

LED Flashlight White Paper

By Henry Schneiker

www.HdsSystems.com

September 17, 2004, updated August 2011

Copyright 2004 - 2011 - all rights reserved

A flashlight (headlamp, torch) designed around an LED can offer significant advantages over the same flashlight designed around an incandescent light bulb or other light emitting device. In this paper, I will cover the luminous, electrical, thermal, optical, mechanical and interface aspects of flashlight design as they pertain to LEDs as well as how the human visual system adapts to different lighting conditions. Although certain aspects of this paper are somewhat technical, even the non-technical person will be able to gain a better understanding of the issues and trade-offs involved in designing an LED flashlight.

By necessity, this paper talks about specific devices. However, I do not mention any manufacturers or part numbers because that information is not important to the points being discussed. As with any project, you should fit your design to the actual parts being used - taking advantage of their capabilities and being careful of their limitations.

The LED

LED stands for Light Emitting Diode. A diode is a device that passes current in the forward direction while blocking it in the reverse direction. When you pass current through an LED in the forward direction, the LED emits light. The color of the emitted light depends on the material and construction of the diode. The amount of light varies roughly linearly with the amount of current passing through the LED, which provides us a convenient way to efficiently change the light output.

LEDs are solid state devices and therefore quite reliable. Mean time to failure is often listed at over 100,000 hours of operation under specified conditions. LEDs are far more reliable than the ubiquitous incandescent light that you normally find in a flashlight. LEDs can also tolerate much higher shock and vibration loads than other common light sources including incandescent lights, florescent lights and other arc lights.

But as good as LEDs are, they not perfect.

The most significant issue from a production point of view is the variability of LEDs.

When LEDs are manufactured, considerable effort is made to get all the LEDs to come out the same. However, even on the same wafer, there is tremendous variation from LED to LED. The variations are so significant that the manufacture must test all the LEDs and sort them into bins. Common bin categories include color tint, light output (flux) and forward voltage.

So to get consistent production results all you have to do is specify the bin you want when you order the LEDs, right? Wrong. To illustrate why this is so, let's take a closer look at what the bin ranks actually mean.

If you take a color bin for a white LED and convert the CIE xy coordinates of the bin corners to CIE L*a*b* values - where each whole number is approximately one distinguishable shade of color or brightness - even a small bin can have 13 distinguishable shades of color diagonally across the bin. A larger bin can easily have 30 or more shades of distinguishable colors diagonally across the bin. This means that if you take two LEDs from opposite diagonal corners of the same bin and view them side by side, you will see that the white light being emitted from each LED is of a different shade of white. Further, people often find some shades of white more pleasing than other shades, with shades closest to the black-body temperature curve being the most pleasing. And people also have preferences as to color temperature - some preferring cooler shades (more blue) and some people preferring warmer shades (more red).

A similar situation results when you look at the light output at a specific current (the flux) and the forward voltage. A typical bin can cover a flux range of n to $1.3n$ and a forward voltage range of v to $1.1v$. When combined in the worst case scenario, this can produce a 40% difference in overall efficiency within the same bin. This difference can be seen easily if you compare the two lights side by side. Now add to this the measurement error tolerance of 10% on both parameters and you get a 51% difference in possible overall efficiency within the same bin. Talk about a quality control nightmare.

If a customer compares two lights and sees an easily distinguishable difference between them, the customer will generally assume that one of the two lights is defective. Educating the customer to the realities of production variations only goes so far. Significant variations will often result in customer dissatisfaction and returns.

There is, of course, a statistical side of this. The actual distribution within the bin is statistical in nature and the probability of any two lights exhibiting a sufficient difference that the customer will notice or be concerned about the difference is also statistical.

The unit-to-unit variation is much higher in LEDs than in most other light sources. Some of these differences can be reduced through a calibration process. However, it might be easier if you have multiple product lines and carefully sort the LEDs for each model - thus reducing the intra-model variations.

LEDs have another problem that is worth noting. LEDs become less efficient as you increase current and temperature. And since LEDs have a negative temperature coefficient, the hotter they get, the more current they want to draw. This makes them unstable and prone to thermal run-away unless careful attention is paid to how they are driven. Further, high temperatures will increase the rate of permanent lumen (light output) loss beyond the typical 8% per 1000 hours.

When I talk about high power LEDs, I am referring to LEDs that are designed to handle one or more watts of continuous power per LED. However, the same principles can often be applied to the low power LEDs or arrays of low power LEDs. The high power LEDs typically provide a fairly low thermal impedance which can be coupled to an effective heat sink to minimize the LED's junction temperature. Every effort should be made to keep the LED cool for maximum efficiency. If the LEDs will be driven at high power levels under conditions where effective heat dissipation is not possible, the LEDs must be thermally protected to ensure reliable operation.

And finally, when you dim an LED by reducing the current, the LED can and often does undergo a significant color shift. The severity of the color shift is dependent on the individual LED. This effect is most noticeable to a human observer with white LEDs. What happens is that as the current drops, the dominant frequency also drops. The amount of frequency shift is quite variable from one LED to the next and a manufacturer seldom tests for it. With white LEDs, the dominant frequency (blue) is used to excite a phosphor that re-emits at the lower end of the spectrum to provide yellow, orange and red. As the frequency shifts down, the blue dominant frequency shifts toward green making less of the required blue frequency available for down-conversion to yellow, orange and red and so the color shifts toward green. The only way to reliably prevent this is to always drive the LED at recommended power levels while dimming.

