Index *Skill Level*

Chapter 1: LEDs

To design a flasher to order it is important to understand how these parts work. LEDs are simple enough, but they have been around for a long time, and have changed quite a bit from their first commercial release. The old parts were fairly dim, and couldn't use much current. It is now possible to buy LEDs that will use over an amp and easily outshine most light bulbs. This article will deal with the dim to medium 5mm type of LEDs, since that is the majority simple ICs can easily power.

LEDs are *current devices*. This means they operate on current once a minimum voltage is provided. Like conventional diodes, they do not limit this current, another component has to do this. Connect an LED to a power source without a resistor and it *will* be damaged, probably burned out. Figure 1.1 shows the conventional scheme to light up an LED.

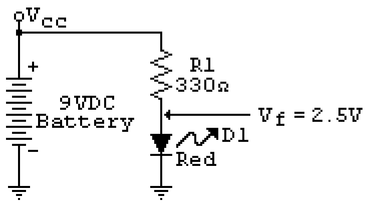

Figure 1.1

The forward dropping voltage, or Vf, of an individual LED is very stable. Go below this voltage and the LED stops conducting. This LED is assumed to be 2.5V, pretty standard for a modern red unit. The target current is 20ma. Going though the math (using Ohm's Law) the resistor is 325Ω. Since 330Ω is the nearest standard resistor value 330Ω it is.

Here is the approximate Vf of most LEDs:

	Older Generation	Newer Generation
Current	10ma	20ma
Red	1.5V	2.5V
Yellow	2.0V	3.0V
Green	2.0V	3.0V
Blue		3.5V
White		3.5V

For the Vf of a specific device you need to refer to the datasheet, and also understand there will be some variation even within a family. Part of the reason LEDs have changed so much is their efficiencies have gone way up. A modern LED at full power can damage your eyes if held directly next to the eyeball with the light shining in. Obviously these are not toys for children. Older LEDs didn't come close to these power levels.

LEDs can also be chained to share the same current to light more than one LED. Since this current is being used twice the apparent efficiency to light these LEDs is increased. Given that the LEDs can vary their Vf it is a really bad idea to parallel LEDs directly. Figure 1.2 shows a fairly typical example of how to do both for increased lighting.

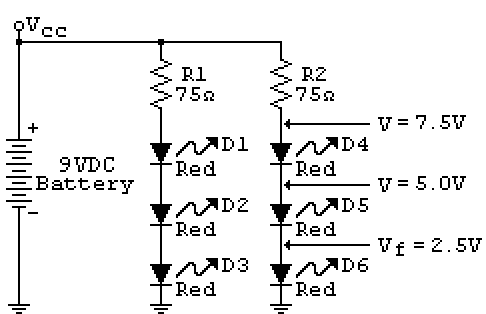

Figure 1.2

The reason it is such a bad idea for parallel LEDs to share their current limiting resistor is normal variations in Vf can cause one leg to draw more current than the other. This can result in the failure of one chain over time, leaving the second chain to absorb all the current. If you have a lot of LEDs in parallel this can lead to a progressive cascade failure, with LEDs popping like corn. You might be able to get by with it, but it is definitely not good design practice.

Chapter 2: Current Limiting

If you are dealing with a stable power supply a resistor is good enough. Be sure to use a resistor that is twice the wattage (or more) than is actually needed. Wattage equals the voltage squared across the resistor divided by the resistance ($P=V^2/R$). This is because some resistors may shift in their values if baked out, or overly stressed.

If the LED current is critical and you need precision, or if the power supply is less than stable, as in the case of automobiles, then better might be needed. A car can vary from 12VDC (battery) to 13.7V when running. This may seem like a small change, but it can create a significant current variation in practice.

The way around this is to use either a constant current source (current regulator) or voltage regulator. Used properly these circuits will stop power supply or LED Vf variations from affecting the design.

The LM317 is an excellent IC for this use. It comes in a wide variety of transistor packages big and small, is easy to use, inexpensive, and has excellent performance characteristics. It can be a voltage or current regulator. It's only downside is it drops about 3 volts. Figure 2.1 shows the two ways of using it's current regulation mode, Vcc can be 5.5V up to 37V, the LED doesn't change its brightness a bit (though the LM317 will get hot, and possibly burn up if not properly heatsinked for extreme voltage). The TO220 case style is shown because it is one of the most available models, and it dissipates heat extremely well.

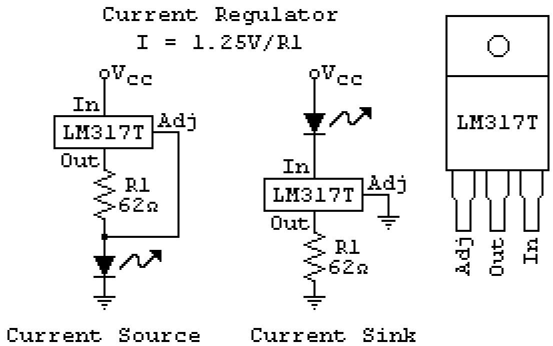

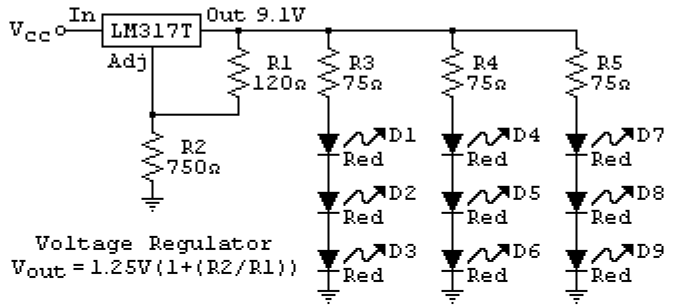

Figure 2.1 *Figure 2.2*

In the figure 2.2 the current is kept constant by keeping the voltage constant. This way one regulator IC can handle many more diodes. The LM317 requires 10ma minimum on its feedback leg, so 120Ω for R1 is pretty much a requirement, though lower values can be used (with an increase in current and no improvement in performance). If there is a long length of wire between the output of the LM317 and its load (the LEDs) you should add a 0.1μF and 10μF capacitor to the input and output pins of the LM317 to prevent the regulator from oscillating.

The 3V drop between the input and output of the LM317 IC can make it unsuitable for some uses. Lets go back to the automotive circuit, where the Vcc can vary between 12VDC and 13.7V. We'll start with this example in Figure 2.3.

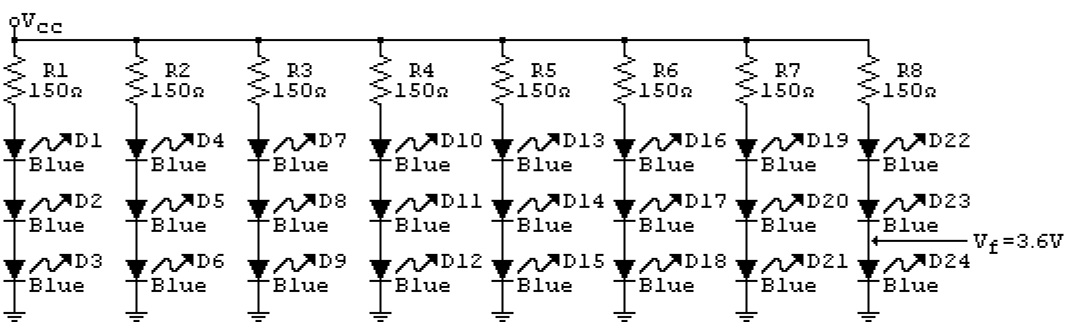

Figure 2.3

Each leg the total voltage drop across all three LEDs is 10.8V. If Vcc is 13.7V, then the current through each leg is 19.3ma. These LEDs were rated at 20ma, so the number matches nicely. However, if the voltage goes to 12V the current in each leg drops to 8ma. Quite a difference, and the LEDs will be a lot dimmer. This would be unacceptable.

If you change the resistors to 56Ω to power the LEDs with 21.4ma at 12V then they would get 51.8ma at 13.7V. Again, this is unacceptable. A regulator is needed. However, remember that the LM317 drops 3V. At 12V it could output 9V, at 13.7 it could output 10.7V. You could remove one of the resistors in the chain, but to use the same number of LEDs the total current would go up by a third.

Being willing to remove an LED

Chapter 3: The LED / Resistor Only Bargraph

LEDs tend to drop a constant voltage when they are conducting. It's not perfect, but it can be used. Take the following schematic in figure 3.1. I've included a schematic for a simple variable power supply using transistors and two 9V batteries if you want to experiment with them.

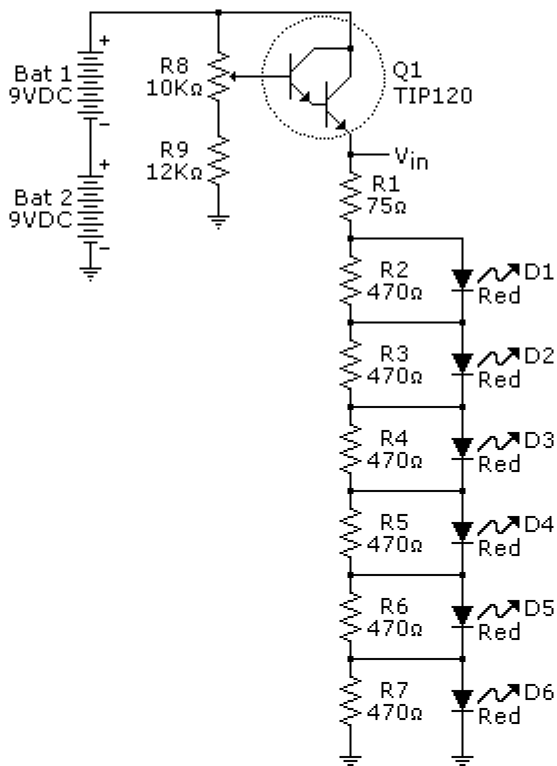

Figure 3.1

As you raise the voltage into the circuit you will see some of the LEDs brighten a little before others. This is due to the variations of the resistors (which are usually ±5%) and the V_f of the LEDs themselves. Ideally they should all come on at the same time.

So what if this effect was cultivated? By using different values for the resistors, we can use the voltage divider effect to turn the LEDs on at different voltages. Once an LED turns on it will stay at the same voltage, and it's matching resistor will not increase its current as the input voltage (V_{in}) continues to rise. The remaining current will be routed through LED, causing the LED to get brighter. Figure 3.2 shows how this would work. You'll note there is an approximate 10% spread between a resistor and the next resistance down.

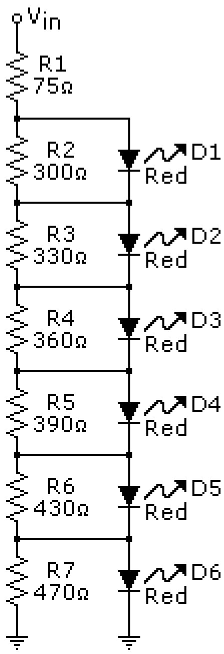

Figure 3.2

Analyzing what voltages will cause which LEDs to light can be tedious, but is predictable. You have to analyze the circuit from scratch every time an LED turns on, and it is critical the Vf used in the calculations match the LEDs. Small errors accumulate in this design. Start by looking at the main current limiting resistor R1. It doesn't interact very much with the bargraph, but it does set the total current. Assuming the V_f of the red LEDs is 2.5V, six LEDs work out to be 15V, and the power supply can go to 16.8V with fresh batteries. I chose an arbitrary current of 25ma, figuring 5ma will go to the resistors. So to figure R1 use:

R1 = (16.8V – (6 X 2.5V)) / 25ma = 72Ω ≈ 75Ω
I_t = (16.8V – (6 X 2.5V)) / 75Ω = 24ma

Next you work through the resistor network. Since R7 is the highest value it will drop the most voltage, turning D6 on first. We are assuming the V_f is 2.5V, so that is the voltage we are interested in for R7, the transition between R7 controlling the voltage and D6. This is a classic voltage divider, so plugging the numbers in looks something like:

R_t = 75Ω + 300Ω + 330Ω + 360Ω + 390Ω + 430Ω + 470Ω = 2,355Ω
I_{R7} = 2.5V / 470Ω = 5.32ma
V_{D6} = 5.32ma X 2355Ω = 12.5V

So at 12.5V the first LED turns on. At this point R7 is not figured as a resistance, but as a constant voltage, and is added to where V_{D6} is calculated. Repeating the procedure to find where D6 turns on:

R_t = 75Ω + 300Ω + 330Ω + 360Ω + 390Ω + 430Ω = 1,885Ω
I_{R6} = 2.5V / 430Ω = 5.81ma

V_{D5} = 5.81ma X 1885Ω = 10.95V + 2.5V = 13.5V

You repeat this procedure for each LED.

R_t = 75Ω + 300Ω + 330Ω + 360Ω + 390Ω = 1,455Ω
I_{R5} = 2.5V / 390Ω = 6.41ma
V_{D4} = 6.41ma X 1455Ω = 9.33V + 5.0V = 14.3V

R_t = 75Ω + 300Ω + 330Ω + 360Ω = 1,055Ω
I_{R4} = 2.5V / 360Ω = 6.94ma
V_{D3} = 6.94ma X 1055 = 7.33V + 7.5V = 14.8V

R_t = 75Ω + 300Ω + 330Ω = 695Ω
I_{R3} = 2.5V / 330Ω = 7.58ma
V_{D2} = 7.58ma X 695Ω = 5.27V + 10V = 15.3V

R_t = 75Ω + 300Ω = 375Ω
I_{R2} = 2.5V / 300Ω = 8.33ma
V_{D1} = 8.33ma X 375Ω = 3.13V + 12.5 = 15.6V

So this bargraph will start at 12.5V and slowly go up and max out around 16V. You'll note it is not very linear (though this can be tweaked), it isn't meant to be. This is not meant for an instrument, but a simple display. There are chips that can do much better at this, such as the LM3914. This chip will do precision displays and a wide range of user options with a minimum of fuss. The schematic shown in figure 3.2 on the other hand will smear, one LED starts to light, but before it is fully illuminated the next one starts to light, so the transition is over 2-4 LEDs. The eye is very good at picking this out however.

While this isn't meant for instrumentation, it has the potential for such. Older LEDs, with their smaller V_f, work much better for this application since each LED is a smaller increment of voltage. Newer isn't always better. You can also improve the predicted values by measuring the real V_f of each LED, and using the real values in the calculations.

You aren't limited to a simple bargraph. Since the resistors choose which LEDs light first you can have several LEDs light up at the same time, or use whatever sequence you choose.

There are other ways to use the fixed V_f of an LED. Someone had a problem where they wanted the LED to go out when the glove box was closed. The catch was they wanted to turn it off when a switch closed (a magnetic reed switch), which is counter intuitive at first glance. Power to this LED would be cut when the key was removed from the ignition, which allowed for this approach.

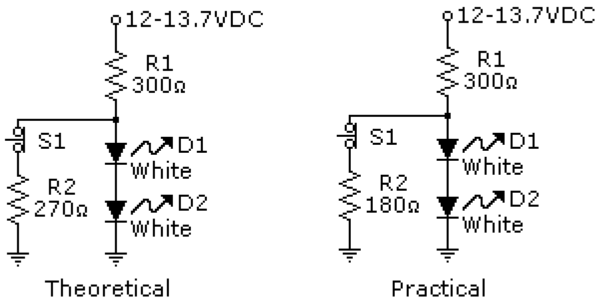

Figure 3.3

R2 was used to control the wattage used by R1 when S1 is closed. If a simple short was used R1 would go over ½ watts, too much heat in a small space. With the addition of R2 this would go down to an eighth watt. Using a similar resistor divider technique I was attempting to get the voltage under the V_f of the combined diodes.

Unfortunately the difference between theory and practice caught up with me, but this was fixed by dropping R2 down to 180Ω.

Chapter 4: The 555 Integrated Circuit

This IC has been around for a long time, over 30 years. The 555 IC could have been designed for LEDs, it is as if they were made for each other. I've written several articles about it, and won't go to the depth I did about the LEDs. Some internals of the 555 IC do need covered, since they relate to LED voltages.

The 555 has a digital output. It is either switched to the positive voltage (high) or the negative (low). An equivalent drawing of it's output would look something like this:

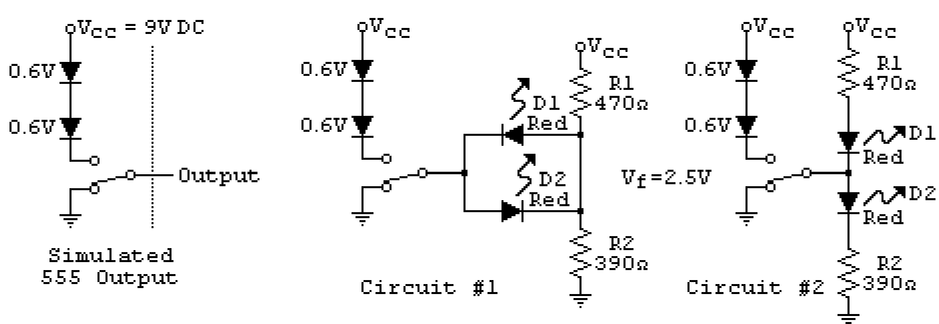

Figure 3.1

Although Circuit #1 and Circuit #2 look different, they are pretty similar in performance. Generally I prefer circuit #2, but #1 will handle some special LEDs that are a red and green LED in the same package. Alternate between the LEDs fast enough, and it appears yellow. In both cases the 555 output shorts one side or the other, leaving the opposite side to light up with full power. The two internal diodes shown (which are actually two base emitter junctions) generate 1.2V, which swamps the LED Vf it is parallel to.

So far I have been showing how to light the LEDs at full power, and how to select the resistor for this. An LED will light up with 1ma and be visible, which will work for a lot of indicator applications. Many cases, such as my experiments, I use a 1KΩ for convenience, and don't worry about it. In the above application this would work out to 6.5ma, which works well enough.

Another issue to be aware of is what the 555 can provide in current. I've already shown it's voltage limitations, but the transistors inside the IC can only provide 200ma before being damaged. There is a general rule in electronics that you should only use half what a component can provide, to make sure the part lasts its expected life. I don't always follow this rule myself, but you need to be aware. The 555 is also rated for 4.5 VDC to 18 VDC, generally this will set the power supply limits of the circuit.

The 555 is a very open ended ICs, and have a lot more uses than just flashers, but for the purposes of this article we'll concentrate on the flasher applications. Shown in Figure 4.2 are two basic configurations that can be used to flash LEDs.

