Control up to four room lights over a twin-wire cable

Ian Thomas

Part 1

It is probably a not-uncommon problem to want to replace the single ceiling light in a room with a more exotic dimmable arrangement only to find that the twin-wire switch cable is concreted in, or otherwise hard to get at without spoiling the room decoration. This project allows you to control up to four room lights — two dimmable and two switched — over the existing twin-wire cable. No need to spoil the wall and ruin the room decor.

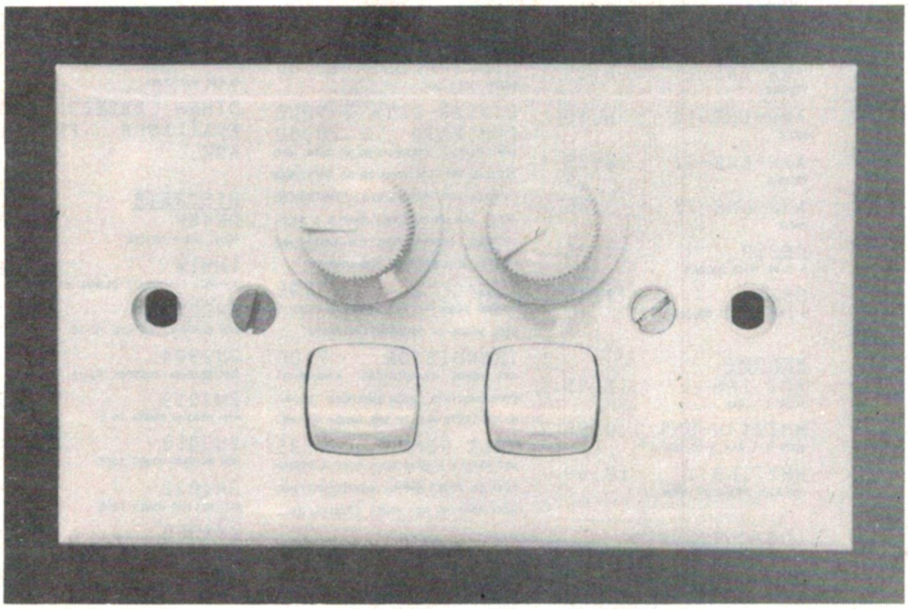

The controls. I modified a standard, commonly available, HPM light fitting (made in Australia) to take the controls. This provided a neat and convenient solution, the switches and dimmer pots snap into the panel, the encoder pc board mounts behind them. The end result is only distinguishable from an 'ordinary' light fitting by the two pc board mounting screws visible here (which should be earthed or covered). If the pc mountings are epoxied in place — who'd know the difference?

IT IS PROBABLY a not uncommon problem to want to replace the single overhead light fitting with a more exotic dimmable arrangement only to find that the control wires to the switch are *concreted in*. Also, in keeping with Murphy's Law, the walls have only been re-papered a year or so ago and the wallpaper is now totally unavailable and to string more wires to the switch would involve a major redecoration of the whole room, so the idea is scrapped.

This project is intended to solve such problems by replacing the single wall switch with the required multiple switch-control pot combination and a very compact printed circuit board. A matching electronics box in the ceiling feeds power to the control board and also recovers the control information to operate triacs to drive the multiple lights.

It was considered essential that the replaced wall controls look exactly like normal 'bought' controls and avoid that 'home made' look so, normal, purchased, wall fittings were used for the controller. The unit in the ceiling merely had to work reliably and could be any size or shape at all.

The design

The first problem to be dealt with in the design was how many controls and what type were wanted. Also, it had to be decided how to multiplex the switch/control pot information onto the single wire pair available. A survey of the many types of controller integrated circuits available (what we were trying to do is not dissimilar to the remote control of a TV or VCR) showed that while there were dozens (literally!) manufactured, precisely none (zip! zero! zilch!) were to be had in Australia unless you wanted to buy 1000 or so.

The only ray of light to be found was from National Semiconductor, and a perfectly satisfactory ray it turned out to be too. National make the LM1871 Radio Control Encoder/Transmitter and its companion LM1872 Radio Control Receiver/Decoder, which are intended to be used as remote controllers for toy cars or model aeroplanes. They provide, in their simplest configuration, two linear analogue control channels plus two separate on-off channels The encoder/transmitter also contains a moderate power high-f_T transistor to be used as an RF transmitter and the receiver/decoder has all the circuitry to construct a (simple) RF receiver so, as I could actually go out and buy some over the counter (good marketing National!) this integrated circuit pair seemed to be the way to go. They would give two completely linear, dimmable lights plus two simple on-off lights which should cover most rooms adequately.

The LM1871/LM1872 integrated circuits operate on a rather peculiar combination of pulse-width modulation plus (once again in their simplest configuration) pulse counting to transmit the analogue and digital information. The receiver/decoder half of the pair contains no memory so the encoder/transmitter must send the information continuously to keep the system 'alive'.

