Univerza v Ljubljani
Fakulteta za elektrotehniko
Tržaška 25, 1000 Ljubljana, Slovenija

Lighting Design with LEDs

Seminar paper for course Electrical Installations and Lighting

André Leston

Professor: prof. dr. Grega Bizjak

Ljubljana, May 2018

1. Introduction

The first thing to know about LEDs is the meaning of the word. LED stands for light-emitting diode.

LEDs are being used in many different situations nowadays, and the number of visual signalling devices that rely on LEDs is increasing, essentially because LED sources have specific characteristics that suit really well almost every situation when there is the need to light something: very quick onset and offset times, straightforward control of luminous intensity through simple adjustments in current and voltage, they can provide high perceived brightness and a relatively long lifetime – all of this at a really fair price.

LEDs can be found almost everywhere: they are replacing our traditional home lighting, whether with the purpose of indeed lighting some area with a good lighting solution or just with decorative means; to backlight LCD panels in TVs, computers or smartphones; headlights and taillights in vehicles; flashlights; traffic lights and signs and even the actual street and road lighting.

The purpose of this seminar paper is to talk about the characteristics of LEDs, how they work, how they are driven, how they behave with temperature, their efficacy and so on.

2. Characteristics of LEDs

2.1. Current instead of voltage

As it was said before, an LED is a diode and, as a diode, it is a current device and not a voltage device. This means that if we have a circuit with a resistor, voltage and current are proportional by the Ohm's law (V = RI). If a voltage is placed across a resistor, a current will flow through this resistor, and this current is I = V/R. If the voltage is doubled, the current will double. Conversely, if a current is pushed through a resistor, a voltage will flow.

When it comes to diodes, Ohm's law cannot be applied, because voltage and current are not proportional – they are exponentially related.

The current through a diode can be modelled by

$$I(V_f) = I_0 e^{k \times V_f} \text{ ,}$$

while the voltage can be modelled by

$$V_f(I) = \frac{1}{k} \times \ln\left(\frac{I}{I_0}\right) .$$

In the above equations, V_f is the voltage from the anode to the cathode, I is the current through the diode, k scales the voltage and I_0 scales the current (these two last are constants).

Typical values for LEDs are $I_0 = 3.2\ \mu A$ and $k = 3.64/V$.

By analysing the first equation, we see that the current is a very strong function of voltage. If a voltage of 2.80V is connected to a diode, a current of 85mA needs to run through it. Adding a quarter volt (to 3.05V) more than doubles the current, to 212mA. Increasing the voltage to just 3.30V (another quarter volt) more than doubles the current, again, to 527mA.

So, this means that it is not smart to control a diode by controlling the voltage but, instead, its current. This is the reason why it is said that the diodes, and LEDs in this particular case, are current devices. The forward voltage is constant, no matter how much current flows through it, so the performance of the device is determined by how much current runs through it. In the same way, since power is V x I , and voltage is constant, power is just proportional to current, so the power into the diode is determined by the current.

The optical output of LEDs is given in lumens/watt. So, depending on how many lumens are needed for a specific purpose, we know how many watts are needed and, therefore, how much current is needed. When a power supply for LEDs is designed, it is designed to drive the LED with a constant current.

2.2. Forward voltage

For white and blue LEDs, typical forward voltages range from about 3.1–3.8V. Yellow LEDs are somewhat higher, and red LEDs are down around 2.2V. It's actually a little easier to compare forward voltages of LEDs than ordinary diodes. It's an industry standard to report the V_f at a current of 350mA, even for devices capable of carrying 1A.

As an aside, the underlying reason that LEDs have much higher forward voltages than silicon diodes is that they aren't made of silicon. Their bandgap is different than silicon, that's how they generate light. And the reason there's such a range of voltages for white LEDs is that there's a number of different semiconductors being used in the industry now. Each has its own bandgap. One of the major areas of research for all LED die manufacturers is how to reduce the forward voltage of the device. Reducing this would increase the light output of the LED per watt.

