the closed or open logic state of the reed switches from each stage to the next stage in sequence.

Two additional sets of contacts are provided in each stage, one set may be used to provide visual indication of the logic state of the stage, the second set may be used to trigger particular operations whenever required.

Reed switches may also be used in many types of coding and decoding systems. A simple decimal to binary encoder is shown in Fig. 31 In this circuit, the input is from a decimal keyboard energizing reed relay coils, while the output is in four-bit binary. Single, double, and triple switch relays are required for this application.

MERCURY WETTED CONTACT RELAYS

The mercury wetted contact relay overcomes the problem of contact bounce that is inherent in the dry reed switch.

The construction of the mercury wetted switch is shown in Fig. 32. It consists of a glass encapsulated reed which has one end immersed in a pool of mercury. The other end of the reed is capable of moving between two sets of stationary contacts. The mercury flows up the reed by capillary action and wets the surface of the fixed and moving contacts. Thus a mercury to mercury contact is maintained whilst the contacts are closed.

The resistance of mercury is very low and contact to contact resistances of mercury wetted switches rarely exceed 50 milliohms. This is somewhat less than if the contacts were permanently soldered together!

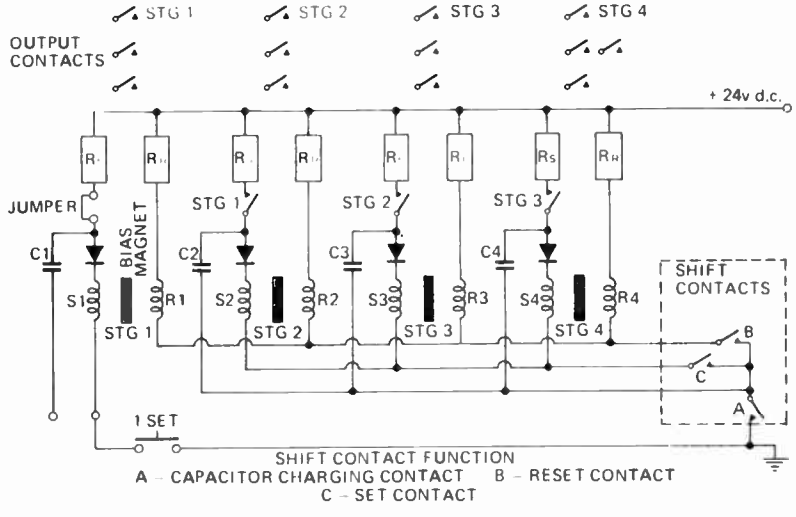

Fig. 30. This four-stage shift register uses reed relays.

Fig. 31. Reed switch decimal binary encoder.

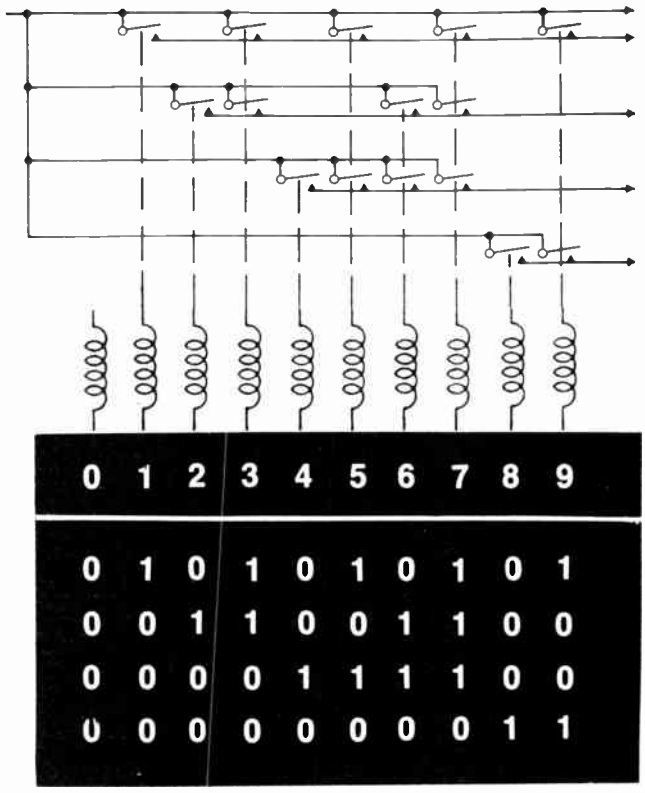

The mercury wetted switch may be opened and closed in a similar fashion to its dry reed counterpart. Operating times are typically 10 milliseconds at normal coil current, falling to three milliseconds at twice the normal ampere-turn rating. The release time is typically four milliseconds under any conditions.

Apart from their high current carrying capacity, mercury wetted reeds have extremely long life since contact erosion is eliminated.

The disadvantages of mercury wetted reeds are poor resistance to shock and vibration, and the need to mount the reed vertically.

FUTURE DEVELOPMENTS

A lot of development work is currently being undertaken – particularly toward the use of cladded reed material.

Nickel-iron reeds combine optimum magnetic characteristics with the high internal damping that is required to minimize contact bounce; but the material is by no means an ideal conductor, and because of this, high resistivity losses within the switch are appreciable at high current loadings.

Cladding with gold or copper substantially reduces many of the undesirable characteristics of the nickel-iron reeds. This cladding reduces the effect of skin resistance – which can be appreciable at high frequencies – and if the cladding is continued right to the ends of the external lead-outs – it virtually eliminates the thermal emfs generated when a copper wire is soldered to a nickel-iron reed in a conventional reed switch.

Another problem currently being investigated is that of reed switch contacts failing to separate, especially after they have been held closed for long periods at high temperatures. This is caused by molecular migration and the resultant metallic bond cannot be broken by the low separating force available. This problem has not yet been completely overcome but current development is toward heat treatment to produce a diffusion of gold into the nickel-iron base, and multi-layer diffusion techniques.

Prices of reed switches are still decreasing, and as they become cheaper, new markets are opening up.

The motor industry in particular is using reed switches in fuel injection and ignition systems. The security industry appreciate the reliable maintenance-free service that can be obtained. Machinery manufacturers are beginning to use reed switches in applications in which adverse environments preclude open switch contacts.

For what other type of switch can remain static for twenty years and then work perfectly the first time that it is actuated? ●

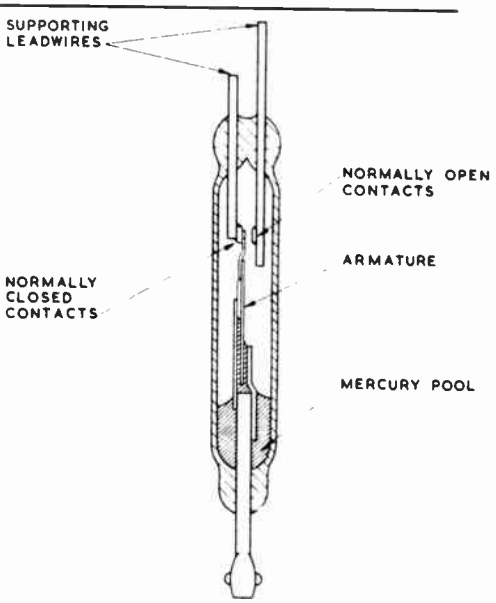

Fig. 32. Construction of a mercury wetted reed switch.

Tone operated switch

Ian Johnston, Mt Eliza Victoria

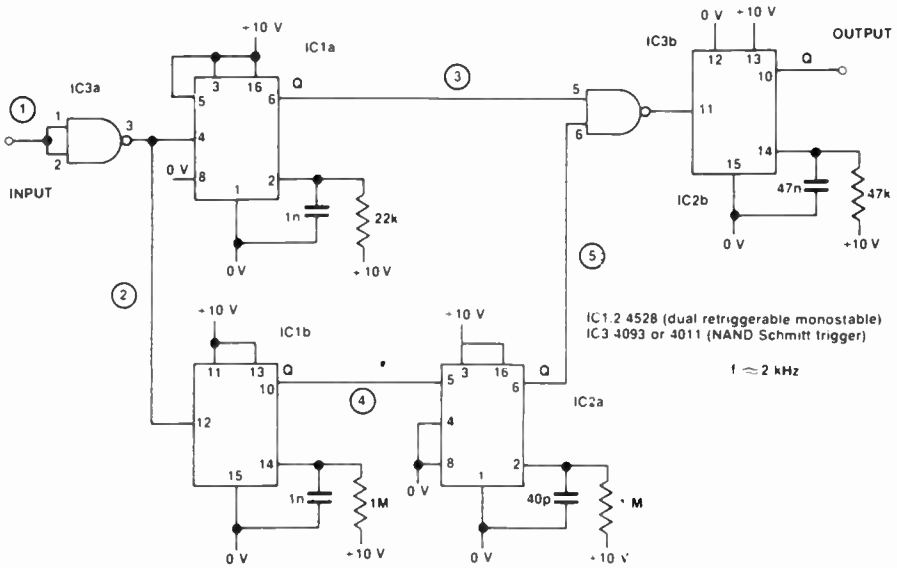

Audio frequency switch. Sine wave input (peak value equal to supply voltage) applied to input will cause output to go high within a certain narrow range of frequency. Response outside this range is zero

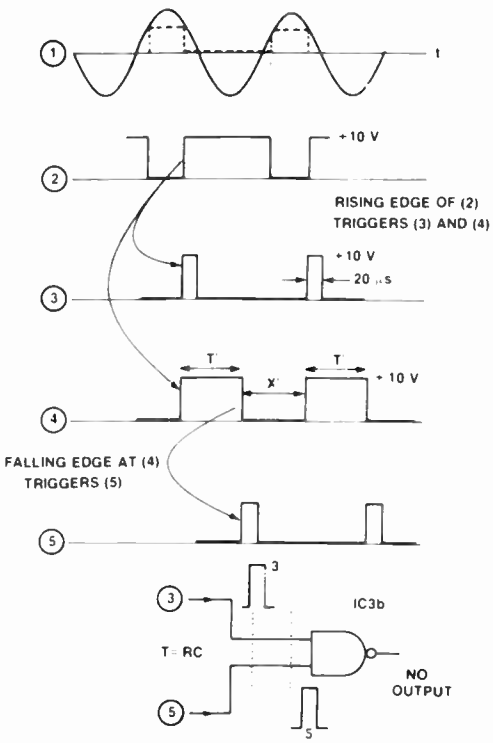

Case 1. Input frequency is below the acceptor frequency

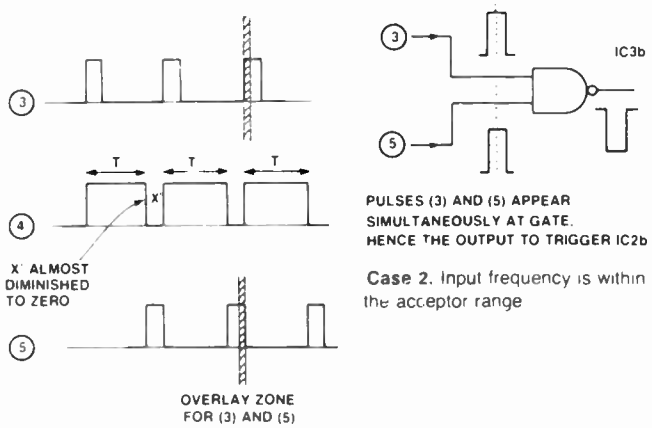

Case 2. Input frequency is within the acceptor range

Case 3. Input frequency is higher than the acceptor frequency. IC1b is retriggered before its timed interval has elapsed. Its output is permanently high and the waveform (5) disappears. Hence no triggering of IC2b is possible.

This is a new approach to the idea of operating a switch by means of a tone, such as a whistled note. It avoids the problems of acceptor filters formed by twin-T networks in the feedback paths of op-amps. They either oscillate or fail to discriminate against loud noises pitched just outside the acceptor frequency.

Gate IC3a squares up the input waveform and triggers IC1a and IC1b on the rising edge of its pulse. IC1a produces a short pulse (approximately 20 μs) which is applied to gate IC3b.

IC1b is set to a time delay that equals the period of the accepted frequency. It forms a reference which measures the period of every cycle of the incoming signal.

IC2a is triggered on its negative edge by the output pulse of IC1b. At the desired frequency the pulse from IC2a (which is short — about 20 μs) will appear at the gate almost simultaneously with the pulse from IC1a. These two pulses overlap, causing an output pulse at the gate.

The output pulses appear at intervals of T (period of the acceptor frequency) and since IC2b is timed longer than T, it is repeatedly retriggered and its output remains high for the duration of the signal.

(Note diagrams of waveforms).

When the incoming signal is too low in frequency, pulses do not coincide at gate IC3b.

When the incoming signal is too high in frequency, IC1b is prematurely retriggered, remaining permanently high at pin 10 (Q), causing elimination of any pulse output from IC2a.

(1) Ideally, C in IC1a and IC1b should be the same value to equalise start up delay upon triggering, but with small values of C this delay is negligible.

(2) Frequency of acceptance is controlled by R and C on IC1b. Bandwidth of acceptance is controlled by R and C on IC2a.

(3) A 4011 can be used in place of the 4093 for IC3, but the latter is probably preferable.

(4) IC's 1 and 2 are 4528 dual retriggerable monostables.

(5) IC2b can be connected to be positive or negative edge-triggered; it does not matter which. IC2a, however, *must* be negative-edge triggered.

(6) The input signal should rise to a peak which is about equal to the supply voltage to ensure reliable triggering.

(7) The circuit will not trigger at one half or one quarter of the acceptor frequency.

I intend to use the circuit in a 'hands-free' whistle-switched intercom in my workshop (where I frequently have dirty hands).

In a breadboard assembly using two complete circuits, the values of C and R must be chosen for IC1b: low note — 1n, 1M5; high note — 1n, 1M. Cheap greencap capacitors were used. The frequencies of acceptance measured were 900 Hz and 2 kHz respectively.

This arrangement can be made to switch on to a few whistled bars of 'I did it my way', and switch off to a few bars of 'Colonel Bogie'. Needless to say, the circuit will also find use in less frivolous applications.

When a sine wave (peak value equal to supply voltage) is applied to the input and output will go high within a certain narrow range of frequencies. Outside this range the response will be zero.

The frequency at which the circuit responds is a function of the time constant set by the capacitor and resistor on IC1b. The 'breadth' of the response range depends on the time constant set by the capacitor and resistor on IC2a.

Optocoupler devices and their applications

Ever wanted to control one circuit with another without having any intermediate electrical connection? Devices that provide coupling via a beam of infra-red light are called 'optocouplers' and they're just perfect for the job. Here's a run down on a host of popular optocouplers and how to use them in practical circuits.

Brian Dance

THE ELECTRONIC circuit designer is often faced with the problem of providing a high degree of isolation between two circuits which must nevertheless be able to pass alternating signals from one part of the circuit to the other.

For example, one may wish to have one part of the circuit completely isolated from the mains, yet use signals from this part of the circuit to control the flow of the mains current through a load. Another example occurs in patient monitoring equipment where the small voltages developed by the beating of a heart can be coupled into mains powered equipment without any danger of the equipment causing a current to flow through the heart.

Optocouplers

Optocouplers use a beam of infra-red radiation (or occasionally, visible light) to convey the signal from one part of the circuit to the other without any electrical connection whatsoever between the two parts. They are sometimes known as photon-coupled devices or as opto-isolators. They may be employed to replace conventional relays when a fast response is required or when sparking at relay contacts must be avoided in an explosive atmosphere.

An optocoupling device consists of an infra-red emitting device or other lamp on its input side and some form of detector for the radiation on the output side, both the emitter and detector being in a light-tight enclosure. The silicon detector itself may be a photo-transistor, a photo-Darlington device, an opto-triggered triac or even a field effect phototransistor.

No matter which of these device types is employed, the silicon detector has its maximum sensitivity at a wavelength quite near to that at which the gallium arsenide device emits with its maximum intensity. In other words, the devices are spectrally well matched so that a small emitter device current can produce a reasonably large response in the detecting device.

Types

A very large number of types of optocoupler have been marketed with the electrical characteristics of both the emitter and the detector having to be specified in every type. Rather than involve readers with a mass of type numbers, this article will concentrate on a limited number of readily-available devices.

The 4N26, 4N28 and MCT2 devices are examples of those using a phototransistor as a detector, the 4N33 has a Darlington output stage, the 6N139 (equivalent to the MCC671) has a 'split-Darlington' fast output device, the MCT6 is a dual device and the MOC-3020 has a triac output for 240 V mains supplies.

Dual-in-line

Although some optocoupled devices are fabricated in circular metal packages, the most common types, including those listed above, are produced in dual-in-line (DIL) packages with a typical construction like that shown in Figure 1. The emitter and detector

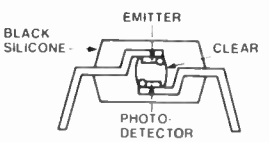

Figure 1. Cross-section through an optocoupler.

are placed fairly close together with a clear insulating material between them. The black silicon body of the device prevents stray radiation from falling on the detector. A circuit symbol is shown in Figure 2.

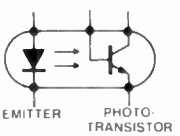

Figure 2. Symbol for an optocoupler having a transistor output stage.

In most DIL devices the radiating emitter is connected to pins on one side of the device, while the detector is connected to pins on the other side. This arrangement provides the maximum possible electrical isolation between the input and output circuits. Many of the simpler optocoupled devices differ from most other dual-in-line devices in that they have a total of only six connecting pins.

The basic internal circuitry of the devices under discussion is shown in Figures 3 to 7 inclusive. The three devices 4N26, 4N28 and MCT2 with a single phototransistor output all have the connections shown in Figure 3. The dual device type MCT6 is housed in the 8-pin package of Figure 4 so that the additional pins required are available.

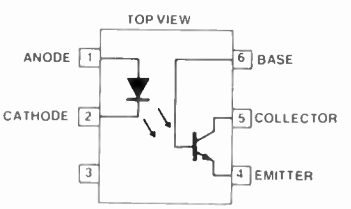

Figure 3. Pinout for the 4N26, 4N28 and MCT2 devices.

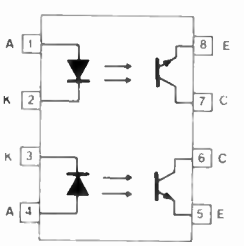

Figure 4. Pinout of the MCT6 dual optocoupler device.

DEVICE TYPE	4N26	4N28	MCT2	MCT6 (dual)	4N33	6N139	MOC-3020
Output device	← Phototransistor(s) →				← Darlington →		Triac
CTR (%) Min.	20	10	—	5	—	400	—
Typ.	50	30	—	50	500	—	—
Isolation (kV)	1.5	0.5	1.5	1.5 (0.5 between channels)	1.5	3	7.5 (max 5 sec)
Isolation resistance (Typical ohms)	10^{11}	10^{11}	10^{12}	10^{12}	10^{11}	10^{12}	—
Isolation capacitance at 1 MHz (pF)	1.3	1.3	0.5	0.5	0.8	0 6	—
Maximum emitter current (mA)	80	80	60	60 per emitter	80	20	50
Typical emitter voltage at 50 mA	1.2	1.2	—	—	1.2	1.5 (at 1.6 mA)	—
Maximum reverse input voltage (V)	3	3	3	3	5	5	3
Input capacitance (pF)	150	150	250	—	150	—	—
Maximum power Total	250	250	250	400	250	—	300
(mW) Input	150	150	200	100	150	35	
Output	150	150	200	150	150	100	
Output transistor:							
BV_{CEO} (min.)	30	30	85	85	30	—	
BV_{CBO} (min.)	70	70	165	—	50	—	
BV_{ECO} (min.)	7	7	14	13	5	—	
h_{FE} (typ.)	250	250	60	—	5000	—	
I_{CEO} (nA typ.)	50	100	50	50	100	1000	
$V_{CE\ (SAT)}$ (typ.)	0.2	0.2	0.24	0.24	1.0	—	
Typical I_c for I_F = 10 mA	5	3	—	—	—	—	
Typical bandwidth (kHz)	300	300	150	150	30	—	
Package	Figure 3	Figure 3	Figure 3	Figure 4	Figure 5	Figure 6	Figure 7

$V_{DRM} = V_{RRM} = 400$ V
Trigger current 5 mA (typ.), 20 mA (max.)
Holding current 100uA (typ.) at $V_{AK} = 3$ V
Maximum output current 100 mA RMS

TABLE 1. Basic data on the types of optocouplers discussed.

The 4N33 with its high-gain photo-Darlington output device is encapsulated in a 6-pin package with the same type of connections as the phototransistor output types of Figure 3; except for the performance differences, these devices are pin-for-pin replaceable.

The 6N139 'split-Darlington' output device has its output transistor base brought out to a separate pin, so the 8-pin dual-in-line package of Figure 6 is employed; this enables the input diode connections to be kept on the opposite side to all of the output connections.

Finally, the MOC-3020 with its triac output stage, is housed in a 6-pin dual-in-line package with the connections shown in Figure 7.

The basic parameters of these devices are listed in Table 1, but it cannot be over-emphasised that some of these values apply only under certain operating conditions stated on the data sheet which cannot all be shown in a table of a reasonable size.

It can be seen that most of the specifications required for the MOC-3020 differ from those of the other devices in their nature owing to the fact that the output triac must be specified in a different way to transistors and Darlingtons.

Which type?

If one wishes to use an input signal to control alternating current from the mains in a load, the MOC-3020 will generally be the best device from those under discussion. This optocoupler will be discussed separately from the others.

If one has to design a circuit which requires two separate control coupling systems, this can be done using the dual MCT6 device provided that phototransistor outputs are suitable for the particular application concerned. Indeed, two of these MCT6 devices can be inserted into a 16-pin dual-in-line IC socket so that one has a quadruple coupling system. (Quad devices in a single package are manufactured, but are not so common as the types under discussion.)

This leaves us with a choice, in the case of single devices, of those using a phototransistor or those employing a photo-Darlington output stage. The types using a phototransistor are most commonly employed, since they provide a fast response and can usually handle input signals with frequencies of over 100 kHz (see Table 1).

Photo-Darlington output devices provide a higher gain, but the bandwidth (or maximum usable signal frequency) is about an order of magnitude less than devices which use a simple phototransistor output; in addition, devices using a photo-Darlington output stage may be priced some 50% higher than those employing a phototransistor output, although this is not always the case.

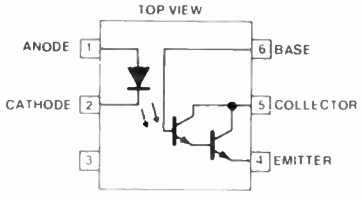

Figure 5. Pinout of the 4N33, an optocoupler having a photo-Darlington output stage.

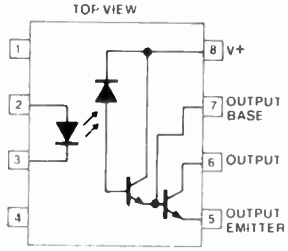

Figure 6. Pinout for the 6N139 (or MCC671) device which has a separate photodiode for maximum speed and a Darlington output for high gain.

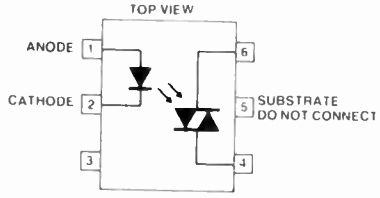

Figure 7. Pinout of the MCO-3020, which has a triac output stage.

The single devices of Figures 3 and 5 (but not the dual device of Figure 4) have the transistor base connected to a separate pin so that suitable circuitry may be used to trade-off gain in order to obtain a better high frequency response. The maximum usable frequency will be obtained when the output phototransistor is connected as a photodiode using only the base and collector connections of Figure 3, but a relatively large input current will then be needed to produce a small output current: the CTR value may be under 0.1%.

Apart from the more limited response speed of a device with a photo-Darlington output stage, it can be seen from Table 1 that the saturation voltage (under high input conditions) is much greater for the photo-Darlington device than for a simple phototransistor. Both the speed of response and the saturation differences are inherent properties of photo-Darlington devices and are not limited to optocouplers.

CTR

In order to understand some of the figures quoted in Table 1, we must first examine the ways in which certain device parameters are specified. The user need not consider any of the internal optical design points, since the manufacturer takes care of such considerations when he is designing the devices concerned.

Optocouplers are supplied as sealed units, although opto-interrupter modules are also manufactured in which there is a slot between the emitting diode and the detector so that a metal vane passing through the slot can interrupt the beam; such opto-interrupters can, for example, be used in car ignition timing systems.

One of the most important parameters of an optocoupled device is its *current transfer ratio* (CTR) which is the ratio of the output current to the input current under certain conditions specified by the manufacturer; it is usually expressed as a percentage and, broadly speaking, may be considered as the 'gain' of the device. It may be noted that devices with a triac or a thyristor output do not have a CTR value.

It can be seen from Table 1 that typical values of the CTR in the case of devices which have a simple phototransistor output stage is of the order of 50% — which means the collector current in the output phototransistor will be about half that to the input diode emitter.

The minimum value in a device of any specified type may be considerably less than that of the typical value. However, in the case of devices with photo-Darlington outputs, a CTR value of 500% is more common — which means the output current is five times the input current.

In some special devices a short light pipe is used to carry the radiation from the emitter to the detector, inevitably with some loss, so the CTR value may be reduced in such devices which may be able to withstand a much higher voltage between their input and output sides.

Unfortunately, the CTR does not have a constant value but varies widely with the diode input current and with the device temperature. Figure 8 shows the typical variation of the CTR value of the MCT2 device (which has a simple phototransistor output stage) with the forward input current passed through the emitter diode.

Each curve is for a different MCT2 device, the wide spread being due to variations in the phototransistor gain, the emitter efficiency and the coupling efficiency between the two internal components. The percentage values quoted on each curve are those for a 10 mA input current.

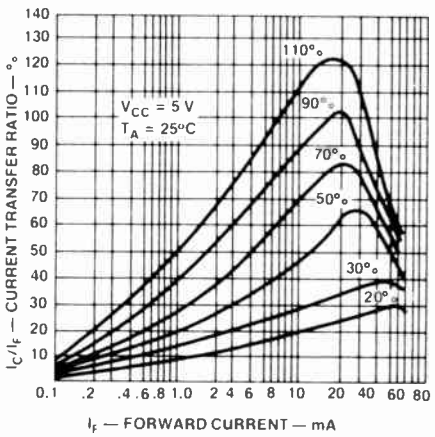

Figure 8. Variation of the current transfer ratio (CTR) with forward current in typical MCT2 devices.

The CTR value of a 4N26 or 4N28 can vary by a factor of about 2.5 between high temperatures (where it is relatively low) and very low temperatures, while devices with Darlington outputs may show variations of double this factor between temperature extremes. Rather smaller variations are more commonly found.

Isolation

Manufacturers of optocoupled device specify a maximum voltage which may be safely applied between the input and output sides of the device. In most devices this is in the range 500 to 8000 V, depending on the device type, but special types can be obtained for higher voltage isolation.

The resistance between the input and output sides of a typical device is often around 10^{11} to 10^{12} ohm. Although this seems very high, if a potential of a few kilovolts is applied across the device, a current of somewhat under 100 nA can flow. This is comparable with the current through the output of a high gain device when the input current through the emitter is under 1 mA.

If an optocoupler fails under a high applied voltage between its input and output sides, a short circuit will normally develop as a track is formed between the emitting and detecting devices. The problem can be reduced by the use of suitable current limiting resistors or protective devices in either the input or output circuit.

The stray capacitance between the input and output circuit of an optocoupler is typically of the order of 1 pF (Table 1). It can provide some unwanted coupling in circuits designed to be able to operate at high speeds, especially when inductive loads are being switched.