The encoder section of the LM1871 generates a series of up to six pulses in each data frame with each pulse separated by a fixed off time. At the end of the up-to-six pulses a much longer synchronisation pulse is sent so the receiver/decoder can lock on correctly. The first two pulses in each data frame are of variable width to convey the analogue information of the two channels and the other, up-to-four, fixed width pulses give the status of the two on-off switches. Figure 1 shows a typical encoder output and indicates how the width of the first two pulses gives the analogue data.

As Figure 1 shows, the analogue information is contained in pulses repeated at the frame rate whose width varies. The values recommended by National give a frame rate of 50 Hz. However, as the analogue information derived from this repetition rate data is to be used to control triacs running from 50 Hz mains the two similar frequencies would be sure to give rise to all sorts of unwanted slow beat effects (you would be amazed how sensitive the eye is to slow beats in, say, light intensity). For this reason I opted to have a frame rate of 75 Hz so any beats would be at 25 Hz and thus unnoticeable.

Also, the National design data is based on bandwidth restrictions and no such restraints exist for us. I also opted to shorten the fixed off time between all pulses to 0.4 ms from the National-recommended 0.5. The analogue channel limits were left the same as in the recommendations at 1 ms minimum and 2 ms maximum. The digital information pulses were left at 0.5 ms to finish defining all the times in the frame sent from the encoder.

The LM1871 is intended to be used as a very low power 27 MHz or 46 MHz (or, by straining and a bit of extra circuitry, 72 MHz) radio transmitter. This seemed like a good idea except that I had no desire to further crud up the ether with lighting control information. Instead, I chose to use the wires carrying power to the controller to carry switched low frequency RF data in the ▶

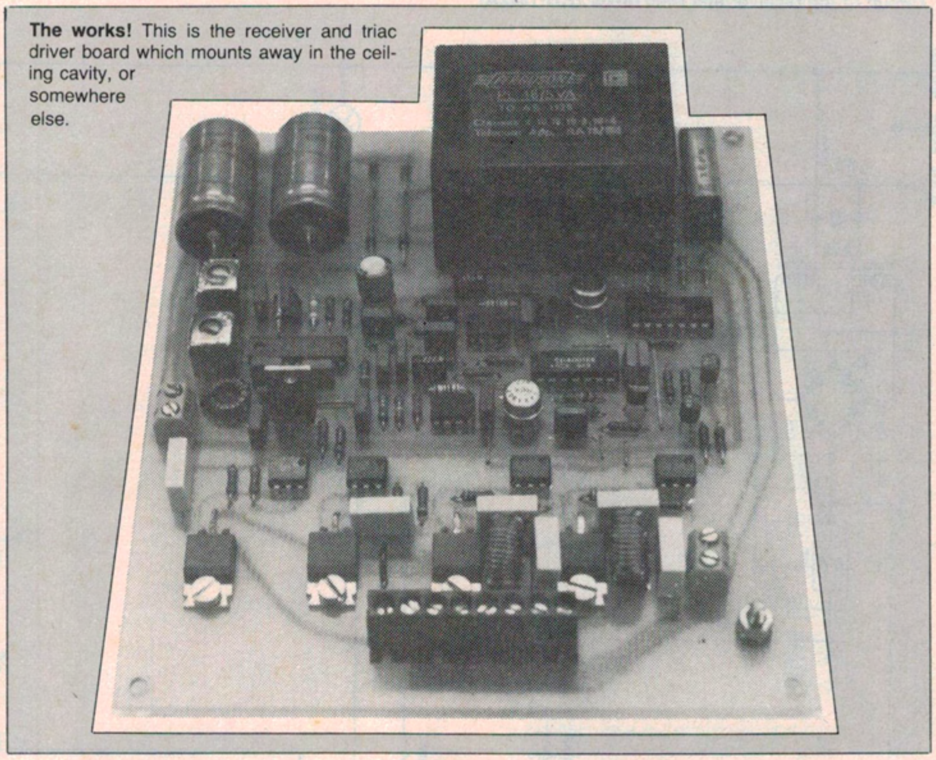

The works! This is the receiver and triac driver board which mounts away in the ceiling cavity, or somewhere else.