2.3. Reverse breakdown

If LEDs conduct in the reverse direction, there's an excellent chance that they've broken. Most of them have a reverse breakdown voltage of only 5 V. This very low voltage can present a serious problem in a practical circuit. Even with several LEDs in series, the breakdown is still only tens of volts. There isn't any room for error. Any glitch or noise in the control loop of the power supply, or a hiccup in the AC line voltage, may be enough to momentarily generate breakdown, and it doesn't take long - diodes can be broken in microseconds under the right conditions.

Some white LEDs can be driven above their rated current by a large amount. Their phosphors saturate, and the light starts to turn blue, but if this does not represent a problem, they can take quite a bit more current. In general, current generation LEDs seem to have plenty of bond wires to carry a lot of current. So, the real limitation is whether the die can take more current without burning up or developing hot spots. This is primarily a thermal question until it gets to very high currents, and very high currents are bad for efficacy.

2.4. Efficacy instead of efficiency

Efficiency usually refers to electrical conversion. A certain amount of energy is put into a power supply; a certain less amount out gets out. The electrical efficiency is defined to be the output power divided by the input power. It is usually expressed as a percentage. Since there are always some losses, it is a number always less than 100 and not less than 0.

Things are not so straightforward for LEDs. Power gets in, but what gets out is not power, but light. Now light, just like electricity, can be measured in watts. And indeed, deep blue LEDs' light output is measured in watts. Efficiency is then straightforward. It is the output power in watts divided by the input power in watts.

For example, for traffic lights, the point is to know how bright red, yellow and green appear to people. Similarly, for light bulbs using white LEDs, the objective is to know how bright people perceive them to be, not how much optical power is coming out. So, for these nonblue LEDs, watts aren't used, lumens are. And so, efficiency is not the right term either. Instead, the term "efficacy" is used.

In the following table it is possible to compare the efficacy of different light sources, as well as some other interesting parameters in which LEDs are obviously better.

Light Source	**Efficacy (lm/W)**	**Life (hours)***	**Correlated Color Temperature (K)**	**Onset Time†**
Tungsten filament lamp	12-20	750-4000	2700-3200	0.1-0.3 s
Fluorescent (incl. compact)	60-100	10,000-30,000	2700-7500	1-60 s
Metal halide	80-110	10,000-20,000	2800-5000	60-300 s
Xenon	30-60	1000-5000	5000-6000	1 µs
Light-emitting diode (white)	90-130	50,000-100,000	3000-8000	10-20 ns

3. Thermal Performance of LEDs

3.1. Thermal Shifts

Temperature has a big impact on LEDs and on their performance. Many important parameters of LEDs have been discussed in the previous topic and now they will be connected with temperature and how it affects them.

A white LED, for example, is composed of some different components: the die (which emits blue light); the phosphor (which converts part of the blue light into white light); the silicone encapsulant (that protects the die and the phosphor) and the package to keep everything together. Each of these components contribute to the thermal performance of the LED and to its aging.

There is an absolute maximum temperature at which the die starts to fail. There is also a runaway temperature at which some of the parts of the die become hotter and also fail, and all of this causes the wavelength of the blue light to shift. Although this shift is small the phosphors absorb light in a very small band, so it can be enough to affect the efficacy with thermal aging.

The encapsulant needs to be mechanically strong and also optically clear and this is a problem. Initially, epoxy was used to build the encapsulant because it was inexpensive, but it had a problem: it turned yellow with age. So, now, silicone is universally used for high-brightness LEDs. Silicone is way better than epoxy, but it is still a polymer, so eventually, it will also turn yellow with heat and time and this, obviously, will eventually affect the colour of the emitted light – and, therefore, affecting the efficacy.

As for the package, if it is not perfectly reflective, it will partially absorb light and, as it also must survive high temperature for thousands of hours, eventually it will gradually turn yellow, all of this affecting efficacy. Cree (2010) says "… the primary mode of degradation for HB LEDs is the package itself". This leads to an engineering dilemma: cost of the package vs thermal aging performance.

