Jan 2015

A Junk-Box Dawn Device

Jamie Davies, MM0JMI

Purpose:

In the winter months, mornings start late this far north and most working people have to wake when it is still dark. This is much less pleasant than the natural experience of waking with the dawn, as we do in spring and autumn.

Katie hinted that the bedroom in her flat would benefit from a device that brings room light up gradually to create an 'artificial dawn' just before she has to get up. Her bedroom already had a standard lamp easily powerful enough to create the right effect, so the challenge was simply to arrange for it to brighten from nothing to full brightness gradually – not as gradually as a real dawn, but over a few minutes – at a time of her choosing every morning.

This is the story of its design and construction, written up in case anyone else finds parts or all of it it useful for a similar project.

Design considerations:

The best design would:

1. Be safe.
2. Be simple to operate (no more difficult than a normal alarm clock)
3. Keep time without a mains connection (Katie's cleaner has a habit of unplugging things)
4. Be silent.
5. Be built from parts to hand in my collection of 'may be useful one day' electronic components, or available on a Sunday in a small Scottish seaside town. I did not want to wait for mail orders to come from specialist electronic dealers. Hence the 'junk-box' approach.

These considerations resulted in the following design:

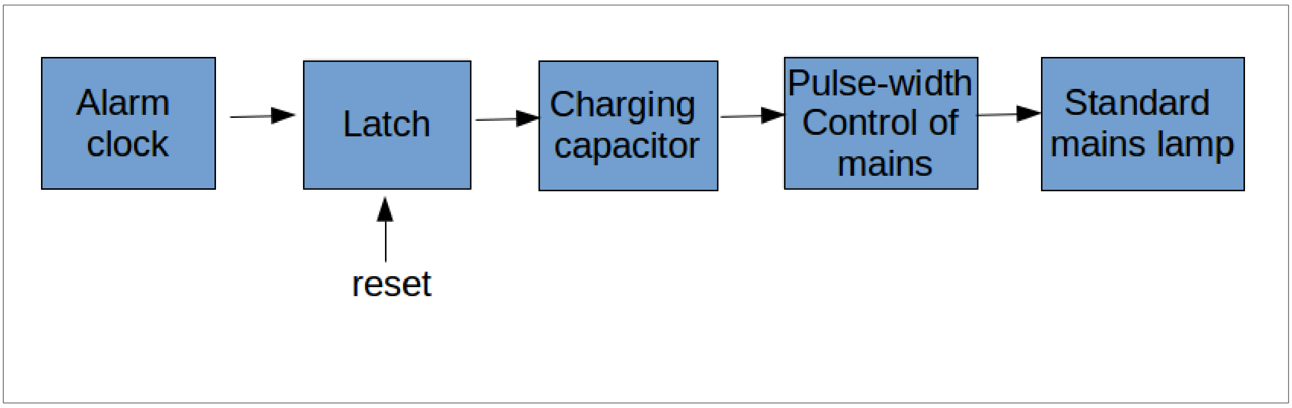

The use of a standard battery-powered alarm clock deals with points 2 and 3, and using a digital one rather than a clockwork one also deals with point 4. Travel alarms are also easy to obtain from my local household hardware shop. The idea of using the rising voltage on a slowly charging capacitor as a brightness control allows the use of 'junk box' components with no need for micro-controllers. The idea of using the existing mains lamp keeps everything simple and avoids extra clutter.

To meet criterion 1 (safety), I decided to use a commercial, CE-marked mains controller module from Kemo, which was spare from a previous project. While I am happy to build high voltage stuff for my own radio ham purposes, having a device connected to the mains in the same room as someone is sleeping raises the safety bar, so a tested professional module seems a better idea.

Buying and re-purposing the clock.