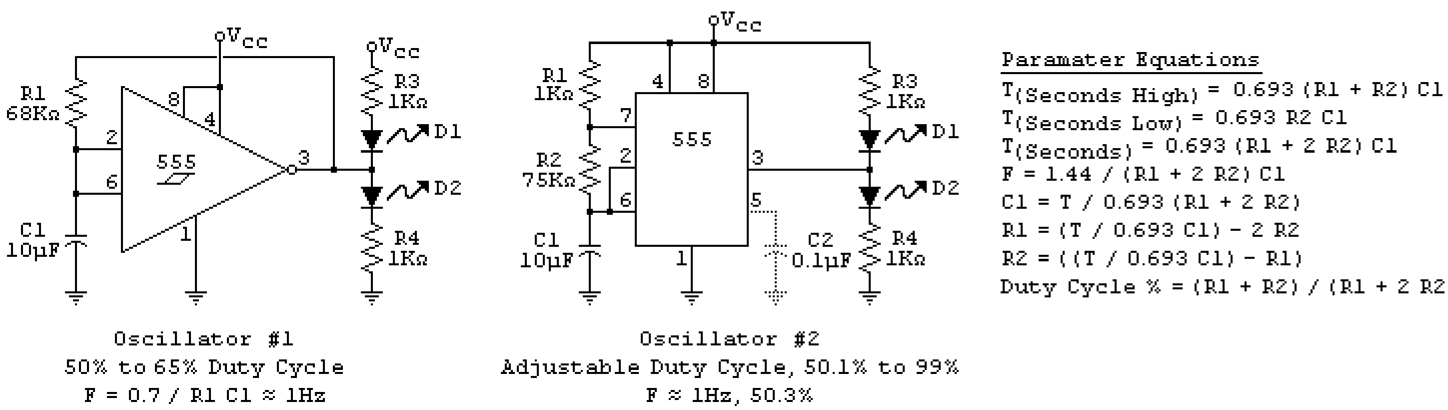

Figure 4.2

Oscillator #1 is in the family of Hysteretic Oscillators, which is usually made with op amps. The 555 version adds some its own twists, since the output isn't quite rail to rail (as shown by the two diodes in the first illustration). Its duty cycle is hard to predict, as it is somewhat dependent on power supply voltage. The higher it's power supply voltage, the closer to 50% it becomes. However, for many applications the duty cycle imperfection is hard to see, so it can be used in a large number of applications with good results. You can even put a potentiometer for R1, which allows the flasher to cover a really large range of rates and frequencies.

Oscillator #2 is straight out of the 555 datasheet. With the addition of a second resistor it overcomes all the problems with oscillator #1, including the 50% duty cycle. For 50% R1 needs to be as low as possible, which is balanced by the fact that at one point R1 is completely across the power supply, thus being one of the components that set the total current draw of the circuit.

C2 is a bypass capacitor. For a single 555 on a battery you don't really need C2 or any other bypass capacitors, which is why I show it as a "ghost" image. There is an exception to this rule, which will be covered in the following article.

So what if you need a single LED that is one only 10% of the time? It is simple, use the D1 side for your LED. If you need 90% then use the D2 side for your LED.

Chapter 5: The 555 and PWM

The 555 has a use that doesn't fall under flasher nor light chaser, but deserves mentioning since it concerns LEDs. That is PWM (Pulse Width Modulation). You could vary the intensity of an LED by varying the current to it, but in many cases this isn't a preferred option, nor is it really linear. PWM allows for truly linear intensity control of a LED.

Shown below in Figure 5.1 is how PWM works. Basically the intensity of the LED brightness is a direct function of how long the power supply is on versus how long it is off, usually expressed in a percentage. This percentage is a direct indicator of LED intensity.

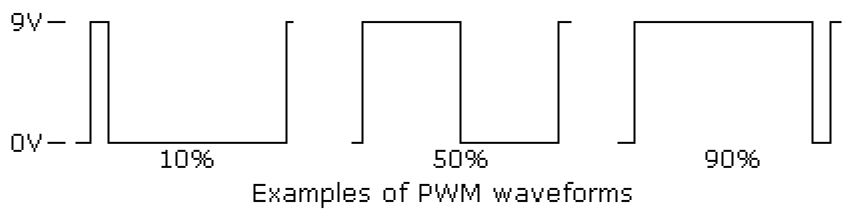

Figure 5.1

Part of the key to how PWM works is speed, the human eye can not perceived changes faster than 30 frames per second (33 Hz), a fact that is used by TV sets the world over. Under this frequency it is possible the on/off of the LED can be seen as a flicker, faster rates than that the average power is seen as a uniform light. The 555 can go much faster, of course, but this sets the minimum.

One of the key features of PWM is that since it is fundamentally digital very little power is used when the light is low or off. There is also a second advantage, LEDs are not a linear device. The intensity of the LED does not vary proportionally to current, but it does vary proportionally using PWM. This makes it a preferred method for adjusting LED brightness.

Figure 5.2 shows several ways to make a quick and easy 555 PWM controller. If you will compare this drawing to Figure 4.2 the resemblances will be obvious. The second drawing is almost the same as the Oscillator 2 in Figure 4.2, since this design has PWM inherent in its design.

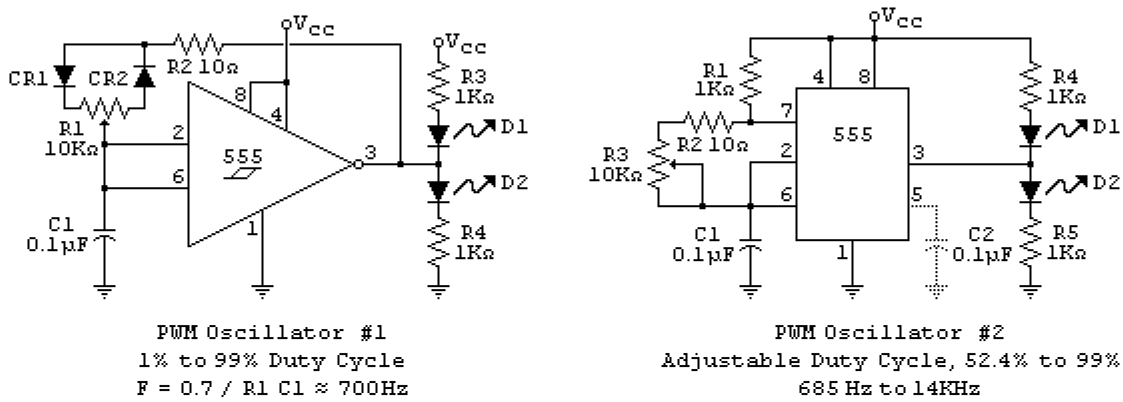

Figure 5.2

There are many ways to generate PWM signals, and they have many uses, such as controlling motor speeds or Class D audio amps. LEDs are only one example of how PWM can be used.

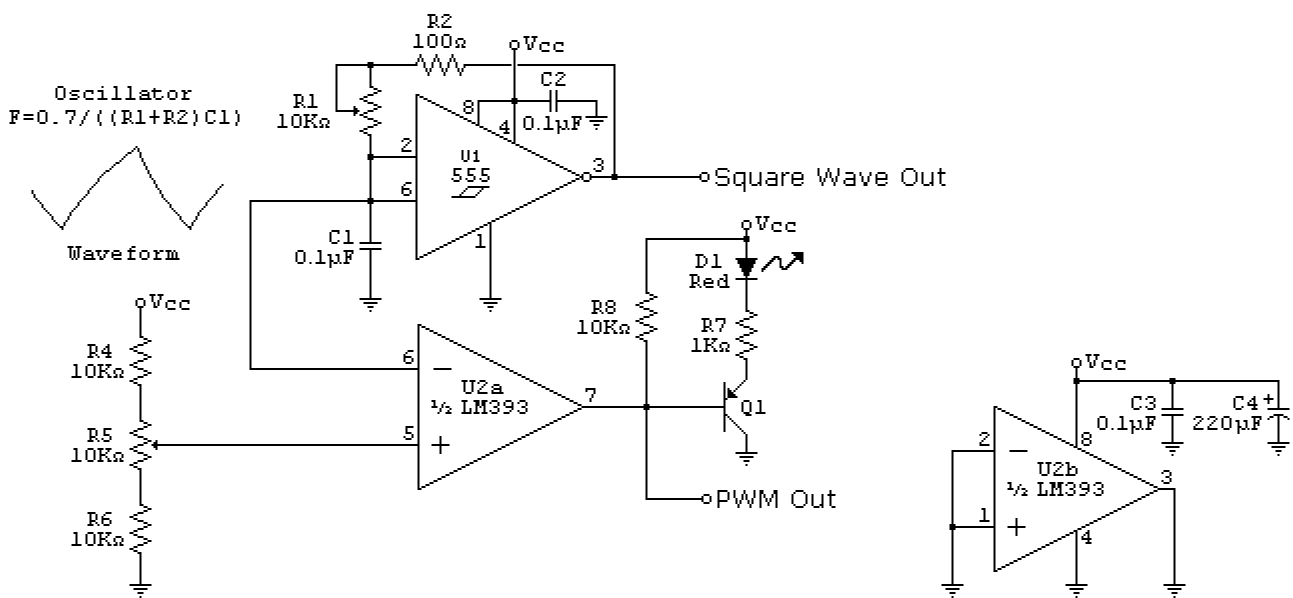

Figure 5.3

This particular circuit, with minor modification, could be used for a Class D audio amp as well as modulate an LED brightness. It has the added ability to adjust the PWM frequency independently of the PWM percentage, which can be very useful. The LM393 dual comparator and the LM339 quad comparator absolutely require a pull up resistor as shown with R7, usually a 10KΩ resistor. Unused comparators need grounded as shown to prevent unwanted oscillations and current surges. It does not matter if the ouput is grounded or not, but grounding it can make for a simpler printed circuit board design. Since the max current from both chips is 16ma, I've added a transistor driver to reduce its load, and R7 can be tweaked for maximum LED brightness.

Chapter 6: Low Power Applications

While the 555 isn't a power hog, it is a product of the 70's. It has 15KΩ resistance, not counting the rest of the circuitry. It will drain a battery very quickly, in days if not hours. Several manufacturers have come out with low power CMOS versions, such as the TLC555 and the 7555. These parts are pretty similar to each other, though not exact. They can both drive an LED going to ground (low), but have about 10% the current capability going to Vcc (high). As the power supply voltage drops the current they can provide radically reduces, so with really low voltages you will have to use a transistor to light an LED to full brightness. On the other hand the CMOS versions draw about one hundredth the current for its internal circuitry, so they definitely have their uses.

Figure 6.1 shows some low power long duration flashers.

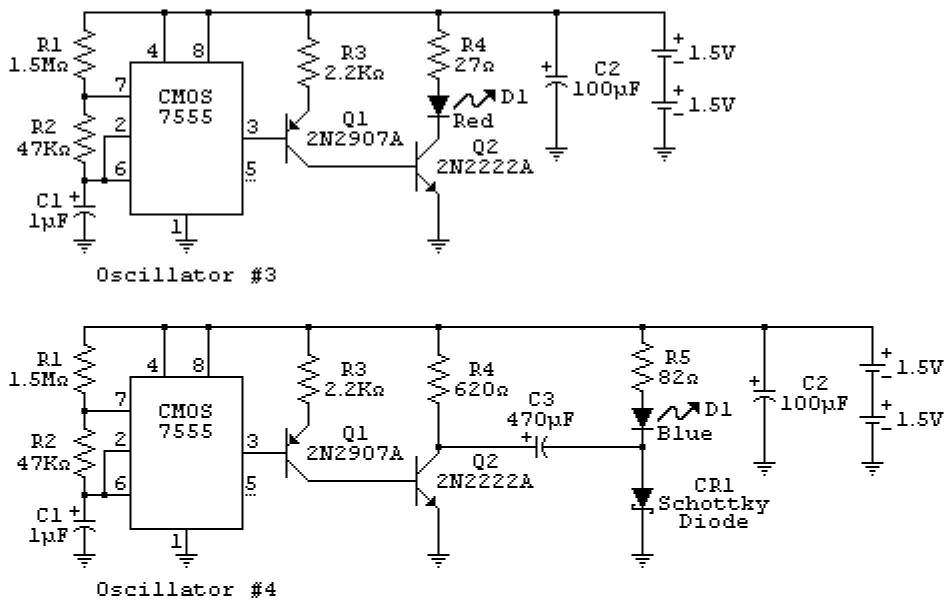

Figure 6.1

Oscillator #4 uses a capacitor voltage multiplication to boost the 3V from the battery to almost double that, enough to drive the 3.5V Vf of the blue LED. The Schottky diode drops a fraction of what a conventional diode does, or a Germanium diode could be used for much the same reason.

Capacitor C2 was added after experimentation showed that it was necessary for maximum life. Without it the circuit basically dimmed and died after two weeks, using AAA alkaline batteries. Adding the capacitor extends the flash life, my test circuit has worked more than 3 months using AAA batteries. This is because the circuit is only on 3% of the time, the remaining 97% the capacitor takes on a charge. I suspect this is a unique case, but it is interesting.

Chapter 7: The Joule Thief

The classic Joule Thief uses transistors. The basic principle, using an inductor to kick the voltage from the battery up until it will power an LED has also been applied to the 555 also. Figure 7.1 is a redrawn schematic, the original source was uploaded on another thread.

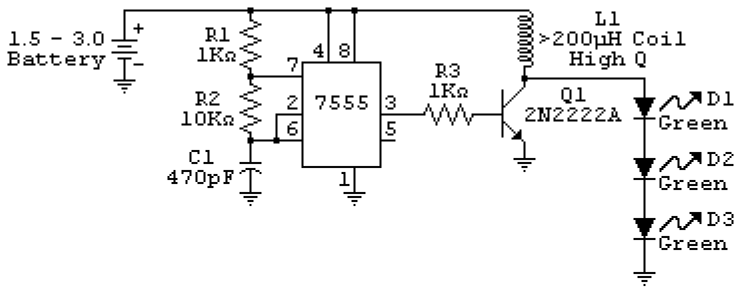

Figure 7.1

The 555 has been so useful over time that a dual version, two complete 555s, have come out. They also have their CMOS versions. I applied this to the following schematic.

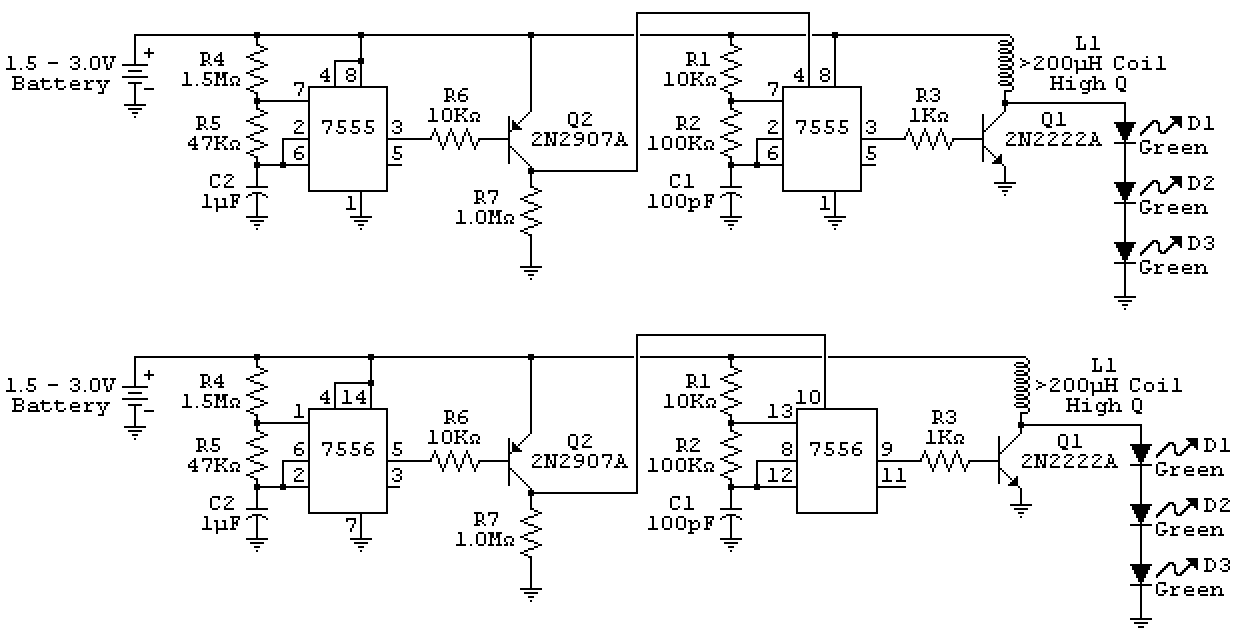

Figure 7.2

These schematics use a feature that hasn't been shown to date. Pin 4 is an Enable pin for the 555, it is possible to use a 555 oscillator to control the second one, the voltage booster. This design works, and should make a battery or two last a very long time, but it could be improved quite a bit. Using two batteries to make 3V improves the brightness of the LEDs substantially. You may notice there is no current limiting resistor. This is because at 3V there simply isn't enough voltage to turn the LEDs on, all the current driving these LEDs is coming from the inductive kick of the coil.

There is a way to flash 20 different LEDs from 4 555 ICs. Each LED would have it's own flash pattern, no two alike (though some are inverted from others), half of the LEDs will be on at any time for a total of 100ma. Basically we're merging Circuit #1 and Circuit #2 together, and using the way the 555s switch on the outputs for this effect. This could be used in a Christmas Tree, or just a light panel for a kinetic sculpture, or some other special effect. The base idea could be expanded even further for more LEDs, however the current draw on the 555s quickly approaches their limit. For 10ma per LED, 5 would be the max (150ma, 30 LEDs). At 6 would be 42 LEDs (210ma). The colors shown in Figure 8.1 were selected at random, and are by way of example.