3.2. Temperature affecting electrical and optical behaviour

The forward voltage drop is probably the most important performance change of LEDs with temperature. Typical decrease of forward voltage runs between -2 and -4 mV/ºC. Practically, this means that if the LED is running at constant current, the V_f goes down as the temperature of the device goes up. Consequently, the power into the LED also goes down with temperature and this, in turn, means that the light output also goes down. A simple calculation can be made to show what this means:

Assuming the 25 °C forward voltage is 3.6V and also assuming (to make the calculation easy) that the dV_f/dT is 3.6mV/°C. If the temperature goes up from 100 to 125°C, then the forward voltage drops by 3.6mV/°C × 100°C = 360mV, to 3.24 V.

This drop is 10% of the room temperature forward voltage and the power into the device is also reduced by 10%. Even if the efficacy of the LED was unaffected by power, which isn't true, the light output would be down 10% at temperature just because the input power is down. This means that, to ensure the light output is not affected by this forward voltage drop is to run the LED at constant power, meaning that it would be necessary to increase the current the keep this power constant. Again, easy calculations to measure this potential increase of current:

Assuming 700mA into the LED at room temperature, a power of 3.6V x 0.7A = 2.52W. At 125ºC the voltage would drop to 3.24V so the current would need to be increased to 2.52W/3.24V = 778mA, an increase of 11%.

Optical effects on LEDs involve all its components. To start, brightness and efficacy of an LED decrease with increasing temperature. Analysing a typical datasheet, it is possible to see a curve of relative brightness vs temperature at constant current (figure below):

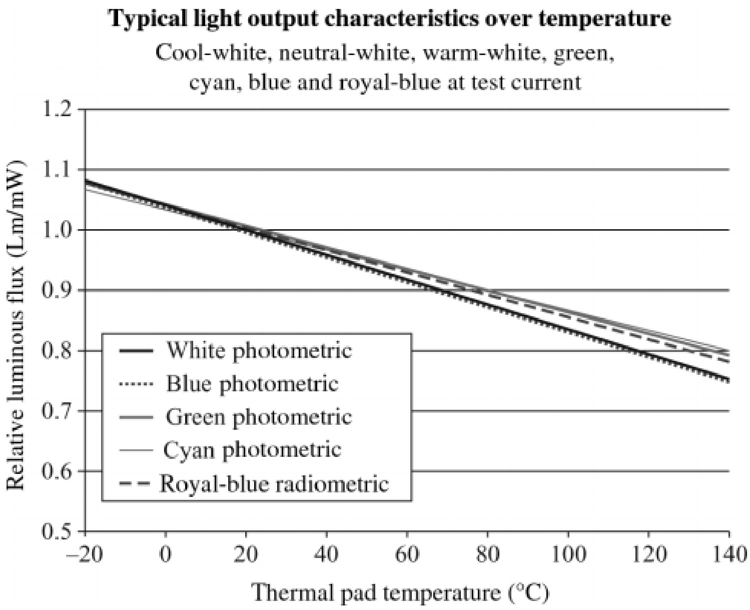

Constant current is the common way to drive LEDs, so this is a useful curve. But, in reality, it mixes two effects. With temperature, the efficacy decreases but the input power also decreases (because it comes with constant current). So, even though this curve shows the brightness decrease, it does not show the efficacy decrease.

So, to find the efficacy as a function of temperature, it is necessary to do some calculations again. The graph shows that from 25 to 125 °C at constant current, the light drops from 1.00 to 0.78. (it is important to note that this is flux, not efficacy. The "Lm/mW" on the y-axis means "lumens or mW," and it is because "royal blue" is measured in watts, and all the other colours are measured in lumens.)

Now, this is a 100 °C delta. The datasheet also shows that the dV_f/dT = 3 mV/°C, so the forward voltage drops from a typical value of 3.15 to 2.85V. This means that the input power drops from 3.15V × 350mA = 1.103W to 2.85V × 350mA = 0.998W. This is a drop to 90% of input power. Efficacy, on the other hand, has dropped from 1.00 to 0.78/0.90 = 0.86, or, in other words, efficacy has dropped 14%. The other 8% of the light drop is, again, due to the constant current drive circuit producing less power for the LED at high temperature.

3.3. LED lifetime: Lumen degradation

The previous sections have talked about temperature effects that can be quickly measured in a lab. Now it's time to talk about very long time thermal effects. As currently defined, an LED lamp is said to have a lifetime that is equal to the time required for half of the lamps to get to 70% of their initial light output. LED lifetime is defined the same way.