Typical application

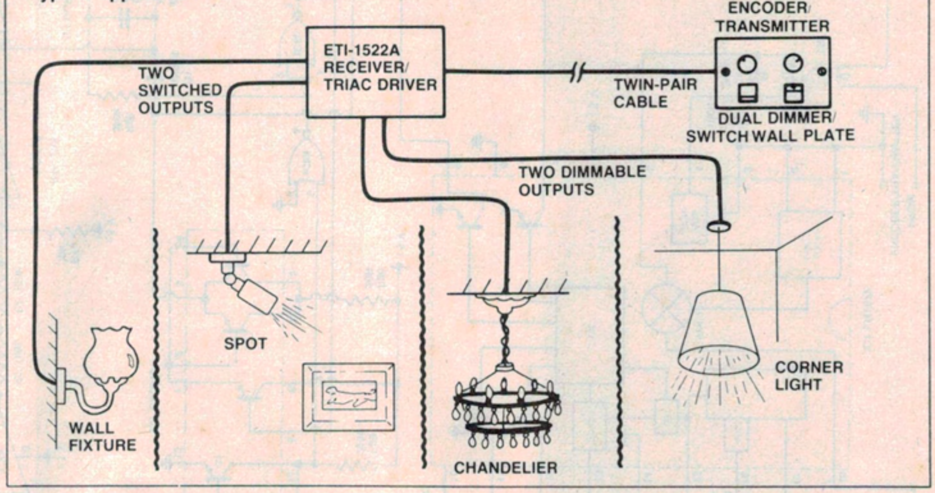

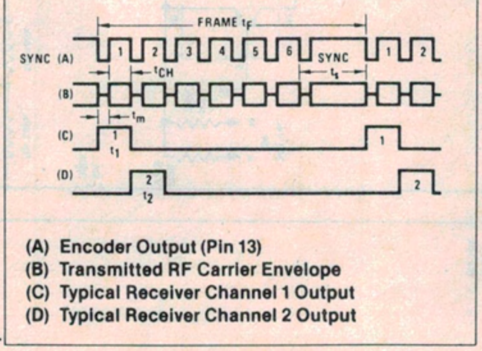

Figure 1. Typical encoder output waveform. The width of the first two pulses gives the analogue data.

Circuit of the receiver and triac driver (ETI-1522A)