The Power Supply

The purpose of the power supply is to supply and regulate power to the LED and ensure a stable operating environment. The power supply can vary in sophistication from a voltage source in series with a resistor - such as a battery with or without an

additional resistor - to a fully regulated constant power system with thermal regulation. The characteristics of the power supply directly affects cost and the stability of the light output.

LEDs have a steep V_f/I_f curve. This means that a slight change in forward voltage (V_f) will have a relatively large affect on the forward current (I_f). Further, this curve is different for each LED. To drive this point home, if you assume you have LEDs that all come from the same bin and the bin provides a V_f tolerance of 10%, the power consumed by LEDs in that bin can vary from w to $2w$ if powered with a constant voltage.

As I mentioned earlier, LEDs have a negative thermal coefficient. This means that V_f decreases with increasing temperature. With a constant voltage power supply, increasing the temperature slightly decreases the V_f and increases the current by a relatively large amount - causing a net increase in power and a further increase in temperature. For any given thermal path, there is a power level at which the heat transfer can no long keep up the the increasing power and falling V_f . Thermal runaway is the result and it will destroy the LED if nothing limits the power to safe levels.

Regulating the current or power through an LED effectively prevents thermal runaway. Regulating power through an LED has the added benefit that it tends to provide better brightness regulation over a wide temperature range. A good regulating circuit will keep the flashlight brightness the same throughout the life of the battery. Both systems can include thermal regulation to prevent dangerous heat build-up if the thermal path is inadequate at a high power setting.

A direct drive flashlight - i.e., batteries in series with an LED, with or without an additional resistor - cannot provide a constant output brightness. The brightness/time graph of this class of flashlight produces a smoothly descending curve. This is because the battery voltage sags quickly and thus the power through the LED drops rapidly. Suffice it to say that in the first quarter of the battery life, the brightness drops significantly. Battery life in this case is often defined as the amount of time it takes for the brightness to drop to 25% of the original value with fresh batteries. Another problem with direct drive flashlights is that power through the LED can be at dangerously high levels until the internal resistance of the battery increases sufficiently to limit the power to safer levels. As a result, direct drive flashlights are usually less reliable due to a higher LED failure rate.

The efficiency of a system is controlled by the efficiency of all the parts. In a typical LED flashlight, you have a battery, a power supply (or regulator) and the LED. The

efficiency of the battery drops as the power increases. Most power supplies and LEDs have an optimum point where they are most efficient and are less efficient at higher or lower power settings. This optimum point is often around 10 to 20% of rated power. In general, the increase in battery efficiency at low power settings more than offsets any decrease in power supply and LED efficiency at low power settings. The net result is that there is little or no net decrease in total efficiency at typical low power settings.

At the higher power settings, a compounding affect can be observed that decreases total efficiency in a non-linear fashion - often faster than N squared. The primary cause of this rapid drop in efficiency is the $I^2 R$ (current squared times the resistance) losses found throughout the circuit - especially where high currents are found. The next biggest loss in a switching power supply is usually coil saturation. And finally, as the LED junction temperature and current density goes up, the LED efficiency drops.

Batteries have a large affect on a flashlight's design because they supply the power. Alkaline batteries develop a high internal resistance as they are used up. Therefore, alkaline batteries are not well suited to high current applications because large amounts of their power will be lost due to the internal resistance - dissipated as heat. Further, alkaline batteries stop producing useful amounts of power by the time the temperature drops below freezing, making them inappropriate for cold environments. Lithium batteries (in many different chemistries) have a fairly low internal resistance throughout their life and work at temperatures well below freezing. The economics of lithium batteries has changed to the point that they are now similar in cost to alkaline batteries, making them the first choice in battery power for most flashlight applications.

You must be careful when using rechargeable batteries not to over-discharge them. Over-discharging rechargeable batteries can damage or destroy them. When a regulating switching power supply is used, the battery can be over-discharged before the light output drops by a noticeable amount. A flashlight should detect rechargeable batteries and prevent them from being over-discharged. Further, lithium rechargeable batteries often contain a special protection circuit that turns off the battery to prevent accidental over-discharge - causing instant darkness. Once the protection circuit turns off the battery, only recharging the battery will turn the battery back on. A flashlight should detect a lithium rechargeable battery and avoid the conditions what will activate the protection circuit and suddenly leave you in the dark.

As your batteries nears the end of their life, a flashlight should signal you and reduce the brightness gracefully. You should be notified that the batteries are nearing the end of their life so you can find a safe place to change your batteries. If you cannot change the batteries, the light should continue to provide light for as long as possible - long

enough for you to return to safety - and fresh batteries. Even if that takes you a day or two. Once you know the batteries are at the end of their life, you should be able to manually reduce the brightness further to make the remaining battery power last even longer.