The emitter

The emitting diode will have a maximum continuous current rating, normally some tens of milliamps as indicated in Table 1. In some devices, pulsed currents above the maximum continuous current are permissable.

A maximum value is also imposed on the reverse voltage which may appear across the emitter diode. The application of a higher reverse voltage can cause it to breakdown and perhaps pass a destructive current; however, this problem is easily avoided by connecting an external diode across the emitter diode as shown in Figure 9.

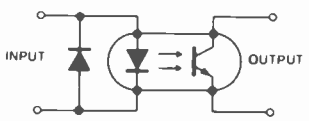

Figure 9. If a reverse voltage is likely to appear across the optocoupler emitter an external diode can be used to 'clip' it.

Although gallium arsenide diodes have been the main type used in optocouplers, there is an increasing trend to employ gallium-aluminium-arsenide types, since the latter not only emit photons more efficiently, but also provide a slightly better spectral match to the silicon detector. Thus an appreciable increase in the CTR value can be obtained.

In many optocouplers one must be careful to observe not only the total power dissipated in the complete package, but also the power dissipated in the separate input and output devices, as indicated in Table 1.

The detector

As with any other phototransistor or photo-Darlington, there is a certain value quoted for the maximum voltage which may be applied between the collector and the emitter with the base unconnected without risk of the device undergoing breakdown; this is BV_{CEO}. Similarly, values may be quoted for BV_{CBO} and BV_{ECO}.

A maximum collector current may also be quoted together with a maximum collector leakage current with base unconnected, I_{CEO}, under specified conditions.

The characteristics of the detector determine the speed of response and the bandwidth, since the emitting diodes are fast. The response time can be reduced by the use of a smaller value of load resistor, but many

manufacturers quote rise and fall times and bandwidths with load resistors which are so small that the circuit would have an inadequate gain for most applications.

The response speed of an optocoupler can be improved by using the circuit of Figure 10 in which the collector load is effectively reduced to a very low value by the virtual earth input impedance of the operational amplifier.

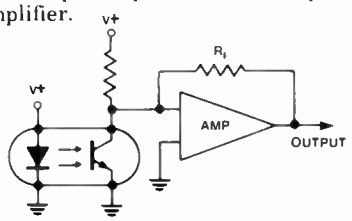

Figure 10. Response speed may be increased by the use of the virtual earth input of an op-amp.

An even simpler way of obtaining a faster response at the expense of a reduced value of the CTR involves connecting a resistor, between the base and emitter of the output transistor. As the value of this resistor is reduced, the response becomes faster until in the limit, when the resistor is a short circuit, one is using the detector as a photodiode.

If one expects to be working with a very small input current, one might expect the use of a high gain device with a photo-Darlington output would be ideal. This is not necessarily true, since the overall efficiency can fall at such currents to the point where a device with a phototransistor would be better.

APPLICATIONS

Optocoupling devices can be employed to replace relays and pulse transformers in a wide variety of applications in which high isolation may be desirable or essential. They provide fast signal transfer with excellent noise immunity. They are suitable for interfacing with TTL and CMOS circuits and can also be used for analogue signal coupling.

Circuits designed for use with single phototransistor output optocoupled devices can generally employ the 4N26, 4N28 or MCT2, but note should be made of the individual differences listed in Table 1.

For example, the 4N28 is limited to applications in which the voltage across the device does not exceed 500 V, while when the other devices are selected, it may be as great as 1.5 kV.

The phototransistors in the MCT2 and in the dual MCT6 outputs are much higher voltage devices than those used in the 4N26 and 4N28.

The bandwidth of the 4N26 and 4N28 is typically greater than that of the other two types, but so is the isolation capacitance between the input and output. However, these points are not likely to be of any great importance in most applications.

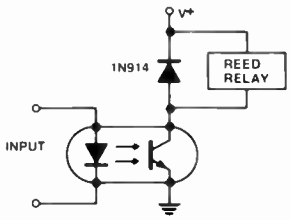

Figure 11. Using an optocoupler to isolate a reed relay.

Relay control

The simple circuit of Figure 11 shows how a small input current may be employed to control a reed relay. The inductive back-emf from the relay coil formed when the current ceases to flow through it is by-passed by the 1N914 diode so that this relatively high voltage pulse cannot damage the output transistor of the optocoupler.

The supply voltage used, V+, should have a value about equal to the voltage required by the relay, but should not exceed the V_{CEO} value of Table 1 for the optocoupler used.

Although the use of a reed relay is suggested so that the output current of the optocoupling device is kept quite small, other types of small relay can be controlled with careful circuit design. Obviously this type of circuit provides better isolation than many types of relay.

The circuit can easily be modified so that the relay does not close until the input has been applied for a short time. One merely connects a capacitor across the input diode and feeds this diode through a series resistor. The delay time before the relay closes will be dependent on the time taken for the capacitor to charge through the series resistor.

Isolated audio

The circuit of Figure 12 shows how an audio output completely isolated from the audio input signal may be obtained. A positive bias is applied to the input signal, V_S, so that the emitter diode polarity is satisfied.

The value of the input resistor R1 should be chosen so as to limit the modulating input current to a maximum of 5 mA. The 100 ohm load resistor of the phototransistor results in rather a low gain, but the 741 stage provides a gain of about ten so that a reasonably large output voltage is obtained.

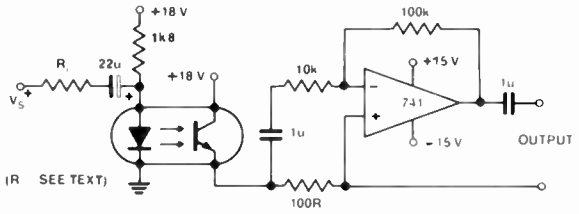

Figure 12. An isolation circuit covering the whole audio range.

The low value of the collector load resistor enables an upper frequency up to 20 kHz to be obtained, while the lower frequency response is determined by the values of the coupling capacitors employed — about 25 Hz in the case of the values shown.

Two separate +18 V supplies are required if complete isolation between the two parts of the circuit is needed. The input resistor R_1 may consist of a variable resistor in series with a fixed resistor if it is required to alter the output signal voltage without any danger of receiving an electrical shock from the output circuit when the latter is at a relatively high voltage.

TTL interface

Optocouplers are widely used in interface logic circuits where the logic signal must be transferred from a circuit at either a high or a low voltage level to a circuit at a very different voltage level.

The circuit of Figure 13 shows how an optocoupling device employing a simple output transistor may be employed to couple the output of a TTL gate to one of the inputs of a TTL 7413 device at a very different voltage level. The 7413 Schmitt circuit provides switching.

A Fairchild report suggests that the base of the output phototransistor of the optocoupling device should be connected to the emitter through a resistor of about 200 kilohm to prevent false triggering of the outputs.

Another logic circuit for coupling an input to a 7413 device is shown in Figure 14, but in this case the 4N33 with its photo-Darlington output device is used.

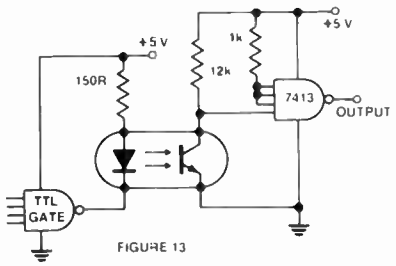

Figure 13. Isolating TTL circuits with an optocoupler.

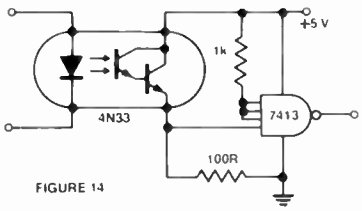

Figure 14. Control of a TTL Schmitt trigger circuit from a 4N33 photo-Darlington device.

It may be noted that in Figure 13 the load resistor (12 kilohm) is much higher than in Figure 14 (100 ohm), but the use of the higher gain of the 4N33 makes up for the lower value of load resistor.

Simple latch

The very simple latching circuit of Figure 15 can employ a pair of 4N33 photo-Darlington output devices. Initially, S1 is open and no current flows through either 4N33. If S1 is then closed, a current flows from the positive supply through the diode emitter in the upper 4N33 and through the emitter in the lower 4N33, the output of the upper device being shorted out by S1 during this time.

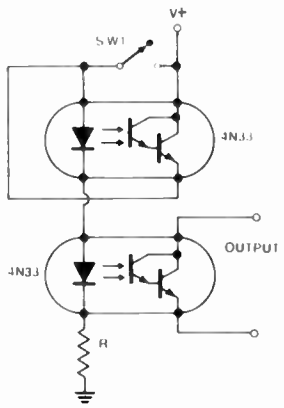

Figure 15. A latching circuit using two 4N33 devices.

When S1 opens, the short is removed from its output circuit, but the response time of the latter is longer than that of the emitter. The current therefore flows through the output of the upper 4N33, through the diode emitter of this same device to maintain the output in its conducting state and through the emitter of the lower 4N33. Thus the output of the lower device remains in its conducting state after S1 has re-opened.

The voltage across the two forward-biased emitting diodes is around 3.5 V and it is convenient to operate these diodes at about 5 mA. Thus, a suitable value for the resistor R is (V − 3.5)/0.005 or about 3.9 kilohm with a 24 V supply.

Bidirectional control

The output current of an optocoupler using a phototransistor or a photo-Darlington device must flow only in one direction, so such a device cannot control alternating current.

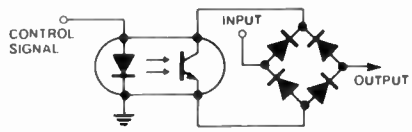

Figure 16. Controlling a bi-directional current using an optocoupler.

This problem can easily be overcome by the use of the circuit of Figure 16, in which the input-to-output current is rectified by a diode bridge circuit before being fed to the output stage of the optocoupled device.

The control signal which switches the output on and off must be unidirectional.

Power supply

Optocoupling devices can be used to isolate the control voltage of a regulated high voltage power supply from this supply line. The basic circuit which may be used is shown in Figure 17.

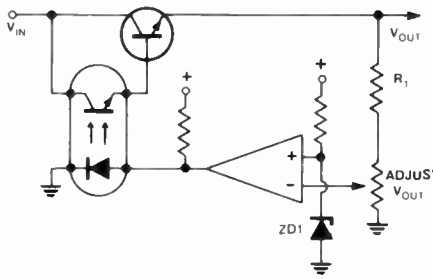

Figure 17. Using an optocoupler in a high voltage series-pass regulator.

A current flows from the stabilised output supply through the high value resistor R1 so that the variable resistor taps off a voltage proportional to the output voltage. This is compared with that across the zener diode D1 using the operational amplifier.

The output signal from this amplifier is fed to the emitter of the optocoupled device which is used to control the series pass transistor and hence to keep the output voltage constant. Thus, the amplifier device output is isolated from the high voltage supply.

A photo-Darlington device may be used in this type of circuit for higher feedback loop gain, but an external pass transistor is always required, since the output devices incorporated into optocouplers can handle only very limited power.

Fast interface

The 6N139 with its 'split photo-Darlington' output device enables the high speed of the separate photodiode to be combined with the high gain of the Darlington connected internal transistors. Although the CTR has a minimum value of 400% at a 500 mA input current, the device output can switch in a few microseconds.

A fast non-inverting logic interface circuit using this device is shown in Figure 18. The maximum switching speed depends on the load resistor, R2, and the input resistor, R1. If R1 has a value of 180 ohm a current of about 17 mA will flow to the output of the TTL input device from the internal emitter diode and the use of a 100 ohm load resistor for R2 will then enable data rates of about 300 kbit/s to be obtained. On the other hand, R1 may be increased to 1k8 for a 1.7 mA diode current with R2 2k2 for a maximum data rate of nearly 50 kHz.

Electrocardiograph amplifier

The use of an optocoupled device to provide complete isolation of a patient from electrocardiography equipment is shown in Figure 19. The electrodes from the patient are connected to the programmable 4250 preamplifier stage which operates from +/−3 V battery supplies, nulling facilities being provided by the variable resistor connected between pins 1 and 5.

The same +3 V battery supply provides the bias for the high gain BC109 transistor which drives the diode emitter of the optocoupling device.

The output phototransistor of the optocoupler receives a base bias so that some current is always passing through its collector circuit. This enables the positive and negative parts of the signal waveform to be obtained at the output.

This is a particularly important application of optocoupled devices, since without the isolation provided by such a device, small currents could be fed into the patient which in certain circumstances could produce death.

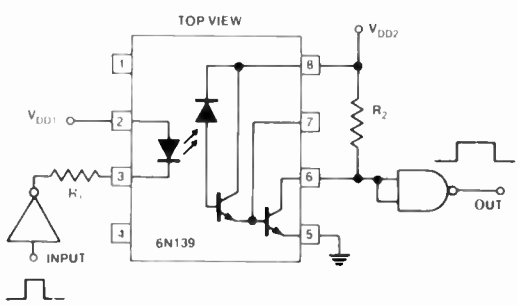

Figure 18. A fast TTL interfacing and isolating circuit using the 6N139.

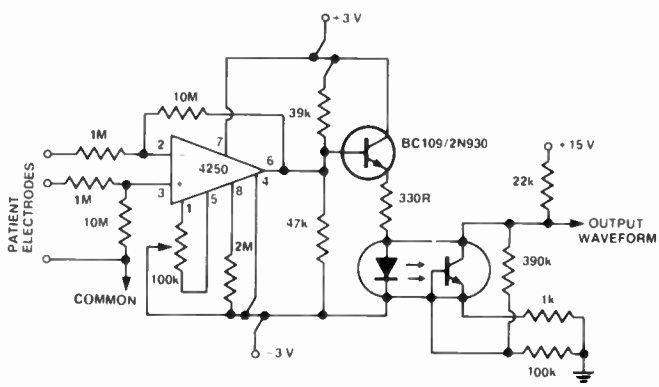

Figure 19. An electrocardiograph preamplifier circuit providing isolation of the patient from the equipment. (Litronix.)

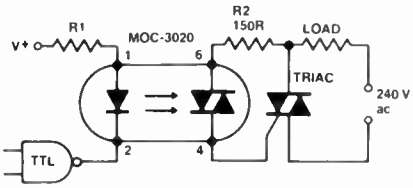

Figure 20. Control of ac power where there is a resistive load, using the MOC-3020.

The MOC-3020

The small triac in the MOC-3020 output can provide a current of up to 100 mA. This is too small for controlling the mains current passing through the load in almost all applications, but is adequate to trigger an additional external triac.

A circuit of this type is shown in Figure 20 in which the output of the TTL gate, controls the emitter current of the MOC-3020 which triggers the internal triac, the latter triggering the external triac.

The latter device should be selected so that it can hold-off the applied mains voltage and also pass whatever current is required by the particular load being used.

Figure 21 shows the use of the MOC-3020 to switch the ac current through a lamp fed from the 240 V mains when the lamp current is less than 100 mA. As the filament of the lamp has a much lower resistance when it is cold, care must be taken to ensure that the initial peak current is not excessive (about 1 A for a very short time is permissible).

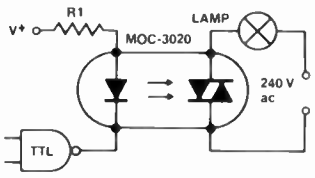

Figure 21. Controlling a lamp on the ac mains using a MOC-3020 (but watch the power rating).

In the circuits of Figures 20 and 21, the load is resistive and conduction of the internal triac ceases when the mains voltage passes through zero during the course of the mains cycle.

In the case of an inductive load (such as an electric motor), however, large back-emf pulses can be generated when the current ceases to flow through the load and this could cause the internal triac of the optocoupler to operate in an improper way.

This problem can be avoided through the use of the type of circuit shown in Figure 22, the values of the components of the 'snubber network' connected across the external triac being dependent on the load inductance and resistance.

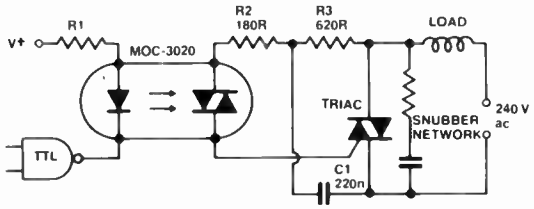

Figure 22. Control of ac power where the load is inductive (i.e: a motor), using the MOC-3020. Note the use of a 'snubber' network. Typical values for the RC network would be R = 180 ohms, C = 220n.

Conclusion

Simple optocoupler devices can be employed in a wide range of circuits from the simplest types to quite complex ones. At prices ranging from under one dollar up to a few dollars, they are excellent value! ●

Power Supplies

Some pointers on the use of monolithic three-terminal voltage regulators from Tim Orr.

THE PRODUCTION OF stable regulated supply voltages has been simplified by the introduction of three terminal voltage regulators. These devices make power supply design relatively simple, but the designer must be aware of other important details that can cause poor results. Consider a simple unregulated supply, fig. 1.

A mains transformer isolates a piece of equipment from the potentially lethal mains voltage and provides a suitable voltage to be rectified, smoothed and applied to a voltage regulator. The secondary voltage of the transformer is measured in volts RMS at a particular loading.

If the transformer is rated at 15 V at 10 volt-amps (VA), the output voltage will be 15 V when the load upon the transformer secondary is 10 VA (10 watts).

If the load is removed the output voltage will rise. The percentage change from load to no load is known as the transformer regulation and is typically 10 – 20%.

If the rectifier is followed by a capacitor-input smoothing circuit, the dc voltage is 1.414 (the square root of 2) times the transformer RMS voltage. Thus a 15 VRMS (loaded) transformer will generate about 20 Vdc when full wave rectified and smoothed, which will rise to about 25 Vdc when the load is removed (assuming 20% regulation, see fig. 1).

When calculating these voltages, remember that each forward biased diode in the current path will drop 0.6 V, so a full wave rectifier will cause a 1.2 V drop.

The smoothing capacitor should be selected to withstand the peak no load voltage from the transformer. Also, make certain that the polarity of an electrolytic capacitor is correct; they can literally explode if wired up backwards!

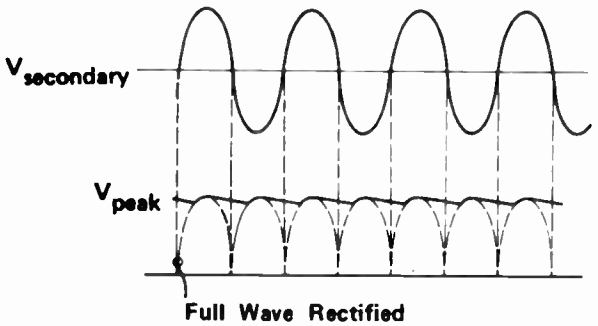

Fig. 1 (above) A simple unregulated power supply, (top) the output (with a load resistor).

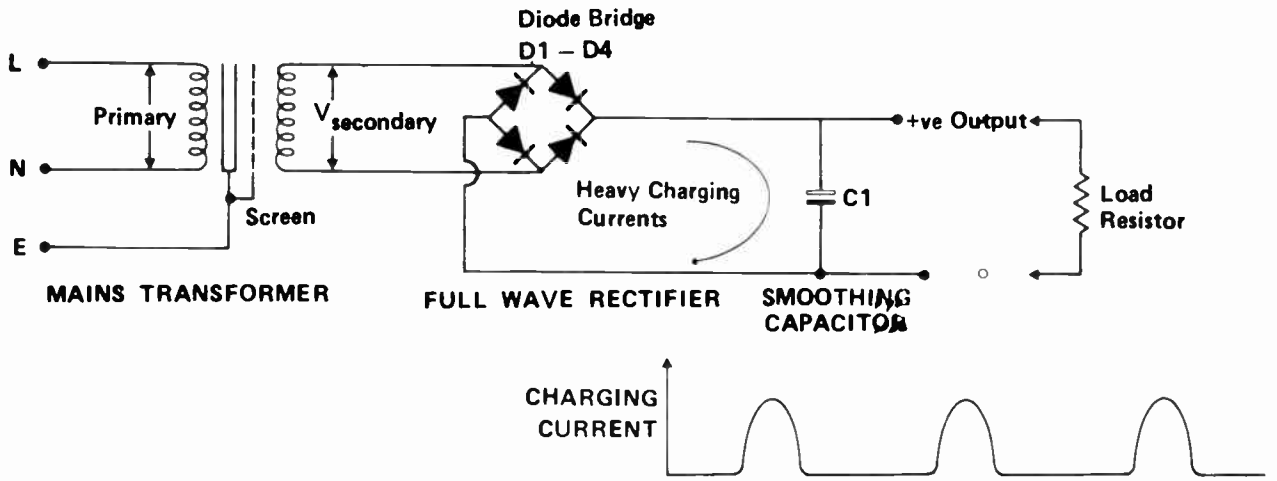

There are three sections, a step down, isolating transformer, a diode bridge and a smoothing capacitor. The transformer is driven from the mains. Some transformers have a copper screen to isolate the primary winding from the secondary windings. For safety, and noise reduction, this should be connected to earth.

Another type of mains transformer uses what is known as a split bobbin; the primary is wound on one bobbin, the secondary on another. Thus the two windings are inherently physically isolated. These two transformer types are generally constructed on what is known as an 'E' core; take one to bits and you will find that it is constructed out of lots of 'E' shaped laminations. These 'E' laminations are butted into 'I' laminations, and clamped together. This butting together of the laminations can cause magnetic field problems. The wider the gap between the 'E' and 'I' laminations, the larger the magnetic field around the transformer.

The magnetic field can generate a significant amount of induced hum in nearby electronics, this can be overcome by using a low leakage toroidal transformer which is constructed from circular laminations. The primary and secondary windings are wound through the centre of the toroid. The toroidal transformer, by virtue of its 'continuous' laminations results in a low stray field and a low profile design, making it ideally suited for audio amplifier applications.

When a load is placed upon the power supply shown above, the output voltage appears as a DC voltage on top of which is a ripple voltage. This can be thought of as two separate periods, a charge period where the capacitor is charged up by the power supply and a discharge period where the load discharges the capacitor.

This charging and discharging generates a ripple voltage which has a period of 10 ms (100 Hz). A load current of 100 mA, and a 100μ

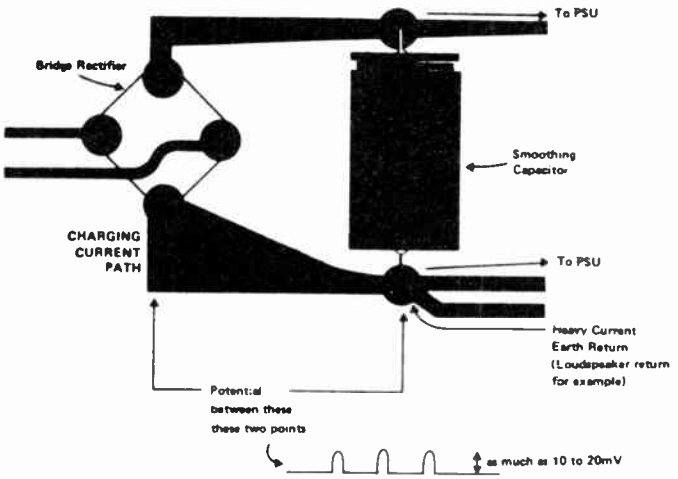

Fig. 2. Correct pcb layout for power supply design. Note the thick short tracks from the bridge rectifier to the filter capacitor.

capacitor will result in a ripple voltage of about 0.7 V p-p.

As a rule of thumb allow 1 to 1.5 V p-p maximum ripple if a voltage regulator is being used. This will generally result in an output ripple of less than 1 mV.

Generally the discharge period is much longer than the charge period. This means that the transformer is only supplying power to the capacitor for short periods. During these periods the smoothing capacitor is rapidly charged, and it is quite common for these current surges to exceed several amps. This can cause mains BUZZ problems when laying out printed circuit board designs for power supplies.

The correct layout is shown in Fig. 2. If the current surge is 1 A and the track resistance is 20 milliohms then the voltage developed between the rectifier and the filter capacitor will be 20 mVpp.

Table One

V secondary at rated load	V peak at rated load	V peak off load transformer regulation 20%
5 VRMS	7.1 V	8.5 V
6 VRMS	8.5 V	10.2 V
9 VRMS	12.7 V	15.3 V
10 VRMS	14.1 V	17.0 V
12 VRMS	17.0 V	20.4 V
15 VRMS	21.2 V	25.5 V
20 VRMS	28.3 V	34.0 V
25 VRMS	35.4 V	42.4 V
30 VRMS	42.4 V	51.0 V
35 VRMS	49.5 V	59.4 V
40 VRMS	56.6 V	67.9 V

Voltage regulators

A voltage regulator takes a varying (unregulated) input voltage and produces a stable (regulated) output voltage. There is a wide range of fixed voltage three terminal regulators to choose from, with a choice of maximum current handling, output voltage and positive or negative operation. The data sheets for these devices contain lots of seemingly complex information and so a glossary of terms is given here.

Ripple Rejection

The ratio of the ripple voltage at the regulator input to that at the output, generally expressed in dB. Typically of the output, generally expressed in dB. Typically of the order of 60 dB (1000 to 1). that is, 1 Vpp of ripple at the input ends up as 1 mVpp at the output.

Temperature Coefficient

The output voltage change for a change in regulator temperature, expressed in mV/°C.

Input Voltage range

The range of input voltages over which the regulator will function normally. For example, a 12 V regulator may work from 14.5 V to 30 V. At 14.5 V the regulator will 'drop out' and lose its regulation. Regulators generally need at least 2 to 2.5 V in excess of their output voltage. At 30 V the regulator will go 'pop' (time to buy a new one).

Output voltage

The voltage at the output terminal with respect to ground. Generally within ±5% of stated value.

Line Regulation

The change in the output voltage caused by a change in the input voltage, typically of the order of 0.2%.

Short Circuit Current

The output current when the output is shorted to ground.

Output Noise Voltage

The RMS noise voltage measured at the regulators output, not including any ripple.

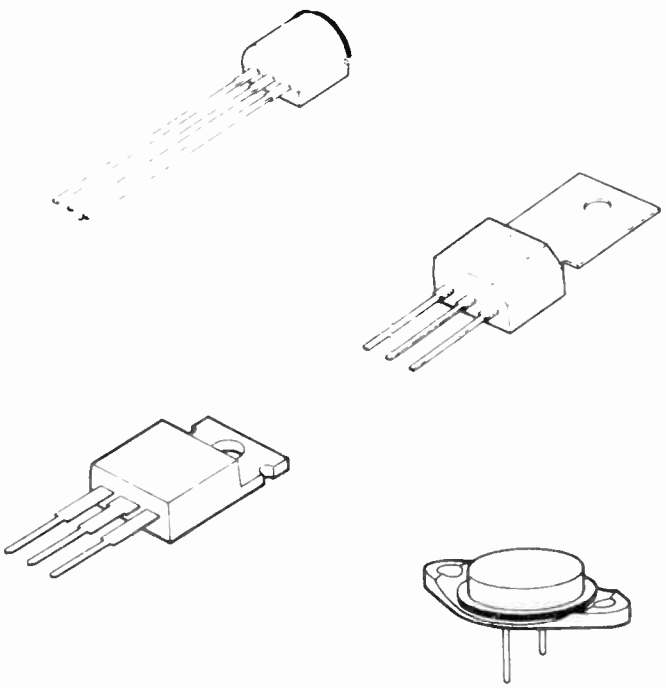

Regulator IC packages. From top: TO-92, TO-202, TO-220, TO-3

Power Dissipation

The maximum power that the regulator can safely generate on a particular heatsink.