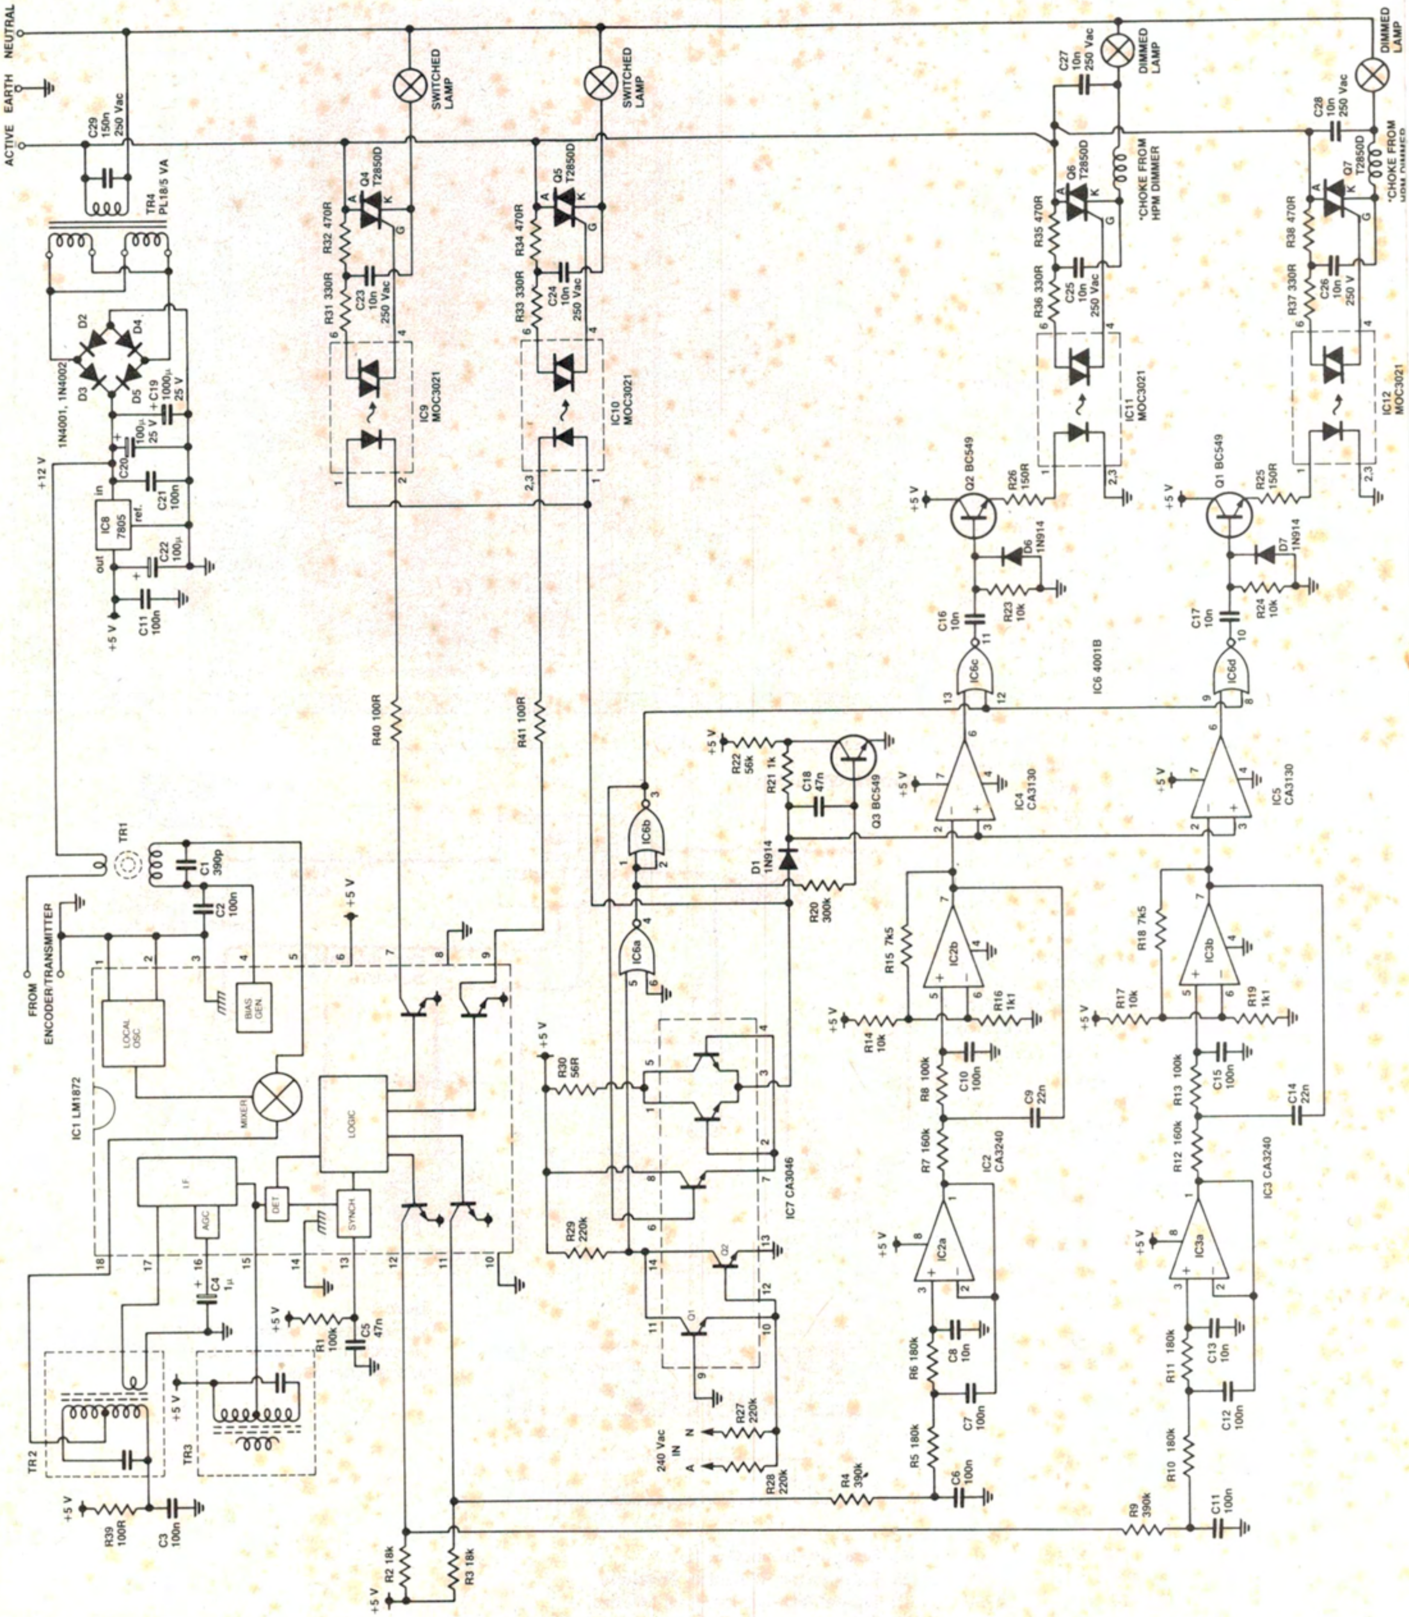

HOW IT WORKS — ETI-1522

The light controller can be divided into two major sections; the transmitter and the receiver and the receiver itself can be divided into several sections which are: the RF receiver, analogue filtering, mains synchronisation and triac drive. First, consider the transmitter.

The transmitter sends its information as to potentiometer position and switch status by sending a series of between three and six data pulses followed by a much longer pulse to synchronise the system (see Figure 1). All the pulses are separated by an off period of exactly the same time. The first two pulses are individually pulse-width modulated to encode the two analogue channels and the next one-to-four pulses contain the switches, encoded by a simple counting system. As the width of the first two pulses and also the number of subsequent pulses can vary, the actual time taken for data transmission varies as the input controls are varied. However, the LM1871 allows for this by varying the length of the synchronising pulse so the length of every data frame is the same.

While it is interesting to go into details of the internal operation of the LM1871 the essential part is to understand how the external components affect the operation of the integrated circuit. The IC consists of several separate RC timing circuits and some combining logic, together with voltage reference generators to give both stable timing and stable RF output level. Overall data frame timing is set by R8 and C10 and this time constant sets the 75 Hz data rate for the system. The time between each individual pulse of the data and the duration of each data pulse is set by capacitor C11, together with different resistors for the inter-pulse time and the actual pulse durations.