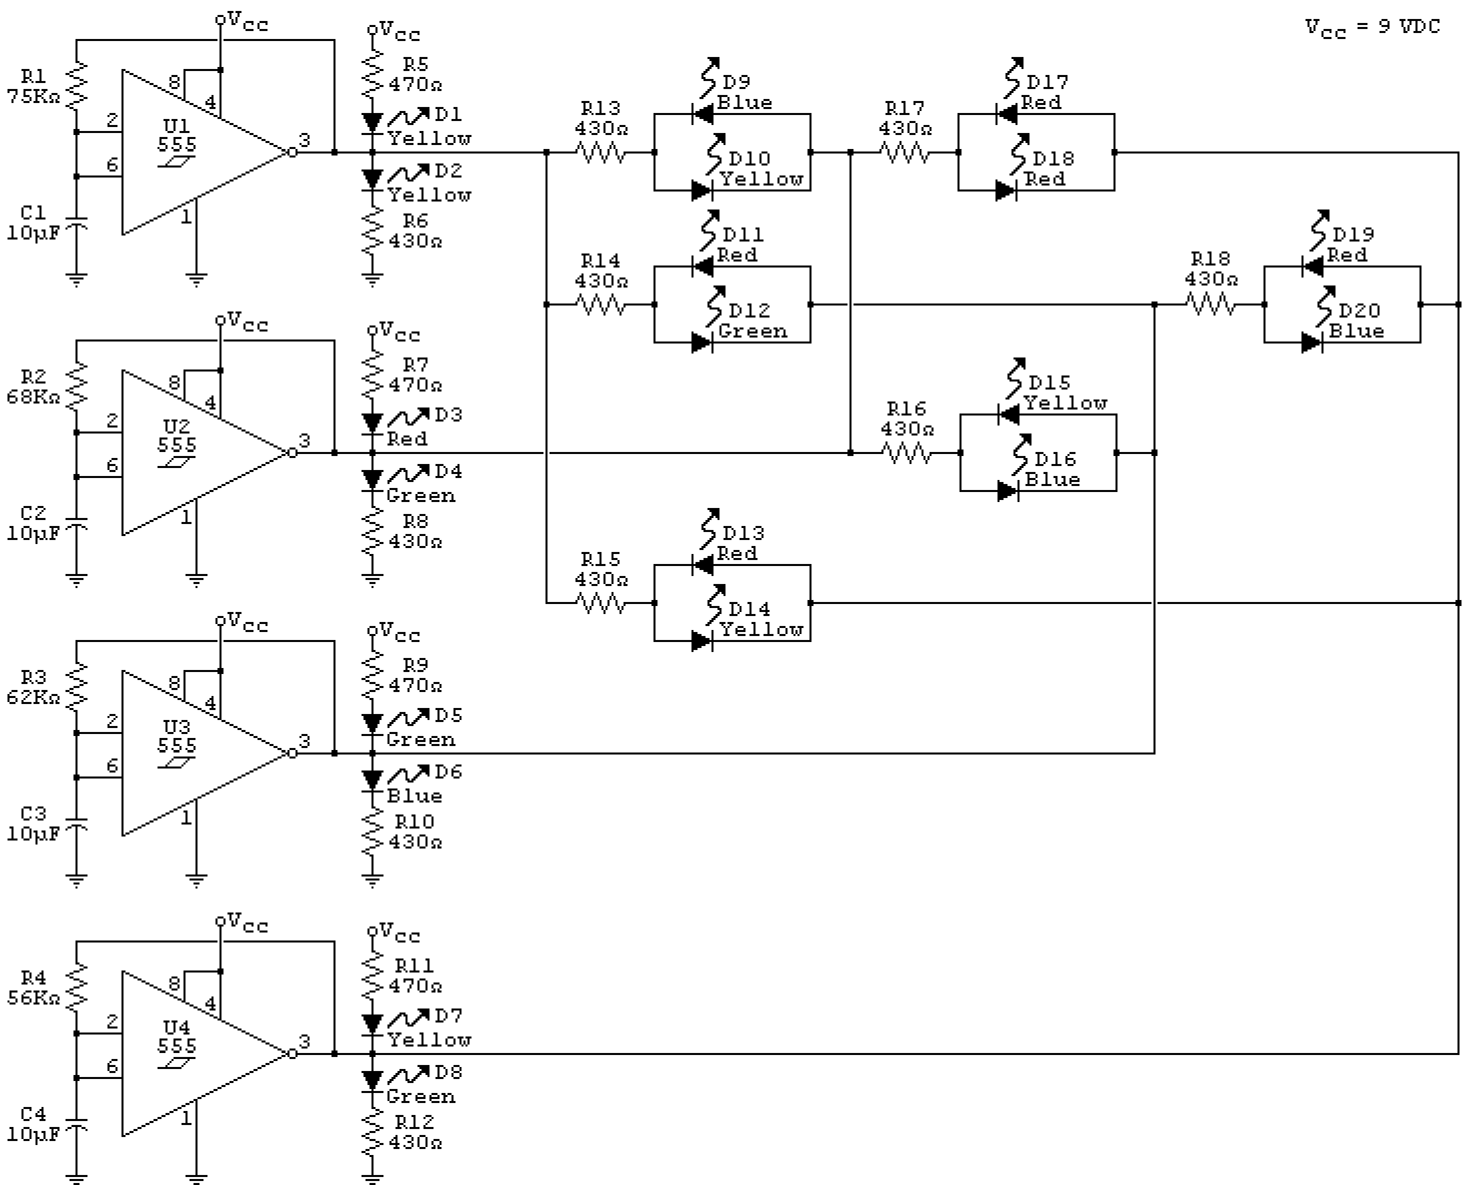

Figure 8.1

Chapter 9: Light Chasers

Light Chasers take a flasher to the next step. Many cases they are done with microcontrollers, small computers, but that isn't really necessary unless some kind of computation for the display is really needed. Two nifty ICs, the [CD4017 and CD4022](), are perfect for this kind of application. They will sequence almost any number of outputs. The data sheet shows how to cascade even more 4017s for more than 10 outputs, and one 4017 can do 2-10 outputs. For CMOS this chip has incredible drive, rated up to 6.8ma best case! I have designed it using 10ma for direct drive of LEDs, though this is definitely not recommended by the manufacturer, and may not work in everyones build.

Figure 9.1 is an old design of mine. This circuit has worked for over 25 years, though not continuously (figure several months on that level). Again, the CD4022 is very stressed, so this isn't a recommended design (but I would use it again in non critical uses).

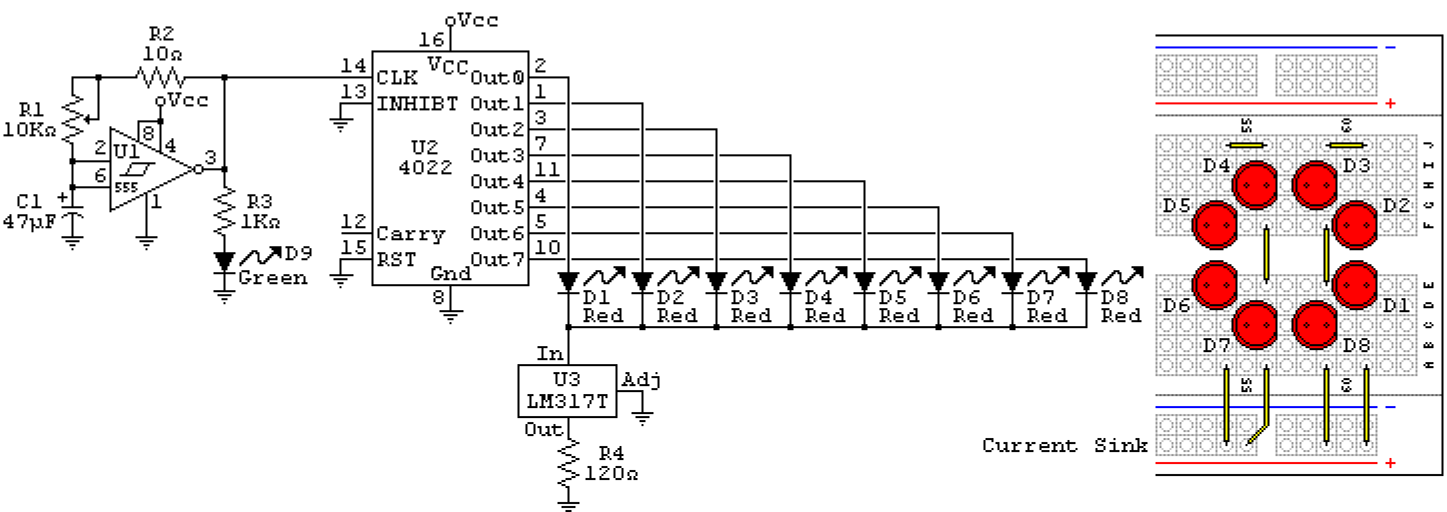

Figure 9.1

The thing to note about this design is it makes absolutely no difference how many LEDs are in each chain, as long as you are under the Vcc/Vf limit (and don't forget the LM317 3V drop). Why is this important? Take the following circuit in Figure 9.2 as an example.

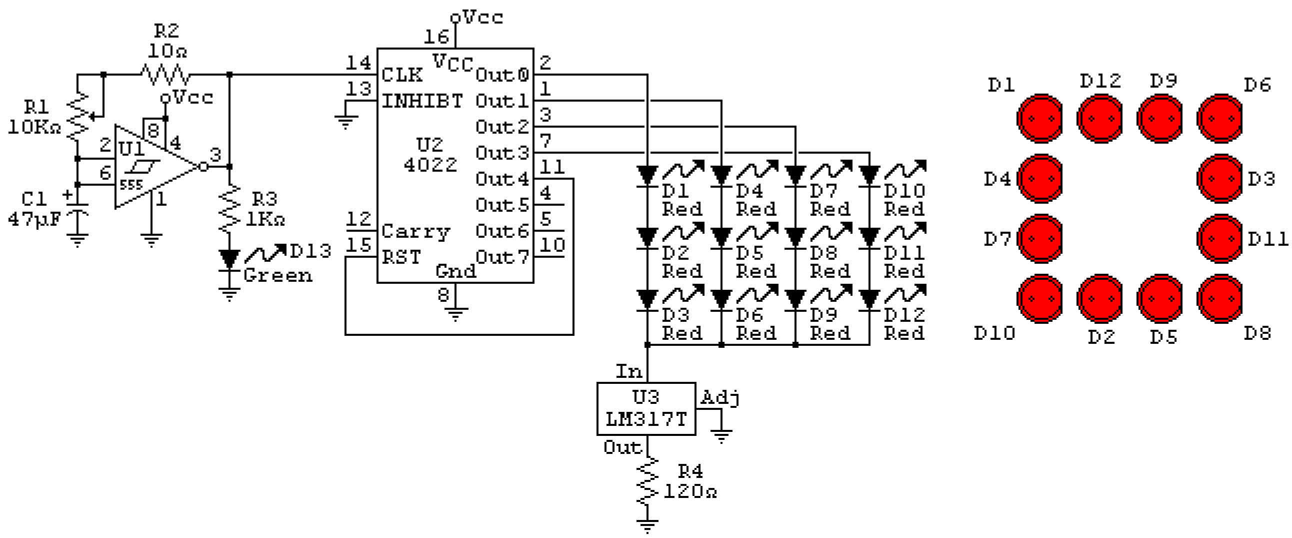

Figure 9.2

With this circuit there are 3 lights apparently chasing around the square. We have all seen variations of this effect on signs and in supermarkets. The thing to remember is this was done by how the LEDs were arranged and wired. It could have as easily been runway lights. I have done this in friends cigarette ashtray with good effect. The arrangement of the lights is more important that the circuit driving them in many cases.

Note how the CD4022 was limited to 4 counts. This is a common theme in using these chips. The 4017 is probably more popular, but it can be limited in a similar way. This is important when you want to generate patterns, which will be discussed later.

Chapter 10: Transistor Drivers

The CD4017/4022 low current output means we have to have some means of increasing this drive. It is easy to become spoiled by the 555, with its relatively huge output currents. It can be fun to cheat a little with something like the 4017, forcing it to go beyond it's ratings, but at some point everything will go permanently dark. These chips can work for decades if kept within their ratings. Fortunately it is easy to use transistors as simple switches, to fully drive modern LEDs. A lot of the schematics have already shown this to one degree or another. Most moderate LEDs seem to focus around 20ma. In some cases much more current is needed, either because the LED requires it or there is a large quantity of LEDs.

BJTs (Bipolar Junction Transistors)

The humble 2N2222A NPN transistor has been around for many decades, as has its compliment, the 2N2907A PNP transistor. They perform admirably as a switching transistor, with a rated max of 0.6A. If we derate it to 0.3A this will still drive a lot of LEDs. If a job comes up that is too big for this part there are many other much higher rated transistors to choose from.

There are two ways of using a transistor. The common collector mode shown previously and in Figure 10.1 is a variation of the voltage regulator. It works because CMOS tends to get quite close to the power supplies rails (the plus or minus voltages). The loading on the CMOS chip is the LED current divided by the gain of the transistor. So if a LED array is pulling 100ma, and the gain of the transistor is 50 (which is pretty low, a minimum spec) the current from the CMOS device is 2ma. This design will generate some heat, since the emitter is 0.6V below Vcc (at a minimum). 0.6V X 100ma is 0.06 watts. In extreme cases the transistor can get a lot hotter.

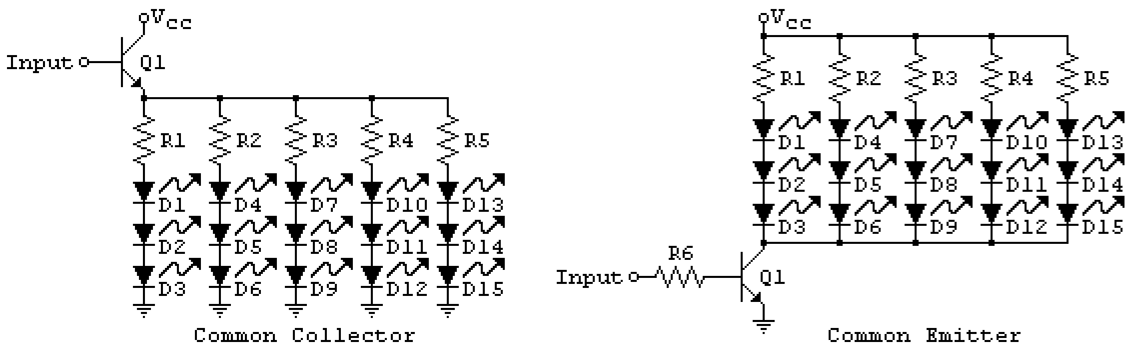

Figure 10.1

The common emitter mode has a different bag of advantages and disadvantages. The transistor acts like a switch because the collector is very close to the emitter voltage, so it generates very little heat. The two most efficient states for any transistor in terms of wattage is when they are fully on (dropping almost no voltage) or fully off (drawing almost no current). Since wattage is voltage times current (V X I), and you have moved one of the variables close to zero the wattage is a very low number. The disadvantage of this configuration is input current, which has to be controlled by R6. A general rule of thumb is the base current should 1/10 the collector current. This isn't always practical, and the collector current should be the base current times the gain of the transistor (Ic = ßIb), but since gain is such a wildly variable number even within a family, the rule of thumb exists.

The way around this is to increase the gain of the transistors. Fortunately this is pretty easy to do with only minor drawbacks. Darlington transistors (aka Darlington pair) and a Sziklai pair. The gain is the two transistors gains times each other, and the only major drawback is the collector emitter will have a minimum of 0.6 volts (as opposed to less than 0.1V for a single transistor in common emitter mode). Shown in Figure 10.2 are examples of the two types in use. In both cases the value of R6 can be increased dramatically.

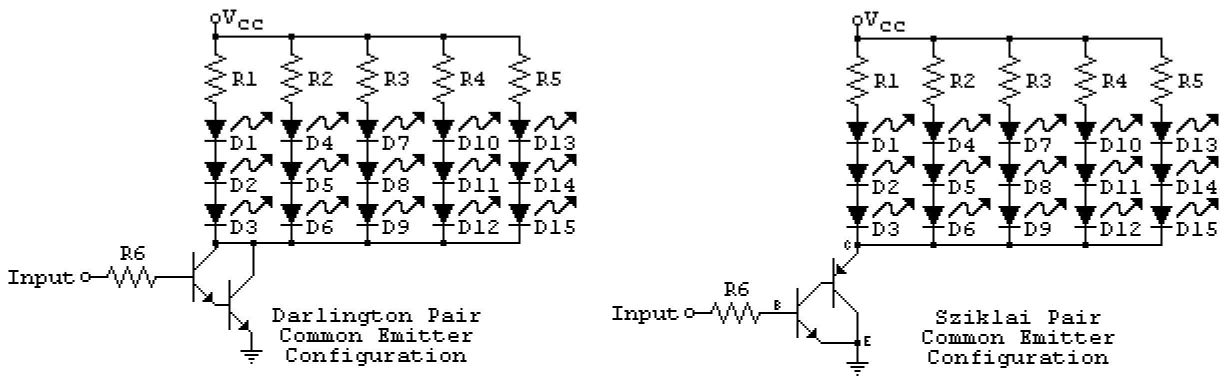

Figure 10.2

<u>MOSFETs (Metal Oxide Semiconductor Field Effect Transistors)</u>

Another common transistor used is a MOSFET, this chapter will deal with N channel enhancement. The N channel refers to polarity while the enhancement designation deals with the transistors internal construction. P channel MOSFETs (or pMOSFETs) are similar in the same way that PNP is to NPN on BJTs.

In many ways this type of transistor is superior to BJTs. Its advantages include almost no switching current (there is an extremely short surge when they switch) and they conduct extremely well when they are on, which means that they rarely get hot. They can conduct extremely high currents, 10 amps or more is not uncommon. Their disadvantages include increased sensitivity to static electricity and at least 10V to switch cleanly, 12V is typical. There is a family of MOSFETs called logic level MOSFETs that can use less voltage to switch (3V to 5V) but they tend to be harder to find.

MOSFETs are voltage controlled devices. They have a large capacitance on the Gate, which is why they have a surge current. You want to switch these devices as fast as you can to prevent them from getting hot, something that only happens when they are partly on. Something like a 555 does this nicely. Unless you have logic level MOSFETs you really need 12V on the 555 due to the 10V limit on the gate.

A gate resistor is needed to prevent something called ringing, which might cause the MOSFET to not switch cleanly. This resistor needs to be as close to the Gate as is practical. Shown below is a typical circuit you might use to power a bunch of LEDs, and a typical package for a MOSFET sold by Radio Shack, an IR510. Not all MOSFETs use the same pin outs, you need to look at the data sheet for each part number.

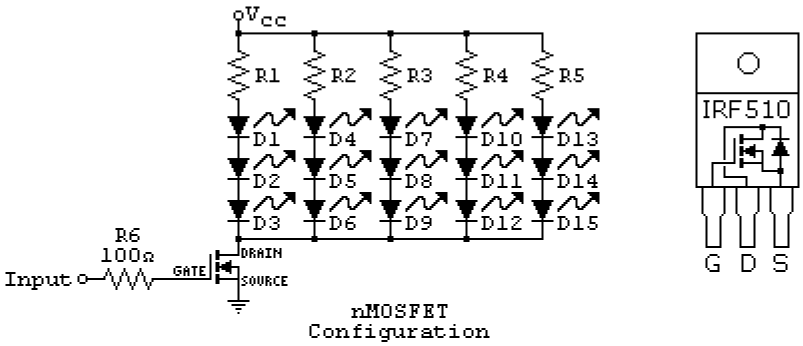

Figure 10.3

Driving A MOSFET

As mentioned, MOSFETs take a lot of short term current to switch cleanly. Their gate looks like a capacitor, with as much as 0.01µF capacitance in extreme cases. If they are being turned on and off infrequently this isn't a problem, once they settle down the current drawn through the gate practically indetectable. However, many of the applications (such as PWM) switches the MOSFET constantly. The large capacitance on the gate can slow the switching speed of the MOSFET, especially if there is a large resistance feeding the gate. The slow switching rates can make the MOSFET run very hot, but if the gate is switched properly the same part will run very close to room temperature.

The way to address the problem is to provide a buffer to help the gate of the MOSFET switch as quickly as possible. Using the PWM circuit shown in Figure 4.3 with a MOSFET and a NPN and PNP transistor (shown in figure 10.4) can be used to correct the issue. Chips are made to do the same job, but this is a quick and dirty method.