As a rule of thumb the regulator case should not exceed about 80°C (which is hot to touch). However, always run the device at as low a temperature as possible. It is thermal ageing that eventually kills electronic devices and for higher temperatures the ageing process is disproportionately faster.

Tricks of the Voltage Regulator Trade

Fig. 1. The conventional arrangement of a three terminal regulator. It is advisable to use a decoupling capacitor connected close to the input terminals. This prevents high frequency instability. If this capacitor is left out then regulation can sometimes be greatly reduced. The decoupling capacitor on the output helps reduce the impedance at high frequencies, where the regulator loses its performance. For best results use a tantalum capacitor.

The internal current limiting of the regulator is shown. When the load current exceeds the current limit, the output voltage drops to almost 0V. This makes the regulator short circuit protected. Another type of current protection is known as 'FOLD BACK' current limiting (shown dotted). This serves to reduce the short circuit current. These devices protect the power supply from abuse. Another type of protection device is the overvoltage clamp, which protects the circuit being supplied from an increase in the power supply voltage. When the supply voltage exceeds a certain level an SCR is triggered on and clamps the rail to ground. This is intended to pop a fuse and so disconnect the faulty power supply (which is better than replacing a $1000 worth of ICs). With foldback the short circuit power dissipated in the regulator is less than that with current limiting.

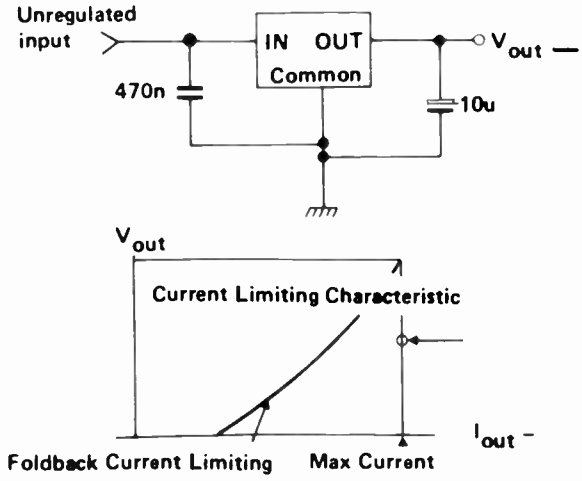

Fig. 2. The output voltage of a regulator can be increased by applying a voltage to the common terminal. This can be done by a zener diode.

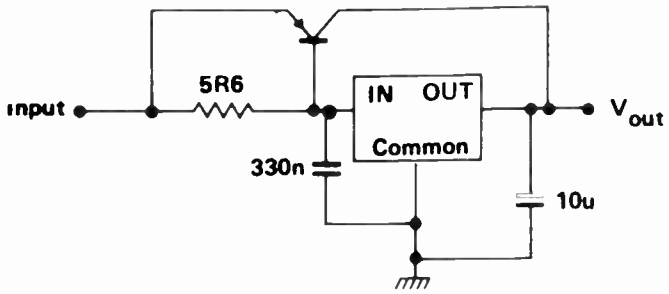

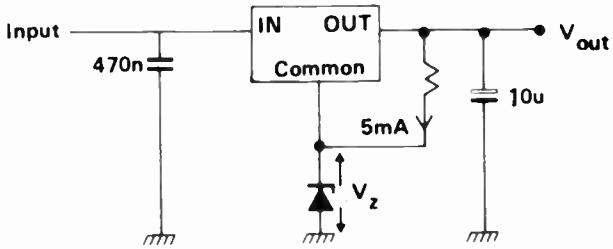

Fig. 3. The output current can be increased by using a bypass transistor. When the current flowing through the voltage regulator exceeds 100 mA (the voltage across the 5R6 being 560 mV), the bypass transistor begins to turn on. This transistor takes all currents in excess of 100 mA and yet the output still remains regulated.

Fig. 4. Sometimes it may be necessary to use a supply which exceeds the maximum voltage rating of the regulator. A simple voltage regulator 2D, and Q1 can be used to overcome this problem. 2D should be chosen so that it is about 6V higher than the regulator output voltage. This technique has the added advantage that the power dissipated in the regulator is less (the rest being dissipated in Q1), and the regulator is presented with a semiregulated voltage, so the output will have better regulation.

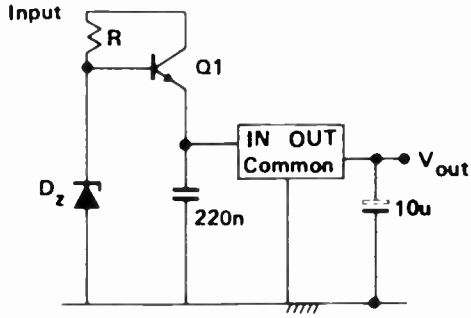

Dual Power Supply

This circuit shows a complete regulated dual power supply. The unregulated rails are obtained from a split secondary transformer, a bridge rectifier and two smoothing capacitors. A positive and a negative regulator have been used to generate the + and - rails. These regulators should be mounted together with insulating washers on heat sinks. The pin out of the negative regulator is different to that of the positive regulator. The two diodes at the output prevent latching up situations (on load) whereby one side starts up faster than the other and forcibly reverse biases it, preventing it from operating.

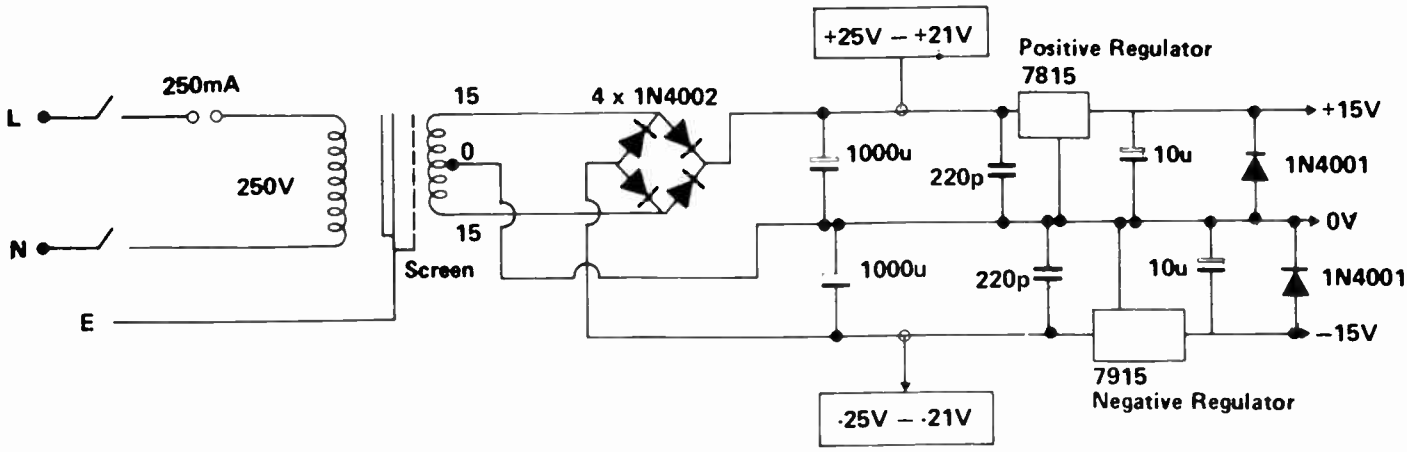

Tracking Regulator

Instead of using a negative voltage regulator to obtain the negative rail, an op amp and a power transistor can be used. The resistor ratio, R1, R2 determines the negative rail voltage. The negative rail is not, however, current limited.

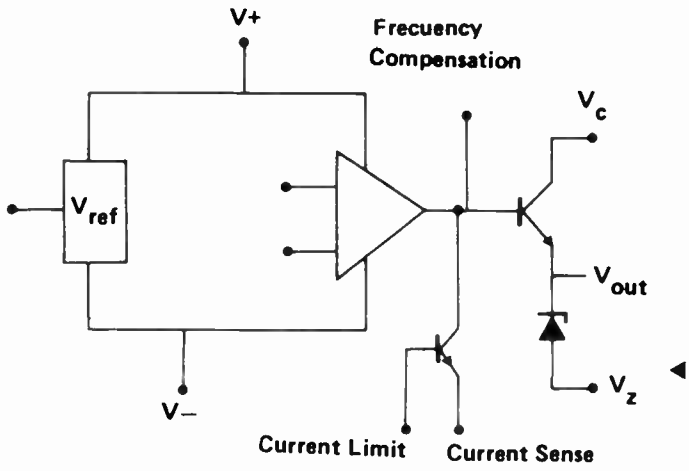

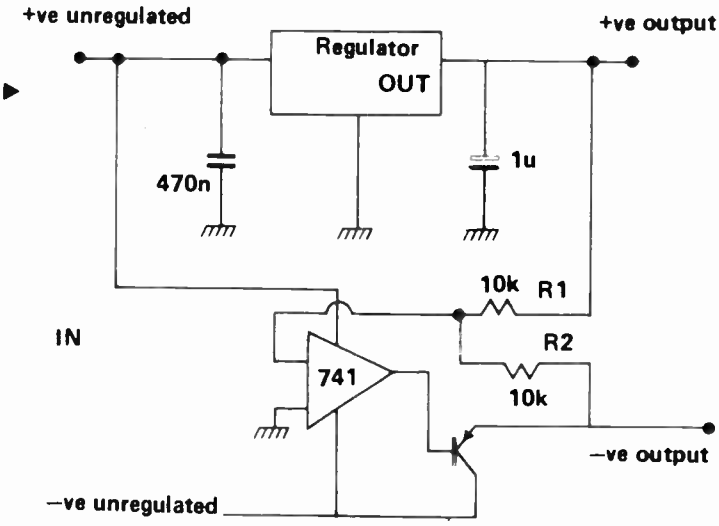

Voltage Regulator

The 723 is an 'industry' standard device. Many manufacturers produce it and the device itself is versatile. It comes in a 10 pin TO5 can or a 14 pin/DIL pack. The device contains a precision voltage reference, with a temperature coefficient of 50 ppm/°C, an error amplifier, an internal transistor capable of handling 100 mA and a current limiting mechanism. By using a few external resistors, a capacitor and maybe an external power transistor, a wide variety of regulator designs can be realised. Right is shown the block diagram of the 723 regulator. As pinouts vary depending upon package, no pin numbers are shown.

Adjustable Positive Voltage Regulator

By using a feedback path (R1, R2), a regulated output voltage can be generated. The voltage reference is connected to the non-inverting input of the error amplifier and the output voltage (via R1, R2) to the inverting input. The error amplifier drives the output transistor on the IC and hence the output voltage is controlled by the feedback voltage. If R1 and R2 are replaced by a potentiometer, the supply can be made variable. A 100 pF capacitor is used to stabilise the device. Rsc is used as a current limit control. When the current through Rsc (the load current) generates a voltage of 560 mV accross it, a current limiting transistor is turned on which in turn shorts out the regulating transistor, causing the output voltage to collapse towards 0 V.

$$V_{OUT} = [V_{REF} \times \frac{R_1 + R_2}{R_2}]$$

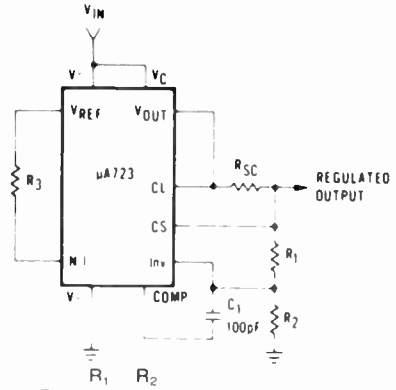

Note: $R_3 = \frac{R_1 R_2}{R_1 + R_2}$ for minimum temperature drift.

R_3 may be eliminated for minimum component count.

Regulated Power Supply

Sometimes it is necessary to make a simple power supply using discrete components when a non-standard voltage is required. The circuit shown uses all the basic elements of a voltage regulator, that is, a reference voltage ZD1, an error amplifier and a series control transistor Q1. The zener diode , ZD1, sets up a reference voltage of 5V1. This diode has a temperature coefficient of -1.2 mV/°C (a 5V6 zener is best at -0.2 mV/°C). The resistor ratio of R3 and R2 sets the output voltage and the op amp provides the error correction (regulation). C1 is used to reduce the output impedance at high frequencies.

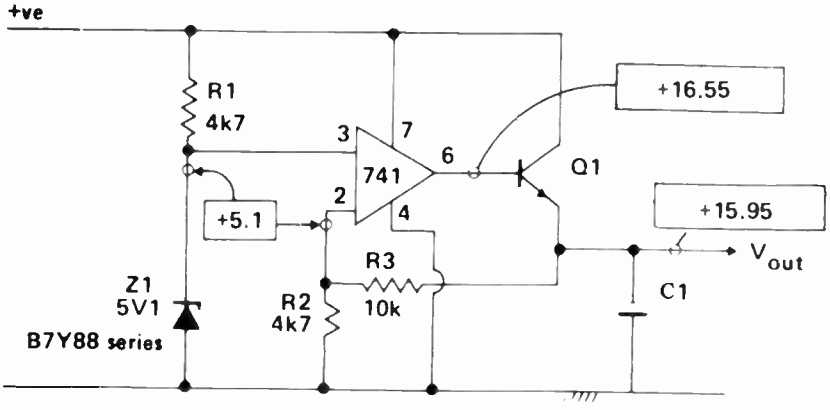

SCRs, triacs and power control

SCRs and triacs are high speed solid state switches specifically intended for use in ac and dc power control applications. Ray Marston explains their basic principles in this edition of Circuit File, to be followed next edition with a stack of application circuits.

Ray Marston

FOR ELECTRONIC switching applications in dc and ac circuits, SCRs and triacs have no equals and wide application. The SCR is like a diode you can turn on and off while the triac is like two back-to-back diodes you can turn on and off (back-to-back SCRs!). That's the easiest way to think of these two very useful, related, devices. But to be able to appreciate their characteristics and how to use them, you need to know about them in somewhat more depth. So, let's look at some basic theory and circuits first.

The SCR: basic theory

The SCR, or silicon controlled rectifier, is a four-layer pnpn semiconductor switching device. It is represented by the symbol shown in Figure 1a. Figure 1b shows the transistor equivalent circuit of the SCR, which takes the form of a complementary regenerative switch in which the base current of Q1 is derived from the collector of Q2 and the base current of Q2 is derived from the collector of Q1. Figure 1c shows the basic connections for using the SCR as a switch in dc power control circuitry. The basic characteristics of the SCR can readily be understood with the aid of Figures 1b and 1c, and are as follows.

(1) When power is first applied to the SCR (by closing SW1 in Figure 1c) the SCR is 'blocked' and acts like an open-circuit switch. This action can be understood by looking at Figure 1b, where it can be seen that, since Q1 base is shorted to the cathode via R1-R2, Q1 is cut off through lack of base current and thus feeds no base drive to Q2, which is also cut off. As both transistors are cut off under this condition only a small leakage current flows between the anode and cathode of the device.

(2) The SCR can be turned on and made to act like a closed switch (or forward-biased silicon rectifier) by simply applying positive gate current by closing SW2 in Figure 1c. This gate current causes the SCR to switch on very rapidly.

If the externally-applied gate current is sufficiently large it will apply base drive to Q1, causing Q1 to start to turn on. As Q1 starts to turn on, its collector current feeds base drive to Q2, causing Q2 to turn on and feed increased base drive into Q1, etc. A fast regenerative action thus takes place, with both transistors switching rapidly into saturation, the total saturation voltage typically being in the range one to two volts.

(3) Once the SCR has been turned on and is conducting significant forward current, the gate loses control and the SCR remains latched on even if the gate drive is subsequently removed. Thus, only a brief pulse of gate current is needed to latch the SCR on. Note from Figure 1b that, because of the presence of R1 and R2, the SCR can *not* be turned off by shorting or reverse-biasing the gate-cathode terminals of the device.

(4) Once the SCR has latched into the on state it can only be turned off again by momentarily reducing its anode current below a value known as the 'minimum holding current'. Since turn-off occurs whenever the current is reduced below this critical value, it follows that turn-off occurs automatically in ac circuits near the zero-crossing point at the end of each half-cycle.

(5) Internal capacitance inevitably

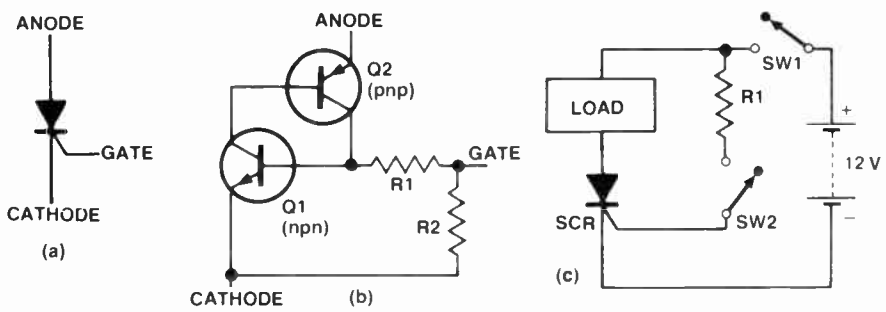

Figure 1. (a) SCR symbol (b) transistor equivalent circuit of the SCR and (c) basic connection for using the SCR as a dc power-control switch.

DATA FILE 7 DEVICE TYPE NO.	PIV RATING	TOTAL CURRENT RATING, RMS/average	V_{GT} (max)	I_{GT} (max)	I_H (max)
TAG 1/100	100 V	1 A/0.64 A	2 5 V	10 mA	25 mA
TAG 1/600	600 V	1 A/0.64 A	2 5 V	10 mA	25 mA
C106D	400 V	4 A/2.5 A	0.8 V	0 2 mA	3 mA
2N3525	400 V	5 A/3.2 A	2 V	15 mA	20 mA
BT109	500 V	6.5 A/4 A	2 V	15 mA	3 mA
IR122A	100 V	8 A/5 A	1 5 V	25 mA	30 mA
IR122D	400 V	8 A/5 A	1 5 V	25 mA	30 mA
C116D	400 V	8 A/ 5 A	1.5 V	20 mA	35 mA
C126M	600 V	12 A/7 5 A	1 5 V	30 mA	35 mA

Figure 2. Basic details of some of the most popular SCPs

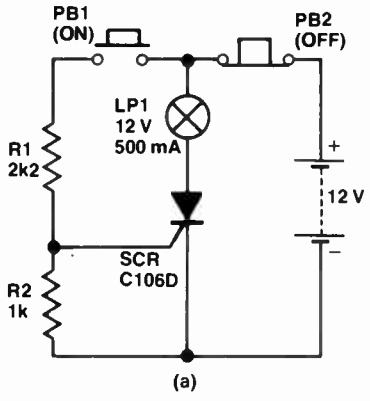

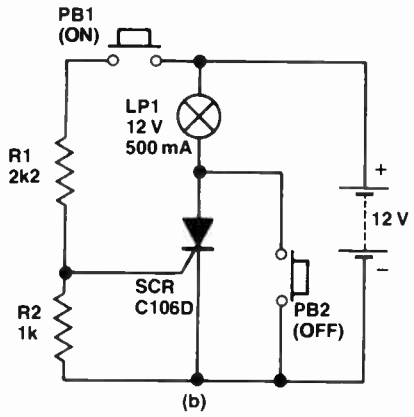

Figure 3. (a) Simple dc on/off circuit and (b) alternative dc on/off circuit.

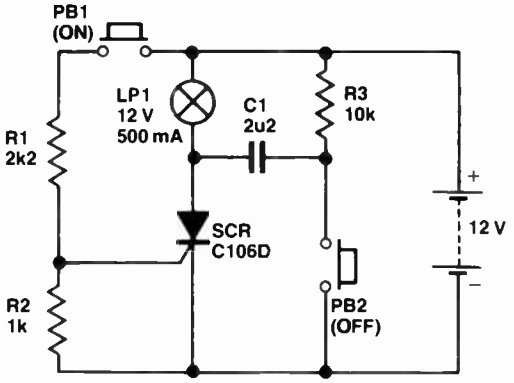

Figure 4. Capacitor turn-off circuit.

exists between the anode and gate of the SCR. Consequently, if a sharply rising voltage is applied to the SCR anode this internal capacitance can cause part of the rising voltage to break through to the gate and thus trigger the SCR on. This 'rate-effect' turn-on can be caused by supply-line transients, and sometimes occurs at the moment that supplies are switch-connected to the SCR anode. Rate-effect problems can usually be overcome by wiring a simple RC 'snubber' network between the anode and cathode of the SCR, to limit the rate-of-rise to a safe value.

These then, are the basic characteristics of the SCR. As you can see, it's a pretty simple device. If you ever need to select an SCR for a particular application, you'll usually find that the most significant parameters are the main voltage and current ratings, plus the gate sensitivity rating and (occasionally) the device's 'minimum holding current' value. The list of Figure 2 gives basic details of a few of the most popular SCRs.

The SCR: basic dc circuits

SCRs have applications in both dc and ac power control circuitry. Let's look first at some basic dc circuits. Figures 3a and 3b show alternative ways of using the SCR as a pushbutton-controlled on/off power switch feeding a 12 volt, 500 mA lamp. In both circuits the lamp and SCR can be latched on by momentarily closing PB1, thereby feeding gate drive to the SCR via R1. Note that the gate is tied to the cathode via R2, to give improved stability. Once the SCR has latched on, it can only be turned off again by momentarily reducing the anode current below the device's I_H value; in Figure 3a this is achieved by momentarily opening PB2; in Figure 3b the turn-off action is achieved by using PB2 to place a momentary short between the anode and cathode of the SCR.

Figure 4 shows another way of achieving SCR turn-off. Here, once the SCR has turned on, C1 charges up to almost the full supply voltage via R3 and the SCR anode, with the R3 end going positive. When PB2 is subsequently closed it clamps the positive end of C1 to ground, and the C1 charge forces the SCR anode to momentarily swing negative, thereby reverse-biasing the SCR and causing it to turn off. The capacitor charge bleeds away rapidly under this condition, but has to hold the SCR anode negative for only a few microseconds to ensure complete turn-off. Note that C1 must be a non-polarised component.

A variation of the capacitor turn-off circuit is shown in Figure 5. A slave SCR is used to replace PB2 of Figure 4 and capacitive turn-off of SCR1 is achieved by briefly driving SCR2 on via PB2. SCR2 turns-off once PB2 is released, since the anode current provided by R3 is lower than the SCR2 holding current.

Figure 6 shows how the above circuit can be modified so that it acts as an SCR bistable or flip-flop driving two independent lamp loads. Assume that SCR1 is on and SCR2 is off, so that C1 is fully charged with its LP2 end positive. The state of the circuit can be changed by briefly operating PB2. SCR2 is then driven on via its gate, and as it goes on it drives SCR1 off capacitively via its anode. C1 then recharges in the reverse direction. The state of the circuit can then again be changed by briefly operating PB1, thus driving SCR1 on via its gate and driving SCR2 off capacitively via its anode. The flip-flop process is repeated ad infinitum.

The dc circuits that we have looked at so far have all used simple resistive 'lamp' loads and have inevitably produced a self-latching

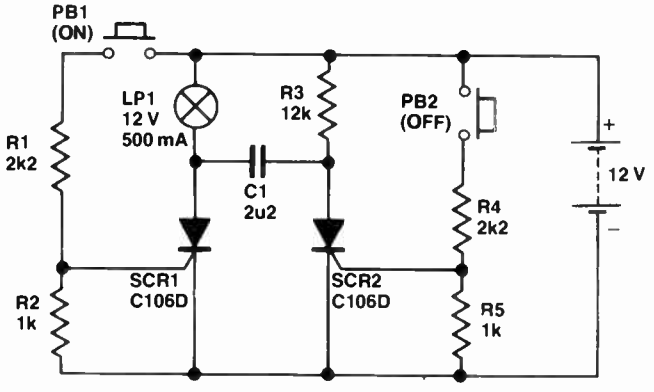

Figure 5. Capacitor turn-off with SCR slaving.

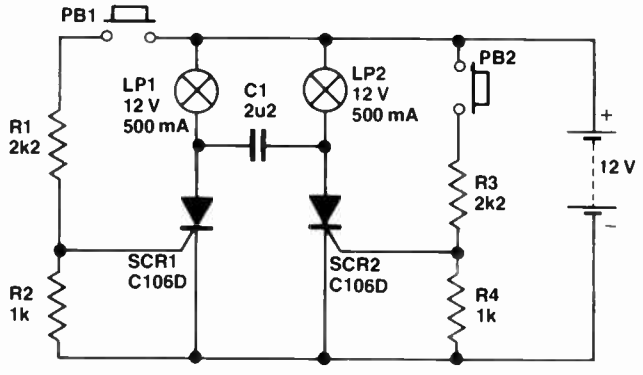

Figure 6. SCR bistable or flip-flop

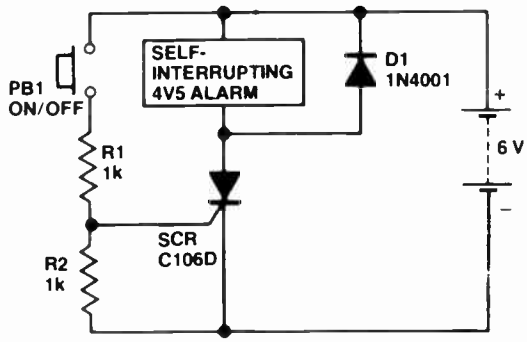

Figure 7. Simple non-latching alarm circuit.

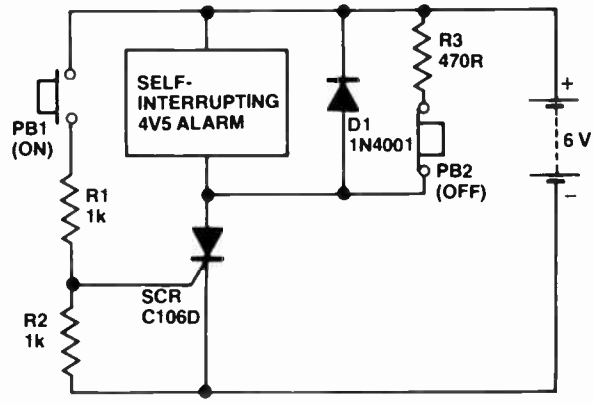

Figure 8. Simple self-latching alarm circuit.

action in the SCRs. Figure 7 however, shows a simple dc alarm circuit driving a self-interrupting load such as a bell or buzzer, and gives a non-latching action.

When self-interrupting devices such as bells or buzzers are connected across a supply, a current flows through a built-in solenoid via a pair of contacts. This current induces a magnetic field in the solenoid and causes a striker to fly outwards and open the contacts, causing the current to fall to zero and making the magnetic field collapse. Once the field has collapsed the striker falls back again and the contacts close, so current is again applied to the solenoid and the action repeats. Consequently, this type of load acts like a switch that repeatedly opens and closes rather rapidly.

When such loads are connected in the Figure 7 circuit therefore, the circuit does not self-latch in the normal way and the alarm operates only so long as PB1 is closed. Because of the inductive nature of such loads, a damping diode must be wired across them when they are used in SCR circuits, as shown in the diagram.

The Figure 7 alarm circuit can be modified to give a self-latching action if required by simply wiring a 470R (or lower) resistor in parallel with the alarm, as shown in Figure 8. In this case, the anode current of the SCR does not fall to zero when the alarm self-interrupts, but falls to a value determined by R3. If this value is in excess of the SCR's 'holding' value, the SCR self-latches. The circuit can be unlatched by briefly operating PB2, enabling the anode current to fall to zero when the alarm self-interrupts.