When the LM1871 is measuring the time between the pulses it is charging C11 from 0.33 V_{ref} to 0.67 V_{ref} by switching in R1. When the LM1871 determines an actual pulse it discharges C11 through the resistor associated with that pulse. For example, when it sets the length of the first pulse in the data frame it switches in the resistance connected to pin 3, which is R7 and RV2. Thus, R7 sets the minimum pulse width of the first analogue channel and RV2 varies the pulse width from that minimum. The same cycle is repeated for the second analogue channel; charging C11 through R1 and then selecting the second discharge resistance, R6 and RV1.

If desired, four more analogue pulse-width modulated channels can be generated by connecting four more variable resistors to pins 1, 18, 17 and 16 but we don't want six analogue channels and also, the LM1871 contains extra circuitry to convert these potential extra channels to two digital channels and at the same time make the receiver much simpler. If pins 1, 18, 17 and 16 are all connected together (as they are here) then for pulses three to six, the same resistor will be used and the last four pulses will have the same length.

Internal logic exists in the LM1871 to operate on information provided on pins 5 and 6 and vary the number of pulses that are sent according to that information. For example, if both pins 5 and 6 are left open (the IC has internal pullup resistors) then only one fixed duration pulse is sent after the two analogue pulses and the rest of the frame time is made up with the synchronising pulse. If both pins 5 and 6 are shorted to ground then a further four fixed duration pulses are sent and the synchronising pulse is proportionally shorter. This is clearly illustrated in Figure 2.

This takes care of the encoding and timing but leaves the problem of the actual transmission. For this purpose the LM1871 provides a single transistor with its emitter internally connected to ground (which removes a few design options but I think they ran out of pins!). As I wanted an oscillator whose frequency was stable but didn't want to spend money on a crystal I chose to use a ceramic resonator as the major frequency selective element. These resonators behave very much like crystals except that they are not quite as accurate or stable and seem to have a whole lot more spurious resonances. In order to control the spurii it is necessary to have the oscillating amplifier selective by itself so that where the spurii occur there is low loop again.

In the circuit chosen, the collector of the oscillator drives a tuned circuit and the power supply and ac earth is connected to a tap almost at the other end of the coil. Thus, the other end of the resonant circuit has a voltage in opposite phase to the collector and of much lower amplitude, suitable to drive the base. The signal to the base is divided down through C6 and C8, together with the ceramic resonator. Thus, the base drive is heavily attenuated for all frequencies except the resonant frequency of the ceramic resonator which controls the frequency of oscillation.

The output from the oscillator is taken from a secondary winding and one end is tied to ac earth via C5. The other side of the winding has a current limiting (and impedance defining) resistor R3, which connects to the +12 volt supply line. Thus, the oscillator signal is coupled at low level to the transmitter supply. R2, C1 and C2 protect the logic circuitry from interference from the transmitter output.

Capacitors C3 and C7 are used to wave shape the data pulses out of the digital section so the RF spectrum of the transmitter output does not "splatter", but this is principally for radio type applications and not strictly necessary here. They are included more for nicety that necessity.

The receiver and data recovery part of the main controller revolves around the National LM1872 which is the companion IC to the LM1871 transmitter. The receiver section of the LM1872 is designed as a superheterodyne receiver, but for this application all that was necessary was a tuned amplifier at 455 kHz with adequate agc as it is indeterminate what attenuation from the transmitter can be expected. For this reason, the local oscillator section of the LM1872 was disabled by grounding its crystal inputs, pins 1 and 2. This makes the mixer of the LM1872 into a simple amplifier (the local oscillator modulates the emitter current of a differential pair and grounding pins 1 and 2 sets a constant emitter current). Thus, what would normally be the RF input becomes an "IF" port and accepts the input from the line.

The actual input signal is coupled in by TR1, a 20:1 step down transformer. The primary is a single turn and the secondary is 20 turns wound around a Philips 3H1 9 mm toroid. This gives a secondary inductance of about 320 μH which resonates with C1, a 390 pF capacitor, at 455 kHz. Thus, the high input impedance as seen at pin 5 of the LM1872 is transformed down by 20^2 to present a few tens of ohms to the 12 volt supply line at 455 kHz and almost no impedance for other frequencies.

The next two stages in the IC provide more than adequate gain for our needs. TR1 couples into the base of one side of a differential pair and TR2 is connected to the collector of the opposite side. The secondary of TR2 couples via diode biassing to the lower section of a cascode pair which is also part of the agc system. TR3 acts as a simple resonator in the collector of the upper transistor of the cascode pair. This collector also drives the agc system and rf detector. Thus, as the agc system has high gain itself, the RF level seen at pin 15 doesn't vary and it is necessary to monitor the dc agc voltage when tuning the receiver.

Once the detector has recovered a binary signal from the RF data bursts the LM1872 must separate the two analogue channels and, by pulse counting, determine the state of the two transmitter switches. The LM1872, like the LM1871, has one master timer which is used to recover frame rate. This is set by R1 and C5. Whenever the LM1872 sees no incoming RF it resets the timer R1C5 and while a continuous carrier is being received the timer is allowed to exponentially rise (see Figure 3).