The clock was the only thing bought especially for this project (though some other things used were new, in the sense of having lurked unused in my bits-and-bobs collection). I purchased a digital alarm clock from a household goods store in North Berwick that tends to sell end-of-line or otherwise 'odd' goods at very low prices. The clock was made by DSL ("Item Number 1713"), and the box looked like this:

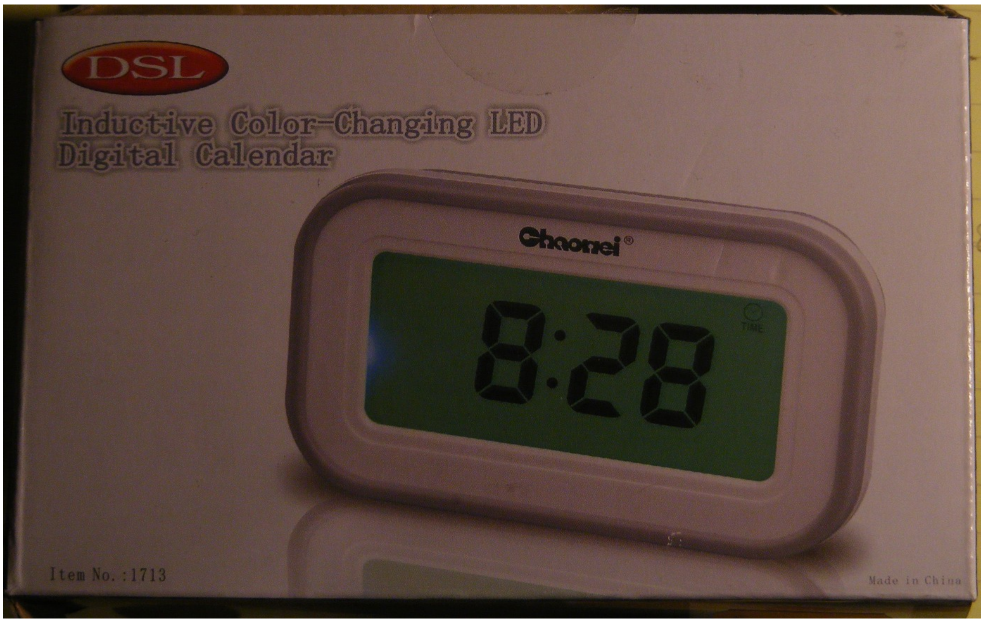

The box says 'calendar' but what was pictured and inside was obviously an alarm clock: the mistake on the box is typical of the things that end up in that shop, and probably the reason for the low prices. Anyway... the clock is probably typical of the breed. I chose this one because the back was held on with screws (thank you DSL! So many others seem to be glued or use clips that are never intended to be released).

The unit runs on 3 AAA batteries, and makes its alarm sound via a small sounder (piezo-electric, I assume) feb by two small wires. I cut these wires and joined them to a twisted pair of red/black hookup wires, soldering the connections, and insulating with heatshrink. I also drilled a small hole in the bottom of the clock, through which the wires emerge.

NB – when you are making the above modification, be aware that *it is not a good thing to let passers-by see, through your window, an alarm clock with home-made wires hanging out of it*: nowadays, your motives may be misunderstood.

Using the alarm's buzzer-feed as a trigger

It turned out that this clock did not simply feed a self-contained buzzer unit with 4.5V dc.. Attaching the wires added above to a bench DVM, to check polarity, and pressing any button on the back of the clock to create the now-silent 'bleep', resulted in a bizarre mix of numbers on the meter and a lot of relay clicking as it changed range. Some readings were far above the 4.5V expected from three 'AAA' batteries. A hand-held DVM showed similar readings. My oscilloscope (not an instrument I expected to be using in this simple project!) showed a bizarre, vaguely-square wave-form of about 75V p-p, with thin spikes reaching well over 100V. The whole thing was very high impedance and easily crushed by even a few hundred ohms across the wires. I assume there must be some kind of multivibrator and coil to generate this inside the clock, and that the clock's sounder is a very high-impedance piezo device.

I wanted a robust low voltage output to trigger the Dawn Device itself. I wasted a lot of time trying to use an LT700 audio transformer, a device once common in early transistor radios that has a step-down of 20:1, to turn the high-impedance, high voltage signal into a low-impedance, low voltage one but even the 'high impedance' side of this was enough to kill off most of the signal.