The first thing to observe is that this is difficult data to collect. Fifty thousand hours is about 8 years of continuous operation. With new generations of LEDs coming out every 6 months, there is no time to measure life before the part is obsolete. So, lifetime measurement has turned to extrapolation from short-time data. In a very general way, everything is expected to age on a logarithmic time scale. If a parameter decreases by 5% in 1000h, it is expected that in the next 9000h it will decrease a further 5%. Unfortunately, this

Arrhenius law applies to a single aging mechanism. In an LED, there are at least four independent aging mechanisms: the die, the phosphor, the encapsulant, and the package. For a given device, there are potentially at least four different logarithmic aging time constants. No one knows how to extrapolate the data yet, so this isn't a current problem, but someday it will have to be addressed.

4. Thermal Management of LEDs

In the previous topic, some of the main effects that temperature has on LEDs were discussed. In this topic the main discussion will be the environment in which LEDs operate and some methods to keep them cool.

4.1. Heat Sinks

When there is the need to reduce the resistance of the thermal conduction path, a heat sink is used, a piece of metal attached either directly or indirectly to the LEDs. Aavid Thermalloy's part number 569000B00000G is an example of what can be done. The big basket in the middle is to allow multiple LEDs to be mounted, the fins are to increase the amount of convective cooling, and the black anodization is to increase the thermal radiation. Together they achieve a thermal impedance of 5.5 °C/W. This will doubtless be the lowest thermal impedance in the system, and so, the design is easy. If the LEDs are dissipating 8W, the temperature rise will be 44°C.

On the other hand, in order for the heat sink to work as advertised, the LEDs have to make proper contact with it. And in this case, proper contact means proper thermal contact. If the LEDs are loosely suspended above the heat sink, of course it won't work properly. The LEDs need to be attached with something that will both hold them in place and have low thermal resistance to the heat sink.

4.2. Fans

Another thermal resistance that can be minimized is the convection. Convection in air is moderately effective at cooling, but it can be dramatically enhanced by a fan. Fans and blowers work by forcing air to move across the hot surface. Fans are typically rated by how much air they move, with units of liters/minute. The reality is that there are a lot of complex motions of the air that determine how much cooling is getting done. The size and shape of the fan, the size and shape of the object being cooled, and its orientation to the fan, as well as other objects in the path of the air flow, all contribute to the actual cooling.

5. DC Drive Circuitry

LEDs need to be electrically driven to emit light. In this topic the discussion will be how to design DC drive circuitry for LEDs. A typical DC source is a battery, as, for example, used in a flashlight. It could also be the output of a switching converter, for example, 12VDC. In any case, what distinguishes DC from AC for practical purposes is that DC typically has a much lower voltage than AC. This means that the regulations governing usage are much easier to comply with. There's no EMI to worry about, the voltages are generally "safe" and the input voltage is generally very steady. On the downside, the currents in a DC drive are higher, and the source impedance becomes an important factor in the design.

The relative values of the supply voltage and the LED voltage dictate what type of converter to use for a DC drive. There are 3 possible cases:

- The supply voltage is higher than the LED voltage;
- The supply voltage is lower than the LED voltage;
- The supply voltage is sometimes higher and sometimes lower than the LED voltage.

5.1. Buck

To cover the first case, where the supply voltage is higher than the LED voltage, a buck converter will be used.

A buck converter can only convert an input voltage to a lower output voltage. There is no way to get a higher voltage out of this topology than the voltage that gets in. This is exactly the way the first case works. As an example circuit, National Semiconductor LM3405 will be used. This part is convenient to use because the switch is integrated inside the IC, and the whole thing fits in a 6-pin SOT, a tiny package.