If the carrier is present for long enough, such as for a synchronising pulse from the LM1871, then a comparator in the LM1872 fires and its decoding cycle is started. The next ON to OFF transition is read as the start of the first analogue pulse and the next ON-to-OFF transition is the end of the first channel pulse and the start of the second. The third transition terminates the second analogue channel pulse and from then on the LM1872 merely counts ON-to-OFF transitions to determine the status of the transmitter switches.

The two digital outputs from the LM1872 are used to drive the cathodes of two optocouplers directly and the anodes are pulsed on zero crossings of the mains. For the two analogue channels the processing necessary is more complicated. The outputs from the IC are set as open collectors and R2 and R3 act as pullups to 5 V. All the timing from the LM1871 is set so that the two analogue channels have pulse widths of between 1 ms and 2 ms (with an allowance for tolerancing — it is less than one and more than two).

As the 75 Hz frame rate gives a pulse repetition rate of 13.33 ms then the two pulse-width modulated outputs have dc components of between 5(1/13.33) and 5(2/13.33) or approximately between 0.4 and 0.8 volts. Superimposed on this is a very strong signal at 75 Hz, plus harmonics. Therefore, the dc component must be dc-shifted and amplified to give an output of between 0 V and +3 V. The two active filters achieve this and provide an attenuation of about 60 dB for the 75 Hz. They are equiripple group delay filters with a phase ripple of 0.5°, which provide an optimum transient and frequency response for this application (see text).

As the dc output is being used to fire triacs, the dc must be related to mains 50 Hz to be of use. IC7, and principally transistors pins 9-10-11, labelled Q1, and pins 12-13-14, labelled Q2, are used as a mains zero cros- ▶

. . . How It Works, continued.

sing detector. The base of Q1 and the emitter of Q2 are both grounded and the emitter of Q1 is connected to the base of Q2 and, via R27 and R28, to the actual 240 volts 50 Hz mains (one of the resistors is connected to the active and the other is connected to neutral so that if active and neutral are reversed, then everything still works). When the mains is at greater than one Vbe above ground Q2 is turned on. Also, when the mains voltage is greater than one Vbe below ground then Q1 is turned on (in fact it must be less than −5 V before Q1 saturates completely). Only for mains voltages between ±0.6 volts are both transistors completely off and their common collectors, pins 11 and 14, at +6 volts. Thus, the two commoned collectors have a positive-going pulse every time the mains voltage passes through zero. This pulse is buffered and cleaned up by two gates in IC6.

The buffered, positive-going pulse is impedance transformed by the rest of the transistors in IC7 acting as a Darlington emitter follower and used to reset a ramp generator formed by Q3 and its associated resistors and capacitors. C18 is forcibly charged to about 2.6 volts through D1 on a mains zero crossing (the other side of C18 is held at 0.7 V by Q3's base). C18 is then charged through R20 and the current is mirrored by collector current through R21. This produces a linear negative-going ramp at the node of C18, D1 and R21. This ramp is applied to the non-inverting inputs of the two analogue channel comparators while the dc signals derived from the filters are applied to the inverting inputs. The two gates of IC6 following the comparators ensure that even when the comparators are always low, which corresponds to lights full on, there is still a pulse on mains zero crossing to fire the triacs. The gate outputs are CR differentiated and buffered, then used to drive the MOC3021 optocoupler input anodes directly.

The four triacs have one terminal connected permanently to the active mains input, the other is connected to the load. This means that when the triacs are off, the driven light fittings are safe. A second, much smaller, triac in the optocoupler is used to trigger the main one on for each circuit. A resistor-capacitor-resistor combination protects the optocoupler triacs from excessively high dV/dt and hence accidental triggering from mains transients. The two triacs that are driven from the dimmer controls have simple LC filters in their outputs to suppress radio frequency interference. The on-off triacs have no need for this as they are always fired on zero crossings.

Circuit of the encoder/transmitter.

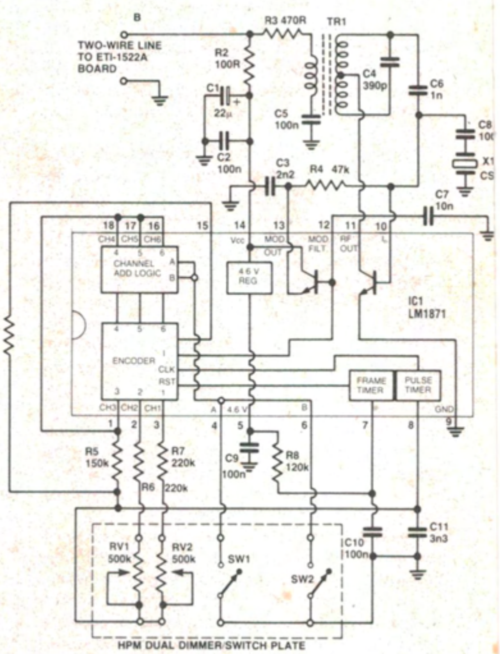

opposite direction. The coil used to operate the transistor in the IC provided as an oscillator also gave an easy way of injecting the oscillator output onto the line. Oscillator frequency stability was provided by using a 455 kHz if ceramic resonator (readily available from Tandy, and others) as part of the feedback.

