

## Constructing the Programmable Christmas Star

The Christmas Star runs through a programmed pattern display sequence, held in EEPROM on a PIC microcontroller. With a suitable PIC programmer, it is possible to load your own custom display sequence into the EEPROM, without affecting the underlying code – without needing to understand PIC programming.

### How it works

The diagram below shows the complete circuit, which consists of little more than the pre-programmed PIC12F683-I/P microcontroller, 20 LEDs and a few resistors.

Using a technique called “Complementary LED Drive”, it is possible to drive 20 LEDs with an 8-pin PIC and five resistors. It relies on two factors:

1. LEDs will only conduct (and produce light) when a high enough forward voltage is present. If the applied voltage is too low, or reversed, they won't light up.
2. The PIC12F683 has tri-state outputs. That is, they can be set high (nearly 3V in this circuit), low (close to 0V), or placed into a high-impedance input state, effectively disconnecting them from the circuit (“off”). Further, the outputs can either source or sink current, up to 25mA.

As an example, consider what happens when the PIC is configured with pin GP5 high, pin GP0 low, and pins GP1, GP2 and GP4 tri-stated (disconnected). Current will flow from GP5 through resistor R1, then LED L19, returning through R5 to GP0. So L19 will light up. Since LEDs are one-way devices, current can't flow through L20, so it stays off.

But there are other paths for current to flow from GP5 to GP0. For example, via LEDs L9 and L10 in series. But these two LEDs in series are in parallel with L19, and L19 is conducting. The forward voltage across a conducting LED is roughly constant; for a red LED it is around 2V. The voltage drop across the series combination of L9 and L10 must be the same as that across L19. So each of L9 and L10 will have a forward voltage of only a half that of L19. If L19 is turned on with a 2V drop, there will be a drop of only 1V across each of L9 and L10 – not enough to make them conduct. So they won't light up.

You'll find many other possible paths for forward conduction; a particularly obvious one is the series combination of L1, L2, L3 and L4. Similar reasoning shows that the voltage across each is only a quarter that across the conducting L19; not enough for them to turn on. Similarly for other paths, such as the non-obvious L13, L16, L17 combination. In fact, with GP5 high, GP0 low, and the other outputs disconnected, only L19 will have enough forward voltage to light up.

Using this technique, it is possible for only five outputs to uniquely address up to twenty LEDs – but only one can be turned on at once. To overcome this limitation, the software uses multiplexing to make it appear as though more than one LED is lit at the same time. The software displays patterns on up to four LEDs which are turned on in sequence, each for 200µs, at nearly 1250Hz, creating the illusion that the four LEDs are on at once.

The remainder of the circuit is very straightforward. Resistors R1-R5 limit current to the LEDs. The current path to a given LED will always flow through two of these resistors, so the effective resistance in series with each LED is 94Ω. Assuming a 3V power supply, and a red LED with a forward voltage drop of 2V, LED current will be 10mA, well within the supply capability of the PIC.

Switch S1 is used as an on/off switch. Resistor R6 holds PIC pin GP3 high until S1 is pressed, pulling the input low. The software polls for this at the end of each display cycle and if S1 is pressed, it puts the PIC into a low-power sleep mode. The PIC is then set to automatically wake up if the switch is pressed again. Debouncing is done in software, so there is no need to external debounce circuitry.



more than one LED light up strongly at once, you probably have one of them in backwards, or perhaps a solder bridge on the board. If one combination doesn't produce any light, while others do, you probably have either a dead resistor or LED, or a soldering problem such as a dry joint.

If all the LEDs check out ok, remove the batteries from the battery back, cut the leads suitably short (15mm or so if using an N cell holder; the leads on a AAA holder will need to be longer), thread them from the back of the board through the hole above C1 and solder to the pads marked + and -, being careful of polarity! If you now reinsert the batteries, nothing should light up; if it does, you have a short somewhere. Next remove the batteries again and use double-sided foam tape to affix the battery holder to the bottom of the board. An N-cell holder fits "vertically" under the battery lead hole; an AAA holder needs to be placed "horizontally" across the board.

Finally, you're ready to insert the microcontroller. Taking antistatic precautions (touch an earthed case first!), be careful to insert the PIC into the IC socket, with the notch on the IC toward the capacitor. Now insert the batteries again and you're finished! At this point, the display may start by itself, but more normally the star will do nothing until you momentarily press the button. The display sequence should now start.

## Operation

Very simple – push the button to start, and press it again to stop. But if you forget and leave the display running, the star will shut itself off after around 3 hours. If this happens, just press the button again to restart.

## Parts list

1	Pre-programmed PIC12F683-I/P
1	100nF monolithic capacitor
5	47R 1/4W resistors
1	10k 1/4W resistor
20	5mm LEDs – any colour, but recommend clear lenses, five colours (four LEDs of each colour)
1	6mm PCB tact switch
1	2-way N or AAA cell battery holder with fly leads
2	N or AAA cell batteries (preferably alkaline)
	Double-sided foam tape (to mount battery holder)