Finally, to complete this section, Figure 9 shows a circuit that can be used to demonstrate the rate-effect turn-on of the SCR, and a method that can be used for rate-effect suppression. Here, the SCR uses a 3 V lamp as its anode load, and is connected across the 4V5 battery supply via SW1. A 4V5 domestic door bell can be connected across the supply via PB1, and enables transient modulation to be applied to the supply line and thus to the anode of the SCR. This modulation can cause rate-effect turn-on of the C106D SCR, which has a critical slew rate of 20 V/us. The network R2 and C1 form a 'snubber' or rate-effect suppression network and can be connected to the SCR via SW2.

To demonstrate the rate-effect, open SW2, close SW1, and then close PB1 so that the bell rings. The resulting supply line transients should be enough to trigger the SCR and turn the lamp on; if not, wire a one ohm resistor in series with the battery. Once the SCR and lamp have been triggered on, they can be turned off again by briefly opening SW1.

Once the turn-on rate-effect has been demonstrated, the effect of the suppressor network can be demonstrated by closing SW2 and SW1 and then operating the bell via PB1. The lamp resistance (plus R2) acts with C1 as a smoothing network that reduces the rate of rise of the anode modulation signal, thereby protecting the SCR against false triggering. R2 is wired in series with C1 to limit the capacitor's discharge currents to safe values when the SCR is triggered on via legitimate signals.

The SCR: basic ac circuits

Figure 10 shows a basic halfwave on/off circuit driving a 100 W lamp from a 120 or 240 Vac power line. With SW1 open, zero gate drive is applied to the circuit, so the lamp and SCR are off. Suppose however, that SW1 is closed. On negative half-cycles, the SCR is reverse biased and gate signals are inhibited by D1, so the SCR is off. On positive half-cycles, the SCR is initially off at the start of each half-cycle, so the full available line voltage is applied to the gate via the lamp and D1-R1; shortly after the start of the half-cycle sufficient voltage is available to trigger the SCR, which turns on. As the SCR goes on its anode voltage falls to near zero, thus removing the gate drive but the SCR remains self-latched for the duration of the half-cycle. The SCR automatically turns off again when the half-cycle ends and the anode current falls to zero.

The Figure 10 circuit gives halfwave operation only. Figures 11 and 12 show alternative ways of obtaining fullwave operation. In these circuits, the ac is converted to rough (unsmoothed) dc via a bridge rectifier and the rough dc is applied to the SCR. With SW1 open the SCR is off, so zero current flows through the bridge and the load. When SW1 is closed the SCR is driven on shortly after the start of each half-cycle of rough dc, so fullwave power is applied to the load. As the SCR goes on in each half-cycle, the gate drive is automatically removed but the SCR stays latched on for the duration of the half-cycle. The SCR switches off at the end of each half-cycle as its anode current falls to zero, so power is removed from the load when SW1 is opened.

Note in the Figure 11 circuit that the load is connected to the dc side of the bridge. A fuse must be placed on the ac side of the bridge, to give protection in the event of a short in the bridge rectifier. In the Figure 12 circuit the load is placed in the ac side of the bridge, which does not need fuse protection since the load itself will limit currents to a safe value in the event of a bridge failure.

A pair of SCRs can easily be wired in

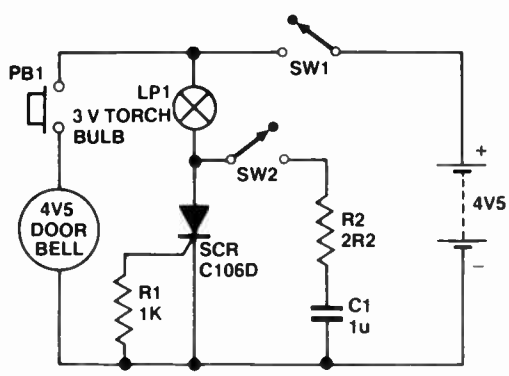

Figure 9. Rate-effect demonstration circuit.

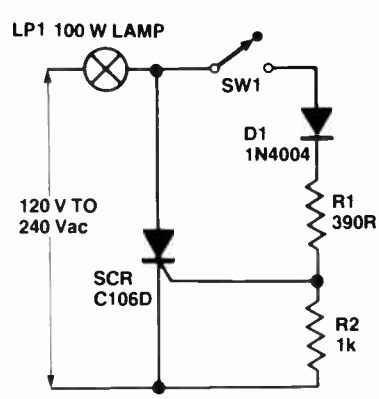

Figure 10. Line driven halfwave on/off circuit.

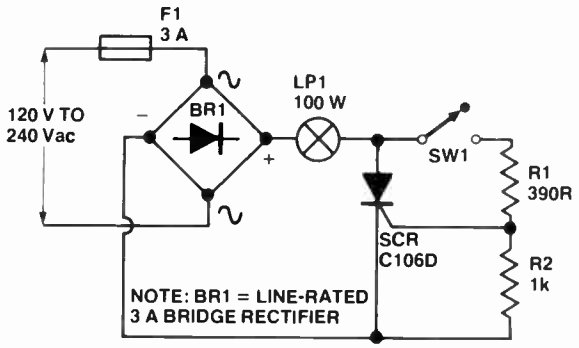

Figure 11. Fullwave on/off circuit.

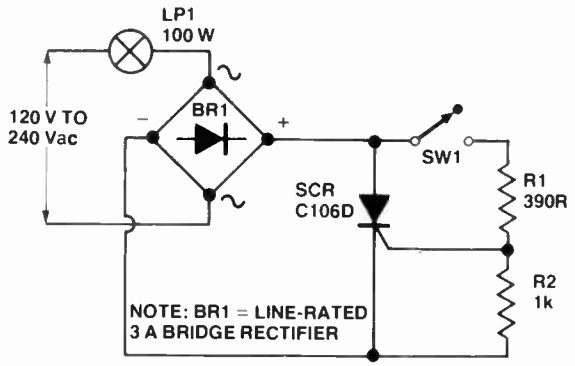

Figure 12. Alternative connection for fullwave on/off circuit.

inverse parallel and used to give fullwave power control without the use of additional rectification. In reality, however, a far more effective way of obtaining fullwave power control is to use a triac in place of the SCRs. Let's now look at triac basics.

The triac: basic theory

A triac can be regarded as being equal to two conventional SCRs connected in inverse parallel within a single three-terminal package, but so arranged that they share a single gate terminal. The triac acts as a solid state power switch that can conduct current in either direction and can be switched from the off to the on state by a gate signal of either polarity.

Figure 13a shows the triac symbol and Figure 13b shows a basic connection for using the device as an ac power switch. The load is wired in series with the triac's main terminals, the combination being wired directly across the ac power line. By closing SW1 dc gate drive can be applied to the triac. Referring to Figure 13b, the basic characteristics of the triac are as follows:

(1) Normally, with no gate signal applied, the triac is off and acts (between MT1 and MT2) like an open circuit switch.

(2) If MT2 is appreciably positive or negative relative to MT1 the triac can be turned on (so that it acts like a closed switch) by applying a gate signal via SW1. The device takes only a few microseconds to turn on. A saturation potential of one or two volts is developed across the triac in the on mode. Once the triac has turned on it self-latches and remains on so long as main-terminal current continues to flow. Only a brief pulse of gate current is thus needed to turn the triac on.

(3) Once the triac has self-latched the gate loses control and the triac can only be turned off again by reducing its main-terminal current below a minimum holding value. When the triac is used as an ac power switch therefore, turn-off occurs automatically near the zero-crossing point at the end of each half-cycle as the main-terminal currents fall to zero.

(4) The triac can be turned on by either a positive or negative gate signal,

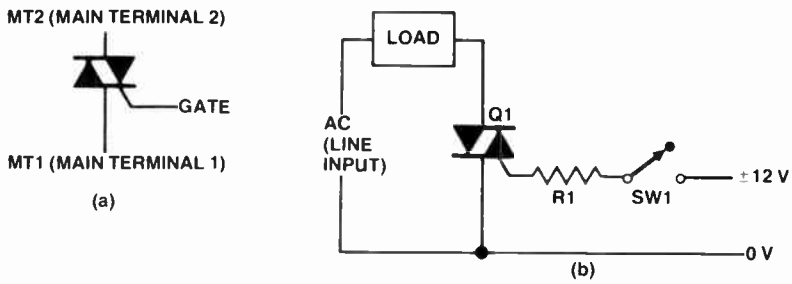

Figure 13. (a) Triac symbol and (b) basic triac circuit with dc gate drive.

irrespective of the polarities of the main-terminal voltages. The device thus has four possible triggering modes or 'quadrants', signified as follows:

I + Mode, MT2 current = + ve, I_{gate} = + ve.
I − Mode, MT2 current = + ve, I_{gate} = − ve.
III + Mode, MT2 current = − ve, I_{gate} = + ve.
III − Mode, MT2 current = − ve, I_{gate} = − ve.

Gate sensitivities in the I+ and III− modes are approximately equal and about twice as high as in the I− and III+ modes.

(5) Triacs can handle very high surge or non-repetitive currents. Typically, a device with a 10 A RMS rating may be able to handle a single-cycle, non-repetitive 50 Hz surge current of 100 amps!

Figure 14 shows basic details of a limited range of popular triacs. In most applications this limited information is sufficient for user needs. Let's now move on and look at some basic ways of using the triac.

The triac: basic circuits

Figure 15 shows the practical circuit of a simple dc-triggered triac power switch in which the dc supply is derived via step-down transformer T1. When SW1 is open, no current flows to the gate of the triac, which is thus off. When SW1 is closed, gate drive is

DEVICE TYPE NO.	PIV RATING	TOTAL CURRENT RATING RMS	V_{GT} (max)	I_{GT} (max)	I_H (max)
C206D	400 V	3 A	2 V	5 mA	30 mA
2N6073	400 V	4 A	2.5 V	30 mA	70 mA
C226D	400 V	8 A	2.5 V	50 mA	60 mA
SC146D	400 V	10 A	2.5 V	50 mA	75 mA
TIC246D	400 V	15 A	2.5 V	50 mA	50 mA

Figure 14. Basic details of some popular triacs.

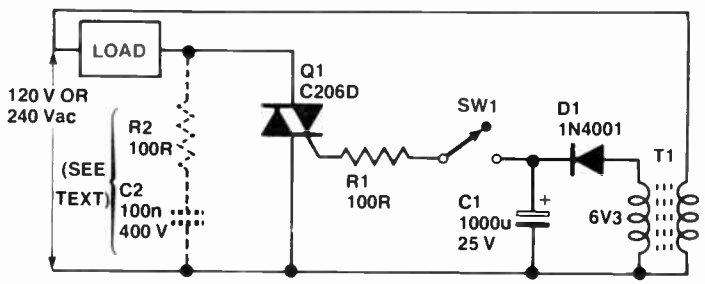

Figure 15. Simple ac power switch with dc gate triggering.

applied to the triac so it and the load are driven on. If an inductive load, such as a motor, is used in this circuit the R2-C2 'snubber' network must be wired in place as indicated to prevent false-triggering by rate-effects.

Note in the Figure 15 circuit that the dc side of the circuit is connected directly to one side of the mains and is thus 'live'. This snag is overcome in the UJT-triggered isolated-input circuit of Figure 16. Here, the UJT (unijunction transistor) operates at several kilohertz and thus delivers roughly 50 trigger pulses to the gate of the triac — via isolation pulse transformer T1 — during each half-cycle of the ac power line waveform. Consequently, the triac is fired by the first trigger pulse occuring in each mains half-cycle, and this pulse occurs within a few degrees of the start of the half-cycle. The triac is thus turned on almost permanently when SW1 is closed and virtually full power is applied to the ac load. The trigger circuit is, however, fully isolated from the mains by transformer T2 and pulse transformer T1.

Figure 17 shows how the triac can be used as a simple line switch with line-derived triggering. With SW1 open, zero gate drive is applied so the triac and lamp are off. Suppose however, that SW1 is closed. At the start of each half-cycle the triac is off, so the full line voltage is applied to the gate via the lamp and R1. Shortly after the start of the half-cycle, enough drive is available to trigger the triac and the triac and lamp go on. As the triac goes on and self-latches it saturates and automatically removes the gate drive until the start of the next half-cycle, thus minimising the dissipation in R1.

Finally, Figure 18 shows how the above circuit can be modified to give either half-wave or fullwave operation. In the halfwave mode, the gate drive is applied via D1, so the triac triggers on positive half-cycles only. In the fullwave mode the triac triggers on both positive and negative half-cycles, as in the case of the Figure 17 circuit.

Phase-triggered power control

The SCR and triac circuits that we have looked at so far have all been designed to give a simple on-off form of power control. These devices can easily be used to give fully-variable power control in ac circuits and are widely used in lamp dimmers and electric motor speed controllers, etc. The most widely used system of ac variable power control is known as the 'phase triggering' system.

Figure 19 illustrates the principle of phase-triggering using a triac as the power control element. Here, instead of the triac being triggered 'directly' from the ac power line, it is triggered via a variable phase-delay network that is interposed between the power line and the triac gate. Thus, if the triac is triggered 10° after the start of each half-cycle, almost the full available line power is fed to the load. If the triac is triggered 90° after the start of each half-cycle, only half of the available line power is fed to the load. Finally, if the triac is triggered 170° after the start of each half-cycle, only a very small part of the available power is fed to the load.

The three most popular methods of obtaining variable phase-delay triggering are to use either a line-synchronised UJT, a special-purpose IC, or to use a diac plus RC network in the basic configuration shown in Figure 20.

The diac can be regarded as a bilateral threshold switch. When connected across a voltage source, it acts like a high impedance until the applied voltage rises to about 35 volts, at which point it switches into a low impedance state and remains there until the applied voltage falls to about 30 volts, at which point it reverts to the high impedance mode. It then stays in this state until the applied voltage rises back to 35 volts, at which point the process repeats.

In the Figure 20 circuit, in each mains half-cycle the R1C1 network applies a variable phase-delayed version of the mains waveform to the triac gate via the diac, and

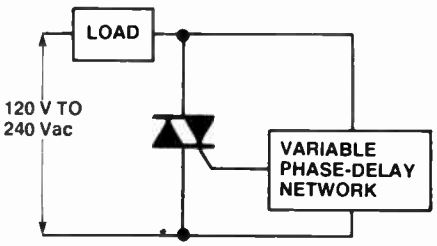

Figure 19a. Variable phase-delay 'switch'.

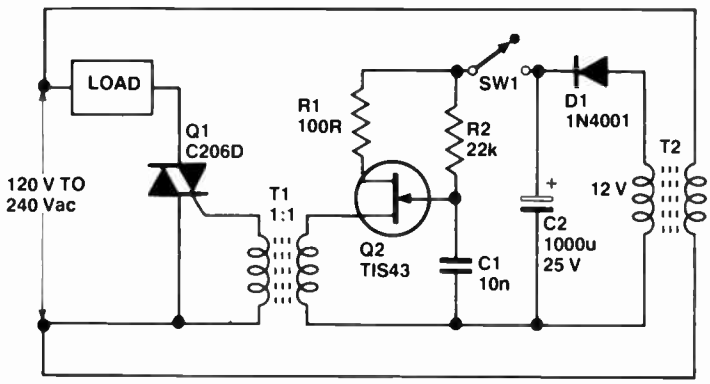

Figure 16. UJT-triggered isolated-input ac power switch.

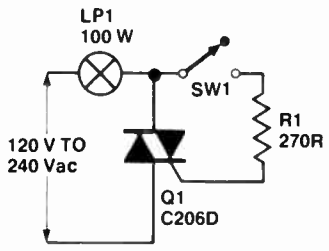

Figure 17. Line-triggered triac switch.

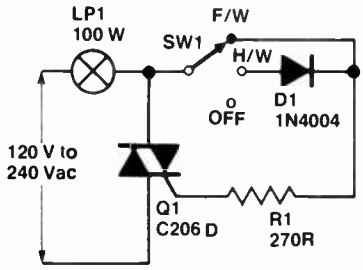

Figure 18. Three-way line switch.

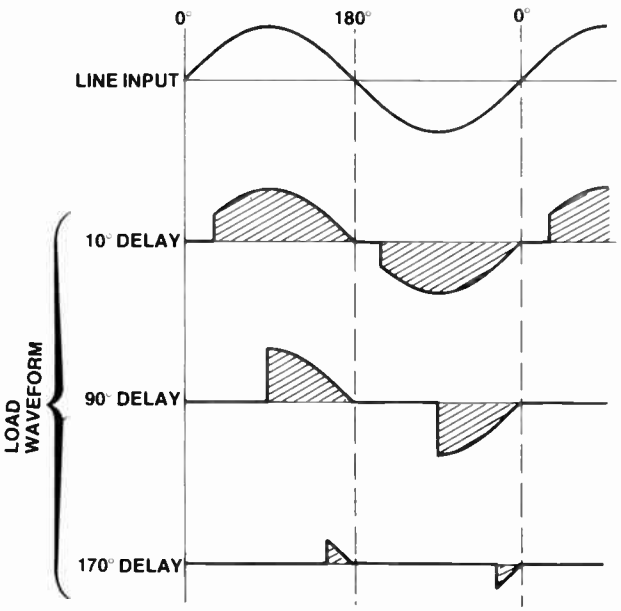

Figure 19b. Variable phase-delay waveforms.

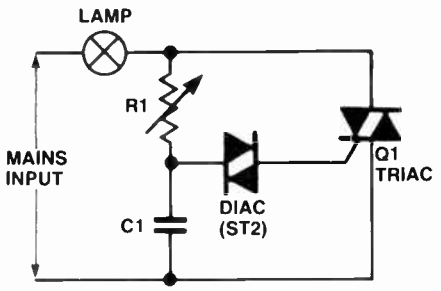

Figure 20. Basic 'diac-type' variable phase-delay lamp dimmer circuit.

each time the C1 voltage rises to 35 volts the diac fires and delivers a trigger pulse to the triac gate, thus turning the triac on, simultaneously applying power to the lamp load and removing the drive from the RC network. The mean power to the load (integrated over a full half-cycle period) is thus fully variable from near-zero to maximum, via R1.

Radio frequency interference (RFI)

You can see from the Figure 19 waveforms that, each time the triac is gated on, the load current changes abruptly (in a few microseconds) from zero to a value determined by the load resistance and the instantaneous mains voltage. This action generates radio frequency interference (RFI). The RFI is greatest when the triac is triggered at 90°, and is least when the triac is triggered close to the 0° and 180° 'zero-crossing' points of the mains waveform.

In lamp dimmer circuits, where there may be considerable lengths of mains cable between the triac and the lamp load, this RFI may be offensive as it will be widely radiated, interfering with radio and television receivers and other appliances. In practical lamp dimmers the circuit is usually provided with an LC RFI suppression network, as shown in Figure 21, to overcome this problem. (Note: in Figure 21 the values in brackets are applicable to 120 V mains operation.)

Zero-crossing techniques

When high power loads, such as electric heaters, are driven from triac circuitry special techniques must be used to minimise RFI. Even if the triac is used as a simple on-off switch in such applications, a 'spurt' of RFI will be generated each time the switch is turned on, and will be of maximum amplitude if the instantaneous phase delay happens to be 90° at the moment of turn-on. RFI problems can be eliminated in high-power applications by using the synchronous or 'zero-crossing' gating technique illustrated in Figure 22.

Here, a low power 12 volt dc supply is generated directly from the mains via R1-D1-ZD1 and C1. A simple zero-crossing detector network (a couple of transistors) is

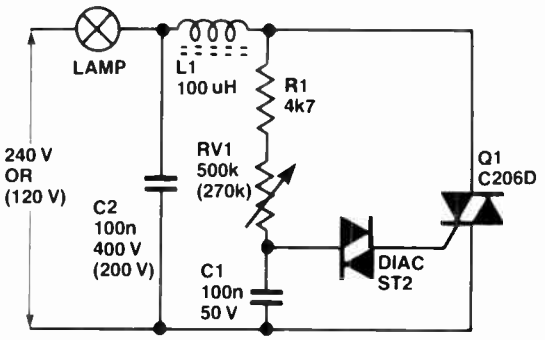

Figure 21. Practical lamp dimmer with RFI-suppression.

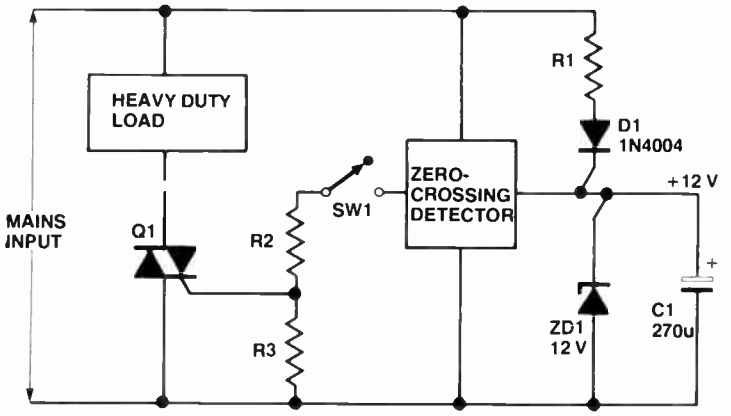

Figure 22. Basic synchronous or 'zero-crossing' mains power switch.

connected directly across the mains and controls the passage of current from C1 to SW1 in such a way that the C1 current is made available for only 5° or so on either side of each zero-crossing point of the mains waveform. Thus, if SW1 is closed, a pulse of gate current is fed to the triac at the start of each half-cycle of mains voltage, at which point the mains voltage is close to zero, so the triac always generates minimal RFI as it turns on.

The 'zero-crossing' technique can be used to provide RFI-free variable power control in high-power loads, such as electric heaters, by replacing SW1 of Figure 22 with a variable mark space-ratio waveform generator so that a variable integral number of complete mains power cycles are alternately fed or not fed to the load.

Figure 23 illustrates the basic principle, in which the total integral period is equal to eight mains cycles. Thus, if the power is alternately switched on for four cycles and off for four cycles, the mean load power is equal to half of the total available power, and if the power is on for one cycle and off for seven cycles, the mean power is equal to only one eighth of the total available power.

In the next edition of 'Circuit File' we'll look at some practical 'zero-crossing' and integral-cycle power controllers, together with a variety of lamp dimmers and motor-speed controllers. ●

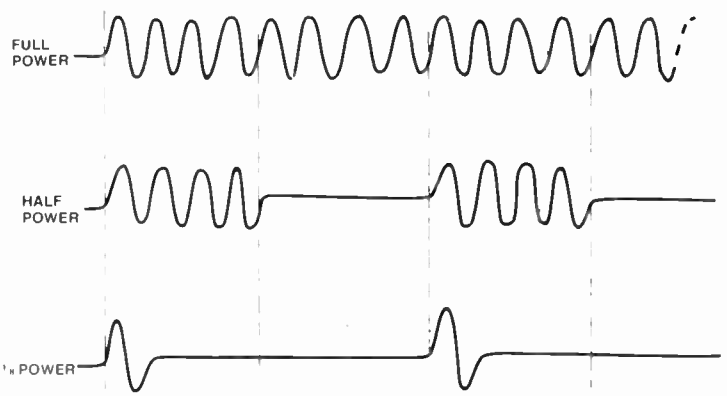

Figure 23. Integral-cycle power control waveforms.

SCRs, triacs and power control

Part two of Ray Marston's short series on these useful devices. In this part, he covers power switch circuits, electric heater controllers, lamp dimmers and motor speed controllers. A whole stack of applications circuits are given — as usual.

Ray Marston

PART ONE of this feature dealt at length with the fundamentals of SCR and triac circuitry, and gave particular attention to the principles of synchronous and non-synchronous triggering. Here, we present a stack of practical circuits for use on 240 Vac mains supplies. In these designs you simply select the triac or SCR rating to suit your own particular application.

Let's start off, then, by looking at some practical triac power switch designs for use in basic on/off ac power line switching applications.

TRIAC POWER SWITCHES

Non-synchronous designs

As was explained in part 1, triacs can be triggered (turned on) either synchronously or non-synchronously with the mains voltage. Synchronous circuits *always* turn on at the same point in each mains half-cycle (usually just after the zero-crossing point), and usually generate minimal RFI. The trigger points of non-synchronous circuits are not synchronised to a fixed point of the mains cycle, and the circuits may generate significant RFI, particularly at the point of initial turn-on. Triac turn-off is always automatically synchronised to the mains, as the device's main-terminal currents fall below the minimum-holding value at the end of each mains half-cycle.

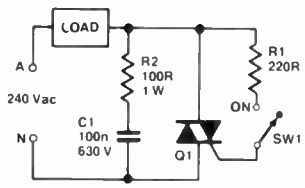

Figure 1. Simple ac power switch, ac line triggered.

Figures 1 to 8 show a variety of non-synchronous triac power switch circuits which can be used in basic on/off line switching applications. The action of the Figure 1 circuit was explained last month, being such that the triac is gated on from the mains via the load and R1 shortly after the start of each mains half-cycle when SW1 is closed, but remains off when SW1 is open. Note, in this circuit, that the trigger point is *not* synchronised to the mains when SW1 is initially closed, but becomes synchronised on all subsequent half-cycles.

Figure 2 shows how the triac can be triggered via a mains-derived dc supply. C1 is charged to +10 V on each positive half-cycle of the mains via R1-D1, and the C1 charge triggers the triac when SW1 is closed. Note that all parts of this circuit are 'live', making it difficult to interface to external electronic control circuitry.

Figure 3 shows how the above circuit can be modified so that it can easily be interfaced to external control circuitry. SW1 is simply

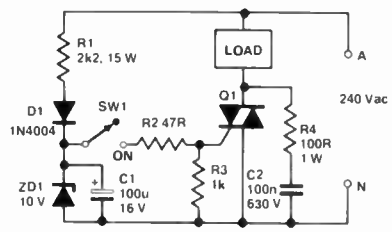

Figure 2. Ac power switch with line-derived dc triggering.

replaced by transistor Q2, which in turn is driven from the 'phototransistor' side of an inexpensive optocoupler. The 'LED' side of the optocoupler is driven from a 5 V or greater dc supply via R4. The triac turns on only when the external supply is connected via SW1.

Optocouplers have typical insulation potentials of 500 to several thousand volts, so the external circuit is fully isolated from the mains, and can easily be designed to give any desired form of remote operation of the triac by replacing SW1 with an electronic switch.

Figure 4 shows an interesting variation of the above circuit. In this case the triac is ac-triggered on each half-cycle of the mains via C1-R1 and back-to-back zeners ZD1-ZD2.

Note that the mains impedance of C1 determines the magnitude of the triac gate current but that C1 dissipates virtually no power. Bridge rectifier D1 to D4 is wired

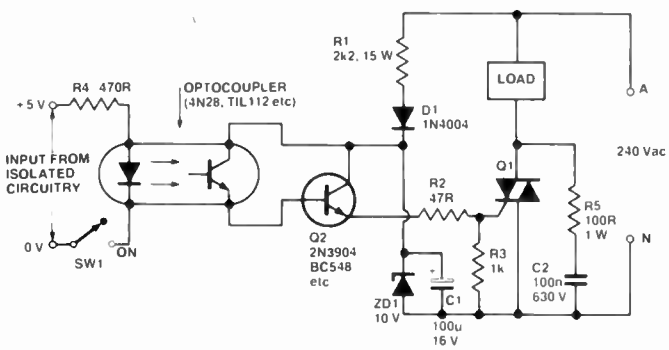

Figure 3. Isolated-input (optocoupled) ac power switch, dc triggered.