As the LM1871 provides all the necessary circuitry to modulate and waveshape the oscillator output from the encoder output, a complete and compact controller was easy to construct.

The actual controls feeding the LM1871 were made from modified light fitting controls purchased for the purpose. I chose to use HPM light fittings as they have a neat and convenient arrangement where a panel is purchased with the required number of holes as there are to be controls and then the controls (switches or dimmers) are snapped into the panel.

In order to keep things looking exactly like ordinary light fittings I bought two light dimmers and removed all the works from them except the 500k pots, which are the dimmer control elements (see How It Works). The end result can only be destinguished by the two mounting screws for the electronics printed circuit board. In order to save space in the controller all capacitors are the smaller 5 mm pin spacing, 63 V metallised film type which are increasingly becoming an industry standard and are available from many manufacturers.

Once the encoder/transmitter had been designed and modelled, the next problem was to design a modified receiver/decoder which would recover the 455 kHz bursts and regenerate the analogue and digital information. The LM1872 contains all the circuitry for a complete superheterodyne receiver which was somewhat more than was strictly necessary for this application. The on-chip local oscillator was therefore disabled and only the IF amplifier section used. Two of the IF transformers from a set carried by Dick Smith Electronics (L-0260) were used and gave satisfactory performance.

The 455 kHz was coupled from the power supply line with a toroidal transformer whose output was approximately resonated with C1, a 390pF capacitor. This means that the input impedance of the RF section is accurately transformed down to about 20 ohms to recover the 455 kHz from the line. The agc for normal applications is filtered with a 100 nF capacitor on pin 16. In this application, where there is no variation in input signal level, the filter capacitor can be increased to 1 μF, or even grater, with no problems and in fact helps by preventing the agc following the input pulses.

The LM1872 IF and detector sections recover the pulses generated by the LM1871 encoder and pass them onto the decoder section whose operation is shown in Figure 2. It can be seen that the decoder simply counts negative-going edges after the synch pulse to determine the status of the two switches in the LM1871 circuit. The timing diagram also illustrates just how the LM1872 derives timing information from the first two pulses after the synch pulse for the analogue channels.

Figure 3 shows how different numbers of pulses are sent from the LM1871 to convey switch status. As the frame timing in the LM1871 is set by a separate resistor and capacitor from the data RC circuit, as the number of data pulses is reduced the synchronising pulse gets longer so the frame rate remains constant at 75 Hz.

In the decoder circuit used, both the analogue and digital outputs were donfigured as open-collector transistors (the LM1872 offers a choice of open-collector or open-emitter drivers for the digital outputs). Processing of the digital outputs, was no problem and the collectors switched the triac drives directly (see "How It Works").

However, the analogue channels needed a lot of processing before they could be used to operate triac drives. As the whole of the LM1872 had a maximum operating voltage of 7 V the whole of the receiver and triac drive circuitry was operated from a regulated 5 V rail. (I'm sure National had a good reason for the voltage restriction but it escapes me).

When resistors R2 and R3 are connected to the output collectors the resultant output is a pulse-width modulated signal with the pulses swinging between about 40 mV (the transistor V_{ce} sat) and 5 V. This signal has a dc component and a somewhat larger ac signal. In order to recover and amplify the dc, and at the same time remove the ac, several methods were considered but finally the simple and straightforward way was chosen; namely — a low pass filter to preserve and amplify the dc and at the same time provide about 60 dB attenuation of the 75 Hz fundamental (and appropriately more for the higher harmonics of the ac signal).

When it comes to filters there must be about as many different types and class of filter as there are applications; and there are an equally large number of realisations for

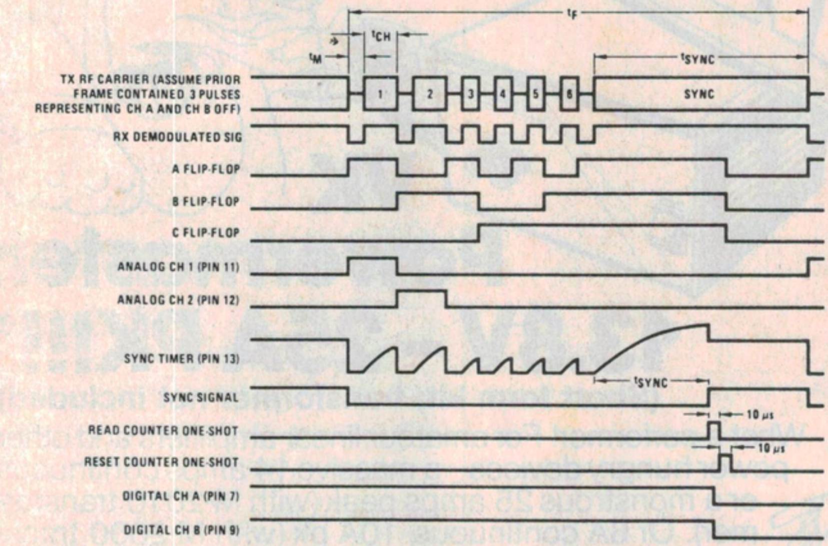

Figure 2(a). Receiver timing waveforms.