Then I remembered a useful property of LEDs: fed even with very high voltage and horribly messy signals at high impedance, they will light and survive (the famous 'RF sniffer' of George Dodds, G3RJV, works this way: it's just a probe, a rectifier, an LED, and the hand capacitance of the person holding it). I found that connecting a normal green LED directly across the leads from the clock caused the LED to glow when any button was pressed. This suggested that an optocoupler may be a reasonable means of converting the unhelpful output of the clock to something more useful. Connecting the LED side of a resistive optocoupler in place of the LED caused the resistance of the optocoupler's output to drop from over 1MΩ to a few kΩ. Resistive opto-couplers get a bad press nowadays because their response times are so slow, hundreds of milliseconds at best. For this application, though, slow response is a really desirable property and a resistive opto-coupler forms an excellent ultra-low-pass filter, as the plodding light-dependent resistor averages the light pulses from the LED that illuminates it. I really like resistive optocouplers for slow control purposes, which is why I usually keep a few in stock.

Having solved the problem of interfacing the clock, the rest of the device could be designed.

The circuit itself.

The circuit for the Junk-Box Dawn Device is shown on the next page. The power supply is conventional. The 12V transformer, rectifier and primary smoothing capacitor were salvaged from an ancient Zip drive power supply. The rectifier and smoothing capacitor were mounted on a small board permanently attached to the transformer itself: I left them there. The 12V is taken down to a regulated 5V by a 7805 regulator: this voltage was chosen to suit the relay I had (see later). The values of C3 and C4 are not critical – they are just smoothing capacitors.

The incoming signal from the clock buzzer feed is changed, via the opto-coupler OPTO1 (the one mentioned above: a Silonex NSL32), into a markedly falling resistance between the +5V line and the base of Q1. The 4k7 resistor is present in series to limit maximum current to Q1's base. D1 is not actually necessary, and is as relic from my attempts to use a transformer and rectifier before I decided on the opto-coupler. The change in Q1's conductivity, when the clock buzzer signal was on, was not enough to switch RLA1 reliably. Adding Q2 in a standard 'Darlington pair' configuration resulted in reliable switching. D2 is present between the emitter of Q2 and the relay to ensure that current does not flow the wrong way into Q2 when Q2 is 'off' but current is flowing towards the relay via D3 (see below). D4 shunts back-EMF from the relay coil, to protect all three transistors and the electrolytics from reverse polarity damage: this is standard practice for hybrid relays and semiconductor circuits.

The latch function is implemented with a relay and diode logic. When the relay is off (before the alarm clock tries to sound its buzzer), the connection between +5V and the anode side of D3 is broken. When Q2 admits enough current to the relay coil through D2 to switch the relay 'on', the contacts close and the anode side of D3 is connected to the +5V line. Current can now flow from this line via D3 to the relay coil, so that the coil remains energised even when the alarm clock falls silent and Q2 stops conducting. The rest of the circuit draws power from the anode side of D3 also, so the effect of the relay-and-diode logic is that a pulse from the alarm clock latches the rest of the circuit 'on'.