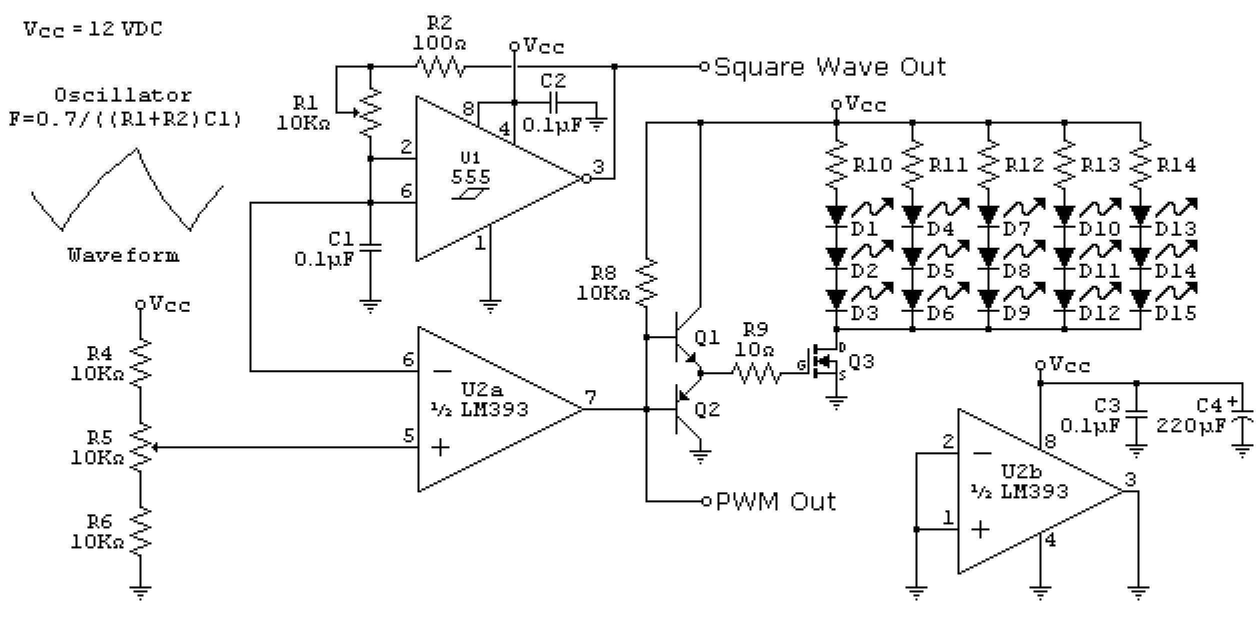

Figure 10.4

The two BJT transistors will convert the low current drive into a high current drive, which is exactly what the MOSFET needs. The current surges through the BJT transistors aren't enough to cause heating. Normal CMOS gates and a CMOS 555 will also drive a MOSFET nicely by themselves, because while they don't have much drive they can handle the surges the MOSFET needs very well. The transistors used inside a CMOS gate is very similar to a MOSFET.

Chapter 11: Making Patterns

Light chasers are cool, but sometimes you want to do more. By steering the current from the 4017 IC sequence other special effects can be created, such as this one.

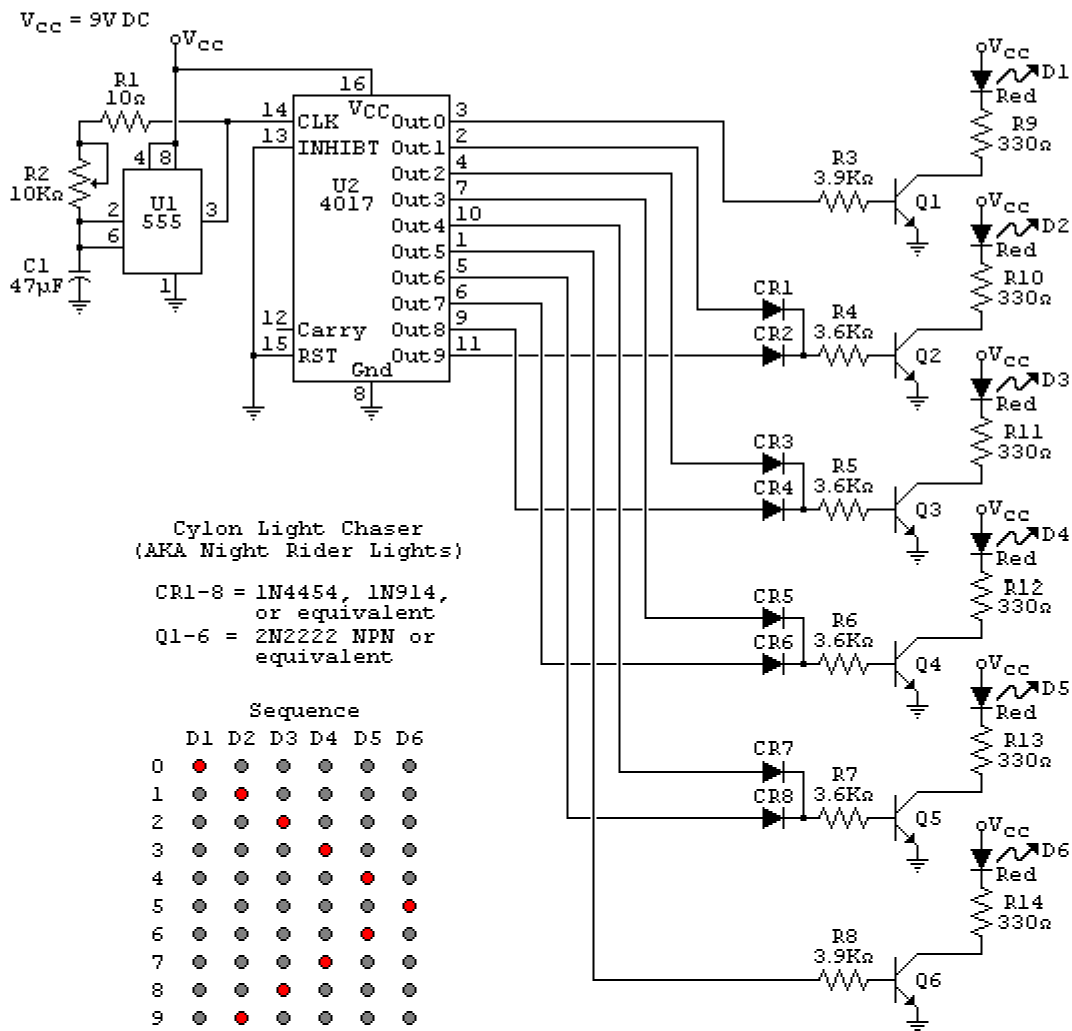

Figure 11.1

This circuit will make the LED light sweep back and forth, a popular Hollywood effect. We have also added transistor drivers that will give the LEDs 20ma without significantly loading U2, which means this particular circuit should last. You may need to add some power supply capacitors, but in general battery circuits are pretty stable without them, as the batteries share some of the same characteristics as capacitors. The voltages from 9V batteries tend to drop fast, down to 7.5 volts, and then stabilize, so be aware. The 555 oscillator will go as low as one cycle every 3 seconds, with the other end being faster than the eye can follow, so it is very open ended for the user.

Another popular design shown in Figure 11.2 is the flasher used in emergency vehicles. This can get you a ticket if you try to use it on a street vehicle, but the basic design is pretty simple.

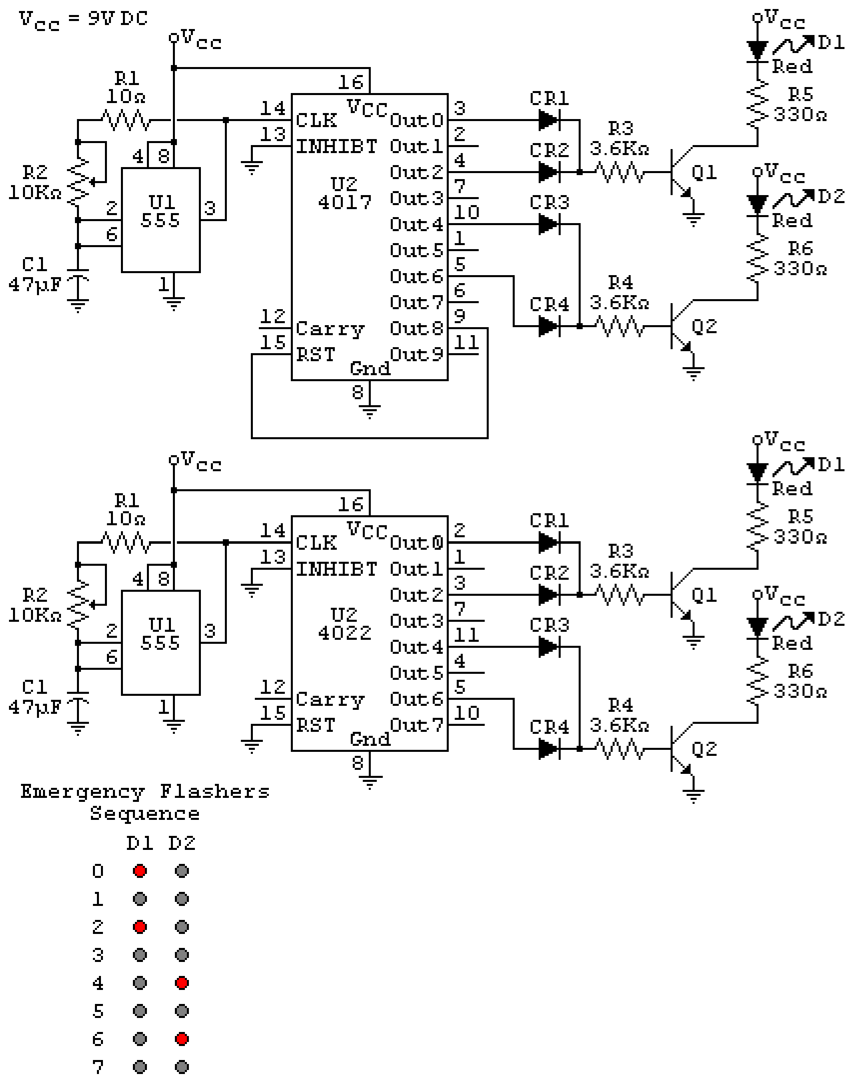

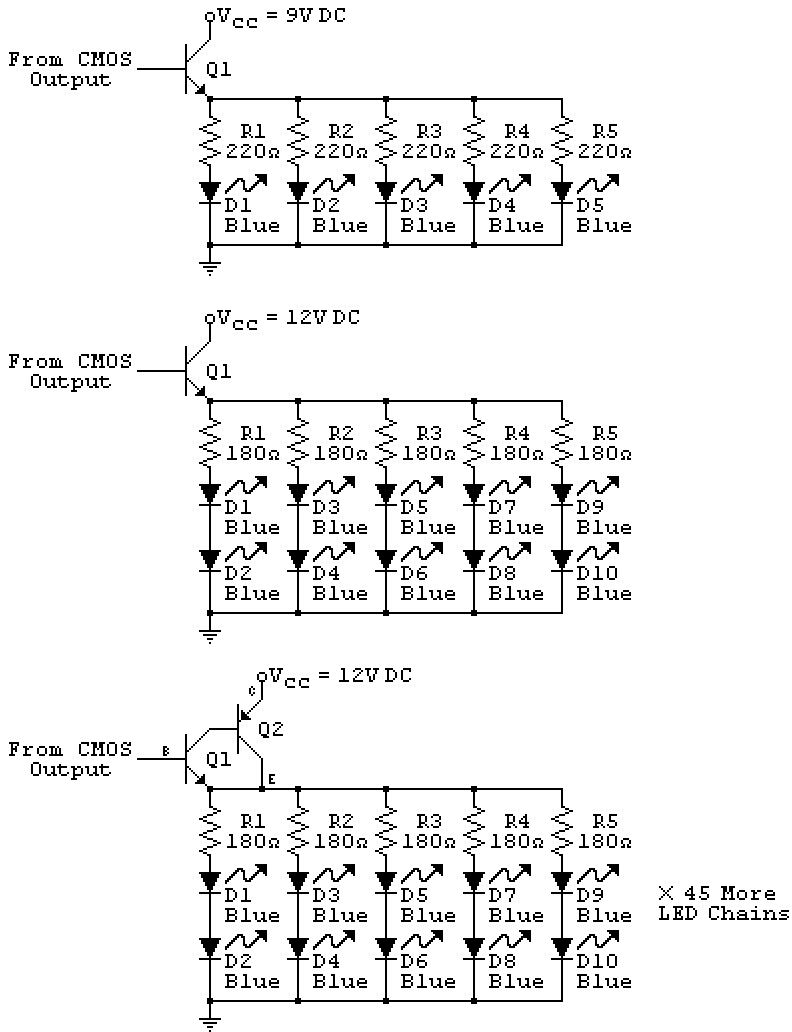

Figure 11.2 *Figure 11.3*

Of course, a design like this practically screams bright lights, so I've shown several options in Figure 11.3. Toys usually use 9V batteries, which can drop as low as 7.5, so this limits what can be done. Some blue LEDs can have a Vf of 3.8V (and 3.8V X 2 = 7.6V). I'd use single transistor in Common Collector configuration shown in Figure 11.3 (also shown in Figure 9.1) to drive individual chains. If the circuit is drawing 100ma for the LEDs, and the transistor has a gain of 50, the current pulled from the CD40XX chips is around 2ma. At 9V and a Vf of 3.8V the LED current is are getting 21ma, if the Vf is 3.5V the LED current is 22ma. At 7.5 and Vf of 3.8 the LED current is 14ma. These calculations show this circuit tries to minimize the current variance for the LEDs. This was covered in the *Current Limiting* chapter.

If you have a stable 12V then the options are more open. Since you can put more LEDs per chain the total current per LED is reduced a bit. The calculated current for this layout is 21ma. If the Vf is 3.5 then the current would be 24ma. Again, the variation is minimized.

But what if you want a lot of LEDs, say 100 of them (50 chains)? This would be a current of 1 amp. A transistor with a gain of 50 would use 20ma through the base, more than the CMOS IC could provide. This would be a good time to use a Sziklai pair as shown. Q2 would definitely have to be a power transistor, but other than that it is pretty straight forward. This would bring the CMOS requirement to 0.4ma, which solves the problem. You could also use the MOSFET shown in Figure 10.3, which is probably the best solution.

I mentioned earlier that the CD40XX ICs could go above their individual counts. The datasheet shows how to do this, as well as Bill Bowden's Website. Figure 11.4 shows how this is done.

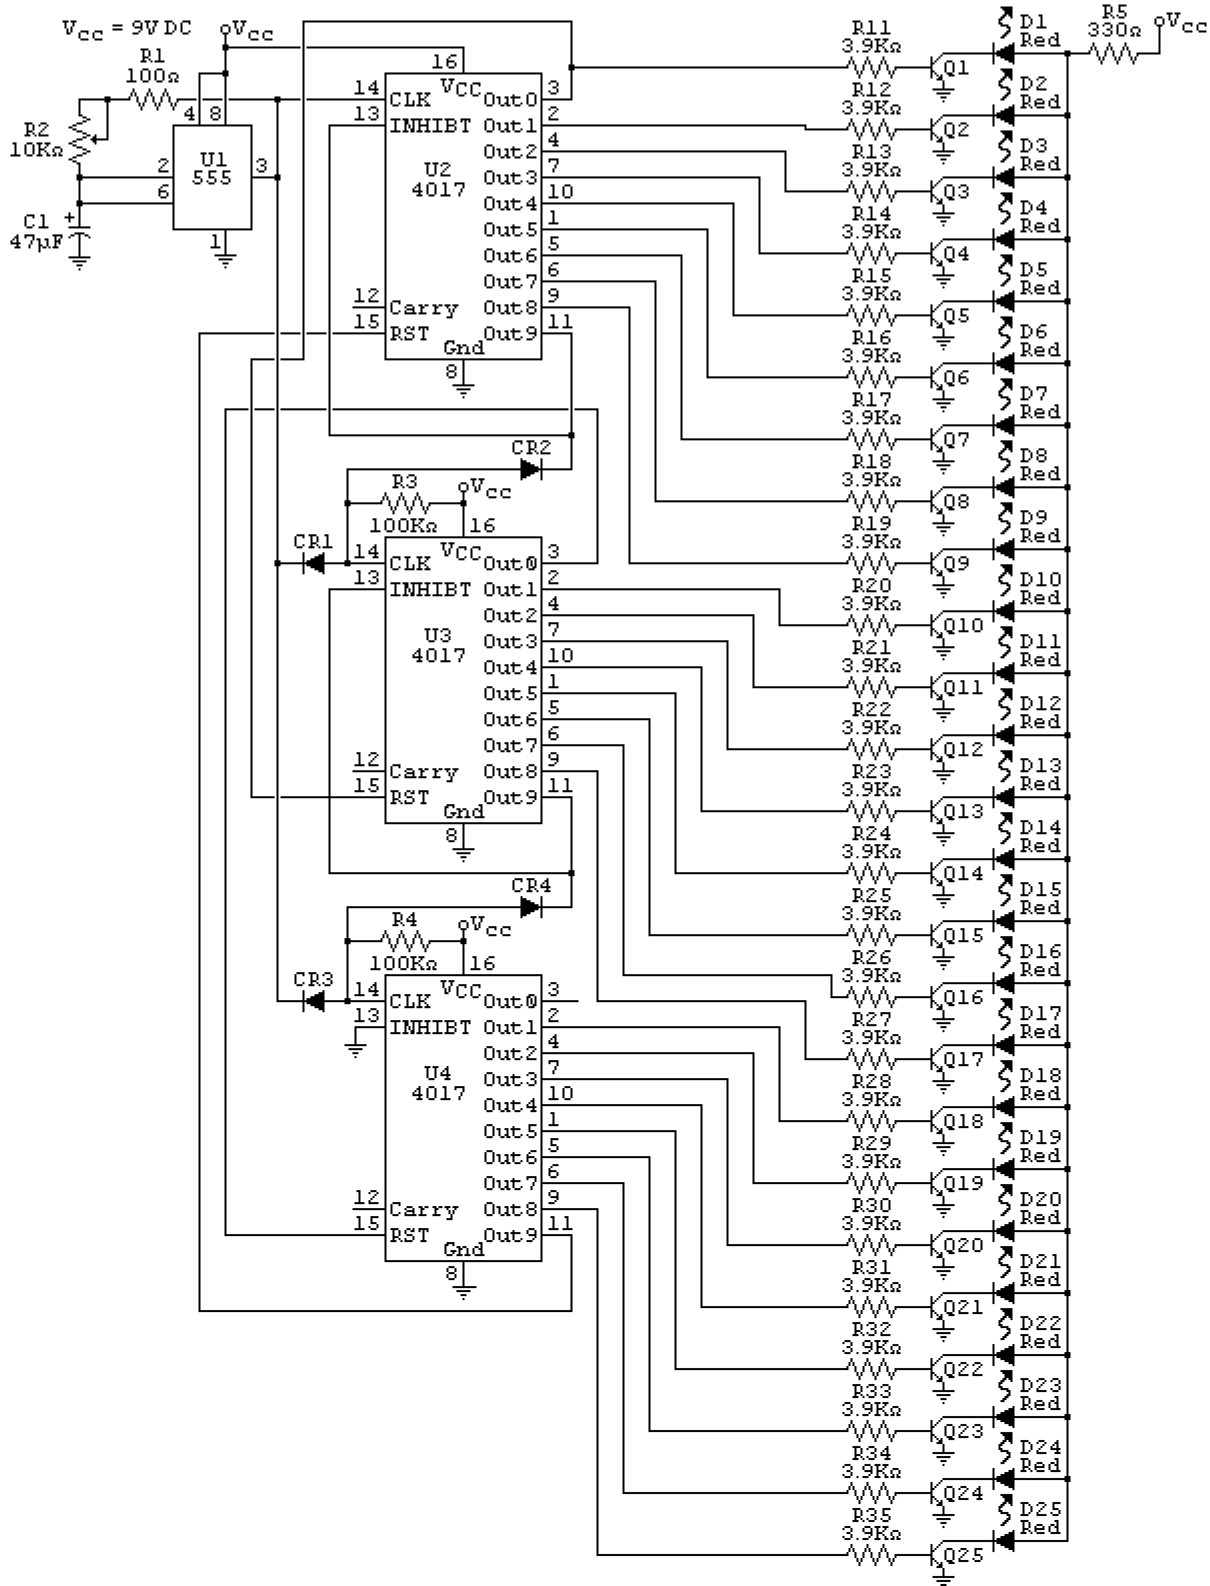

Figure 11.4

The number of transistors and resistors used makes the method shown in Figure 9.1 to drive LEDs more appealing, doesn't it? U3 can be repeated for even more counts, if need be. R5 can be shared as long as you are only turning one LED at a time and using the same color (same Vf). I will be showing some special effects in a later chapter where this won't work.