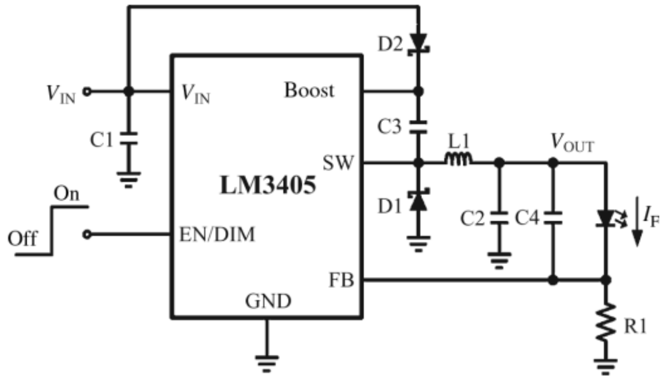

The input voltage can be as high as 20 V on this device, and so the IC is powered directly from the input on the V_{IN} pin. The transistor for the buck is inside the IC, between the V_{IN} and the SW (switch) pins. The LM3405 switches the transistor on and off at a constant frequency of about 1.6MHz. This connects and disconnects the SW pin to the input voltage. When the transistor is on, the voltage is positive across L1, and when the transistor is off, the voltage on SW swings down to a diode drop below ground, turning on D1. The current in L1 thus increases and decreases. The high switching frequency of the LM3405 means that the value of L1 can be very small, typically around 10µH.

The current in L1 is smoothed by C2 and is fed to the LED. C4 is basically in parallel with C2, but also has some function for stability of the circuit. The current is sensed by R1, which produces a voltage proportional to the current and is fed back to the FB (feedback) pin. The IC controls the duty cycle of its internal switch to produce a voltage of 205mV on the FB pin. This voltage is intentionally low in order to avoid significant power loss in R1 (1A x 205mV = 205mW, which can be handled by a ¼W resistor). This may possibly require an RC filter between R1 and the FB pin, to reduce noise fed back to the IC. Other ICs have even lower feedback voltages. You should generally avoid using a normal IC used for producing a fixed output voltage. For example, a 2.5V feedback at 1A would lose 2.5W in the resistor.

The LM3405 can be controlled on and off with a digital signal on the EN/DIM pin (1.8V on, 0.4V off). The final elements are D2 and C3. These are used as a charge pump to produce a voltage on the BOOST pin that is higher than the input voltage. There is a MOSFET inside the IC between VIN and SW. When the MOSFET is turned on, the voltage on SW needs to be as close as possible to the input voltage, as otherwise there is a voltage drop across the transistor, leading to power loss. But to turn on an N-channel MOSFET, the gate voltage has to be higher than the source voltage. The BOOST pin provides that higher voltage.

This circuit can be used, for example, in a flashlight design to convert the USB output of a computer to drive a single power LED.

5.2. Boost

In the second case, the input voltage is always lower than the output voltage. The topology to use is called a boost. As the name suggests, it takes the input voltage and boosts it up to a higher voltage. And just as the buck is incapable of producing a higher voltage output than its input, the boost is unable to produce an output that is (significantly) lower than its input. As an example circuit, the Fairchild FAN5333 will be used. This part is convenient to use because the switch is again integrated inside the IC, and the whole thing fits in a 5-pin SOT.