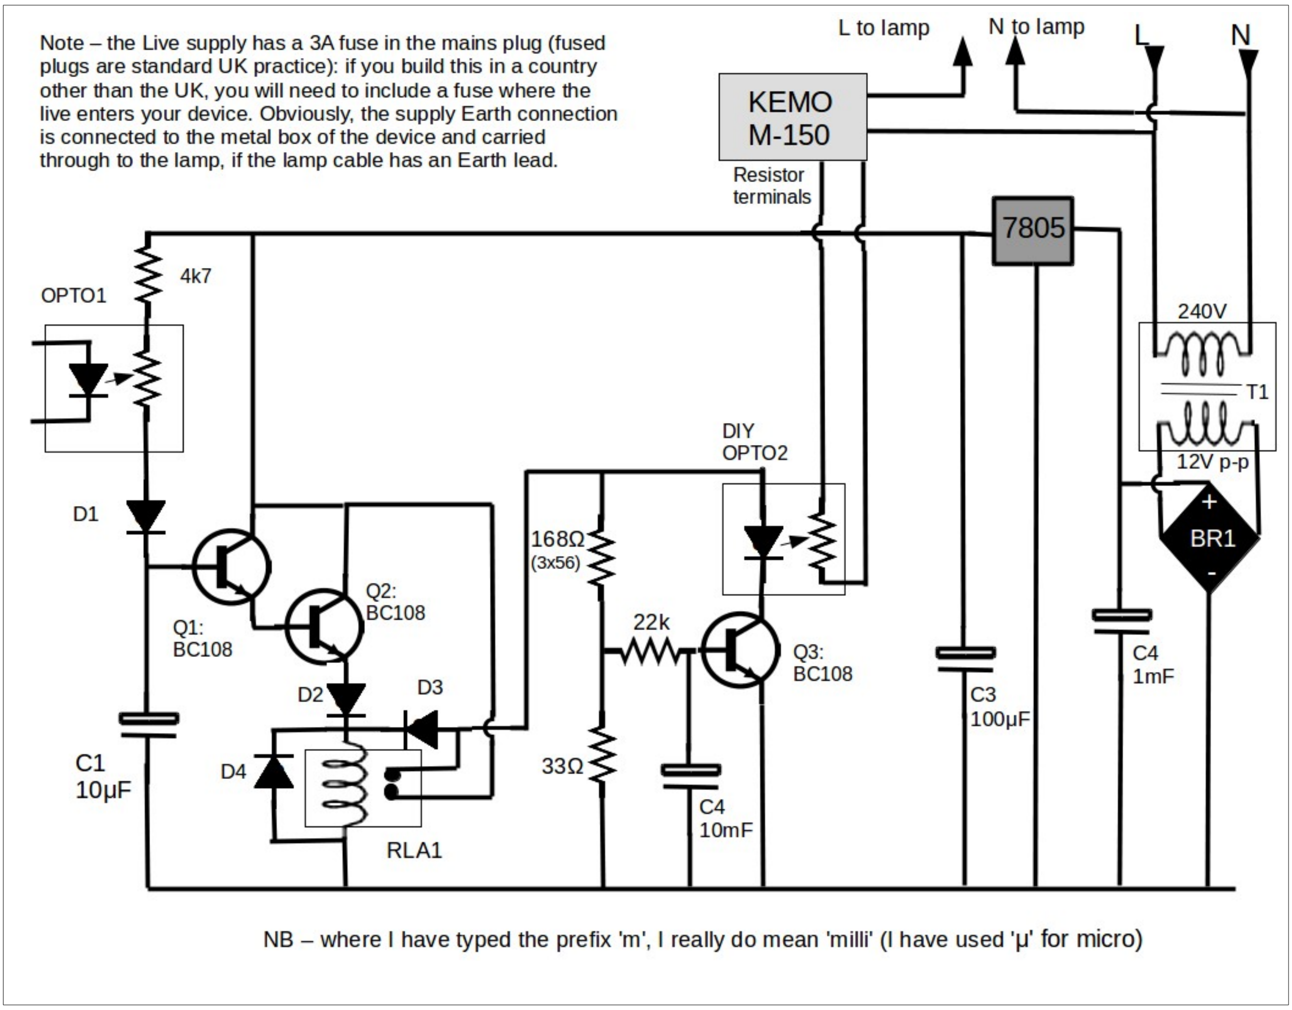

As soon as the anode side of D3 connects to 5V, current starts to flow into the circuit around Q3. The slow rise in brightness of the lamp comes from the slow charging of capacitor C3. Obviously, this voltage cannot be used to drive the relatively low load of the optical isolator OPTO2 directly without OPTO2 discharging the capacitor: instead, the capacitor needs to be connected to a relatively high-impedance input of an amplifier. The obvious candidate would be a JFET (they have very high input impedance and are directly sensitive to voltage) but I did not have any lying about. I did have plenty of BC108 bipolar transistors, a ubiquitous general-purpose NPN type. Bipolar transistors are sensitive to base current, rather than voltage, but the current admitted to the base is related to voltage, with a lower cut-off when the base is about 0.5V above the emitter, and full switch-on current (enough to drive the collector-emitter current to its maximum) when the base is about 0.65V above the emitter; another 0.15V can be added to this for a 'safety margin' of assuring a fully 'on' state, given that individual transistors vary and I tend to buy them in 'bargain bags'. This 0.5V → 0.8V is quite a close range; there is therefore no point in charging C3 right up to 5V. C3 was therefore connected, via a resistor to limit its charging rate, to a point in a simple potential divider that would hold at a maximum of 0.8V with no load being drawn from it. That way, C3 and therefore the base of Q3 will never reach more than 0.8V.