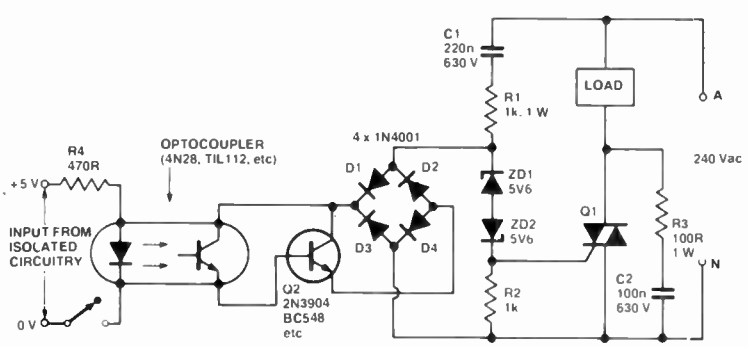

Figure 4. Isolated-input ac power switch, ac triggered

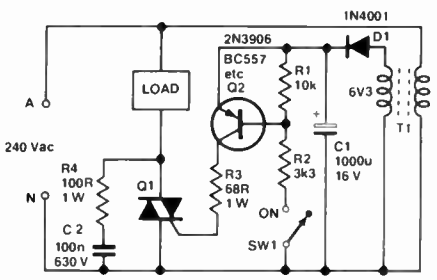

Figure 5. Ac power switch with transistor-aided dc triggering.

Figures 5 and 6 show a couple of ways of triggering the triac via a transformer-derived dc supply and a transistor-aided switch. In the Figure 5 circuit, the transistor and the triac are both driven on when SW1 is closed, and are off when SW1 is open.

In practice SW1 can easily be replaced by an electronic switch, enabling the triac to be operated by heat, light, sound, etc. Note however, that the whole of the Figure 5 circuit is 'live'.

Figure 6 shows how the circuit can be modified for optocoupler operation, so that it can be activated via fully isolated external circuitry.

Synchronous designs

Synchronously-triggered triac circuits *always* turn on at the same point in each mains half-cycle. Usually, the trigger point occurs just after the 'zero-crossing' point at the start of each half-cycle, in which case the triac generates absolutely minimal RFI.

Figures 9 to 18 show a number of on/off power switching circuits that use this form of triggering.

Figure 9 shows the practical circuit of a 'transistorised' synchronous line switch that is triggered near the zero-voltage crossover points of the mains. The triac gate trigger

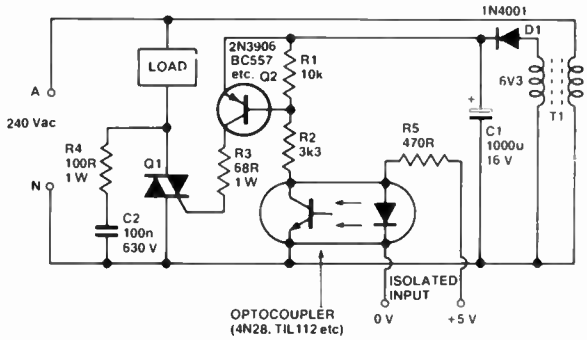

Figure 6. Isolated-input ac power switch with dc triggering.

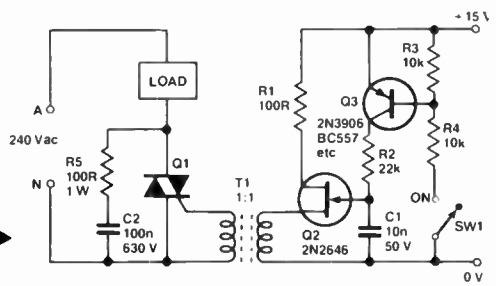

Figure 7. Isolated-input (transformer-coupled) ac power switch.

across the ZD1-ZD2-R2 network and is loaded by Q2. When Q2 is off, the bridge is effectively open and the triac turns on shortly after the start of each mains half-cycle; when Q2 is on, a near-short appears across ZD1-ZD2-R2 inhibiting the triac gate circuit, and the triac is off.

Transistor Q2 is actually driven via the optocoupler from an isolated external circuit, so the triac is normally on but turns off when SW1 is closed.

CONSTRUCTION OF T1, FIGS 7, 8

The core is a 30 mm long piece of 9.6 mm dia. ferrite aerial rod. The primary and secondary are each 30 turns of 0.4 mm dia. enamelled wire (26 B&S) closewound on the centre 15 mm of the core. Use two layers of plastic insulation tape between the two windings and cover complete unit with a further two layers of tape. Bring the primary and secondary leads out opposite ends of the core. Mark the starts of each winding (spots on circuit).

Finally, to complete this section, Figures 7 and 8 show a couple of alternative ways of obtaining triacs triggering from a fully isolated external circuit. In these two circuits the triggering action is obtained from a unijunction (UJT) oscillator (Q2) which operates at a frequency of several kilohertz and has its output pulses fed to the triac gate via pulse transformer T1, which provides the desired 'isolation'.

In the Figure 7 circuit, Q3 is wired in series with the UJT's main timing resistor so the UJT and triac turn on only when SW1 is closed. In the Figure 8 circuit, Q3 is wired in parallel with the UJT's main timing capacitor so the UJT and triac turn on only when SW1 is open. In both of these circuits, SW1 can easily be replaced by an electronic switch.

current is obtained from a 10 Vdc supply that is derived from the mains via R1-D1-ZD1 and C1, and this supply is switched to the gate via Q5, which in turn is controlled by SW1 and zero-crossing detector Q2-Q3-Q4.

The action of Q5 is such that it can only turn on and conduct gate current when SW1 is closed and Q4 is off. The action of the zero-crossing detector is such that Q2 or Q3 are driven on whenever the instantaneous mains voltage is positive or negative by more than a volt or two (depending on the setting of RV1), thereby driving Q4 on via R3 and inhibiting Q5.

Thus, gate current can only be fed to the triac when SW1 is closed and the instantaneous mains voltage is within a few volts of zero. The circuit thus provides minimal switching RFI.

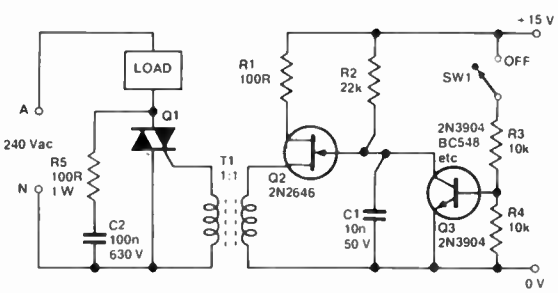

Figure 8. Isolated-input ac power switch.

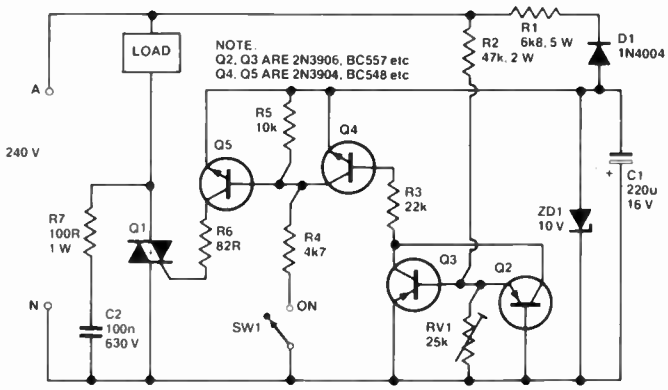

Figure 9. Transistorised synchronous line switch

Figure 10 shows how the circuit can be modified so that the triac can only turn on when SW1 is open. Note in both of these circuits that, since only a narrow pulse of gate current is sent to the triac, the *mean* consumption of the dc supply is very low (1 mA or so). Also note that SW1 can easily be replaced by an electronic switch to give automatic operation via heat, light, etc, or by an optocoupler to give fully isolated operation from external circuitry.

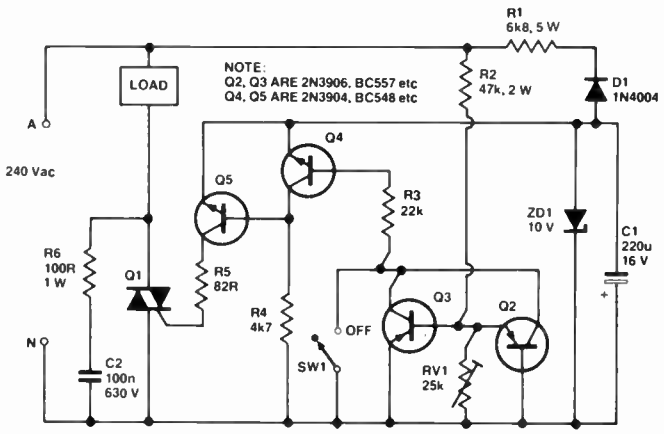

Figure 10. Alternative version of the 'transistorised' line switch.

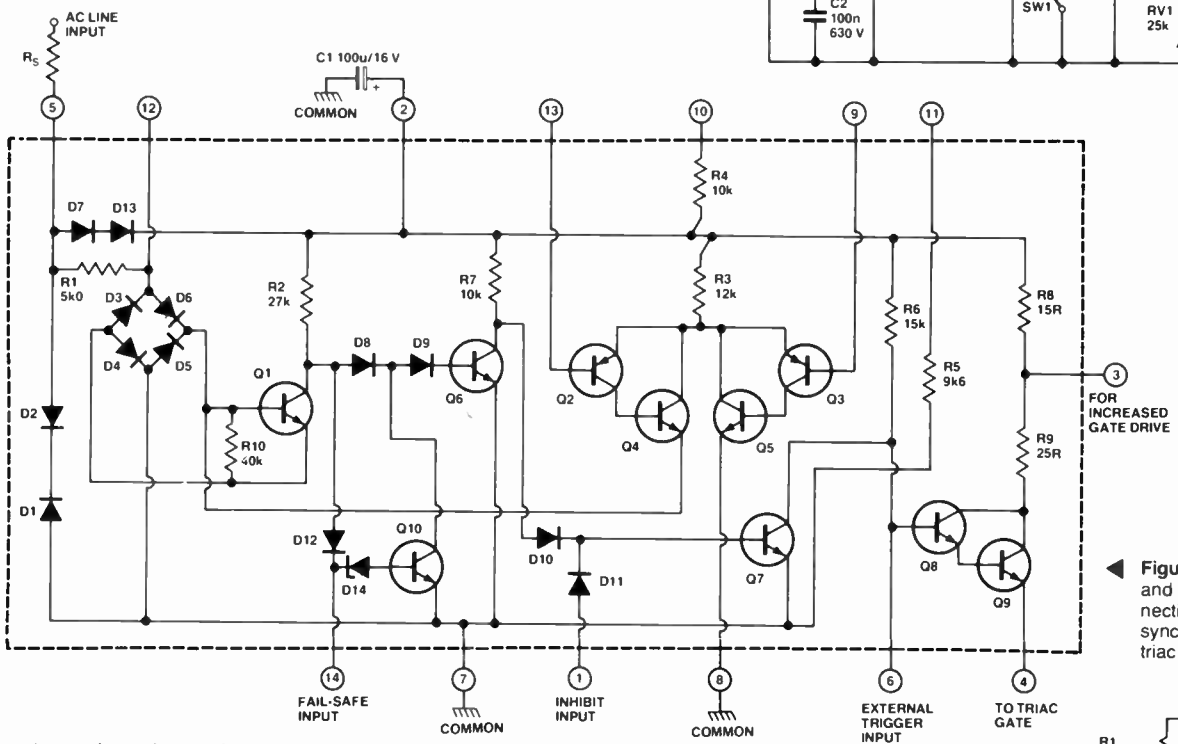

Figure 11. Internal circuit and minimal external connections of the CA3059 synchronous 'zero-voltage' triac driver.

A number of special-purpose synchronous zero-crossover triac-gating ICs are available, the best known examples being the CA3059 and the TDA1024. These devices incorporate mains-derived dc power supply circuitry, a zero-crossing detector, triac gate drive circuitry, and a high gain differential amplifier/gating network.

Figure 11 shows the internal circuitry of the CA3059, together with its minimal external connections. Mains power is connected to pins 5 and 7 via limiting resistor R_S (22k, 5 W or three 68k, 1 W resistors in parallel).

Diodes D1 and D2 act as back-to-back zeners and limit the pin 5 voltage to ±8 V. On positive half-cycles D7 and D13 rectify this voltage and generate 6.5 V across the 100u capacitor connected to pin 2. This capacitor stores enough energy to drive all internal circuitry and provide adequate triac gate drive, with a few milliamps of spare drive available for powering external circuitry if needed.

Bridge rectifier D3 to D6 and transistor Q1 act as a zero-crossing detector, with Q1 being driven to saturation whenever the pin 5 voltage exceeds ±3 V.

Gate drive to an external triac can be made via the emitter (pin 4) of the Q8-Q9 Darlington pair, but is available only when Q7 is turned off. When Q1 is turned on (pin 5 greater than ±3 V) Q6 turns off through lack of base drive, so Q7 is driven to saturation via R7 and no triac gate drive is available at pin 4. Triac gate drive is available only when pin 5 is close to the 'zero-voltage' mains value. When gate drive is available, it is delivered in the form of a narrow pulse centred on the crossover point, with pulse power supplied via C1.

The CA3059 incorporates a differential amplifier or voltage comparator, built around Q2 to Q5, for general purpose use. Resistors R4 and R5 are externally available for biasing one side of the amplifier. The emitter current of Q4 flows via the base of Q1 and can be used to disable the triac gate drive (pin 4) by turning Q1 on.

The configuration is such that the gate drive can be disabled by making pin 9 positive relative to pin 13. The drive can also be disabled by connecting external signals to pin 1 and/or pin 14.

Figures 12 and 13 show how the CA3059 can be used to give manually-controlled 'zero-voltage' on/off switching of a triac. These two circuits use SW1 to enable or disable the triac gate drive via the internal differential amplifier of the IC. Remember, the drive is enabled only when pin 13 is biased above pin 9.

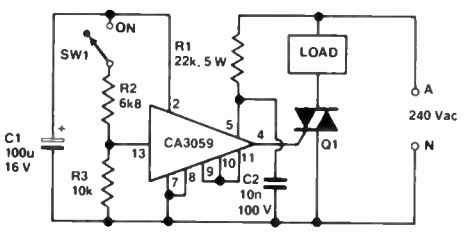

Figure 12. Direct-switched IC-gated 'zero-voltage' line switch.

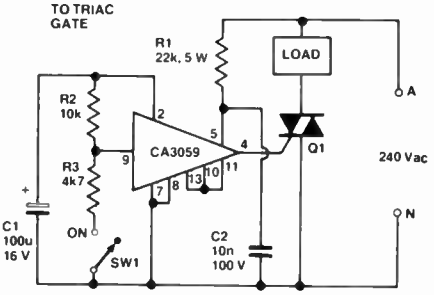

Figure 13. An alternative method of direct-switching the CA3059 IC.

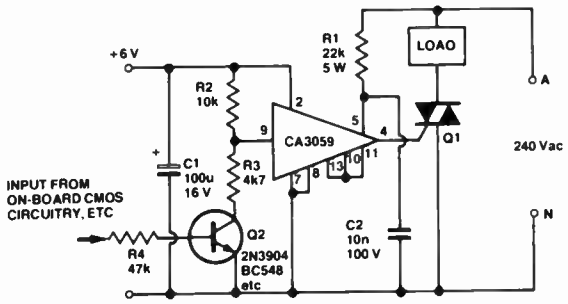

Figure 14. Method of transistor-switching the CA3059 via on-board CMOS circuitry, etc.

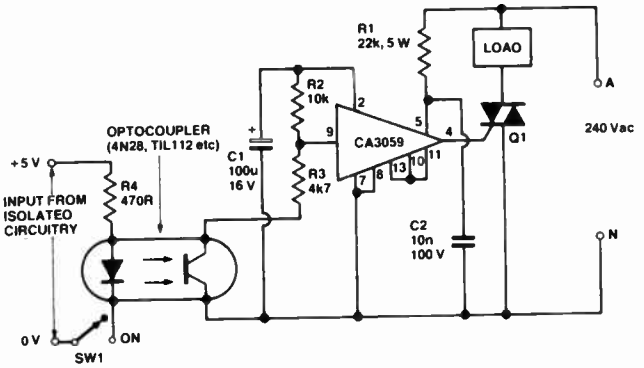

Figure 15. Method of remote-switching the CA3059 via an optocoupler.

In the Figure 12 circuit, pin 9 is biased at half-supply volts and pin 13 is biased via R2-R3 and SW1, and the triac turns on only when SW1 is closed.

In Figure 13, pin 13 is biased at half-supply and pin 9 is biased via R2-R3 and SW1, and the triac again turns on only when SW1 is closed. In both of these circuits, SW1 handles a maximum potential of 6 V and maximum current of only 1 mA or so.

Note, in these designs, that capacitor C2 is used to apply a slight phase delay to the pin 5 'zero-voltage detecting' terminal, and causes the gate pulses to occur slightly after (rather than to 'straddle') the zero-voltage point.

Note in the Figure 13 circuit that the triac can be turned on by pulling R3 low or turned off by letting R3 float. Figures 14 and 15 show how this simple fact can be put to use to extend the versatility of the basic circuit.

In Figure 14, the triac can be turned on and off by transistor Q2, which in turn can be activated by on-board CMOS circuitry (such as one-shots, astables, etc) that are powered from the 6 V pin 2 supply.

In Figure 15, the circuit can be turned on and off by fully-isolated external circuitry via an inexpensive optocoupler, which needs an input in excess of only a couple of volts to turn the triac on.

Alternatively, Figure 16 shows how the TDA1024 can be used in place of the CA3059 to give either directly-switched or opto-coupled 'zero-voltage' triac control.

Finally, to complete this section, Figures 17 and 18 show a couple of ways of using the CA3059 so that the triac operates as a light-sensitive 'dark-operated' power switch. In these two designs the built-in differential amplifier of the IC is used as a precision voltage comparator that turns the triac on or off when one of the comparator input voltages goes above or below the other.

Figure 17 is the circuit of a simple dark-activated power switch. Here, pin 9 is tied to half-supply volts and pin 13 is controlled via the R2-RV1-LDR-R3 potential divider.

Under bright conditions the LDR has a low resistance, so pin 13 is below pin 9 and the triac is disabled. Under dark conditions the LDR has a high resistance, so pin 13 is above pin 9 and the triac is enabled and power is fed to the load. The precise threshold level of the circuit can be preset via RV1.

Figure 18 shows how a degree of hysteresis or 'backlash' can be added to the above circuit, so that the triac does not switch annoyingly in response to small changes (passing shadows, etc) in ambient light level. The hysteresis level is controlled via R3, which can be selected to suit particular applications.

ELECTRIC-HEATER CONTROLLERS

Non-synchronous circuits

Triacs can easily be used to give automatic room-temperature control by using electric heaters as the triac loads and either thermostats or thermistors as the thermal feedback elements.

Two basic methods of heater control can be used, either simple on/off power switching or fully automatic proportional power control. In the former case, the heater switches fully on when the room temperature falls below a preset level and turns off when the temperature rises above the preset level.

In the latter case, the *mean* power to the heater is automatically adjusted so that, when the room temperature is at the precise preset level, the heater output power self-adjusts to balance the thermal losses of the room.

Because of the high power requirements of electric heaters, special care must be taken in the design of triac controllers to keep RFI generation to minimal levels. Two options are open to the designer, to use either continuous dc gating of the triac, or to use synchronous pulsed gating.

The advantage of dc gating is that, in basic on/off switching applications, the triac generates zero RFI under normal (on) running conditions. The disadvantage is that the triac may generate very powerful RFI as it is initially switched from the off to the on condition.

The advantage of synchronous gating is that no high-level RFI is generated as the triac transitions from the off to the on condition. The disadvantage is that the triac generates continuous very-low-level RFI under normal (on) running conditions.

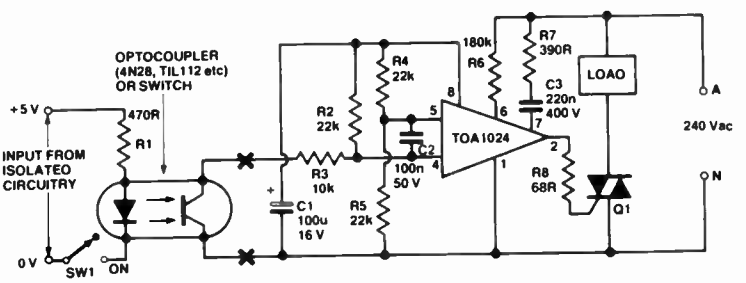

Figure 16. The TDA1024 used to give either directly switched or optocoupled 'zero-voltage' triac control.

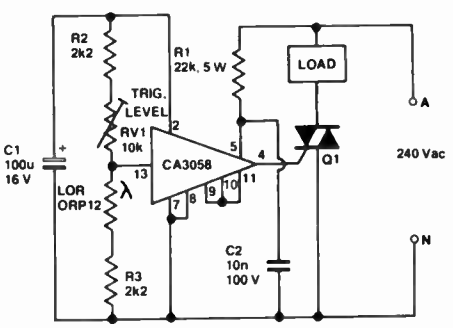

Figure 17. Basic 'dark-activated' zero-voltage switch.

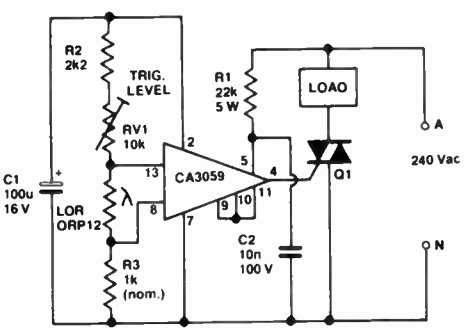

Figure 18. Dark-activated zero-voltage switch with hysteresis provided via R3.

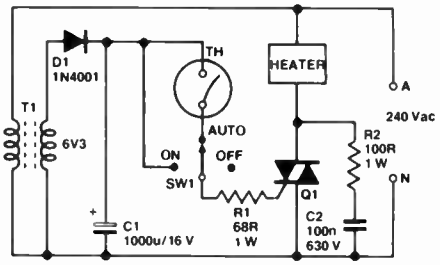

Figure 19. Heater controller with thermostat-switched dc gating.

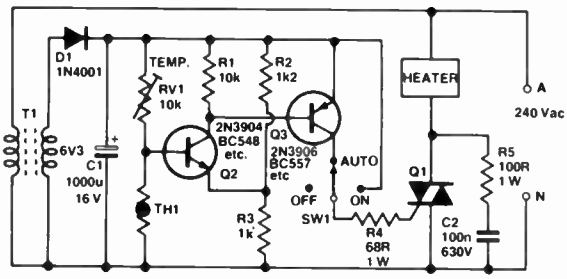

Figure 20. Heater controller with thermistor-switched dc gating.

Figures 19 and 20 show a couple of dc-gated heater-controller circuits, in which the dc supply is derived via T1-D1 and C1, and the heater can be controlled either manually or automatically via SW1. The Figure 19 circuit is auto-controlled via a thermostat.

The Figure 20 circuit on the other hand, is controlled by negative temperature coefficient (NTC) thermistor TH1 and transistors Q2-Q3, and calls for some explanation. RV1-TH1-R2-R3 are used as a thermal bridge, with Q2 acting as the bridge-balance detector. RV1 is adjusted so that Q2 just starts to turn on as the temperature *falls* to the desired preset level. Below this level, Q2-Q3 and the triac are all driven hard on, and above this level all three components are cut off.

Note, in the Figure 20 circuit that, since the gate-drive polarity is always positive but the triac main-terminal current is alternating, the triac is gated alternately in the I+ and III+ modes (or quadrants) and that the gate sensitivities are quite different in these two modes.

Consequently, when the temperature is well below the preset level Q3 is driven hard on and the triac is gated in both quadrants and gives full power drive to the heater, but when the temperature is very close to the preset value Q3 is only 'gently' driven on, so the triac is gated in the I+ mode only and the heater operates at only half of maximum power drive. The circuit thus gives fine control of temperature.

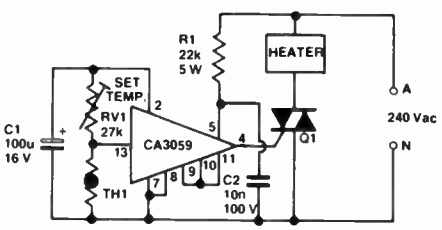

Figure 21. Heater controller with thermistor-regulated zero-voltage switching.

Synchronous circuits

Figure 21 shows how a CA3059 can be used to make an automatic thermistor-regulated synchronous electric heater controller. The circuit is similar to that of the 'dark-activated' power switch of Figure 17, except that NTC thermistor TH1 is used as the feedback sensing element.

The circuit is capable of maintaining room temperature within a degree or so of the value via RV1.

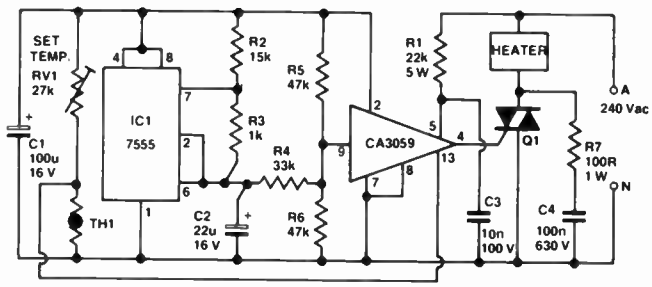

Figure 22. Heater controller giving integral-cycle precision temperature regulation.

Finally, to complete this 'heater controller' section, Figure 22 shows the circuit of a proportional heater controller which is capable of maintaining room temperatures within 0.5°C (depending on sensor placement). In this circuit a thermistor controlled voltage is applied to the pin 13 side of the CA3059's comparator and a repetitive 30 ms ramp signal, centred on half-supply volts, is applied to the pin 9 side of the comparator from CMOS astable IC1.

The action of the circuit is such that the triac is synchronously turned fully on if the ambient temperature is more than a couple of degrees below the preset level, or is cut fully off if the temperature is more than a couple of degrees above the preset level.

When the temperature is within a couple of degrees of the preset value however, the ramp waveform comes into effect and synchronously turns the triac on and off (in the 'integral cycle' mode) once every 300 ms, with a mark/space ratio that is proportional to the temperature differential.

Thus, if the mark/space ratio is 1:1, the heater generates only half of maximum power, and if the ratio is 1:3 it generates only one quarter of maximum power.

The net effect of this action is that the heater output power self-adjusts to meet the room's heating requirements. When the room temperature reaches the preset value, the heater does not switch off, but generates just enough output power to match the thermal losses of the room, giving very precise temperature control.

LAMP DIMMER CIRCUITS

Triacs can be used to make lamp dimmers, which vary the brilliance of incandescent lamps, by using the phase-triggered power control principles described in part 1. The triac is turned on and off once in each mains half-cycle, the mark/space ratio controlling the mean power fed to the lamp. All such circuits require the use of a simple LC filter in the lamp feed line, to reduce RFI problems.

The three most popular methods of obtaining variable phase-delay triggering are to use either a diac plus RC phase delay network, or to use a line-synchronised variable delay UJT trigger, or to use a special purpose IC as the triac trigger.

Figure 23 shows the practical circuit of a diac-triggered lamp dimmer, in which R1-RV1-C1 provide the variable phase delay. This circuit is similar to that described in part 1, except for the addition of on/off switch SW1 which is ganged to RV1 and enables the lamp to be turned fully off.