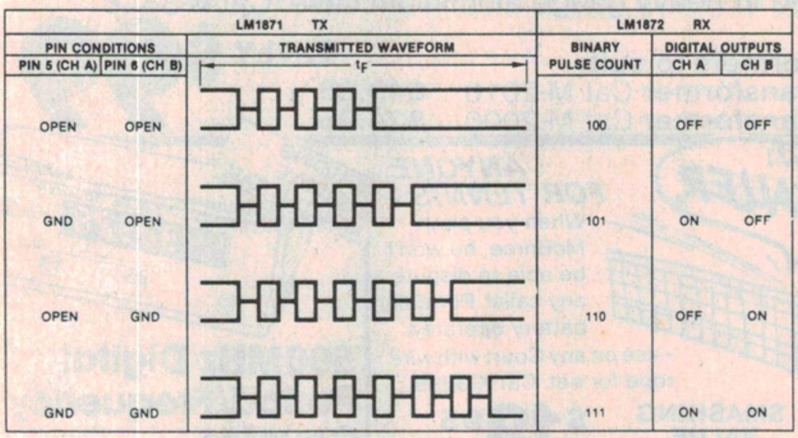

Figure 2(b). Digital channel encoding and decoding via pulse-count modulation.

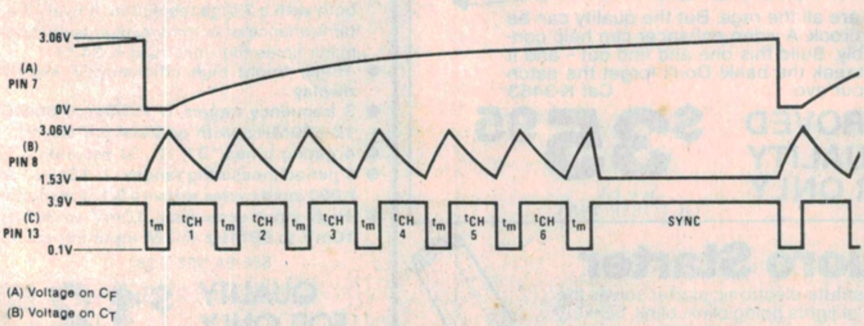

Figure 3. Simplified waveforms for the encoder timing circuits.

each class. Filter types vary from *ye olde Butterworthe* (whose chief merit lies in the simplicity and regularity of its transfer polynomial) up to the more exotic *Cauer-Chebychev* filters which give optimum cutoff response but fierce group delay distortion (and hence ring like a bell).

For this application we really need optimum step response, as if the filter rang then, when the lighting intensity was suddenly changed, it would brighten and fade several times before settling down. This would suggest a Bessel filter which has an optimised step response (it doesn't overshoot or ring at all) but a rather slow and soggy frequency response. However, we can do a little better than Bessel's filter by accepting a tiny amount of overshoot (only a few per cent) and using an equi-ripple group delay filter where the filter phase response is allowed to vary a small amount around the ideal linear (Bessel) response and a slight improvement in frequency performance obtained.

The filter actually used is a 5th-order, 0.5° equi-ripple group delay filter which sounds all gosh! gee! wow! fantastic!! but nonetheless actually works just fine. (Linear phase or equi-ripple phase filters also have the advantage that they don't have any very high Q sections and so are a bit more forgiving on component tolerances).

The first section of the filter is a 3rd-order block with the capacitor values shifted to preferred values and the resistor values varied accordingly (this involves fun things like solving a hextic equation (you beat it to death with a calculator!)), but it means you can buy the components in this country. The second section is only a 2nd-order block but it has all the gain necessary to bring the pulse-width modulated dc up to about 0-to-3 volts.

At this stage in the circuit the fact that the information was transmitted at 75 Hz has been removed and all that remains is to relate the dc signals to the mains 50 Hz. In order to do this, IC7, a transistor array, is used to generate pulses on zero crossings of the mains and these pulses are used to generate a ramp which is compared with the dc from the filters (once again, see "How It Works"). After suitable gating to ensure the triac trigger pulses occur when wanted and don't when not (for a pot fully anticlockwise it is essential that the lights be completely out and when fully clockwise they should be full on — care was taken here with end-limit tolerances) the processed analogue outputs are fed to the triac drives.

All four triacs are triggered through opto-isolators to ensure that all the electronics is safe to work on **but beware!! there is 240 volts on the board!!** The two triacs that are switched from digital outputs have no output filtering as they are triggered on zero crossings, but the two dimmer triacs are filtered to suppress RFI. ●

. . . to be continued, next month.