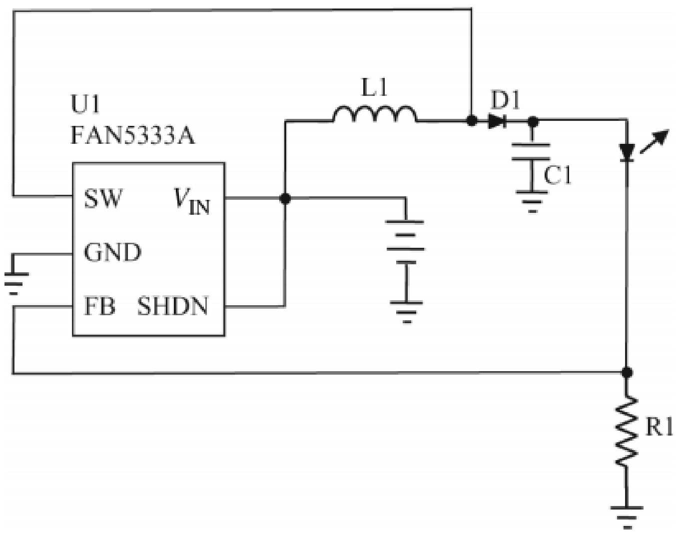

Analysing the figure helps to understand how a boost works. Once again, the input voltage goes to the V_{IN} pin, this time limited by the manufacturer to 6V. In this case, the transistor, diode, and inductor are still present, but are connected together differently. The input voltage is always attached to the inductor. The reason the inductor current doesn't continue to increase is because the voltage on its other side is higher than the input voltage (that's why the boost has to have higher output than input voltage). The transistor is operated by periodically pulling the SW pin to ground. When the transistor is on, SW is at zero volts, and so the inductor current increases. When the transistor is off, SW goes up to a diode drop above the output voltage, and so the inductor current decreases.

The purpose of the diode D1 is thus clear. When the switch is on, the SW pin is pulled to ground. If the diode wasn't there, that would pull the output to ground as well. The current from the inductor goes through D1 and into C1 and the load. The current is sensed by R1, producing a feedback signal at the FB pin. The FAN5333A controls the duty cycle of its internal transistor to produce a regulated voltage at the FB pin (110mV), corresponding to the desired current through the LED. When the switch is on, the inductor current is going to ground. To keep the LED on during this time, there has to be enough energy stored in capacitor C1 to ensure that the voltage across the LED doesn't change significantly.

In this specific schematic, the shutdown pin, SHDN, is connected high, so that the IC is always on. The only other aspect to note is this design is that no charge pump circuit is required. The internal MOSFET goes from SW to ground, so the input voltage is high enough to run its gate.

This circuit can be used, for example, to design a flashlight boosting the output of 2 D alkaline cells to drive a 3W LED.

5.3. Buck-Boost

The buck converter can only be used when the output voltage is less than the input voltage. Conversely, the boost converter can only be used when the output voltage is greater than the input. What should be used when the input is sometimes higher and at other times lower than the output? This situation arises because battery voltages are dependent on a number of factors, including their state of charge and temperature. So a fresh battery may have a voltage greater than the LED it is driving, but when it's been mostly discharged, its voltage may be lower.

The traditional circuit for this power supply is called a buck-boost. As the name suggests, it is a buck and a boost stuck together. This involves two transistors and two diodes, one set for each. Conceptually, it works by turning off the buck and running the boost when the input voltage is low, and by turning off the boost and running the buck when the input is high. For many LED drivers, however, this may be too expensive. Four external power devices eat up quite a bit of board space and money; using integrated switches costs even more.

This can be used, for example, to convert the output of a car battery to run an LED taillight.

6. AC Drive Circuitry

When there is the need to run an LED light off the electrical grid, AC drive circuitry is used.

AC power conversion is considerably more complicated than DC. Not only is it technically harder, there are also government regulations concerning EMI emissions.

6.1. Rectification

The first step in using AC power is rectification.

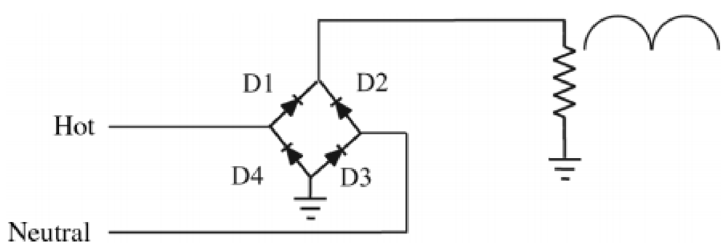

Rectification converts the sinusoidal AC line to DC with some ripple. When hot is positive with respect to neutral, current flows through diode D1, through the load (here a resistor), to ground, and then through diode D3 from ground back to the neutral. When it's negative, neutral is positive with respect to hot. Then the current flows through D2, through the load in the same direction, through D4, and then back to hot. Since current flows through the load in the same direction regardless of the polarity of the AC line, this is a DC supply. Of course, the actual voltage the load sees still runs from zero to peak each half-cycle.

6.2. EMI

One of the things that make design of AC ballasts much harder than DC is electromagnetic interference (EMI). Governments require that devices that attach to the AC line not produce more than a specific amount of electrical noise. And switch-mode power supplies, since they switch at high frequency, generate a lot of noise. There are two types of noise, and one of them has two varieties. The design must meet all of the noise requirements for all of these.