A defect of the simple Figure 23 design is that it suffers from considerable control hysteresis or backlash. If the lamp is dimmed by increasing the RV1 value to 470k, it will not go on again until RV1 is reduced to about 400k, and it then burns at a fairly high

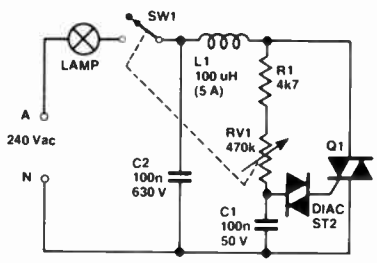

Figure 23. Practical circuit of a simple lamp dimmer.

brightness level. This 'backlash' is caused by the diac partially discharging C1 each time the triac fires.

The 'backlash' effect of the Figure 23 circuit can be reduced by wiring a 47R resistor in series with the diac, to reduce its discharge effect on C1. An even better solution is to use the gate slaving circuit of Figure 24, in which the diac is triggered from C2, which 'copies' the C1 phase delay voltage. But here, R2 protects C1 from discharging when the diac fires.

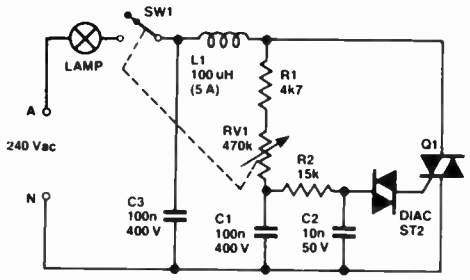

Figure 24. Improved lamp dimmer with gate slaving.

CONSTRUCTION OF L1, FIGS 23 TO 26

The core is a 30 mm long piece of 9.6 mm dia. ferrite aerial rod. Wind two layers of 20 turns, closewound, using the centre 15 mm of the core, with 0.63 mm dia. (22 B&S) enamelled wire. Cover with two layers of plastic insulation tape.

Finally, to complete this 'lamp dimmer' section, Figure 26 shows how a dedicated IC, the Siemens S566B 'Touch Dimmer' chip, can be used as a 'smart' lamp dimmer that can be controlled by either touch pads, push-button switches, or via an infra-red link.

The action of this chip, which gives a phase-delayed trigger output to the triac, is such that it alternately ramps up (increases brilliance) or ramps down (decreases brilliance) on alternate operations of the touch or pushbutton inputs, but 'remembers' and holds brilliance levels when the inputs are released.

The IC incorporates 'touch conditioning' circuitry, such that a very brief touch or push input causes the lamp to simply change state (from off to on, or vice versa), but a sustained (greater than 400 ms) input causes the IC to go into the ramping mode, in which the lamp power slowly ramps up from 3% to 97% of maximum and then down to 3% again, and so on.

The touch pads used with this circuit can be simple strips of conductive material; the operator is safely insulated from the mains voltage via R8 and R9.

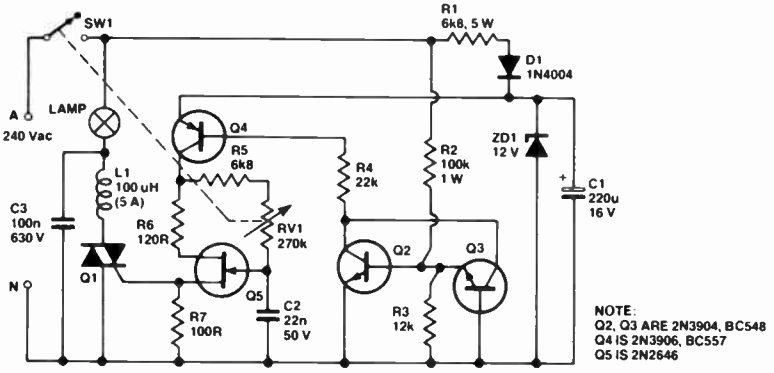

Figure 25. UJT-triggered zero-backlash lamp dimmer.

If absolutely zero backlash is needed, the UJT-triggered circuit of Figure 25 can be used. The UJT is powered from a 12 Vdc supply derived from the ac line via R1-D1-ZD1-C1. The UJT is synchronised to the mains via the Q2-Q3-Q4 zero-crossing detector network, the action being such that Q4 is turned on (applying power to the UJT) at all times other than when the mains is close to the zero-crossover point at the end and start of each mains half-cycle.

Thus, shortly after the start of each half-cycle, power is applied to the UJT circuit via Q4, and some time later (determined by R5-RV1-C2) a trigger pulse is applied to the triac gate via Q5. The UJT resets at the end of each half-cycle, and a new sequence then begins.

UNIVERSAL-MOTOR CONTROLLERS

Domestic appliances such as electric drills and sanders, sewing machines and food mixers, etc, are almost invariably powered by series-wound 'universal' electric motors (so called because they can operate from either ac or dc supplies).

When operating, these motors produce a back-emf that is proportional to the motor speed. The *effective* voltage applied to such motors is equal to the true applied voltage minus the back-emf. This fact results in a degree of self-regulation of the speed of the motors, since an increase in the motor loading tends to reduce the speed and back-emf, thereby increasing the effective applied voltage and causing the motor speed to return towards its original value.

Most 'universal' motors are designed to give single-speed operation. Triac phase-controlled circuits can easily be used to provide these motors with variable speed control. A suitable 'diac plus phase-delay' circuit is shown in Figure 27. This circuit is particularly useful for controlling lightly-loaded appliances such as food mixers, sewing machines, etc. However, you only get a limited range of control.

Electric drills and sanders are subject to very heavy load variations, and are not really suitable for control via the Figure 27 circuit. Instead, the variable speed-regulator circuit of Figure 28 should be used.

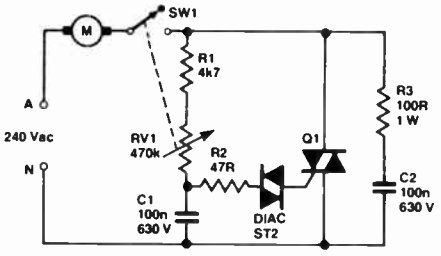

Figure 27. Universal-motor speed controller for use with lightly-loaded appliances (food mixers, sewing machines, etc).

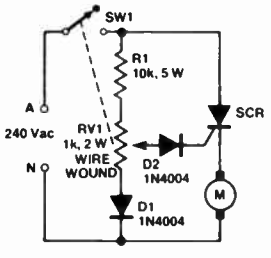

Figure 28. Self-regulating universal-motor speed controller for use with electric drills and sanders, etc.

This circuit uses an SCR as the control element and feeds half-wave power to the motor (this results in only a 20% or so reduction in available speed/power), but in the off half-cycles the back-emf of the motor is sensed by the SCR and used to give automatic adjustment of the next gating pulse, giving some speed regulation. The R1-RV1-D1 network provides only 90° of phase adjustment so all motor power pulses have minimum durations of 90° and provide very high torque.

At low speeds the circuit goes into a 'skip cycling' mode, in which power pulses are provided intermittently, to suit motor loading conditions. The circuit provides particularly high torque under low-speed conditions, but the motor 'chatters' somewhat. Like the previous circuit, only a limited range of control is provided. ●

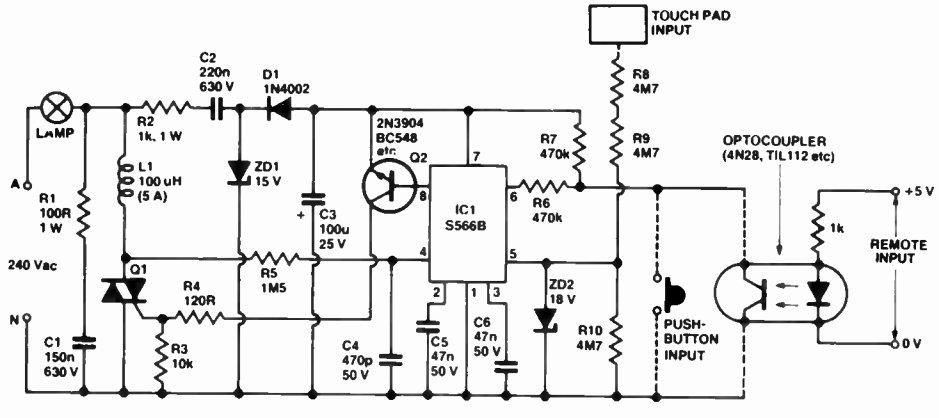

Figure 26. 'Smart' lamp dimmer controlled by a dedicated IC.

Circuit source guide

Here is a collection of circuits selected from the voluminous files of your Editor, Roger Harrison and some other sources. From this anthology you should be able to derive other circuits or assemble a system from a variety of 'blocks' to suit a particular application or solve a circuit problem. Applications covered range from audio to RF, timing to dc control, measurement to musical, etc. You may have seen some of these ideas before, but there are bound to be plenty you haven't.

This feature is intended for the experienced experimenter, and construction details are not given. While the circuits have been checked for accuracy and feasibility, they have not necessarily been built and tested. We are unable to answer queries on individual modifications or construction techniques.

OPTO

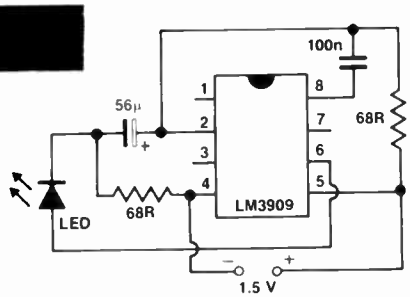

LED Light Booster

The LM3909 LED flasher IC is well known. It can be used to boost the brightness of ordinary LEDs by providing them with high current pulses at around 20 kHz — too fast for the eye to see flashing — giving an impression of increased brightness.

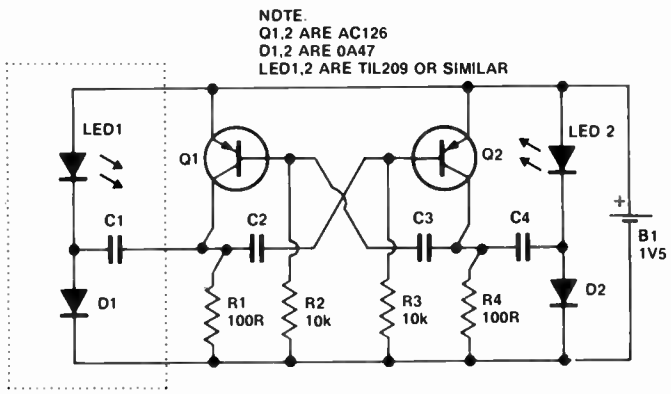

LEDs on 1.5 V Battery

As most LEDs require a forward voltage between 1.6 and 2.3 volts, it's difficult to power them from a 1.5 V battery. This circuit is an astable multivibrator and voltage doubler that boosts the voltage across the LEDs. To make the LEDs appear to be on continuously, C1 and C2 should be 47n, C3-C4 10µ. To make the LEDs flash alternately, C1-C2 should be about 100µ, C3-C4 should be about 10 times that. To operate a single LED, omit LED1, D1 and C1.

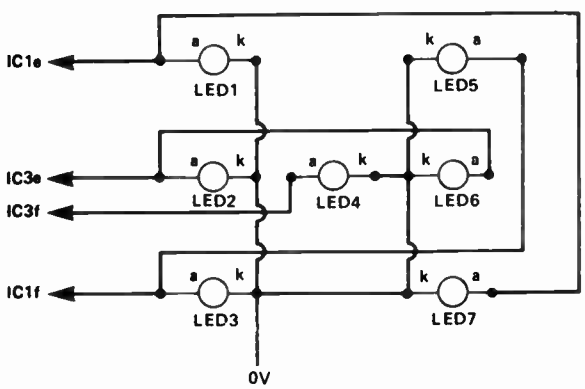

Die/dice Roller

A dice is more than one die — a six-sided solid with faces numbered 1 to 6 (i.e: a cube with numbered faces — you've all seen at least one). This is the electronic version. It comprises an oscillator and a counter with a 1 to 6 display output to LEDs. The LEDs can be arranged in the traditional die face, as shown.

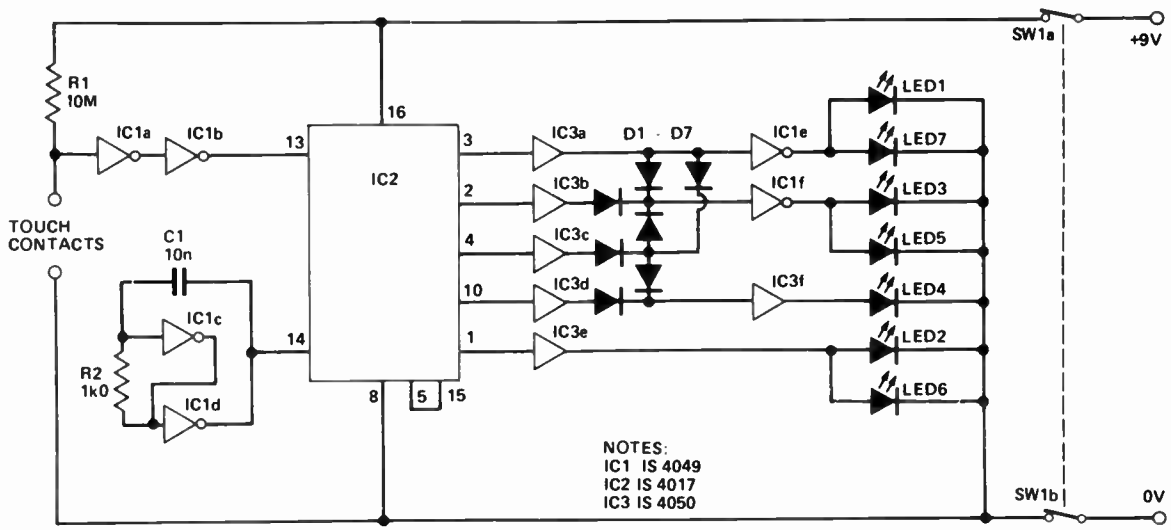

Wideband Amp with dc Feedback

For high, stable gain, a wideband amp requires several stages with multiple dc feedback paths. The two left hand and two right hand transistors here form common emitter common collector pairs, the common collector devices providing a high impedance load for the previous transistor and a low source impedance for the following stage. This reduces internal capacitive feedback. A CA3018 transistor array IC is used. The lower cutoff is determined by the capacitor values. Use low self-inductance metallised poly low voltage types.

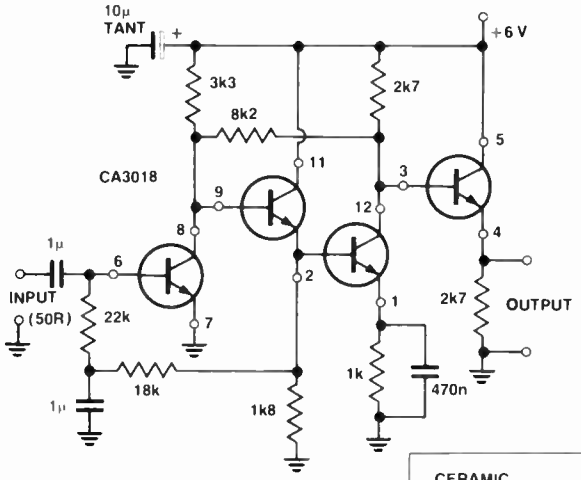

MIDBAND GAIN APPROX 50 dB
BANDWIDTH (−3 dB) 1 kHz-30 MHz.
MAX INPUT APPROX. 4 mV RMS

Wideband Power Amplifier

Well, not all that much power — in the milliwatt range, really. This uses the same configuration as the common-base/common-collector amp opposite, but power gain is provided and the circuit will deliver between about 10 and 180 mW output, depending on power supply voltage. The common collector stage provides current gain, transferring power to the load. This stage has to operate at a fairly high quiescent current. In fact, to get the bandwidth, both stages have to be operated at high quiescent current levels.

mon base stage compensates for the input capacitance of the common collector stage. The RFC in the emitter of the output device gives 'lift' at high requencies, helping broaden the bandwidth. R1 can be used to match the input to the source driving impedance. Two of these amplifiers can be cascaded to deliver around half a watt output with little decrease in overall bandwidth. Increase the 1µ capacitors to decrease lower cutoff.

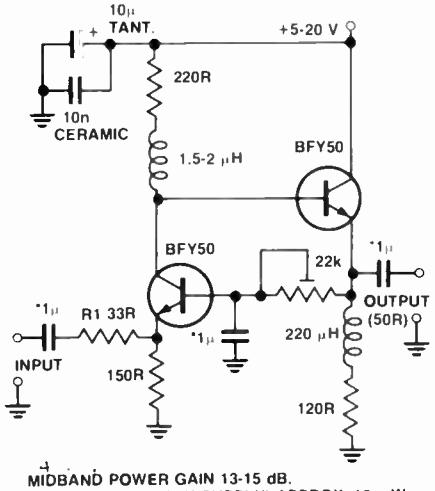

MIDBAND POWER GAIN 13-15 dB.
POWER OUTPUT (5 V SUPPLY) APPROX. 10 mW
(20 V SUPPLY) 180 mW; BANDWIDTH APPROX.
50 kHz-15 MHz.
* 50 V METALLISED POLY OR MONOBLOCK
(LOW INDUCTANCE) TYPES.

100 kHz — 10 MHz Crystal Oscillator

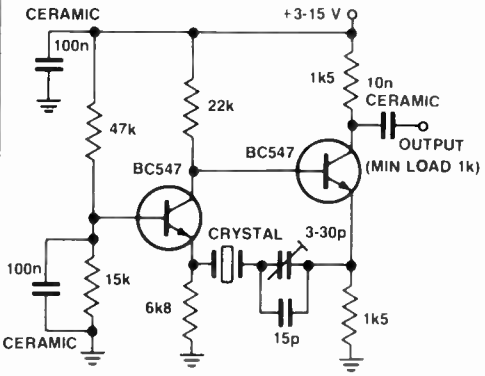

This untuned crystal oscillator will operate reliably over the range from 100 kHz to 10 MHz. Output level depends on supply rail, generally around one-third (peak-to-peak). Output is basically sinewave, but you can 'square up' the output by decreasing the value of the emitter resistor of the output transistor. Dropping it to, say, 220 Ohms gives good harmonics beyond 30 MHz from a 100 kHz crystal. While BC547 transistors are recommended, many types can be used, providing their gain-bandwidth product is 250 MHz or above, e.g: 2N5777, 2N3563 4, 2N3642, 2N2222, BF115, BFY90 etc. For crystals below 100 kHz, transistors with good LF gain are recommended — such as the 2N3565. Supply voltage can range from 3 V to 15 V. The trimmer capacitor is to set the crystal on frequency if necessary. If not, use a 100p ceramic.

IC Crystal Oscillator

This low frequency crystal oscillator provides an essentially squarewave output. Upper frequency limit is below 10 MHz. The output voltage swings virtually from rail to rail (5 Vp-p here).

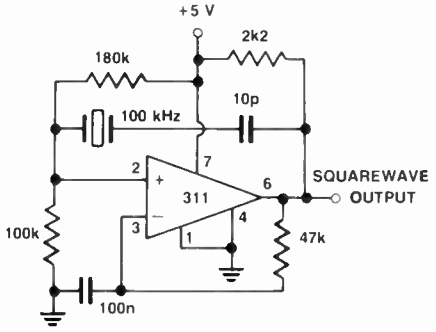

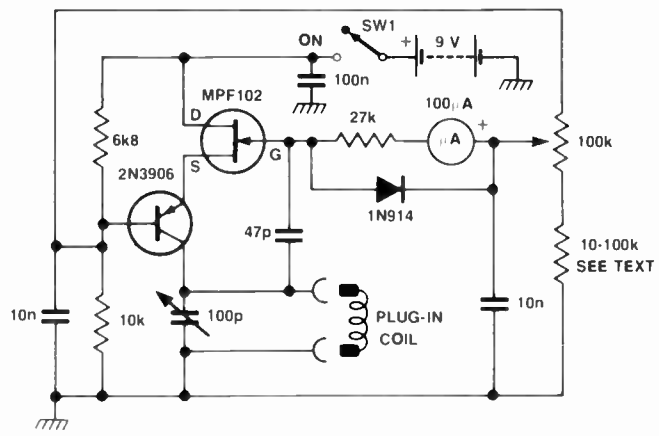

Solid-State Dipper

This circuit employs a grounded-base oscillator to provide tuning via a single-gang variable capacitor where the rotor plates can be grounded. Simple two terminal plug-in coils can be used too. The circuit will oscillate over quite a wide frequency range, the 100k pot adjusts the oscillation amplitude to suit the meter and battery voltage. The series resistor to common should be adjusted to cut-and-try so that the pot operates over its range. At upper VHF, the 47p capacitor should be reduced. All capacitors should be ceramic types and short leads are recommended in construction.

Wideband Cascade Amplifier

The major feature of the cascade amplifier configuration is the isolation between input and output. This makes for very good stability regardless of the load reactance. The 'lower' transistor is connected as a common emitter amplifier, the 'upper' one operating in grounded base. This circuit employs two transistors as a common substrate from a transistor array IC (CA3046), each having h_{FE} of 110 and f_T of 450 MHz. Upper cutoff determined is by R1. Increase all capacitor values to reduce lower cutoffs. A CA3018 may be used or discrete transistors such as 2N706, 2N2369, 2N3607, MPS3646 or 2N5769.

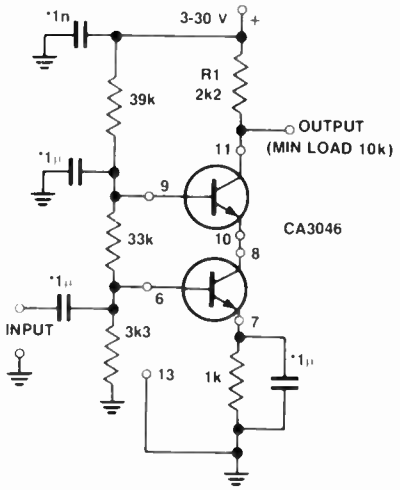

MIDBAND GAIN (LOADED) APPROX. 32 dB.
BANDWIDTH (−3 dB) APPROX. 5 kHz-4 MHz.
ALWAYS GROUND PIN 13 OF CA3046
* 50 V METALLISED POLY TYPES
OR TAG TANTALUMS (WATCH POLARITY)

Common Base/Common Collector Wideband Amp

This amplifier configuration has the advantage of low input and low output impedances — around 100 Ohms here. Two devices from a transistor array IC are used. Use dipped tanatalum capacitors and keep leads short for best results. You can match the input impedance to the impedance of the source by varying R1, which varies the emitter current of the common base stage and hence the input impedance. Linearity of this configuration is very good, but gain is not high.

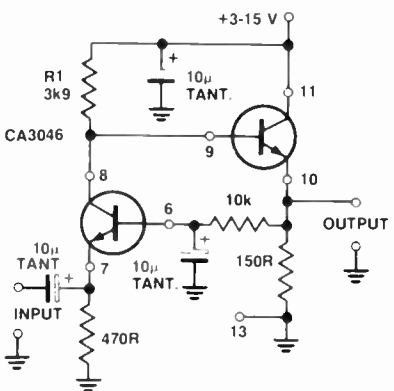

MIDBAND GAIN APPROX. 17 dB.
BANDWIDTH (−3 dB) APPROX. 150 Hz −3.5 MHz.
MAX. INPUT AT 5 V SUPPLY APPROX. 40 mV RMS
ALWAYS GROUND PIN 13 OF CA3046

TYPICAL PERFORMANCE CURVES
(MHW710-2)

FIGURE 1 INPUT POWER, EFFICIENCY AND VSWR versus FREQUENCY

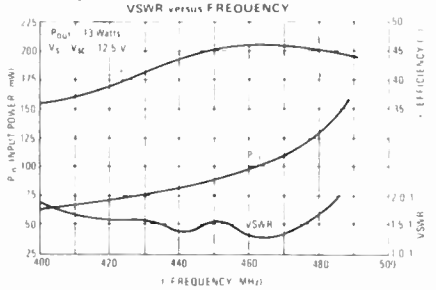

FIGURE 2 OUTPUT POWER versus INPUT POWER

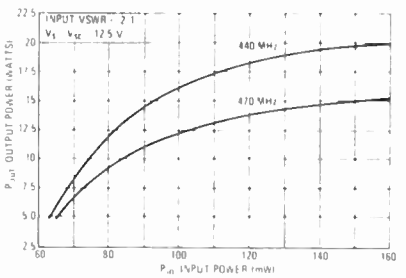

FIGURE 3 OUTPUT POWER versus VOLTAGE

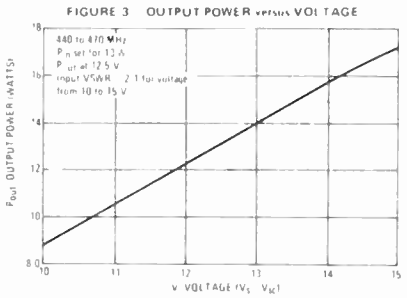

FIGURE 4 - OUTPUT POWER versus GAIN CONTROL VOLTAGE

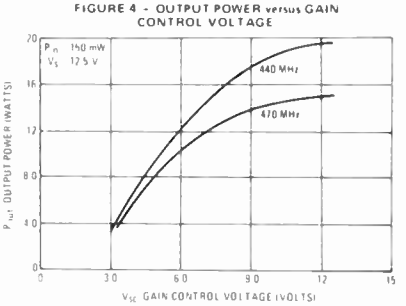

Amateur TV Modulator/Output Stage

According to Les Jenkins VK3ZBJ, high level modulation for amateur TV applications has distinct advantages over low level modulation followed by a string of linear amplifiers. The main problem with the latter approach is getting the linearity, particularly with solidstate stages. This circuit (from Les) does the job, only requiring about 100 mW of RF drive at 426.25 MHz and standard level composite video input from the camera. Peak RF output is 10 watts which will put quite a respectable signal on the air.

The Motorola MHW710 UHF 'gain block' is readily available and comparitively cheap. Good UHF construction practices should be used. The MHW710 is made for stripline termination to the pins. It should be bolted directly to a heatsink with double-sided pc board (glass fibre or teflon-glass) butted beneath the pins with appropriate tracks to provide termination. C1 and R1 should have virtually no leads. The output filter need not be anything fancy, a stripline or coaxial bandpass type should do the job. Remember, you're dealing with video bandwidths, keep it broad.