The potential divider was made with a 33Ω bottom resistor and a 168Ω top resistor, this '168Ω' resistor actually being made from 3 x 56Ω resistors in series: 5V x 33/(33+168) ≈ 0.8V. If I needed the light to begin to come on the instant the device triggered, it would have been necessary to tie the bottom end of C3 not to ground but to say 0.5V, so that C3's voltage varied just through the transistor operating range of 0.5V → 0.8V. In practice, it does not matter if C3 is charged all the way from 0V – it just means that the charging time from 0V – 0.5V would be 'dead time' when the lamp would do nothing, and brightening would occur only after this. As it happens, 0.5V is almost exactly 63% of 0.8V, and the time constant τ of an RC circuit is defined as the time taken to reach 63% (that is, $1-e^{-1}$) of final charge. Therefore, if the bottom end of C3 is just grounded, the lamp will remain dead for one time-constant, and will brighten over the next one. The time constant of an RC circuit is RxC; with C3 = 10mF (again, using what was to hand) and a desired time constant of about 4 mins, R comes out as 22k (rounded to the nearest value of easily available resistor). So, with these values in the circuit, we can expect a delay of 4 minutes between clock alarm time and the first brightening of the lamp, and a further 4 minutes or so (depending on the exact characteristics of the individual transistor used for Q3) of gentle artificial sunrise before the lamp remains fully on until reset. Here are the real data:

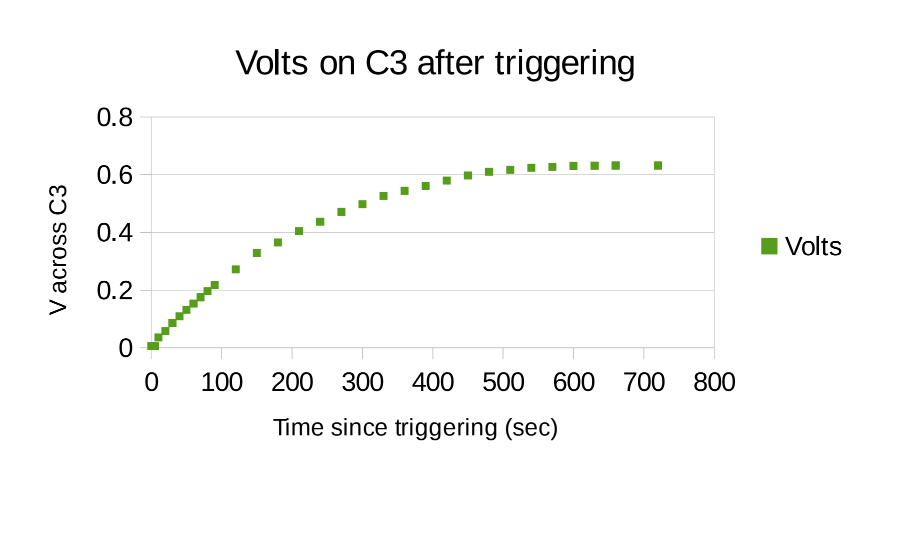

The voltage rises asymptotically towards a plateau. The time constant (time to 0.63x plateau height) is about 4 mins, as expected. The plateau is lower than expected for no-load conditions, 0.63V not 0.8V, presumably because of the current drain into the base of Q3.

As the voltage rises, modest currents begin to flow into the base of Q3, and therefore the collector current rises and the LED in OPTO2 begins to glow. As the LED becomes brighter, so the resistance of the light-dependent resistor (LDR) in the resistive optocoupler OPTO2 falls. Here are the real data from the final design of optocoupler:

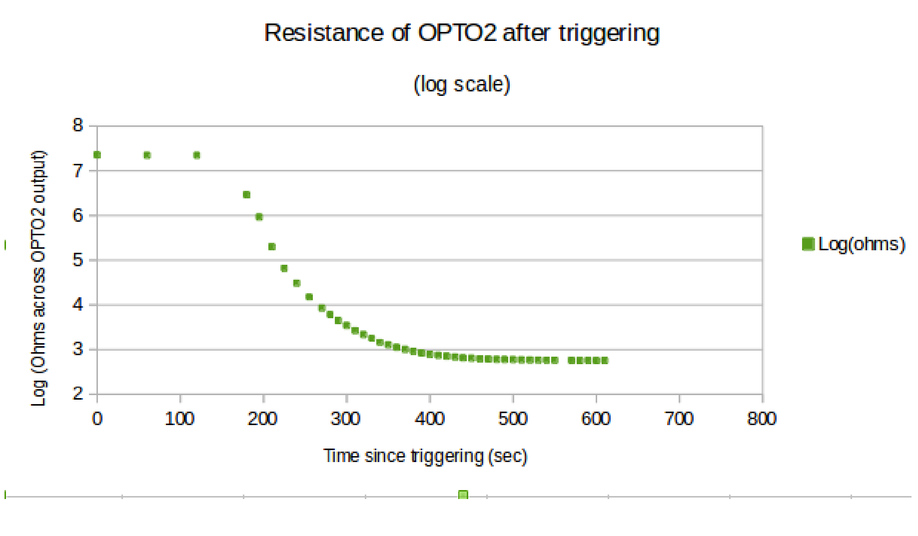

Significant changes happen after about 1 time-constant (again, as expected), and continue for about 1.5 more time constants.

The optocoupler drives the Kemo M-150 mains module, which expects a 470k (lin) potentiometer between two of its terminals. There is a problem with the M-150, though: those terminals for the resistor have a potential across them of hundreds of volts ac. The data sheet does not tell you this. I assume that the black box of the module contains a Triac, the gate of which is connected, probably via a Diac, to one of the resistor terminals, the other of which is connetced to the Live terminal. This voltage is too high for a Silonex NSL32 optocoupler. The data sheet says max 60V: so did the smoke pouring out of one when I tried it in the circuit before reading the data sheet carefully. I had no other ready-made resistive optical isolators, and transistor ones would be useless for high voltage ac. I therefore made a DIY optical isolator by gluing an LDR (light dependent resistor) into one end of a 1cm length of thin black drinking straw, and gluing an ultra-high brightness 3mm white LED into the other. Breadboard experiments with this confirmed that it showed a roughly 300kΩ → 300Ω swing in the resistance of the LDR as the LED was powered up from dark to 10mA (achieved, for the purposes of test, by placing a 1k resistor in series with it and connecting the pair across a 9V battery). The whole assembly was stabilized in heat-shrink tubing.

Final connection to the mains was made via a switched plug (these are available via various traders on ebay: the plugs take a conventional 3/5/13A fuse and have a large switch and neon indicator moulded into their top). This switch provides a simple reset (just switch the whole thing off for a few seconds – long enough for C1 to discharge is enough. It does not matter if C3 is still somewhat charged because that cannot switch the latch back on; C3 will gradually discharge via the base of Q3 and via the 22k and 33Ω resistors anyway).

Building

The circuit was built on some spare Veroboard (probably the last time I will ever use the stuff – I use Manhattan construction for most things now). I built everything into a die-cast aluminium box, bolting the transformer and mains module into the box, and bolting the circuit-board in on stand-offs. The Earth connections (from the mains plug and to the lamp) were made by ring terminals to a separate bolt through the case, with a locknut ensuring this safety-critical assembly stays tight.

Mains connections were by spade connectors that were insulated (shrink-tube) for my safety when making measurements. My version of the Kemo module had spade terminals – I think other versions come with permanently attached leads. To the back of the case I bolted a pipe clamp, so that the box could be mounted at the base of the metal stem of the standard lamp. The clock was glued to the top of the case, the 'buzzer feed' cable passing through holes where the two boxes met so that it is invisible from the outside and no danger to anyone.

Remember, when choosing a position for your clock, that you will need to be able to reach the batteries to change them one day, and also to reach the programming buttons to set alarm time (in this clock, the order of buttons on the back, seen from the front side, is down-up-set-mode). Also, remember you will need a Tungsten or a Halogen bulb in the lamp: most low-energy types of bulbs do not tolerate triac-type dimmers.

<End of file>