The first type of noise to consider is conducted noise. This is noise that comes out of the design and is conducted along the input power wires. It is typically caused by devices that turn on and off quickly and carry power, in particular the MOSFET and the power diode.

There are two types of conducted noise: normal mode and common mode.

Normal mode noise is due to different current in the hot and neutral power wires.

Common mode noise is due to the current in the hot and neutral wires with respect to the ground wire.

The second type of noise is radiated. This is typically caused by current loops, that is, current that runs in one direction and then returns, with a physical separation between the two paths.

6.3. Power Factor Correction

Many devices above 50W are now required to be power factor corrected (PFC). Additionally, new regulations seem to require PFC for lighting devices regardless of power level. So, the design may well need to be power factor corrected.

7. Example of Design

To make a complete design, all of the various components (the LED, the driver, the dissipater) need to be integrated. A printed circuit board for the LEDs and for the ballast needs to be done. A design for optics to get the light from the LEDs to go where it should, and to have the characteristics that it should needs to be done, and this design must survive in the customer’s environment.

The first step in creating a good PCB (printed circuit board) is to make a good schematic. A good schematic is more important than it seems, because a good schematic is easy to look at and understand. This is important for anyone else who has to look at the schematic, for example, for a consultant or for a design review.

Below there is an example of a good schematic.

The first reason for it to be a good schematic is in the IC pin-out. It shows the pins out of numerical order and with some pins on all four sides. Probably this is to minimize the visual complexity of the schematic. It also shows the IC pins in their actual order. This makes it obvious that the IC connections to the gate and source of the MOSFET are on the same side as VIN and GND. The pins thus have to either be routed around the IC (possibly a noise problem with the gate drive) or the IC could be rotated 180° or be mounted on the backside of the PCB. Conscientiously putting the pins in the right place on the schematic will make it obvious if connections need to cross over other connections, a frequent occurrence with ICs.

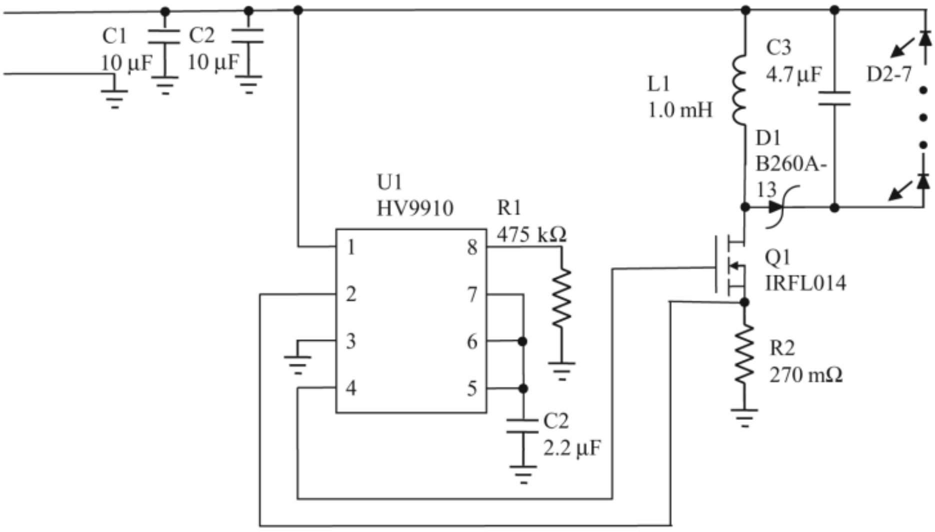

8. Bibliography

- Ron Lenk, Carol Lenk: Practical Lighting Design with LEDs;

- https://www.nfpa.org/-/media/Files/News-and-Research/Resources/Research-Foundation/foundation-proceedings/2012-SUPDET/13Bullough-et-al-paper.ashx?la=en&hash=16B7D062601F6EE5DF5F16A8EFB330F6F918D8A4

- https://electronics.stackexchange.com/questions/173283/what-destroys-a-led-in-the-reverse-direction

- https://www.sciencedirect.com/science/article/pii/S1359431111001359