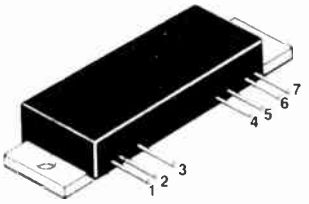

MHW710 PINOUT

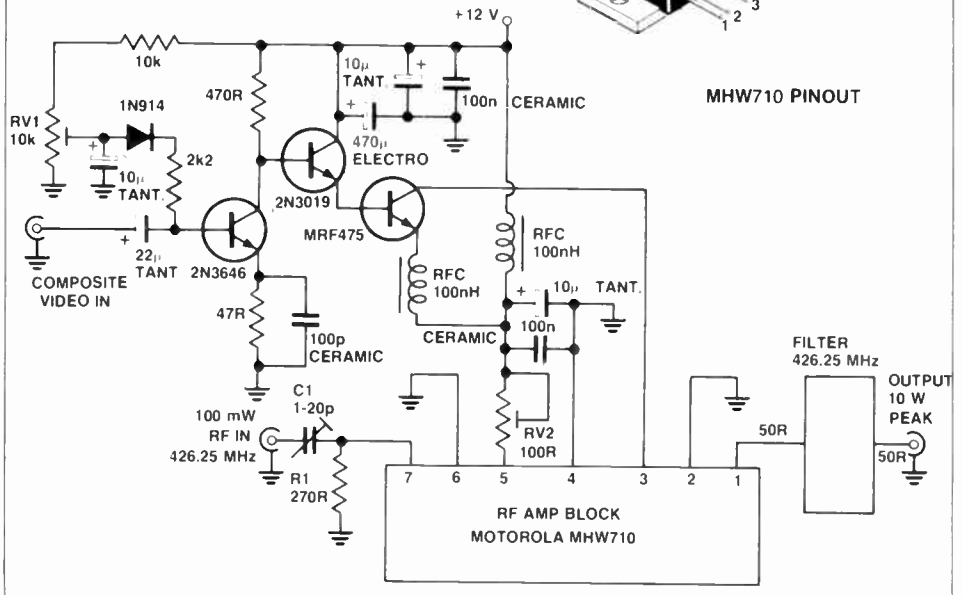

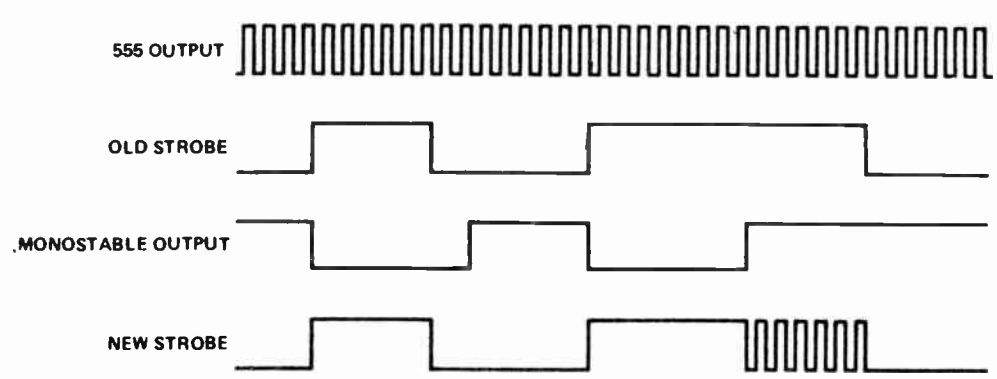

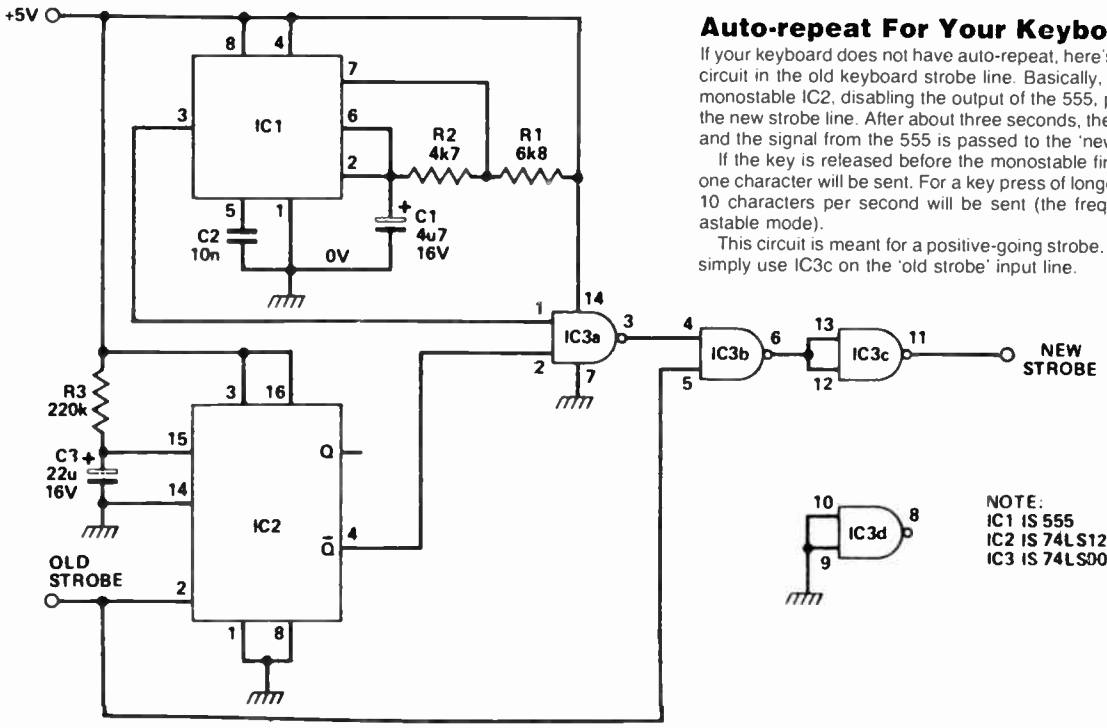

Auto-repeat For Your Keyboard

If your keyboard does not have auto-repeat, here's the solution. Just insert this circuit in the old keyboard strobe line. Basically, the strobe line activates the monostable IC2, disabling the output of the 555, preventing its pulses reading the new strobe line. After about three seconds, the output of IC2 changes state and the signal from the 555 is passed to the 'new strobe' line.

If the key is released before the monostable finished its timing period, only one character will be sent. For a key press of longer than three seconds, about 10 characters per second will be sent (the frequency set by the 555 in its astable mode).

This circuit is meant for a positive-going strobe. For a negative-going strobe, simply use IC3c on the 'old strobe' input line.

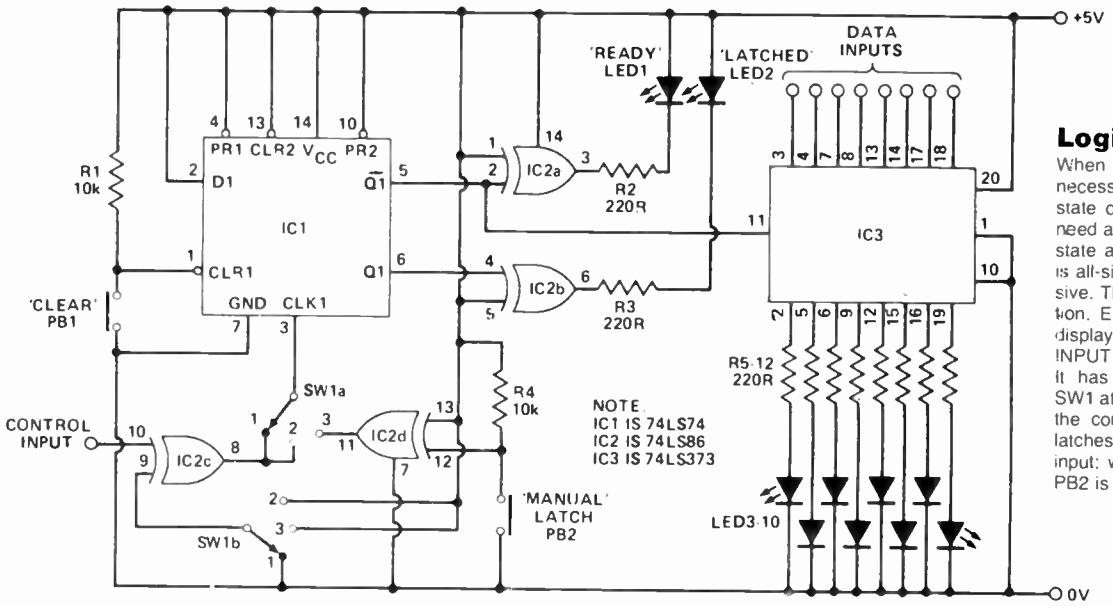

Logic Analyser

When testing logic circuits it is often necessary to know the simultaneous state of several nodes. To do it, you need a bunch of logic probes or a logic state analyser. The 'professional' tool is all-singing-and-dancing, and expensive. This circuit is a simpler implementation. Eight input bits are latched and displayed, controlled by the CONTROL INPUT from the circuit being analysed. It has three triggering modes: with SW1 at 1, it latches on a rising edge on the control input; with SW1 on 2, it latches on the falling edge control input; with SW1 on 3 it latches when PB2 is pressed. PB1 clears the latch.

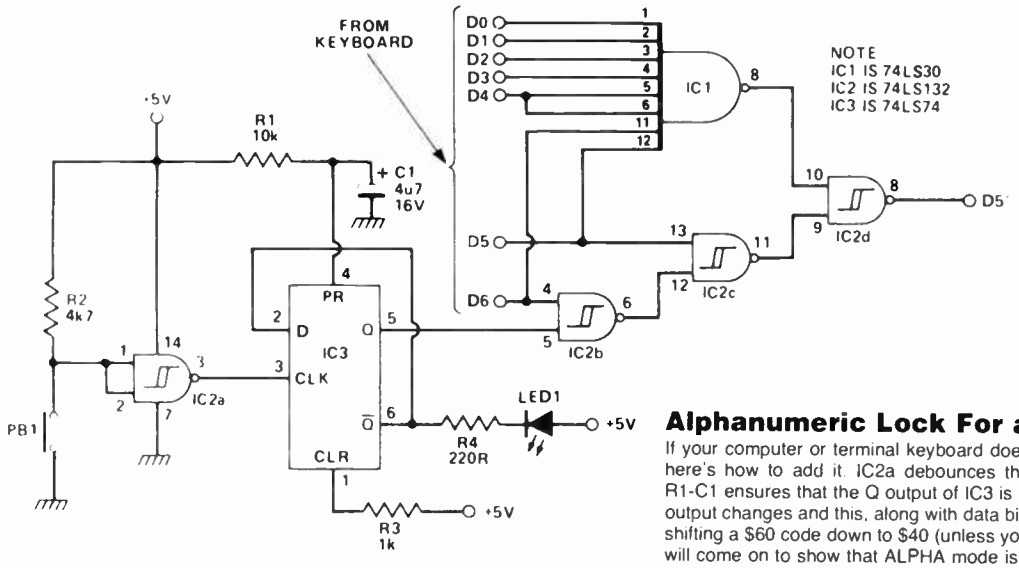

Alphanumeric Lock For a Keyboard

If your computer or terminal keyboard does not have an alphanumeric lock key, here's how to add it. IC2a debounces the alpha-lock key PB1. On power-up, R1-C1 ensures that the Q output of IC3 is high. Each time PB1 is pressed, the Q output changes and this, along with data bit 6 being high, will make D5 a low thus shifting a $60 code down to $40 (unless you pressed the DELETE key). The LED will come on to show that ALPHA mode is selected when PB1 is pressed.

To insert this into your keyboard it is only necessary for the original D5 line to pass through this circuit (D5 to D5') and add a keyswitch for PB1.

HCMOS Crystal Oscillator

This circuit works well with crystals in the MHz region, up to 15 MHz, perhaps more. It starts reliably and can be operated on any supply rail between 5 V and 15 V. The 22p and 33p capacitor values may need to be increased for lower frequency crystals — but maintain the ratio.

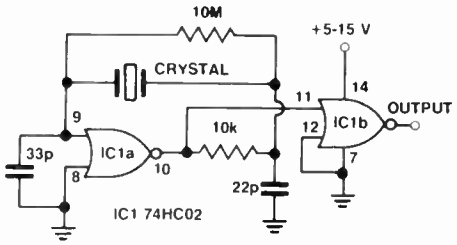

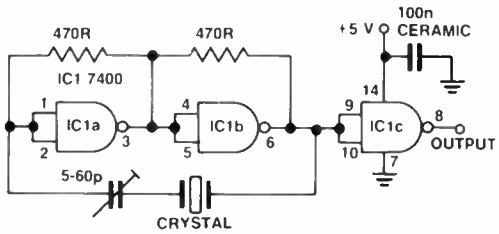

Reliable TTL Crystal Oscillator

Many TTL crystal oscillators have been published over the years, and many exhibit unreliable oscillation — particularly on startup. This one has none of those problems. This will work with crystals in the 1 MHz to 20 MHz range as well as with 'Ceralock' ceramic resonators. The trimmer is to set the crystal to frequency, if necessary. If not, substitute a 100p-1n ceramic capacitor.

GENERATORS

Quadrature Sinusoid Oscillator

This oscillator provides two sinusoid outputs precisely 90 apart. All capacitors should be 1% types. R1C1 = R2C2. The oscillation frequency is given by:

$$f = \frac{1}{2\pi\sqrt{C2R2C3R3}}$$

(R1C1 = R2C2)

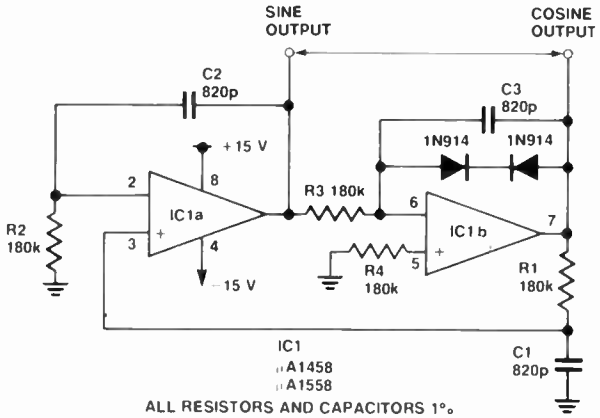

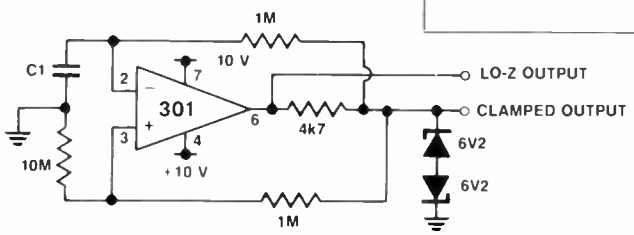

Low Frequency Squarewave Generator

The output frequency of this generator is determined by C1 and the 10M resistor. A low impedance output, that swings almost from rail to rail, can be taken from pin 6 of the 301, or a clamped output — limited by two 6.2 V zeners — gives a 6.8 V peak-to-peak output. Maximum supply voltage is ±18 V.

EHT Meter

This simple dc meter circuit will read 0-5 kV and 0-10 kV. A high impedance input op-amp is used as a meter amplifier, driving a 100 µA moving coil meter calibrated 0-5 and 0-10 to read kilovolts. The input divider comprises a 1000M resistor, made up of 10 x 100M resistors or a 1000M EHT probe, and a 100k resistor. These should be 5% types at least, preferably 2% types. A Class 2 or Class 2.5 (common type) meter movement should be used which will provide 2% or 2.5% full-scale accuracy. The resistive divider need not be any more accurate. The two zeners on the input provide over-voltage protection. Calibration is simple. Short the input and adjust RV1 to zero the meter. To calibrate the 10 kV range, set SW1 to 2 then apply 1.00 volts across R11 and adjust PR1 so the meter reads full scale. For 5 kV, set SW1 to 3 then apply 0.50 volts across R11 and adjust PR2 to read full scale on the meter. Resistors R1 to R10 should be mounted in a 'string' and covered in heatshrink tubing to prevent arc-over between their ends at peak voltage. Use ½W or 1W resistors for their voltage rating.

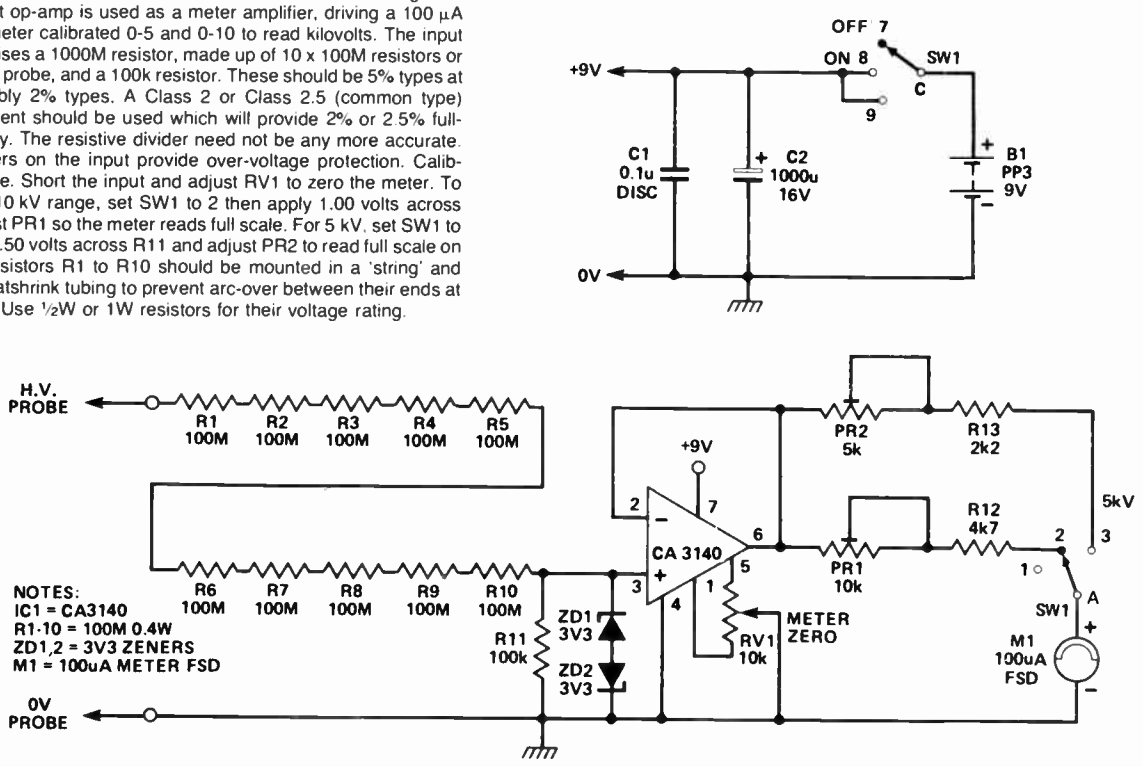

mV-nA Meter Amplifier

This meter amplifier can be calibrated to read either 10 mV or 100 nA full-scale. Zero set is provided by the 2k trimpot and full-scale calibration by the 5k trimpot. The back-to-back diodes on the input provide input overrange protection.

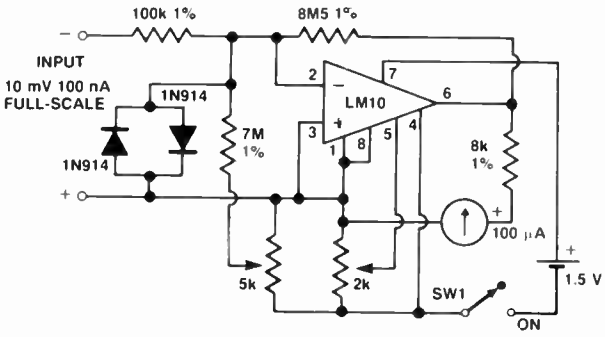

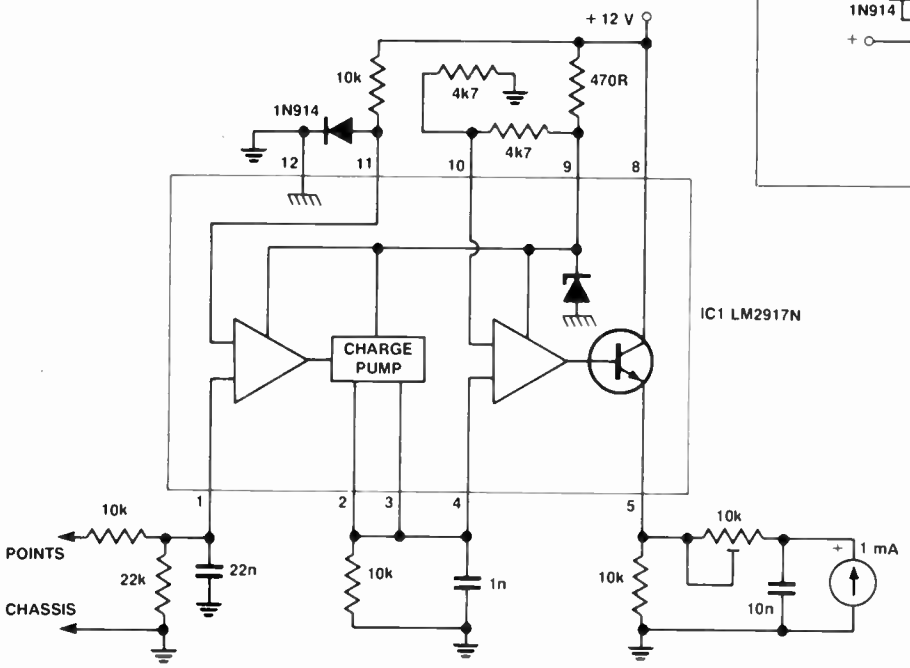

Dwell Meter

The LM2917 is a frequency-to-voltage converter, used here to measure the 'points closed' period or engine 'dwell'. Construction is non-critical. Calibration is by means of the 10k trimpot. Use a 50% duty cycle square wave of a few volts to calibrate the meter — 45 is half-scale for four cylinder engines.

0-100°C Thermometer

A precision temperature sensor, the LM134, provides an output of 10 mV/°K. The LM10 provides 'normal' temperature offset and amplification so that a moving coil meter movement can be used to indicate temperature.

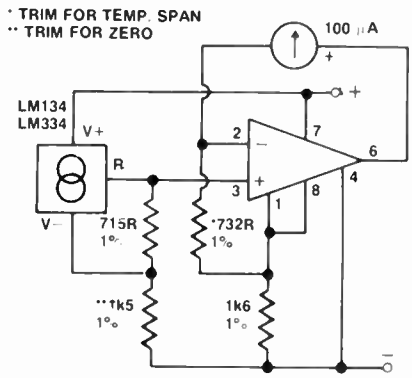

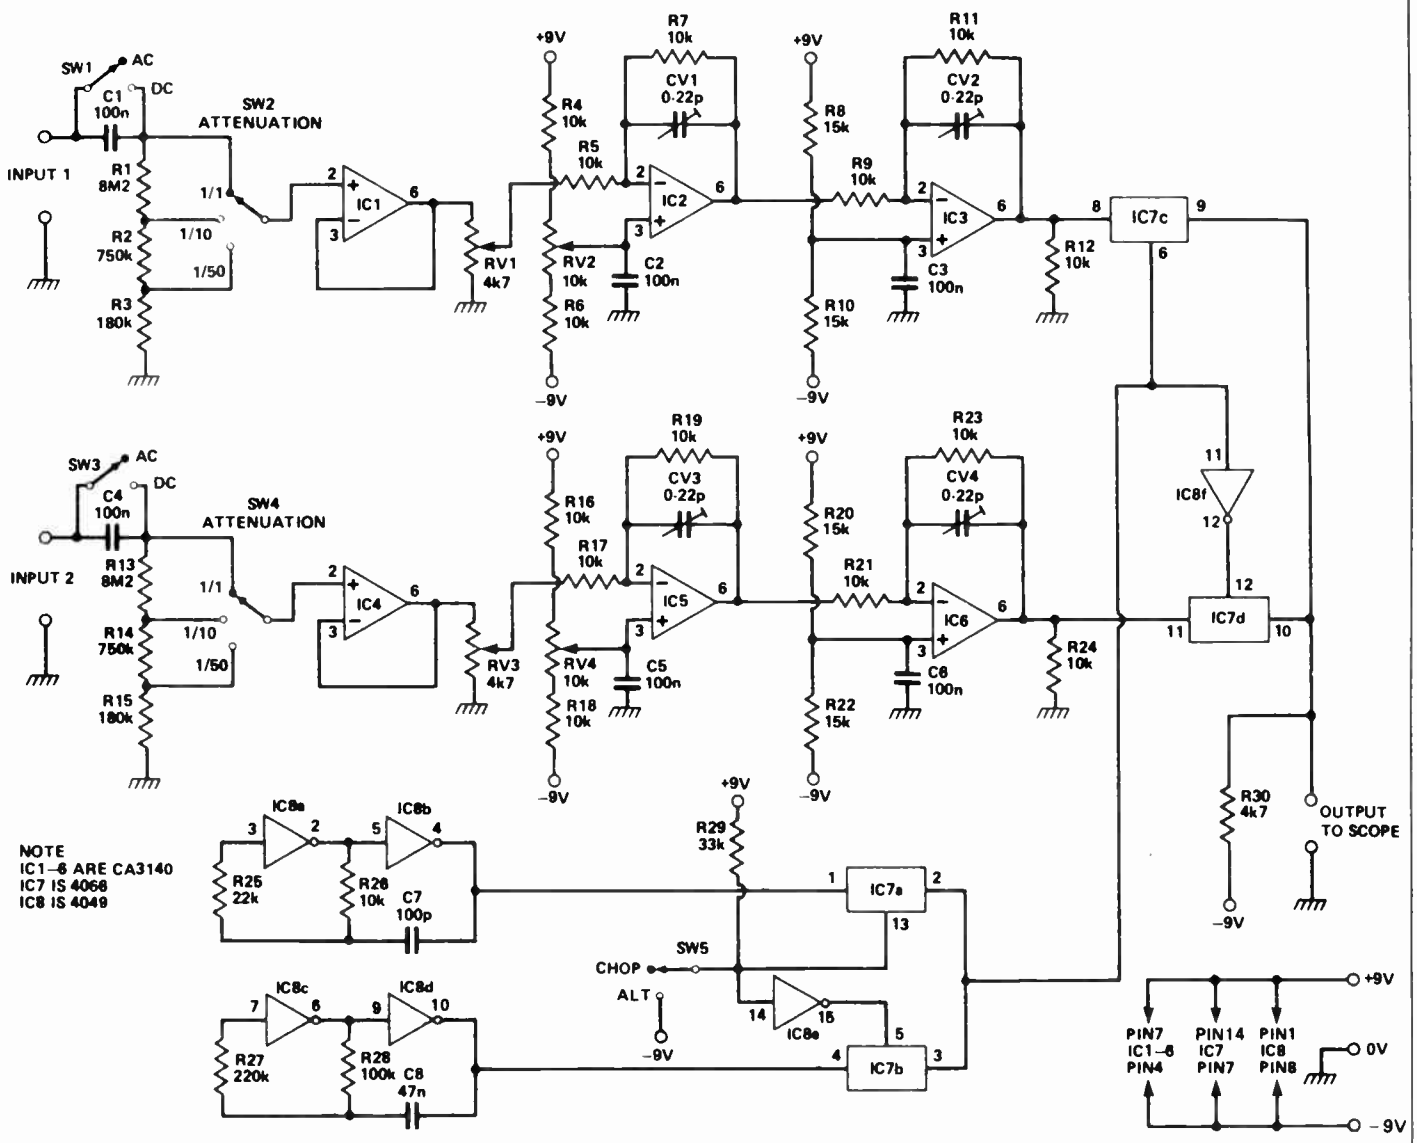

CRO Dual-Trace Switch

This permits a single-channel CRO to be turned into a dual-trace unit. Two modes of display are provided: 'chop' and 'alternate'. When examining frequencies between dc and about 15 kHz 'chop' mode is used as trace speed across the screen (timebase speed) is relatively slow. About 15 kHz, 'alternate' mode is used when the timebase makes separate sweeps of the screen for each trace.

Two preamps are used, each having a bandwidth extending from dc to about 2000 kHz. The outputs of the channels are switched to the CRO input by an electronic changeover switch comprised of IC7c-d and IC8f, driven by an oscillator. Two oscillators are used, one for the chop mode, one for the alternate mode, as different switch speeds are necessary.

Setting CV1-CV4 is critical. Procedure is identical for both channels. Inject a 1 Vp-p 100 kHz squarewave into Input 1 and set SW2 to the 1/1 position. Setting of SW1 is unimportant. Set RV1 to maximum and RV2 to mid-position. Connect a CRO to the output of IC2 and adjust for the best squarewave response. (No overshoot and least rounding). Then move to IC3's output and adjust CV2 for best response. Repeat the procedure for the other channel.