12. Special LED Effects

Throbbing LEDs

A common theme with LEDs is to slowly turn them on and slowly turn them off. Some people have called this throbbing, and there are quite a few ways to do it. It can be done simply with a 555 and a couple of transistors. Figure 12.1 shows how this is done.

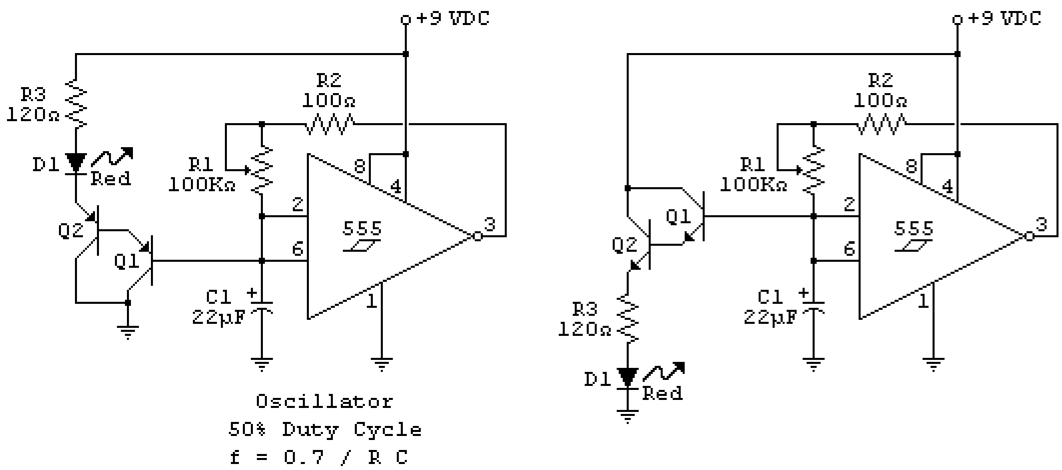

Figure 2.1

While you can buy Darlington transistors prepackaged you don't want to use them in this case, because they also have a built in resistor that will interfere with the circuit. Instead use two separate transistors, such as 2N2907's. It also requires 9V as a matter of course, due to the way the voltages work out. You can also flip the LED/Transistor upside down and use NPN (2N2222) transistors as shown.

The voltage needs of the circuit above is a key point against it. You can do something similar with PWM. With PWM the voltage can be less than a volt above the Vf of the LED, a major advantage. It also can handle a lot more power without anything getting hot, another major advantage. Figure 12.2A shows the basic setup for doing this. Note the similarity of the schematic to Figure 4.3. Due to the low power supply voltage a CMOS 555 was used. I also used a LM393, which is a dual comparator similar to the quad LM339 shown earlier. Both chips require a pull up resistor on the output, and in the case of the MOSFET, a more robust driver.

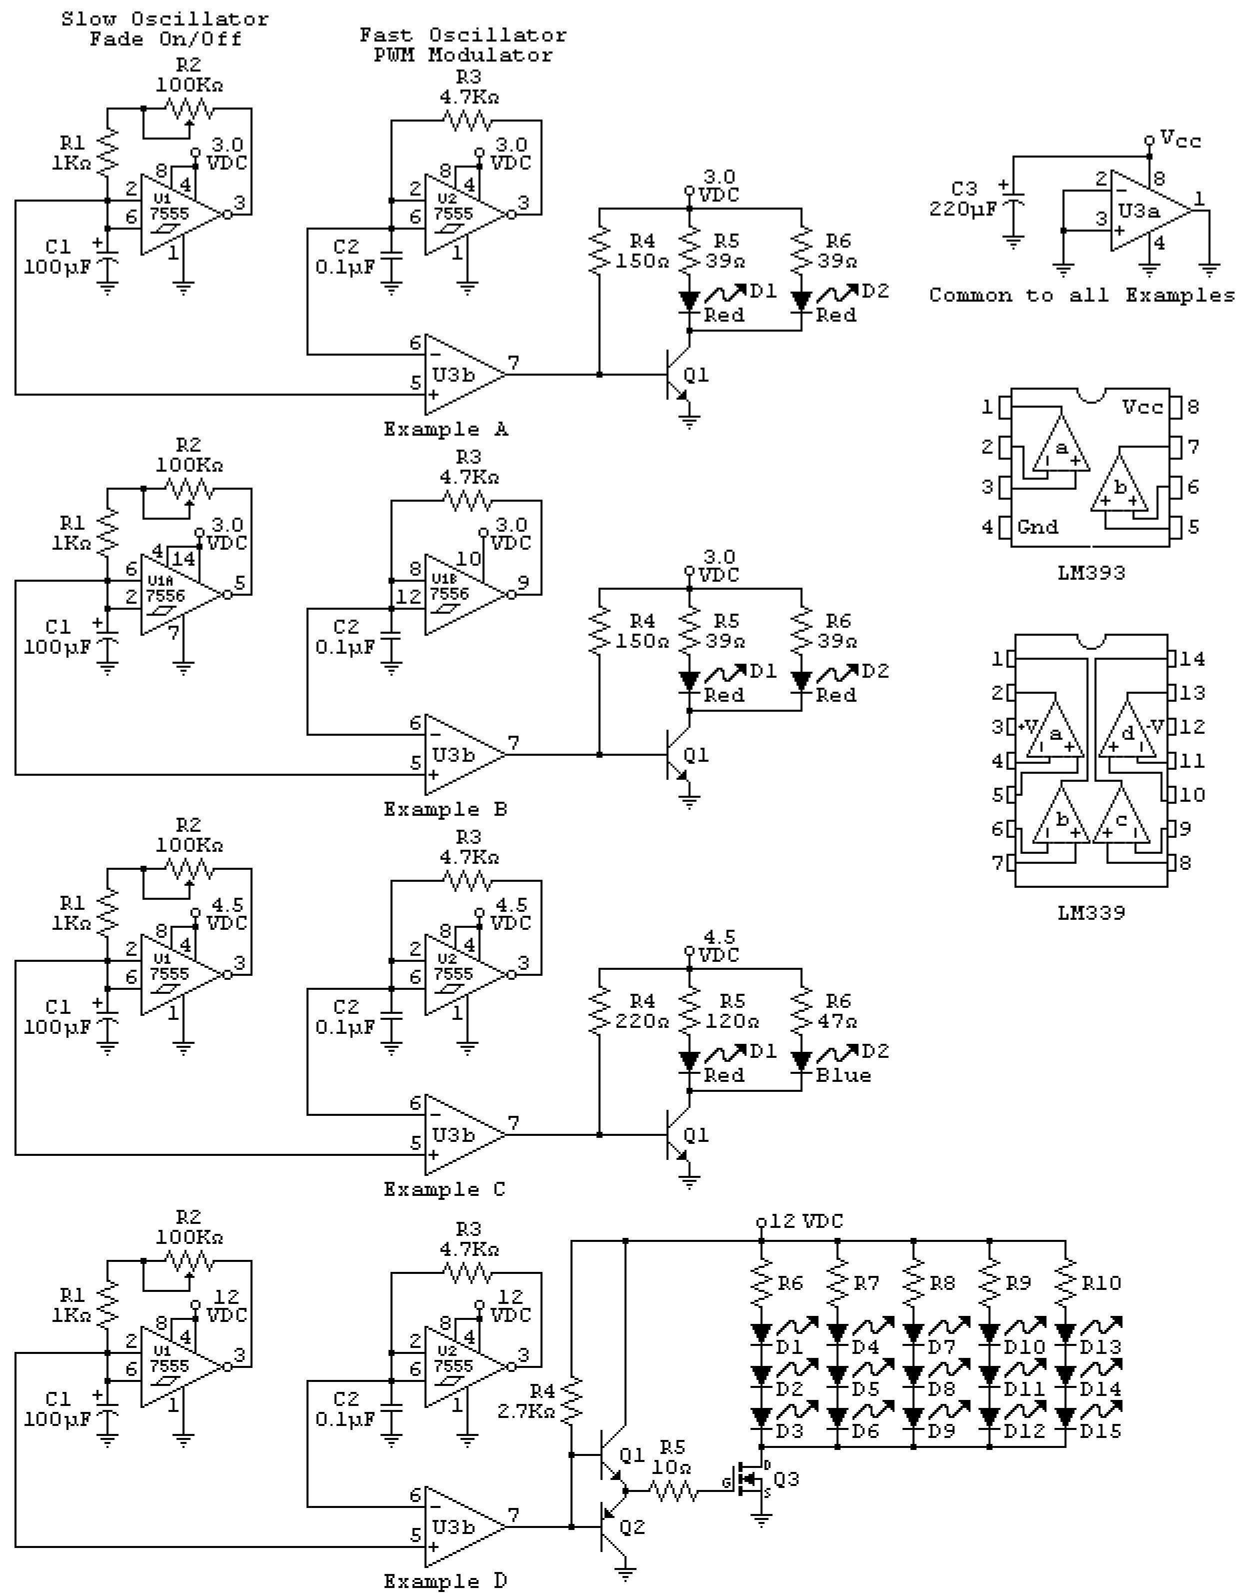
Slow Oscillator
Fade On/Off
Fast Oscillator
PWM Modulator
R1
1KΩ
R2
100KΩ
R3
4.7KΩ
C1
100µF
C2
0.1µF
3.0
VDC
U1
7555
U2
7555
U3b
R4
150Ω
R5
39Ω
R6
39Ω
D1
D2
Red
Q1
Example A
U1A
7556
U1B
7556
Example B
4.5
VDC
220Ω
120Ω
47Ω
Blue
Example C
12
VDC
12 VDC
2.7KΩ
10Ω
Q2
Q3
R7
R8
R9
R10
D3
D4
D5
D6
D7
D8
D9
D10
D11
D12
D13
D14
D15
Example D
C3
220µF
Vcc
U3a
Common to all Examples
Gnd
LM393
LM339

Figure 12.2

Since this circuit requires two 555's as a matter of course I've also included how you would wire a 556 (a dual 555) to do the same job in Figure 12.2B. I've assumed that D1 and D2 Vf is 2.2V, different dropping voltages would require adjustments in R5 and R6. What if D2 were a blue LED that dropped 3.6V? The power supply would have to be increased, Figure 12.2C shows how the end result might look.

Another side feature of this design is the LEDs will be in complete sync, they will become dark and hit their maximum brightness together, even though the individual LED legs are quite different.

I mentioned that this design could achieve much higher power levels. Figure 12.2D shows how this is done with a 12VDC power supply. If used in a car be careful how you implement it, as it could get you a traffic ticket for having emergency lights without a license.

Fading LEDs (AKA Comet Trails)

The sweeping lights shown in Figure 11.1 are interesting and dramatic, but they can be improved if the lights fade slowly, creating a trail behind the sweeping pattern. This mode has also been called a comet trail, and can be accomplished by using the circuit shown in Figure 12.3A, which is attached to the outputs of the 4017 shown in Figure 11.1. There the emitter following transistor circuit work with the capacitors to delay the LED going out just a little, creating a trail of fading LEDs behind the LED actively being lite.

A friend of mine who goes by the handle of Fenris did this in one of his projects, and posted the details at Project: Knight Rider style sweeping light. It has schematics as well as some videos of the final product.

The first example shown works, but charging the capacitor can overload the 4071's feeding it, so Figure 12.3B shows how this loading could be reduced, putting the brunt of the current on Q2 to charge the cap. If you need more than one input (which sweeping lights will need) you can use Figure 12.3C. Note that the extra transistors are also replacing the diodes in the circuit.

Figure 12.3D shows how you would use this concept in a total circuit. It is pretty similar to the schematic shown in Figure 11.4, but I simplified it somewhat to reduce the size of the schematic.

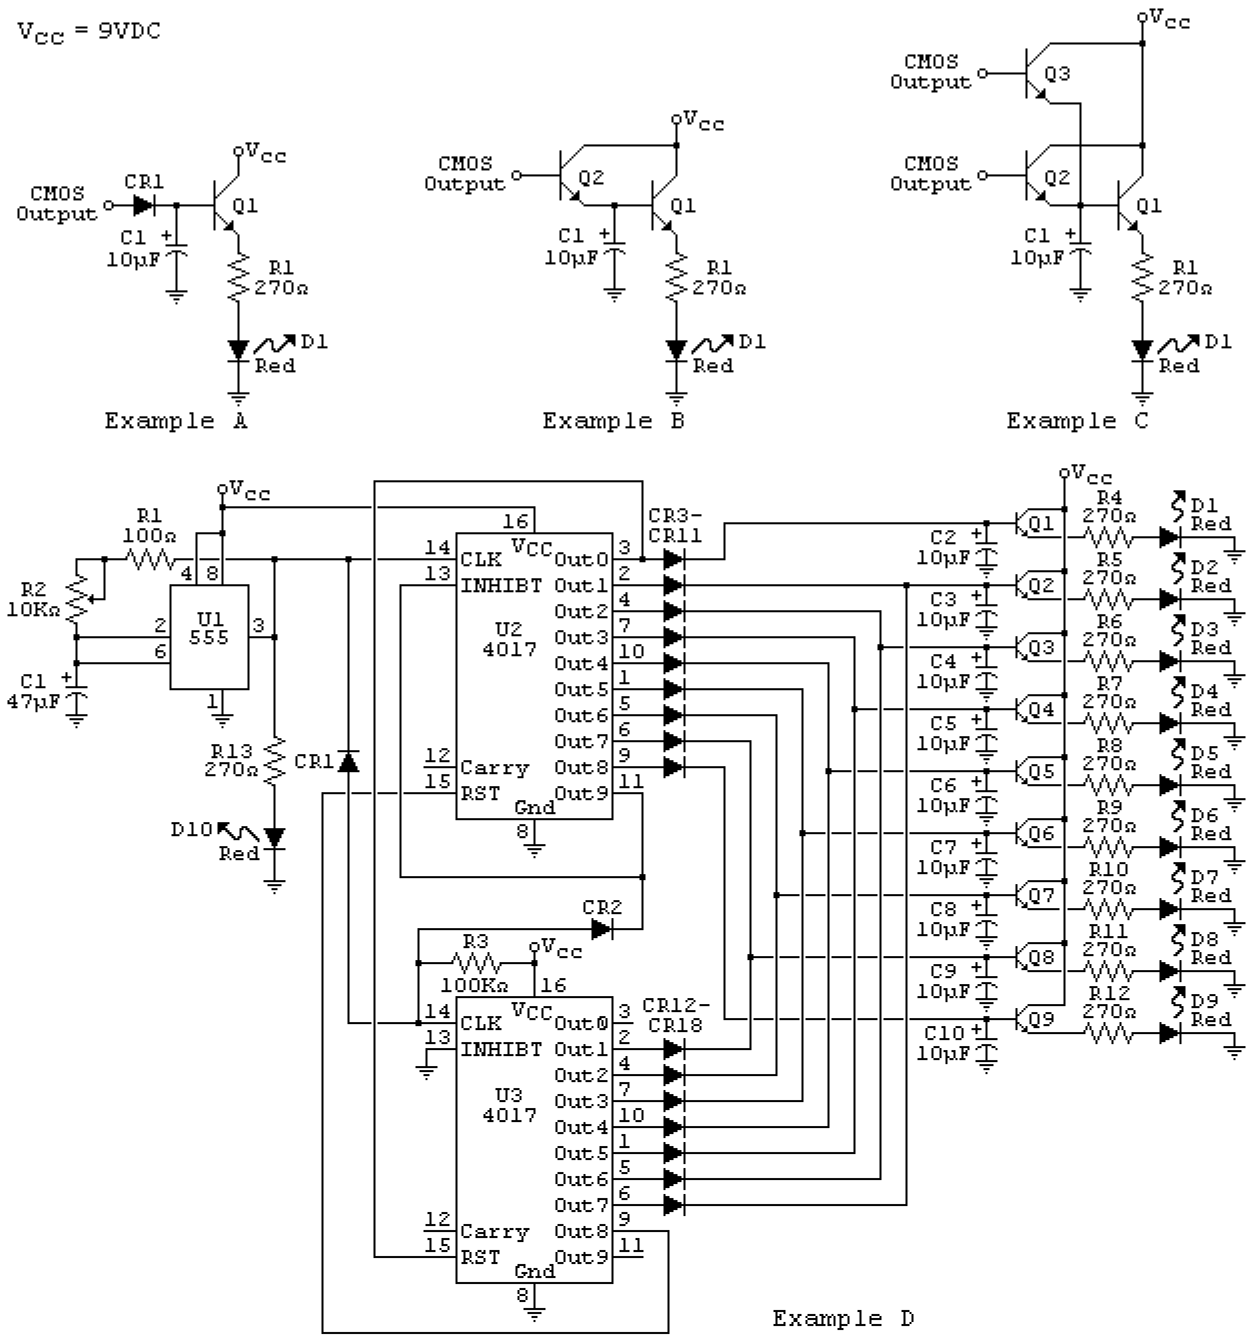

Figure 12.3

<u>Flickering LEDs (Fire!)</u>

Another effect that is often wanted is flickering LEDs. The flicker is harder than it sounds, it isn't a simple circuit. MicroControllers often use tables and other tricks to create a pseudo random number that controls the brightness of the LED, but there is an older technique that works quite well.

Back in the day of tube radios model train enthusiasts would use a small light bulb connected to an AM radio in place of the speaker. The bulb would glow a dim red and flicker, almost like a camp fire, when the radio was tuned to a station.

Modern portable radios don't have anything like the voltages older tube radios used, but you can still use the base concept. It doesn't matter if it is AM or FM. Feed the earphone output of the radio into either one of the following circuits.

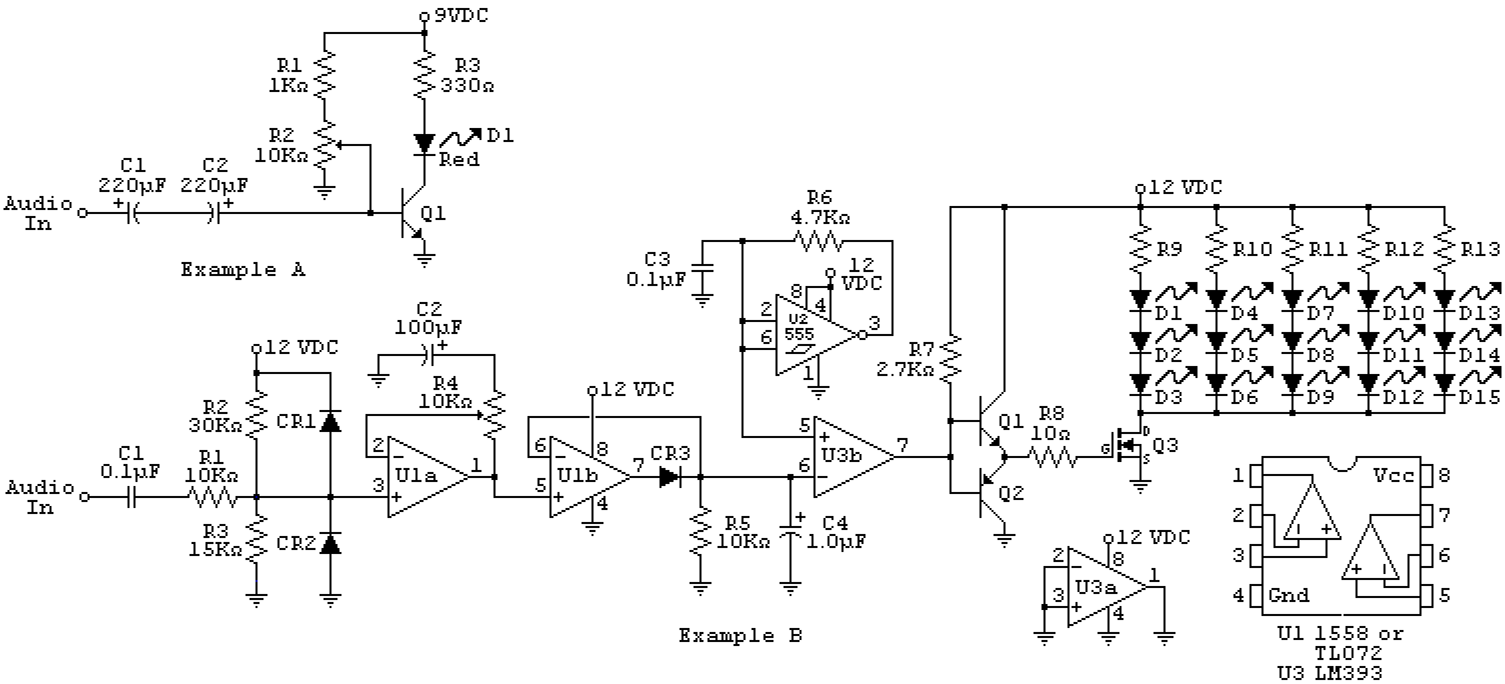

Figure 12.4

If you need a small light the simple transistor circuit in Figure 12.4A would work, while the Figure 12.4B would be for a bigger flame, such as you might see in a Halloween Cauldron that is often commercially available, the kind that uses fans, ribbons, and lights to create a large flame effect.

If the source is a stereo radio then you could have two banks of LEDs almost in sync with each other. This would make for an even better effect.

This circuit is basically a single channel color organ. Add audio frequency filters and it would be straight out of the sixties.

RGB LEDs (Millions of Colors)

Now we come to a fairly new device, a LED that has 3 colors (and 3 LEDs) built in. They can be common anode or common cathode (in other words, the direction of the LEDs can be different), but they are usually 4 lead devices, with one of the leads being shared by all 3 LEDs.

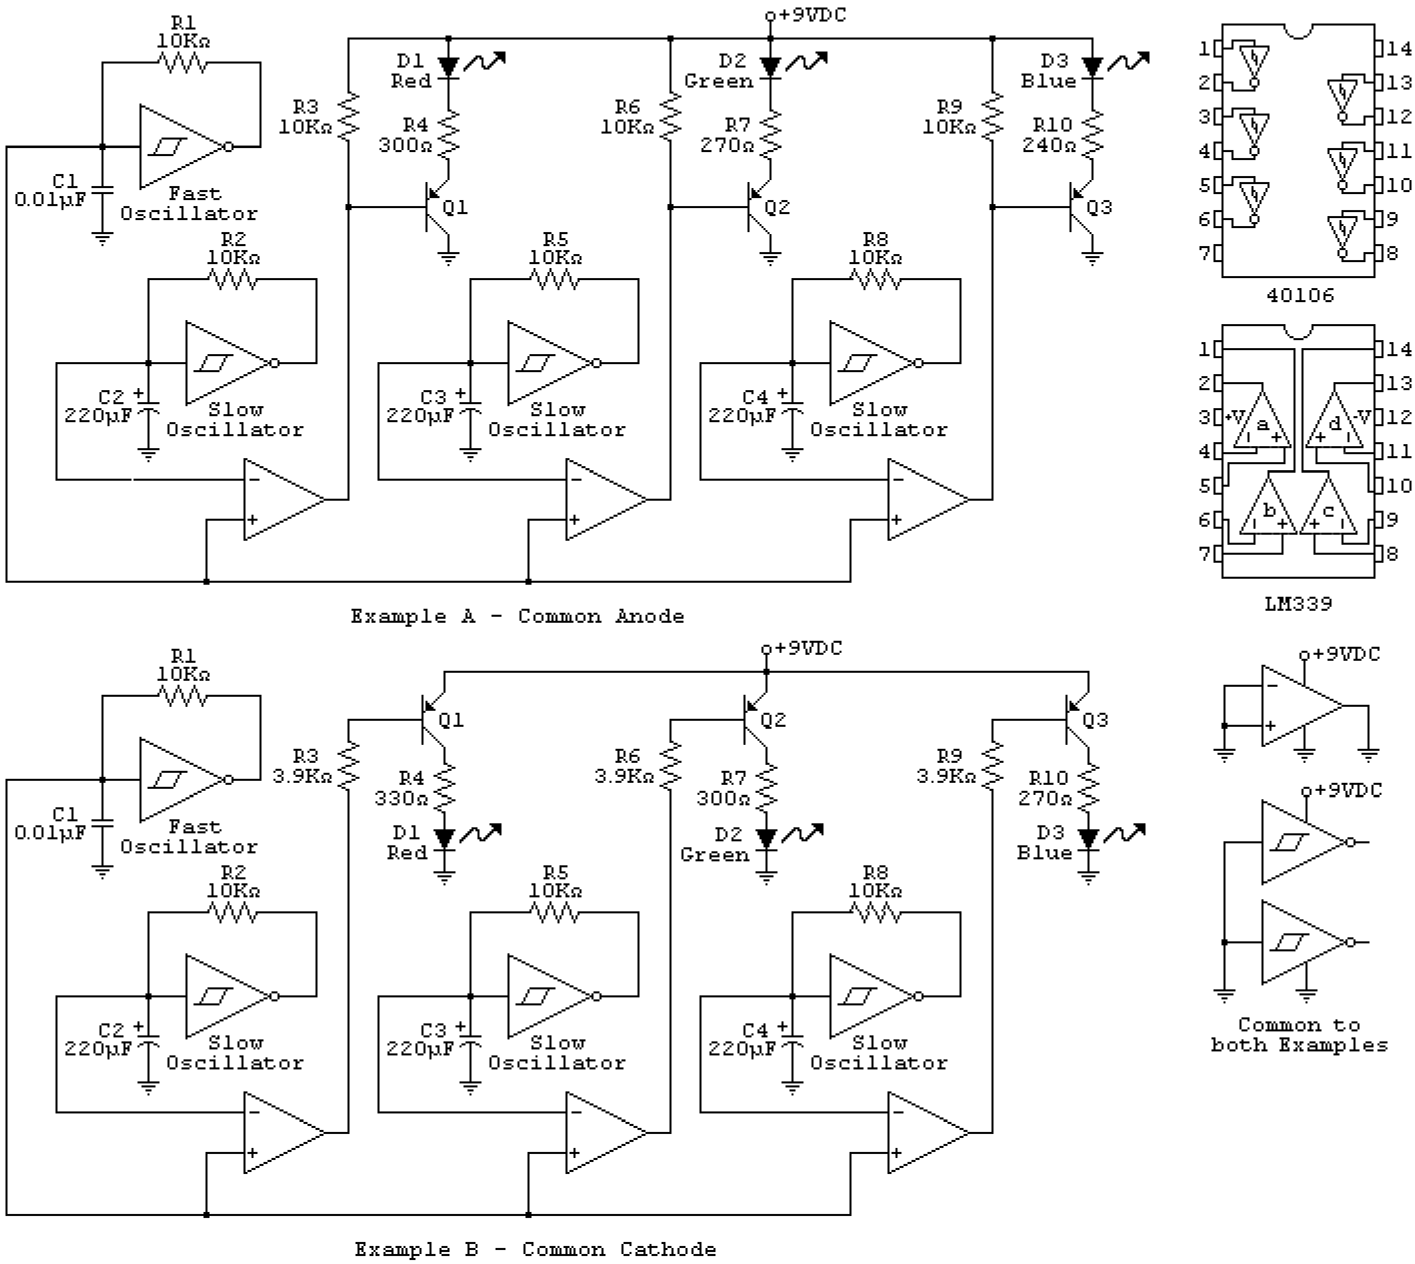

Figure 12.6

While a 555 or 556 would work well for this design there is another way that will dramatically cut parts count. The CMOS chip 40106 has 6 Schmitt Trigger inverters in one package with a power supply spec of 3V to 15V, and since we need four of them this one chip can replace two 556 chips (or four 555's). As usual the unused gates are tied off to prevent oscillations and other unwanted problems.

This design has 3 independent oscillators, which will all fade the LEDs all the way on and all the way off. You will probably want to vary R2, R5, and R8 a little to make them oscillate at slightly different frequencies, but even simple variations in tolerances of the components will cause them to vary. Over time, every possible intensity of RGB will be displayed, each with its own color (much like a color TV). It will cover the entire visible spectrum, baring imperfections in the LED itself.

Conclusion

LEDs are among the more fun circuits to build. They are easy to construct, give instant feedback when they work, and can be tweaked in many different ways. Your imagination can take these basic ideas even further. The field is still advancing very quickly, future models of LEDs will probably replace light bulbs, and we'll be building circuits to make them do wild and crazy things right along side. If you are interested in the history of these LEDs, I would recommend the online LED Museum.

Appendix

People have been coming up with schematics based off of the work here. I'll link or attach the schematics for easy access. As they come out with them I'll post them here. Many cases the schematics will be slightly modified as to fit their requirements, but a working PCB layout is always easier to wire than prototyping.

Figure 11.4

25 LED Open Ended Sequencer

An approach from nerdegutta

http://forum.allaboutcircuits.com/sh...t=49520&page=5

Important discosion(LED) : http://forum.allaboutcircuits.com/showthread.php?t=18613

Source: http://forum.allaboutcircuits.com/showpost.php?p=174485&postcount=26

Why Drive High power LEDs with Constant Current?

Abstract: When applying high power LEDs for display backlighting or other illumination applications, there are two reasons to drive them with constant current:

1. *To avoid violating the Absolute Maximum Current Rating and compromising the reliability.*

2. *To obtain predictable and matched luminous intensity and chromaticity from each LED.*

This application note describes the characteristics of a range of typical LEDs and circuits which achieve the necessary constant-current drive.

The forward current vs. forward voltage of six random high power LEDs (three from each of two manufacturers) is shown in **Figure 1**. In this case, driving these six LEDs with 3.4V, for instance, will cause their forward current to vary from 10mA to 44mA, depending upon the LED.

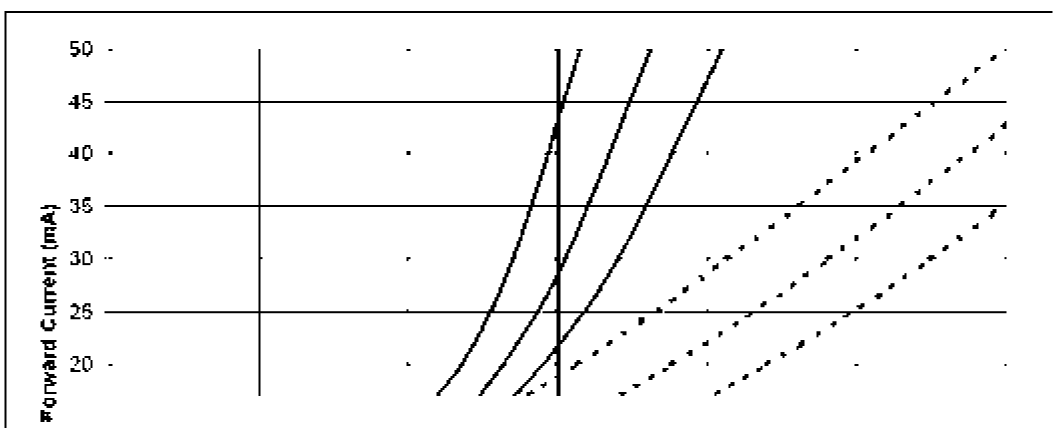

Figure 1. Forward Current versus Forward Voltage for six random high power LEDs (three from each of the two leading manufacturers). Note that the forward current at any given voltage varies widely—10mA to 44mA at 3.4V, for instance.

For reliability, it is important to not violate the LEDs absolute maximum current rating. A typical value might be 30mA abs. max., but as shown in **Figure 2**, the maximum current is de-rated to handle increasing ambient temperature. It is common to limit the current to 20mA for use up to 50 degrees centigrade. By combining the information in Figures 1 and 2, it is obvious that driving high power LEDs with constant voltage is not a reliable solution.

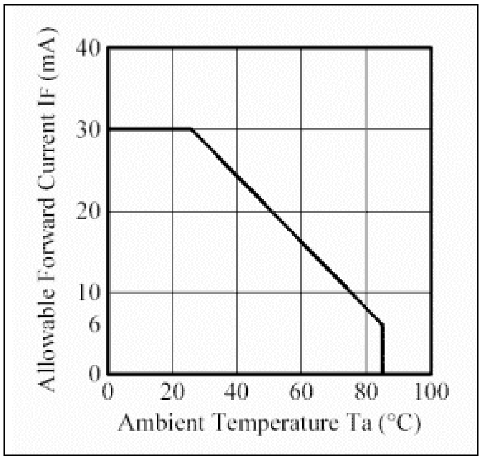

Figure 2. The absolute maximum forward current rating of a high power LED is typically de-rated as ambient temperature increases. (Courtesy Nichia Corporation.)

Furthermore, it is important to note that a high power LED's luminous intensity and chromaticity (color) are tested and best controlled by driving it with constant current. A typical high power LED specification can be seen in **Figure 3**.

(2) Initial Electrical/Optical Characteristics (Ta=25°C)

Item		Symbol	Condition	Min.	Typ.	Max.	Unit
Forward Voltage		V_F	I_F=20[mA]	-	3.6	4.0	V
Reverse Current		I_R	V_R= 5[V]	-	-	50	μA
Luminous Intensity	Rank T	Iv	I_F=20[mA]	720	860	1000	mcd
	Rank S	Iv	I_F=20[mA]	500	600	720	mcd
	Rank R	Iv	I_F=20[mA]	360	430	500	mcd

∗ Measurement Uncertainty of the Luminous Intensity : ± 10%

Color Ranks (I_F=20mA,Ta=25°C)

	Rank a0			
x	0.280	0.264	0.283	0.296
y	0.248	0.267	0.305	0.276

	Rank b1			
x	0.287	0.283	0.330	0.330
y	0.295	0.305	0.360	0.339

	Rank b2			
x	0.296	0.287	0.330	0.330
y	0.276	0.295	0.339	0.318

	Rank c0			
x	0.330	0.330	0.361	0.356
y	0.318	0.360	0.385	0.351

∗ Measurement Uncertainty of the Color Coordinates : ± 0.01

Figure 3. For a typical high power LED, the entire electrical specification is tested at IF =20mA. Therefore, to obtain predictable and matched luminous intensity and chromaticity (color), constant current drive is recommended. (Courtesy Nichia Corporation.)

Figure 4 shows four common power-supply circuits to drive high power LEDs. **Figure 5** shows the corresponding regulation accuracy when regulating our six random LEDs. In figure 5, the regulator's output load-line characteristic is plotted on top of the LEDs' Vf curves. Where the lines intersect is the regulation point for each LED.

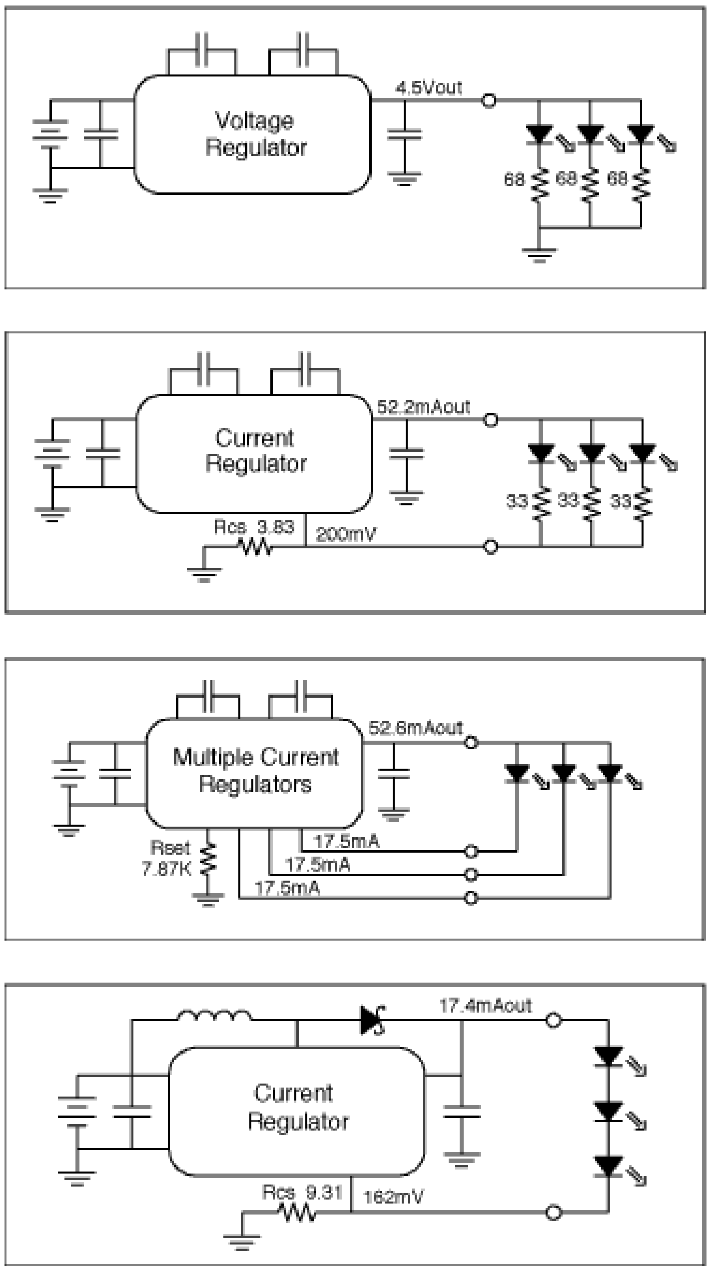

Figure 4. High power LEDs are commonly powered in four different ways: (a) a voltage source and ballast resistors, (b) a current source and ballast resistors, (c) multiple current sources, and (d) a current source with the LEDs in a series connection.