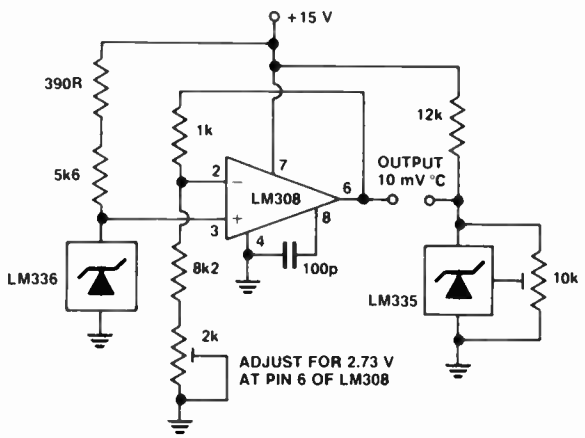

Centigrade Thermometer

This simple circuit can be used to measure temperature in Centigrade degrees and will read out directly on a standard digital multimeter. The LM336 is a precision 2.5 V voltage regulator, the LM335 a precision temperature sensor. The LM308 output provides a 'stiff' reference scaled up from absolute zero. The 2k trimpot is set to provide 2.73 V between pin 6 of the 308 and common. The 10k trimpot is set so that you get 2.982 V across the LM335 at 25 C. The LM335 can be used for contact temperature measurement.

AUDIO

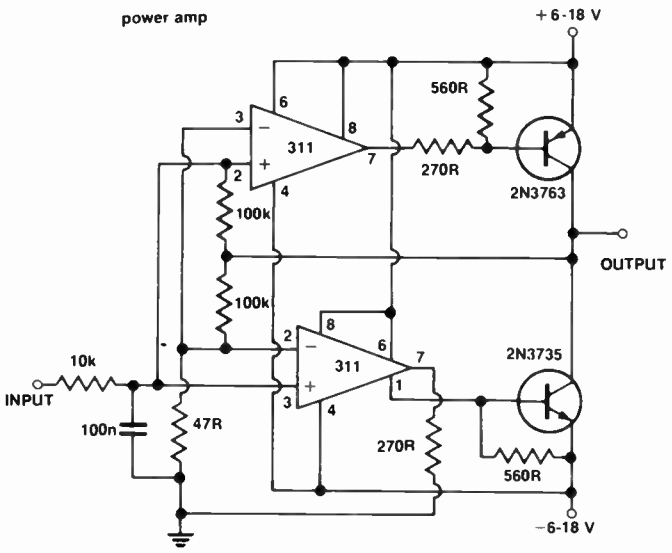

Switching Power Amp

Two 311 comparators used in a class B power amp. The output devices only switch on when alternate input cycles exceed a small voltage threshold. Feedback largely reduces the crossover distortion.

FILTERS

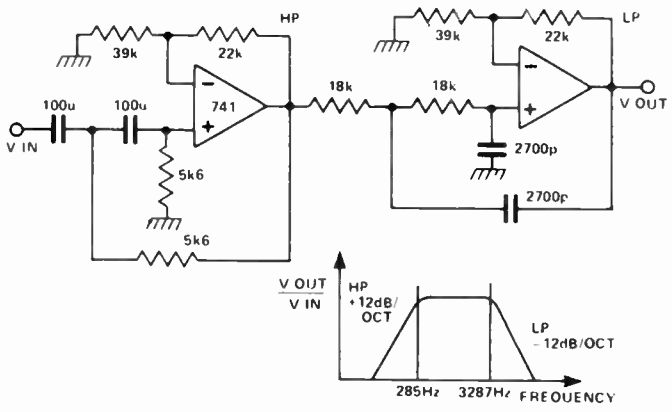

Speech Filter

For any purely speech audio system application — communications receivers, transceivers, dictaphones etc. — it is desirable, for best intelligibility, to limit the audio bandwidth to provide relatively steep rolloff below about 300 Hz and above about 3 kHz or so. Most speech information is contained between these limits. This circuit shows a simple bandpass filter sytem with 12 dB/ octave rolloff below 285 Hz and above 3287 Hz. One unusual application of such a filter would be in a 'light show' system so that the display varies with voice variations of an announcer or performer.

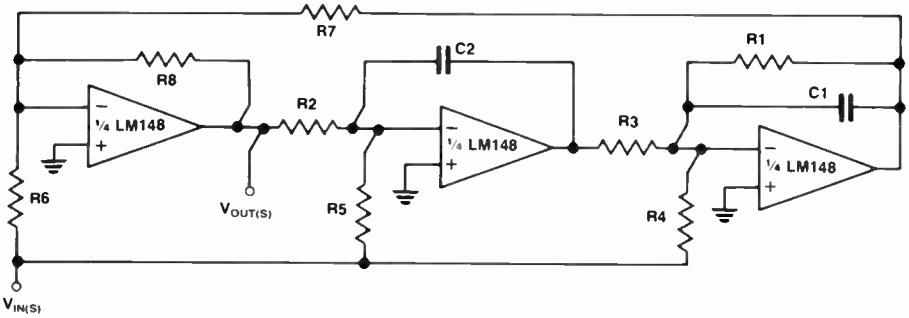

$$Q = \sqrt{\frac{R8}{R7}} \times \frac{R1C1}{\sqrt{R3C2R2C1}} \quad f_o = \frac{1}{2\pi}\sqrt{\frac{R8}{R7}} \times \frac{1}{\sqrt{R2R3C1C2}} \quad f_{NOTCH} = \frac{1}{2\pi}\sqrt{\frac{R6}{R3R5R7C1C2}}$$

Necessary condition for notch: $\frac{1}{R6} = \frac{R1}{R4R7}$

Bi-quad Audio Notch Filter

This notch filter provides good predictability and better noise performance than the state-space approach. Three op-amps from a quad package are employed. For a notch of 3 kHz, R1 is 270k, R2-R3-R5 are 20k, R4 is 27k, R6-R8 are 10k, R7 is 100k, C1-C2 are 1n (use good quality poly or mica capacitors).

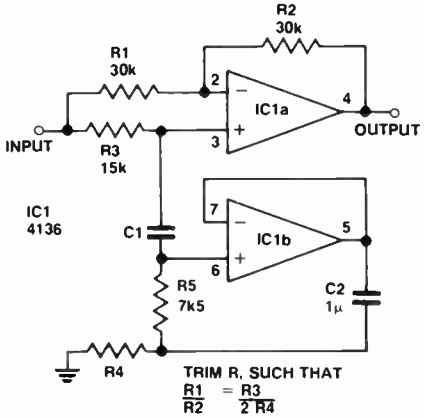

Gyrator Audio Notch Filter

Using two-op-amps from a 4136 quad package, you can implement a very effective notch filter with common component values, the notch frequency ranging from 20 Hz to 2 kHz. Capacitors should be good quality, low voltage metallised poly types.

Notch Frequency as a Function of C1

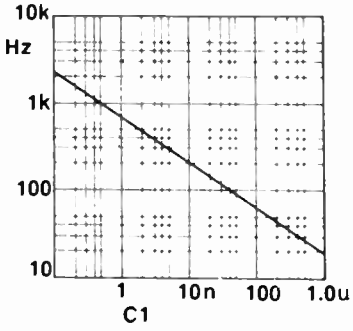

All-pass Notch Filter

This circuit provides adjustable notch attenuation ('depth'). IC1 and IC2 are arranged as 'all pass' filters. They have a flat freqency response but the phase changes with frequency. Overall maximum phase shift is 360°, a phase shift of 180° (reversal) occurring at a frequency of 1/2CR Hz. By mixing the phase-shifted signal with the original, cancellation will form a narrow 'notch' in the frequency response.

The 1k preset varies the amount of phase-shifted signal mixed with the original so the notch attenuation can be varied. The notch frequency, f_C, can be varied by varying the two resistor marked R. e.g: for a notch at 50 Hz, R should be 220k.

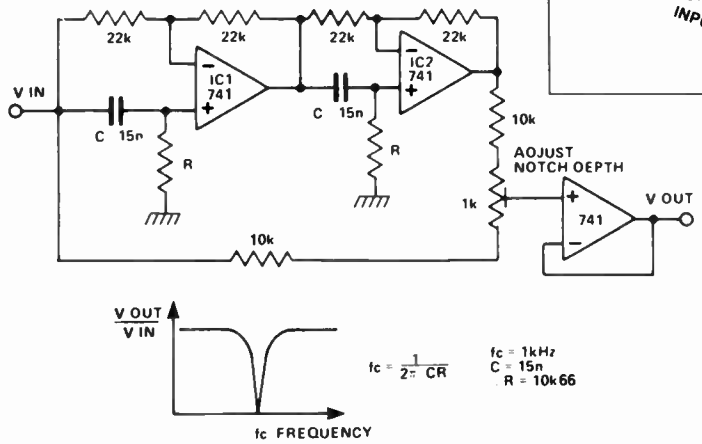

$f_c = \frac{1}{2\pi CR}$

f_c = 1kHz
C = 15n
R = 10k66

Tunable Audio Filter with Adjustable Selectivity

This filter can be tuned from 20 Hz to 1 kHz and features adjustable bandwidth. RV2 sets the frequency, RV1 sets the selectivity. This is a positive feedback control. Advancing it beyond a certain point causes oscillation. Set it just below the point of oscillation for minimum selectivity. The capacitors should be low voltage metalised poly types, 5% or better.

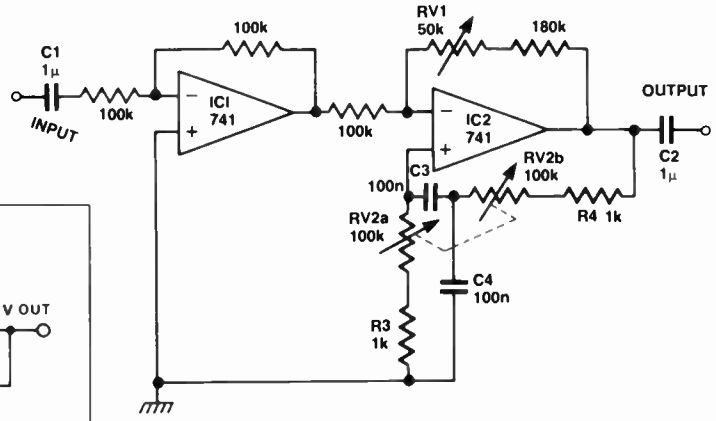

Active 1 kHz Bandpass Filter

This circuit provides a 40 Hz-wide 'window' at 1 kHz. Two op-amps from a quad op-amp IC are used. For best results, the two 10n capacitors should be metallised poly, low voltage types, with a 2% or 5% tolerance — matched values if possible.

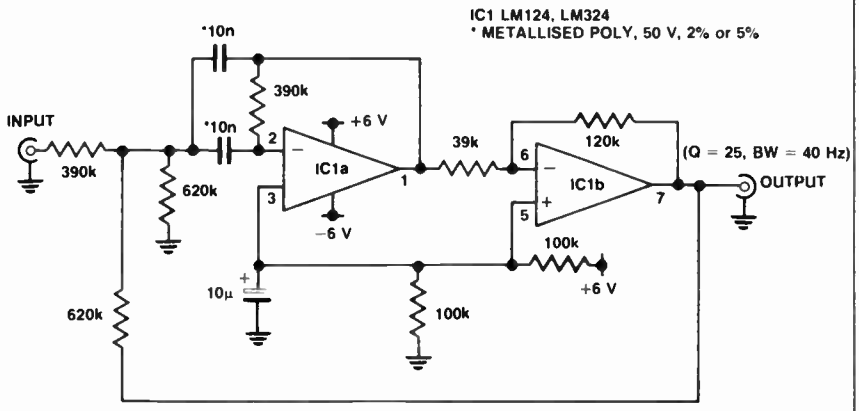

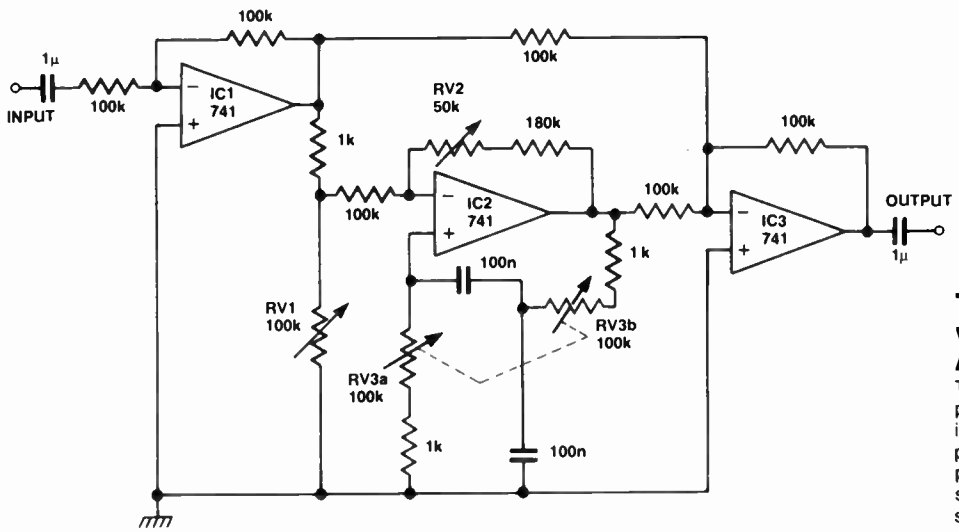

Tunable Audio Notch Filter with Variable Selectivity, Attenuation

This is a modification of the previous circuit to provide a notch, rather than a peak. The original input and the filtered signal are summed in anti-phase at the input of IC3. Hence, the filter circuit provides attenuation at the filter frequency. RV1 sets notch depth, RV2 sets notch width, RV3 sets notch frequency.

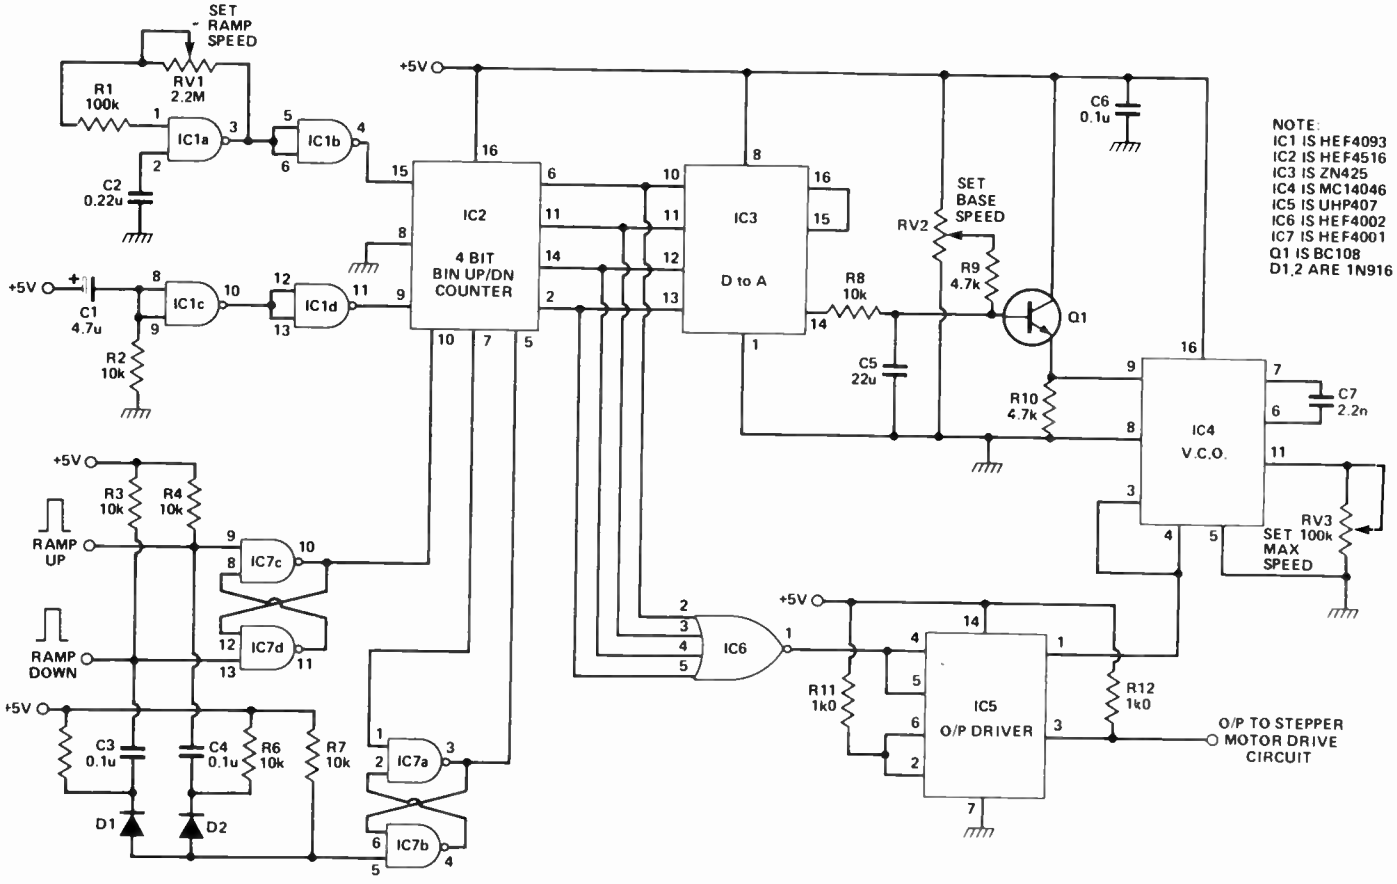

Ramped Stepper-Motor Controller

This has great applications for robotics. When the motor is at rest, the receipt of a RAMP UP pulse will cause the motor to start at its base speed then ramp up to maximum speed and run for as long as required. On receipt of the RAMP DOWN pulse, the motor will slow down to base speed then stop.

IC1a b provides clock pulses (ramp speed) to a four-bit binary up down counter, IC2. On receipt of the RAMP UP pulse IC7c d sets count up and IC7a b enables the counter. Unless a RAMP DOWN pulse is received the counter will reach its maximum count (max. speed) and hold at this until the RAMP DOWN pulse is received. The latter will set count down and enable the counter which then counts to zero (base speed) and hold again until the next RAMP UP pulse.

The counter output drives the D-to-A converter IC3, the ramping output of this controlling the VCO, IC4. The lower frequency of the VCO (base speed) is set by the bias adjustment of Q1. Upper frequency (max. speed) is set by the 100k pot, RV3. IC5 provides open-collector drive for the output pulse train and also the on off gate, controlled by IC6, when the counter is set to zero. IC1c d provides a set zero pulse to IC2 to ensure that the output, at pin 3 of IC5, is off each time the generator is switched on.

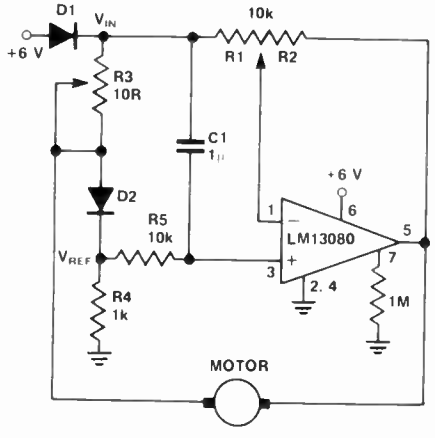

Speed Control for dc Motor

The LM13080 power op-amp is used here to provide a simple speed control for small dc motors requiring less than 0.5 A starting current. This circuit operates by impressing the multiple of a reference voltage across the motor then varying the reference by means of quasi-positive feedback to change the voltage across the motor whenever the load on the motor changes.

It works as follows: D1 brings V_{IN} within the common-mode range of the op-amp. A reference voltage is established by the combined voltage drop through the 10R pot R3 and diode D2 and is applied to the non-inverting input of the LM13080. Resistor R4 is used to bias D2 on. The 10k speed adjust pot is two resistors in one — R1 R2. R1 is the input resistance. R2 is the negative feedback resistance. Thus, the voltage impressed across the motor is given by:

$$V_{MOTOR} = \frac{(V_{REF} + I_3R3)R2}{R1} + V_{BE}$$

The positive feedback is developed as a change in the voltage across R3 due to the change in motor current caused by a variation in the motor's load. Resistor R3 is shown as a pot so that the amount of positive feedback can be adjusted to smooth operation of the motor (no 'hunting' or 'cogging'). Capacitor C1 and R5 serve as a filter for the reference voltage. Use a polyester low voltage capacitor.

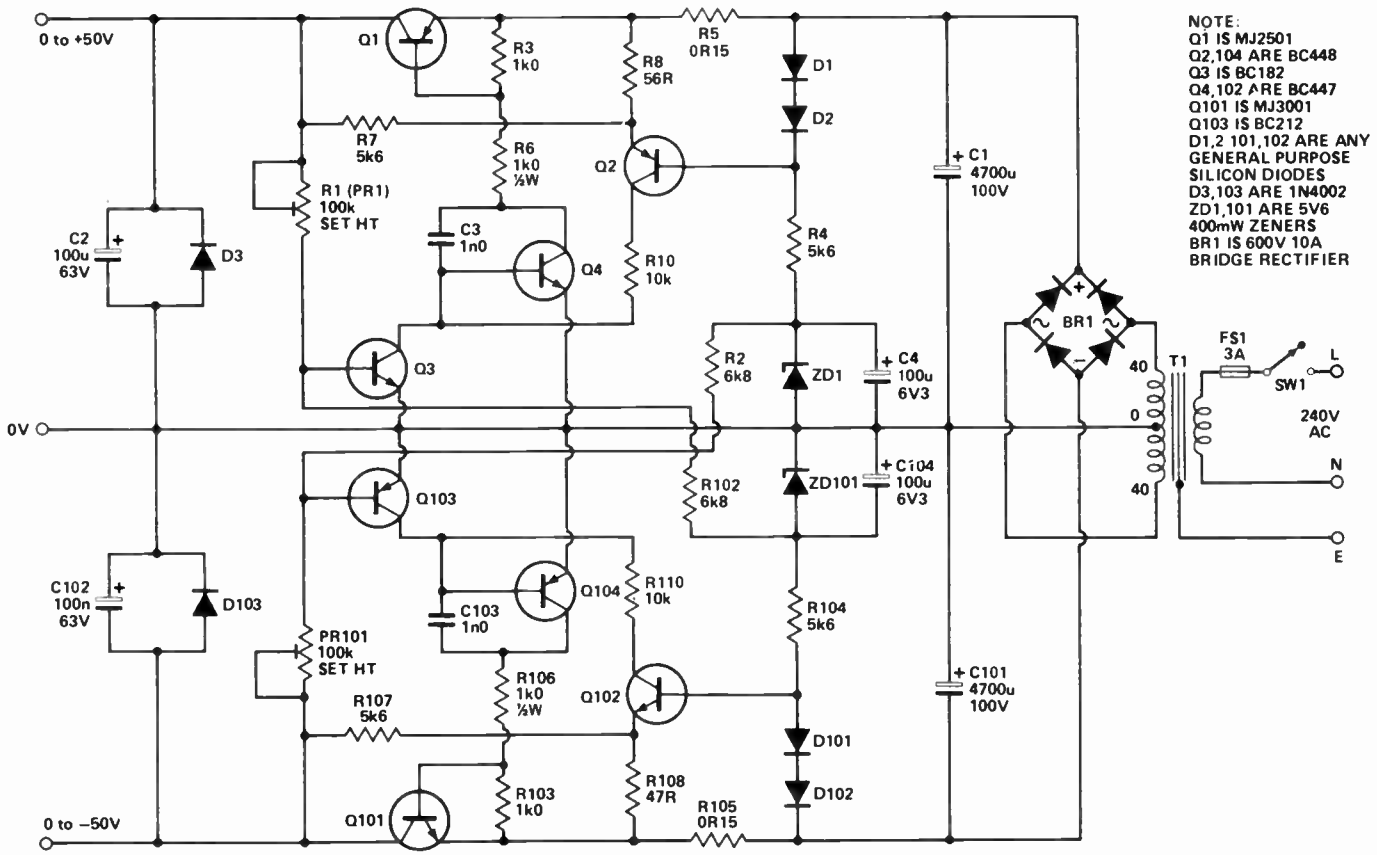

NOTE:
Q1 IS MJ2501
Q2,104 ARE BC448
Q3 IS BC182
Q4,102 ARE BC447
Q101 IS MJ3001
Q103 IS BC212
D1,2 101,102 ARE ANY GENERAL PURPOSE SILICON DIODES
D3,103 ARE 1N4002
ZD1,101 ARE 5V6 400mW ZENERS
BR1 IS 600V 10A BRIDGE RECTIFIER

Twin Stabilised Power Supply

This circuit will deliver up to 50 V on the two output rails at current up to 3 A. The secondary of T1 must be rated at 4.5-5 A to deliver these output currents. Full current limiting protection is provided. This circuit is ideal for use with audio power amp modules requiring supply rails up to 50 V. Note: R8 should be 47R.

50 V/1 A Laboratory Power Supply

A series-pass regulator is employed here with separate circuits for voltage reference and current limiting. The input voltage should be at least 55-60 V, well filtered. The LM10 has an internal reference, which is made use of here.

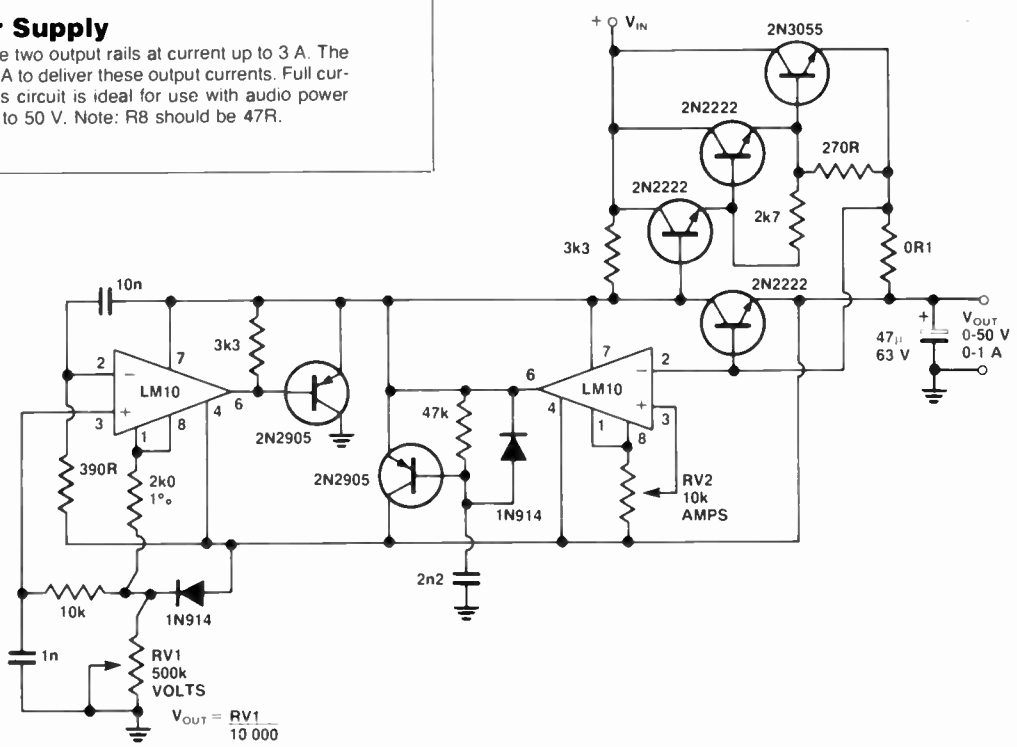

ISBN: 0 86405 045 1