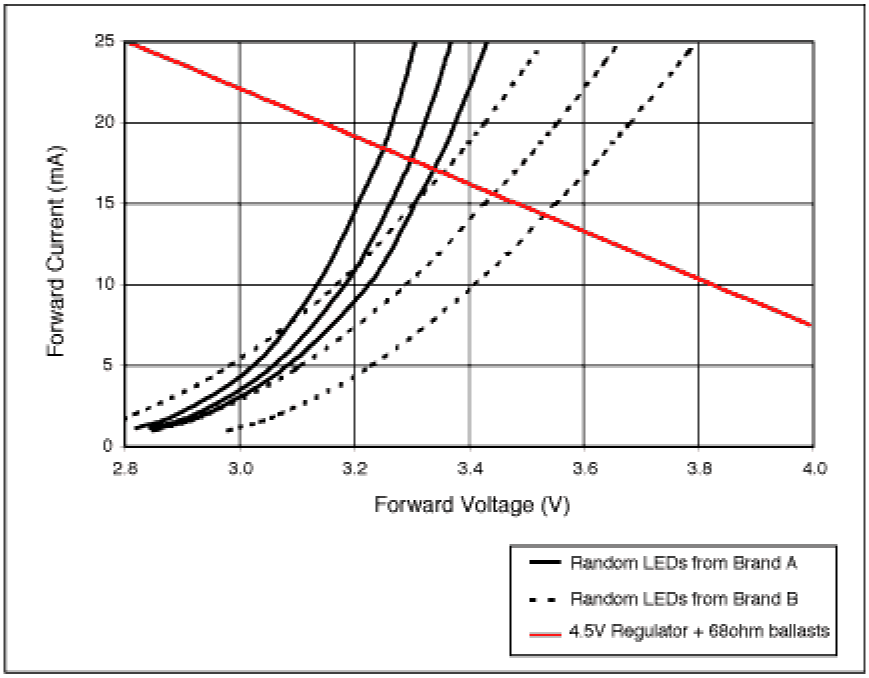
25
20
15
10
5
0
Forward Current (mA)
2.8
3.0
3.2
3.4
3.6
3.8
4.0
Forward Voltage (V)
Random LEDs from Brand A
Random LEDs from Brand B
4.5V Regulator + 68ohm ballasts

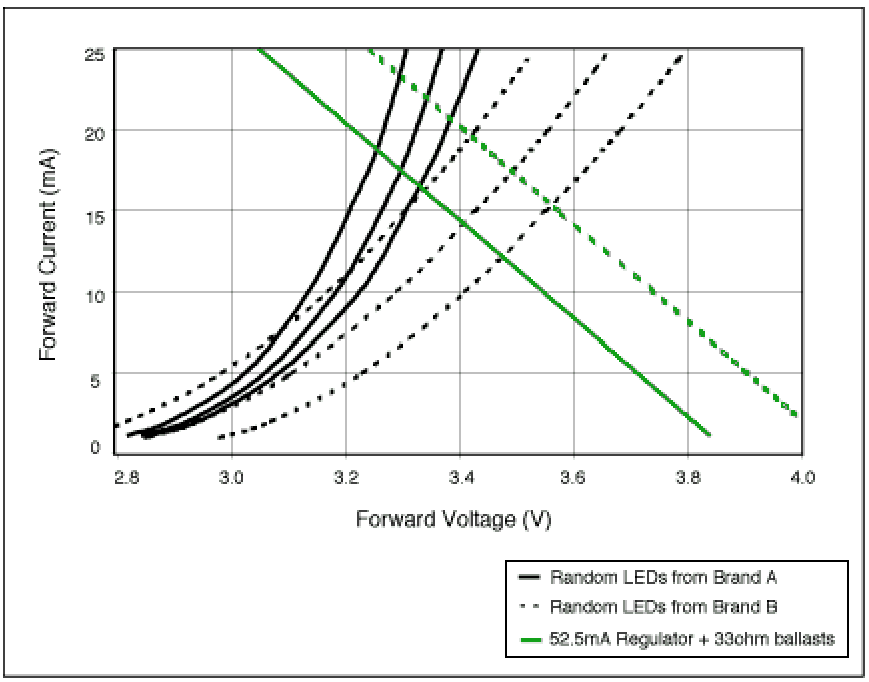
25
20
15
10
5
0
Forward Current (mA)
2.8
3.0
3.2
3.4
3.6
3.8
4.0
Forward Voltage (V)
Random LEDs from Brand A
Random LEDs from Brand B
52.5mA Regulator + 33ohm ballasts

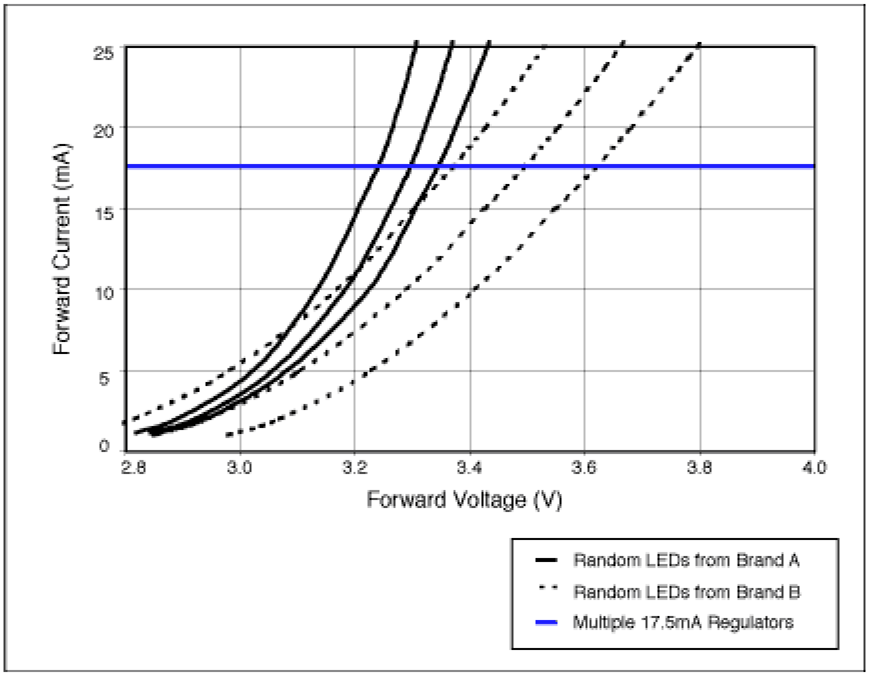

Figure 5. The variation in white-LED forward voltage (Vf) influences the current regulation accuracy differently depending upon the regulation scheme: (a) a voltage source and ballast resistors, (b) a current source and ballast resistors, (c) multiple current sources or a current source with the LEDs in series. The Vf curves of six random high power LEDs (three from Lot A and three from Lot B) are shown. Where the regulator's output Load-Line curve intersects the LED's Vf curve is the regulation point.

The circuit of Figure 4a shows how to use a voltage regulator and ballast resistors to control the LED current. The advantages of this approach are that a great variety of voltage regulators may be applied this way and that only one terminal connects the regulator to the LEDs. The disadvantages are that the efficiency is not very good due to power lost in the ballast resistors and that LED forward current is not accurately controlled. Figure 5a shows that the current in our six random LEDs varies from 14.2 to 18.4mA, while brand A runs about 2mA brighter on average than brand B.

The circuit of Figure 4b regulates the total current applied to the LEDs, while ballast resistors are still used for LED-to-LED matching. The MAX1910 is an example of this type of current regulation. The circuit capitalizes on the fact that the LEDs are better matched within one manufacturer's lot and that most variation is from lot-to-lot or brand-to-brand. Because of this, the ballast resistors can be reduced to waste half as much power while providing similar current control as in the previous circuit. Figure 5b shows that the current in our six random LEDs varies from 15.4 to 19.6mA; however, brand A variation is even less and both brand A and brand B run at the same average current of 17.5mA/LED. The disadvantages are that there still remains significant power loss in the ballast resistors and the LED currents do not match perfectly. Nevertheless, this circuit represents a good compromise between performance and simplicity.

The circuit of Figure 4c regulates the individual current in each LED and requires no ballast resistors. Current regulation accuracy and matching is controlled by the accuracy of the individual current regulators. The MAX1570 is an example of this type of current regulation and exhibits 2% typical current accuracy and 0.3% typical current matching. Because the current regulators are low dropout, the efficiency can be very high. Figure 5c shows that the regulated current is a constant 17.5mA for all six randomly selected high power LEDs. The lack of ballast resistors saves board space, but four terminals are required between the regulator and LEDs. This type of circuit represents a high-performance solution that can readily compete with inductor-based solutions.

The circuit of Figure 4d is a high-efficiency inductor-based boost converter configured to regulate current. A low feedback threshold minimizes the power wasted in the current-sense resistor. Because the LEDs are arranged in series, the LED current matches perfectly under all conditions. Current accuracy is determined by the regulator's feedback threshold accuracy and is independent of the LED's forward voltage variation. The MAX1848 and MAX1561 are examples of this type of current regulation and exhibit 87% (3 LEDs) and 84% (6 LEDs) efficiency (P_{LED}/P_{IN}), respectively. Some other advantages include only two terminals between the regulator and LEDs and that the series arrangement of the LEDs is independent of which particular boost converter is used, giving designers much flexibility. The disadvantages include the inductor's size (especially height), cost, and radiated EMI.

....

Overcurrent protection for high-power LEDs

High power LEDs:

The recent availability of high-power LEDs provide an alternative to Lasers for high intensity light sources that can be easily modulated. Because the only practical means of modulating most light sources is through intensity (or "amplitude") modulation, these modulation schemes invariably the varying of the current through the LED to achieve the desired modulation.

As with any electrical device, there are some practical current limits that must be observed in order to avoid destruction of the device. One of these is excess voltage (from static discharge, for example) but most high-power LEDs are rather well protected in this regard.

Excess current, however is another story!

The Luxeon III LEDs are rated for a (nominal) 3 watts of dissipation, but some of the devices in the series (the Red, Red-Orange, and Amber) are actually rated for continuous operation at a bit over 4 watts - somewhat higher power than the other colors in the Luxeon III product line.

For the remainder of this discussion, we will be assuming the use of the Red, Red-Orange, and Amber Luxeon III devices with the higher power capability.

Given an "infinite" heat sink, the absolute maximum continuous current permitted for these devices is 1.544 amps with a peak pulse current of 2.2 amps although the parameters of that pulse (e.g. width and duty cycle) are not noted.

Through testing, I have determined that a properly heatsinked Luxeon III will *probably* tolerate continuous operation at 2.5 amps for a short period (30 seconds or so) without damage Other experience has shown that at a current of somewhere between 3 and 8 amps, the bond wire on the negative lead of the Luxeon will fuse (open) and render the LED inoperable.

Having accidentally made several Luxeon III's "inoperable" I decided that all optical transmitters will contain circuitry to limit the current to a value that will be safe enough to protect the LED during short-term accidents.

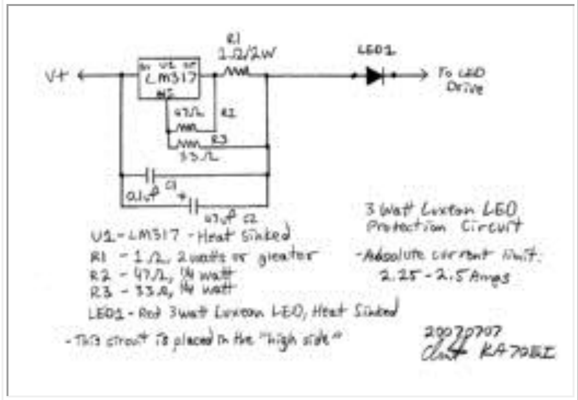

Figure 1:
LM317-based current limiter circuit
Click on image for a larger version.

A few notes concerning use with Luxeon LEDs:

- Red, Red-Orange, and Amber **Luxeon I** LEDs have a much lower peak current rating - about 550 mA. At this lower current, R1 would need to be above 2.5 ohms and R2/R3 would be adjusted accordingly. If a precise value of R1 (around 2.2 ohms) can be selected, R2 and R3 can be eliminated and the "adjust" terminal can be connected directly to the "load" side of R1. This lower current/voltage also means that the total drop of this circuit will be in the 3.5-4 volt range at the limiting current.

- **Philips is apparently phasing out the Luxeon I, III, and V lines in favor of the lower-power Luxeon Rebel devices.** Since I have not used those other devices, the techniques described here may not directly apply unless the maximum current is appropriately adjusted. For the time being, however, the Luxeon III devices are still available from various sources. The Luxeon Rebel devices have a peak current of 700 milliamps, so a 1.8 ohm resistor in place of R1 (with R2 omitted and the "ADJ" terminal connected to the "load" side of R1) would be appropriate.

- With "ultra high-power" LEDs - that is, those with maximum peak currents in the 10's of amps - the following methods may not directly apply as the LM317 and similar 3-terminal regulators simply can't handle the current! In those cases, current-limited supplies, a slightly more elaborate overcurrent protection circuit and plain, ordinary care and common sense are your best defenses against accidentally "killing" an LED!

Current limiting circuitry:

The simplest current limiter is a series resistor. While simple, it's somewhat costly in terms of power dissipation and its effectiveness (both in preventing damage and in allowing normal operation) is somewhat dictated by the available supply voltage: Too high a value and the LED cannot be driven to full current. Too low a value and the LED is not adequately protected. If the supply voltage is lower or higher than expected, the effects are similar.

A solution to this problem is the use of a common 3-terminal voltage regulator, the LM317. This device is commonly used to regulate voltages, using a simple resistive divider to set the output voltage, but it can also be used as a fairly precise current source.

Even though the "official" rating of the LM317 in the "K" (metal TO-3) or "T" (TO-220 "tabbed) packages is 1.5 amps, it is perfectly capable of operating at somewhat higher currents than this. One of the features of the the LM317 is its built-in current limiting and a quick peek at the data sheet will reveal that its internal ("room-temperature") current limit is typically 2.2 amps - but it could be anywhere from

1.5 to 3.5 amps. In testing dozen or so devices from different manufacturers made over the past two decades or so I observed that they typically exhibited an absolute current limit in the 2.25-2.75 amp area and none limited at less than 2 amps when the device was at room temperature.

It is important to note that this "maximum current limit" is only valid when the device is at room temperature. As the device warms up the "thermal protection" begins to take effect, reducing the maximum available current with increasing temperature. I also observed that at "touchable" temperatures (e.g. temperatures at which one could touch the device for 5 seconds or so before deciding that it was too hot) that the reduction in the maximum current was minimal: It wasn't until the device got well above about 150 degrees F (about 65 C) that it really started to limit its current.

Figure 1 shows such a circuit. This is a modification of the typical LM317-based current source in that R2 and R3 have been added, but its operation is the same:

- As the current increases, a voltage appears across R1.
- When the voltage on the "Adjust" terminal is about 1.245 volts below that of the "Output" terminal, the current is reduced.

Normally, the adjust terminal is connected directly to the "load" side of R1: When the voltage drop across R1 exceeds about 1.245 volts, the current reduction occurs. In testing, I built such a circuit, using only R1, but I found that in order to set the limiting current within the 2.25-2.5 amp range, I needed a resistor of about 0.6 ohms - not a common value.

While it would have been possible to synthesize such a resistance by paralleling resistors or just using a piece of known-length and gauge of small magnet wire, I decided to construct a circuit that used more common resistor values and **Figure 1** is the result. In this case, R1 is set for a slightly higher resistance to effect a higher voltage drop while R2 and R3 divide that voltage such that at 2.25-2.5 amps, the Adjust-to-Output differential is in the 1.245 volt range.

Connected in parallel with this circuit is bypass capacitance. If the LED is modulated with waveforms that contain high frequency components (like those in a PWM circuit, high frequency subcarrier, or video modulation) the bypass capacitance (C1 and C2) allows high frequency components to pass around the LM317 while still offering protection to the LED from a DC current fault. The use of two capacitors is recommended as the 0.1uF capacitor will be transparent to the highest frequency components, while the larger electrolytic will better-pass the lower frequency components. It is recommended that 105C (high temperature) and high-frequency electrolytics (such as the low-ESR types designed for switching power supplies) be used - particularly if the capacitor is exposed to the heat from the LM317.

Total voltage drop of the circuit in **Figure 1** at 2.5 amps is around 4.5-5.0 volts. This, in series with a Luxeon III LED would imply that there is a total voltage drop of as much as 7.75-8.25 volts - but this is still enough headroom for many circuits that operate from a 12 volt source - but some circuits may need to be modified. As mentioned before, if R2/R3 are omitted and the Adjust terminal is connected to the load side of R1, the value of R1 can be reduced to an experimentally-derived value in the 0.6 ohm area, reducing the voltage drop at 2.5 amps by a volt or so.

Constructing the circuit:

As can be seen from the pictures in **Figure 2** this circuit can be constructed in a number of ways. It is essential that the LM317 have at least some heat sinking. In normal operation, the average LED current of a Luxeon III will be about 1 amp, peaking to 2 amps or so at 100% positive modulation. Under these conditions, the LM317 will have to dissipate about 2 watts of heat while R1 will be dissipating around 1 watt.

Even though this is a fairly small amount of heat, it is important that this circuit ***NOT*** be mounted on the same heat sink as the Luxeon as a quick check of the Luxeon's specifications will show you that is it warms up, its efficiency drops, so it is best not to add another device that will warm it up! As shown in these pictures, the LM317 is mounted separately from the Luxeon's heat sink - preferably off to the side so that convection will not be likely to cause one heat sink to heat the other.

It worth noting that it is also a good idea *not* to ***excessively*** heat sink the LM317. A good example of this is shown in the bottom picture of **Figure 2** where the LM317 is simply soldered to a piece of copper circuit board material. This provides adequate heat sinking for normal operation where the LM317 is dissipating about 2 watts. If, however, a fault develops and, say, a full 12 volts is placed across the LED drive, this will cause the current to exceed the 2.25-2.5 amp threshold. When this happens, the voltage drop across the LM317 necessarily increases and this will dramatically increase the amount of heat being dissipated by the LM317. In this state, the current will start to drop as the LM317 gets hot, soon reducing the current to well below the 2.25-2.5 amp threshold. It is under these conditions that a very large heat sink is less-desirable as it would prevent the LM317 from heating up and reducing the operating current.

Operational use:

As mentioned before, there must be enough voltage headroom in the power supply voltage and modulator circuit to accommodate the added voltage drop of this protection circuit. In the modulators that I have built, power MOSFETs are used along with 1 ohm current-sense resistors, adding as much as 2.5 volts of additional drop - but these run will drive the LED to over 2 amps with an 11.5 volt supply - the voltage of a discharged lead-acid battery.

Under normal conditions, this circuit is completely transparent: Other than an addition I*R voltage drop across the sense resistor, there are no effects at all. When the current exceeds the preset threshold, however, the LED drive is clipped at that level and could cause some distortion.

A quick look at the schematic will also reveal that I have placed it in the "high" side of the LED supply. This was done in the event that the "low" side of the LED (which, in my case, goes to the modulator) gets shorted to ground: A high-side current limiter will still protect the LED in this case.

Note that because this circuit sits inline, there is no real voltage limit to it aside the voltage differential that would occur under fault conditions when the LM317 is current-limiting. What this means is that if you were to run a string, say, 5 Luxeons in series from a 24-36 volt supply, you could: This circuit would happily protect them all.

Other types of circuits:

When I was deciding how to protect the Luxeon, I was considering other types of circuits such as comparators driving MOSFETs or other transistor-based limiter. While these would work - and some have the advantage of having lower voltage drops - none of them were as simple as the LM317-based circuit shown here and some of them require their own supply voltage to operate. This circuit has the advantage of being cheap, simple, and it need only be placed in series with the current flow, not needing a ground or voltage reference.

Why not a fuse?

There is also the obvious question as to why not use a fuse? If you look closely at a Luxeon III emitter *(see **Figure 3**)* you'll note that it ***already*** has a fuse disguised as a bond wire connecting one of the terminals to the center of the emitter die. If you use an external fuse, it *must* be one of a lower rating *and* faster than that bond wire! Clearly, this rules out slow-blow fuses as well as thermistor-type PTC fuses. It is also worth noting that there is quite a difference between the "hold" current of a fuse and the current at which it will blow quickly, and a current at which the fuse will fail over a longer period of time - not to mention a fairly wide range of variability of fuses of the same ratings - even those from the same manufacturer. Finally, if one uses fuses, it is imperative that spares be kept onhand.

Final comments:
Despite the voltage drop which is something that can be worked around in most circuits, usually through a minor change, the described protection circuit works very well to keep peace of mind - and it could prevent the sudden, expensive end of testing should some sort of fault or accident occur that might open up the LED.

LED DESIGN FORUM: Avoiding thermal runaway when driving multiple LED strings (MAGAZINE) ***A current mirror and monitor can save a parallel string from failure.*** An increasingly common method of increasing the light output from a high power LED cluster is to run parallel strings of LEDs from a single constant current source. But this option is not without its hazards.

A typical high power 350mA white LED has a forward voltage (Vf) of about 3.3V, so if an cluster of 10 LEDs were required in an application, connecting of the LEDs all in series would require a driver capable of delivering at least 33V. If the supply voltage is 24VDC, then an expensive boost converter would be required with all the attendant EMC problems it creates.

Connecting the LEDs as two strings of five wired in parallel requires a higher current 700mA constant current source but only 16.5V output voltage. Thus a low cost buck converter running from 24VDC can be used. The circuit (similar to that shown in **Fig. 1**) can be found in many manufacturers' datasheets.

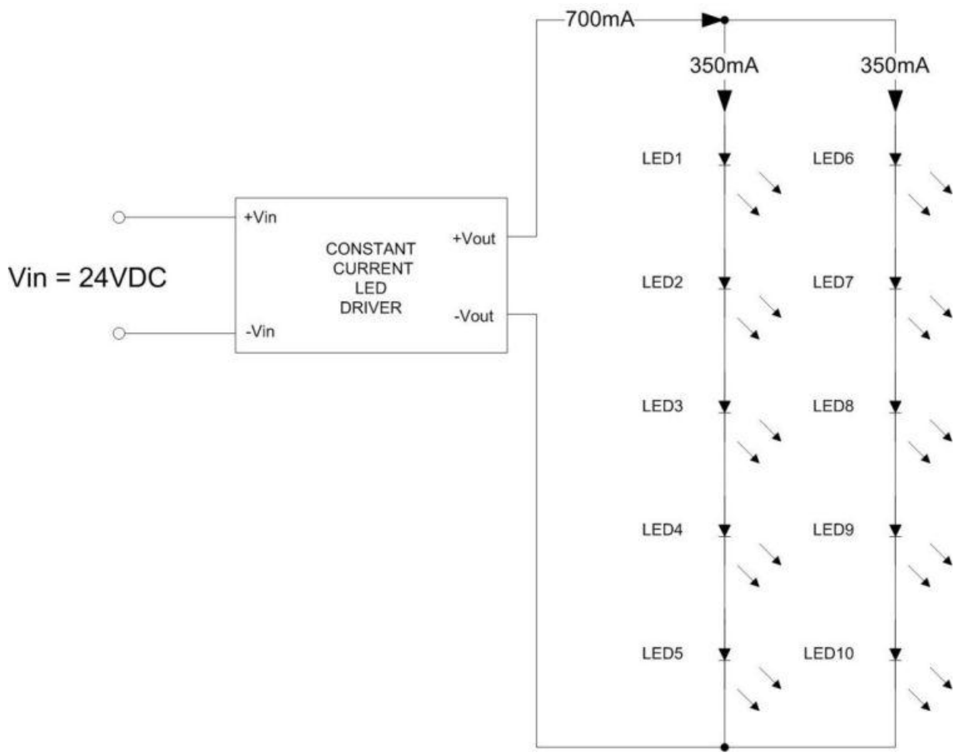

Fig. 1. Driving multiple LED strings

The basic assumption made with this circuit suggestion is that the 700mA regulated current from the LED driver will be shared approximately evenly across both strings of LEDs, i.e. each string of LEDs will see 350mA of current. However, this is rarely the case.

How LEDs become overdriven
Even if the LEDs are all from the same production batch and sequentially manufactured, the Vf of individual LEDs still has a ±20% tolerance. The tolerances mean that the total forward voltage for each string can be very different and therefore the current mismatch significant (see **Fig. 2**).

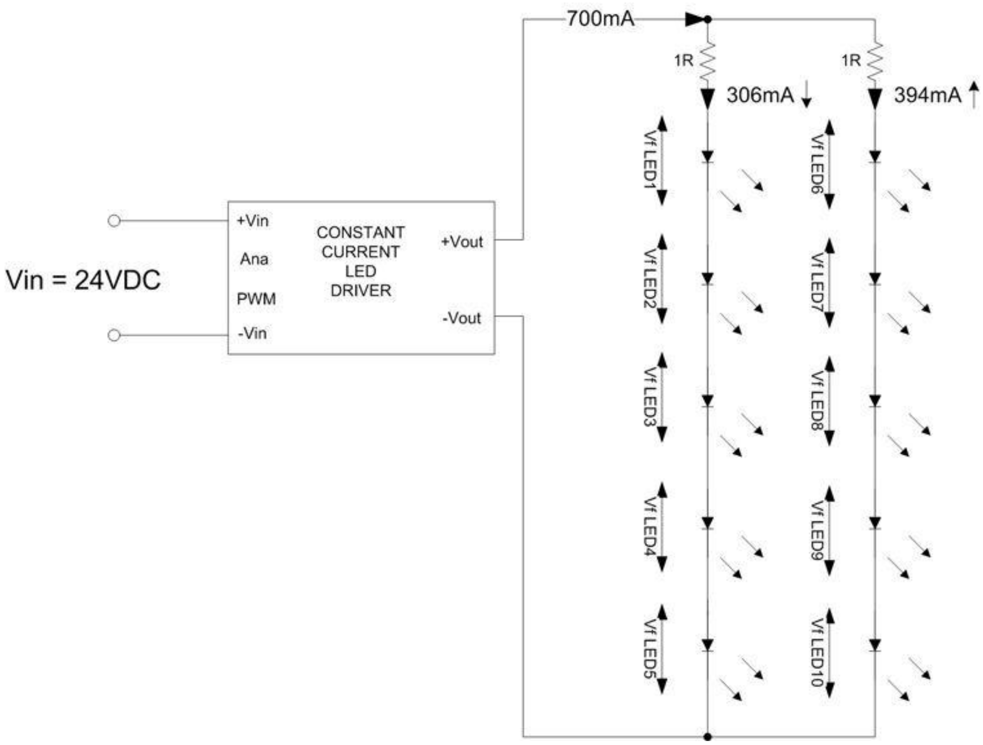

Fig. 2. Real life situation

In a test using identical SMD LEDs from a single production batch and using 1 Ohm resistors to help balance out the forward voltages, the currents flowing in each string were measured to be 306mA and 394mA. The LED driver was still doing its job of correctly limiting the current to 700mA, but the over-current flowing through the second string was seriously overdriving the LEDs.

Worse, as the LEDs started to get warm, the combined forward voltage of the higher current string started to decrease. This increased the imbalance and more current started to flow through the already over-driven string. The current through the other string of LEDs reduced as the constant current driver compensated, so they started to cool down and their forward voltage increased.

The net result was thermal runaway with the majority of current flowing through one string only, even though the LEDs were mounted on a large metal heat sink. The test was stopped when the current imbalance was 600mA to 100mA. Obviously, if this situation was allowed to continue, the over-driven string would eventually fail and then the entire 700mA would flow through the remaining intact string and destroy that as well.

And this circuit is often given as a recommended application example!

Using a current mirror
What is required is a way of balancing the currents flowing through the two strings to ensure that they remain approximately equal, even if the combined forward voltages are mismatched. The balancing circuit must also continually compensate for changes in forward voltage caused by changes in the operating temperature and by aging of the LEDs.

Fortunately, there is a very simple transistor circuit that will do this job admirably. It is called a current mirror and “reflects” the current flowing through one reference transistor onto the current flowing through a second transistor (see **Fig. 3**). As long as the transistors are reasonably well matched in terms of their Vbe values, the currents will also be reasonably well matched.

In tests using Recom's 700mA LED driver and two strings of 350mA Osram LEDs, the currents flowing through the two strings were matched to an accuracy of about 87% over the entire input voltage range of the driver from 16VDC to 36VDC. The LED currents were stable as the LEDs warmed up and no thermal runaway was observed. It is important that the two transistors are both at the same temperature so a copper clamp was used to thermally connect both transistors together to keep their Vbe voltages stable.

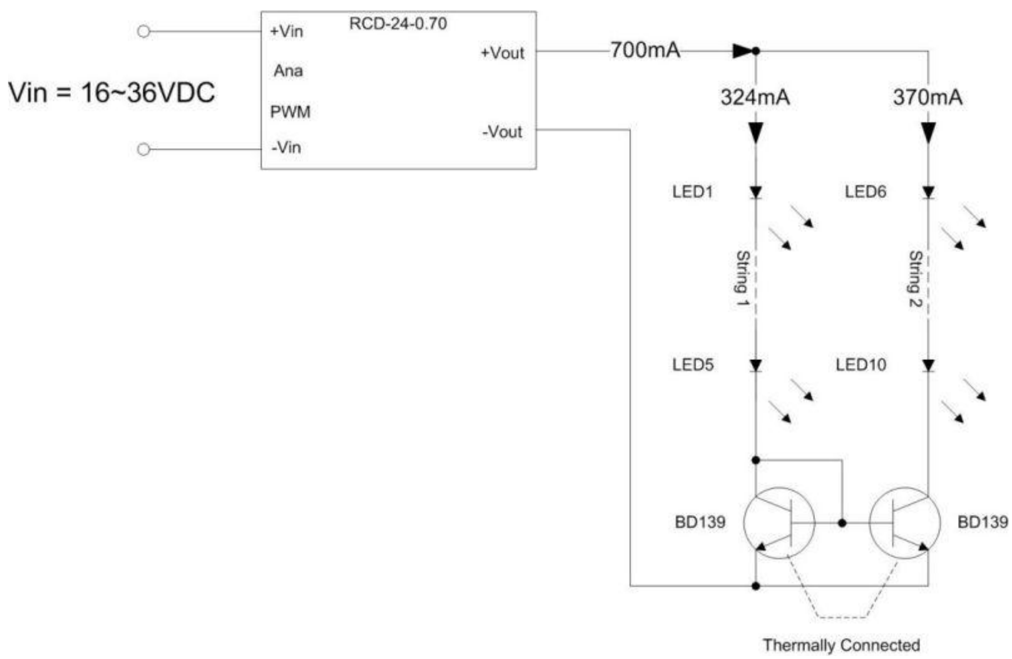

Fig. 3. Using a current mirror

In addition, if any of the LEDs in String 1 fail, the current to ALL of the LEDs is disconnected. Thus the LEDs on String 2 are automatically protected against being over-driven.

However, this circuit is still not ideal. The currents are not perfectly matched and if any of the LEDs in String 2 fail, then all of the 700mA source current will still flow through the first string and destroy it.

Over-current protection

Fig. 4 shows the final version of the current balancing circuit. The addition of 1.5 Ohm resistors in the emitter paths makes the circuit less sensitive to small Vbe changes and balances the currents in the two strings to 99% accuracy.

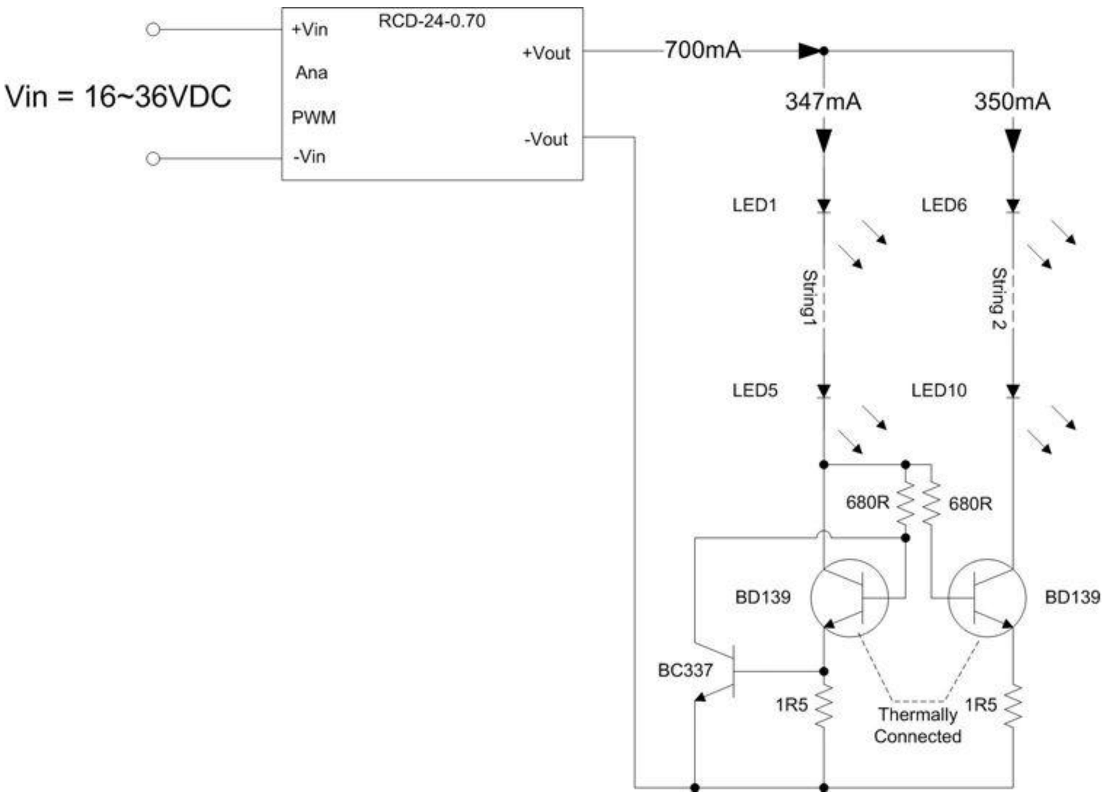

Fig. 4. Over-current protection

The addition of a small signal transistor as a current monitor protects the LEDs from being overdriven in the case of any LED failures. If LED1-LED5 fail open circuit, then the current in the second string falls to zero as before. However, if LED6–LED10 fail, then the current increases in the first string until the voltage developed across the 1.5 Ohm emitter resistor reaches around 0.7V, thus turning on the BC337 transistor and pulling the base voltage of the power transistor to ground and limiting the current. With the component values given in the circuit, the measured current limit was 445mA with String 2 open circuit.

The circuit suggestion given in Fig. 4 can theoretically be extended to any number of LED strings by adding an NPN transistor and emitter resistor to each additional string and tying the transistor bases together. The current flowing through the reference transistor will be faithfully mirrored by all of the other transistors.

However, considering that LEDs are high reliability illumination sources and the driver and associated components need to be equally reliable to get the maximum lifetime out of the system, it is recommended that the circuitry be kept as simple as possible and restricted to only one or two strings per driver.

....

Current source with LM317

Current source A current source is an electrical or electronic device that delivers or absorbs electric current. A current source is the dual of a voltage source. Current sources can be theoretical or pratical. I will handle only the practical with the use of a LM317.

Why using a current source instead of just a simple cheap resistor? In many situations a resistor will be enough but also some devices need a constant current irrespective of the voltage: ex. a 20mA current loop transmitter. Also LEDs are current-driven devices that require current limiting when driven from a voltage source. In most applications, it is desirable to drive leds with a constant current source. The current source is used to regulate the current through the led regardless of powersupply voltage variations or changes in forward voltage drops(Vf) between LEDs.

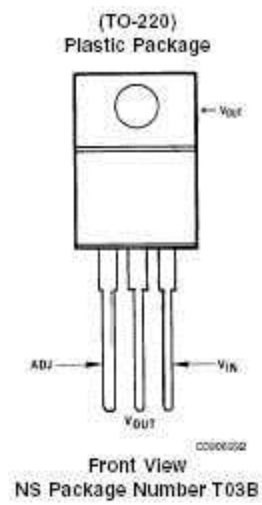

LM317

The LM317 is a monolitic integrated circuit. It's a 3 terminal positive voltage regulator designed to supply more than 1.5A of load current with an output voltage adjustable over a 1.2 to 37V. It employs internal current liniting, thermal shut-down and safe area compensation. The LM317 is cheap, thermal protected, up to 1.5A and easy available.

The schematic

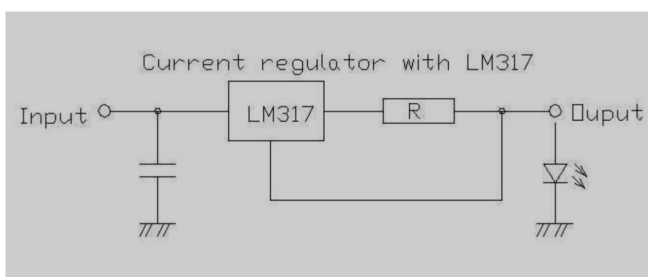

Power dissipation

Because the LM317 is a linear regulator and needs a voltage drop of about 3V, the dissipated power will be the voltage drop over the LM317 multiplied by the current of the circuit.

- P=power loss
- U=supply voltage
- Uf=voltage drop device
- I=current

$$P = (U - Uf) * I$$

Tip: When you have a device with a specific voltage drop and a relative high current(ex. a white lumiled: 3.2V, 0,35A 1W) keep the input voltage as low as possible, but keep in mindthat the LM317 needs a voltage drop of 3V.

<u>Example with a white lumiled:</u> The led needs a forward voltage of 3.2V. To minimize the power losses we will connect the led to a voltage of +-7V (Voltage drop LM317+forward voltage led+1V reserve)

<u>A bad example with a white lumiled:</u> What is the power loss when we connect the led to 11.7V: Well when the the forward voltage of the led is 3,2V and the current is 350mA by a power voltage of 11,7V then with the law of ohm we can find a power loss of +-3W: 3 times the led power (certainly not economical).

I'm sure the LM317 will become very hot. The LM317 devices have internal shutdown to protect from overheating but in all working conditions the junction temperature must be within the range of 0 to 125 deg celcius. So a heatsink maybe required at maximum power loads and maximum ambient temperature.

$$P = (U - U_f) * I$$

$$P = (12V - 3.2V) * 0.35A$$

$$P = 3W$$

When it's impossible to lower the voltage of the powersupply and the current is high maybe a switched current source will be better.<> I will explain this in the future.

<u>Current adjustment:</u>

Now we know the needed input voltages but still don't have a constant output current. For this we gone abuse the voltage regulator. We place a resistor in series with the LM317 and the output device (ex. a led) and connect the adj pin over the resistor. Because the LM317 will regulate the voltage on the adj input allways to 1.25V, we become a constant current through the resistor and connected device

How it works: Over the resistor there is always a voltage present of 1.25V

This means when the current decrease, normaly the voltage over the resistor will be lower also but what happens now: the regulator lets increase his output voltage to adjust a constant voltage over the resistor of 1.2V

So we can calculate with ohms law what resistor is needed to get a specific current.

$$R = U / I$$

$$R = 1.25V / I$$

<u>Example:</u> We will supply 3 lumileds 1W power rated in serie with a 12V battery. The nominal current for the leds is 0,35A We can find the proper resistor value with the formula above:

$$R = 1.25V / I$$

$$R = 1.25V / 0{,}35A = 3.57 \text{ ohm}$$

Thus we need a resistor of 3.57 ohm but will not find one with this value. To solve this problem take a value that's higher. Here in the example we will take one of 3,9 ohm.

The real current will be then I = U / R = 1,25V / 3.9ohm = 0,32A what not will be a problem (it extends the lifetime of the leds)

Power rating of the resistor:

It's easy because:

$$P = U * I$$

$$P = 1{,}25 * 0.32A = 0.4W$$

In real we take a resistor with a 10% higher power rating: here we find 1/2 Watt.