The Optics

The LED chip emits light in all directions. An optical system is used to gather the light and shape it into a beam. The LED itself has a built-in lens. This lens is limited in its ability to gather and shape the light into an acceptable beam and the LED manufacturers only offer a limited number of lens options. An optical system consisting of just the built-in lens in the LED may produce acceptable results for inexpensive flashlights, but for higher quality results, a more advanced optical system is required.

A parabolic reflector can theoretically generate a beam of perfectly parallel light rays the same size as the exit aperture (the large open end of the parabola) given a point source of light at the focal point. Some percentage of the light from the point source will not be incorporated into the beam because it never contacts the reflector surface. The solid angle that intersects the parabola relative to total spherical solid angle is used to define the reflector capture efficiency - the smaller the exit aperture solid angle, the higher the capture efficiency and the brighter the beam - all other things being equal.

In a real reflector you have imperfections that tend to scatter the light beam. First, your light source is never a point source - it is always a complex 3-dimensional surface. Second, the reflector's surface is never perfectly smooth or exactly the correct shape. This leads to angular errors that tend to spread the beam. By designing in and controlling the angular errors you can shape the light distribution (beam pattern) to suit your purposes. Reflectors tend to provide the highest quality beams and the highest overall efficiency.

A transmissive lens, such as a common magnifying glass, uses refraction to achieve a similar result. Instead of light being reflected off a surface, light is bent at the air/lens interface. Transmissive lenses have three main problems. First, they have a relatively low capture efficiency - any light that falls outside of the aperture is lost. Second, as the angle between the lens surface and the light source deviates from perpendicular, an increasing amount of light is reflected off the surface and lost instead of being bent into the lens. Third, the lens projects an image of the emitting surface - as a result, the focal point is normally chosen to be different from the emitting surface so the image is blurred.

There are also compound reflector-refractor such as Total Internal Reflection (TIR) optics that combine properties of both systems. These optics often use the high reflectivity of the air/optic interface to form the reflector surface. These optics tend to be more efficient and have better beam quality than a pure transmissive lens but are not as good as pure reflectors. However, these optics have the advantage that they tend to be more compact than reflectors.

An interesting property of lenses is that they are all about angles. Since the same angular error will produce the same result regardless of size, the only limits to small size are the emitter size and surface tolerances. As a result, a smaller emitting surface can use a smaller optical system and produce the same beam pattern. Or, you can increase the size of the optical system to reduce the angular errors and improve the beam pattern definition. The variations are endless.

The Human Visual System

If you look at typical flashlight marketing literature, you might think that brighter is always better. But allow me show you why always brighter is not always better.

Your eyes can adapt to a wide range of light intensities. At noon on a summer day the sun can illuminate a surface to 120,000 lux while an overhead full moon might illuminate the same surface to only 0.1 lux. Your eyes use three main methods when adapting to changing light levels. Under bright conditions, your eyes exclude excess light by closing the pupils. As the light level falls below the point where the pupils are fully open, the eyes undergo a chemical change to increase the cone (color) cell sensitivity. Finally, as the light level falls below the capability of the cone cells, the rods (no color) cells take over and become maximally sensitive. It takes roughly 20 minutes after you leave a brightly illuminated area to become fully dark adapted.

So how much light do you actually need to perform a task? 1000 to 2000 lux is recommended for fine detail work where maximum visual acuity and color recognition is required. 100 to 200 lux is appropriate for most office work. By 10 lux, colors are less vibrant while visual acuity is good. By 1 lux, unsaturated colors cannot be distinguished and visual acuity is only acceptable. By 0.1 lux, no colors are visible and visual acuity is poor. By 0.01 lux, visual acuity is lousy and you can see objects better if you don't look directly at them.

The forgoing illustrates that providing 50 to 200 lux of illumination is plenty for most tasks that require good color recognition and good visual acuity. Much less light is needed for tasks that only require object shape identification.

It takes one lumen - that's one candela per steradian - to illuminate one square meter (about 10 square feet) with 1 lux. Note that luminous flux - in lumens - is mapped against the color sensitivity curve of the human eye. However lux can also be defined in terms of radiometric flux at a single frequency and is equivalent to 1.46mW of radiant electromagnetic power per square meter at 555nm (the color to which you are most sensitive - a shade of green).

The amount of light illuminating an object depends on how bright the light is and the distance the light is from the object. If you know the amount of light illuminating an object at one distance, you can calculate what the amount of light illuminating an object will be at another distance using the inverse square law. For example, if you move the light twice the distance from an object, the illumination drops to one quarter (2 squared is 4 and the inverse of 4 is 0.25 or one quarter). As another example, if you have a light that generates 500 lux at 1 meter (just over 3 feet), it will generate 100 lux at 2.2 meters (about 7 feet) or 10 lux at 7.1 meters (about 23 feet) or 1 lux at 22.4 meters (about 73 feet). The drop in illumination corresponds to an increase in the illuminated surface area. Put another way, you need 4 times the light to see twice as far with the same illumination level.

The eye responds to light in a logarithmic manner. That is, for the eye to indicate a significant increase in brightness - i.e., an easily recognizable increase in brightness - the amount of light must double. For dealing with the large variations of light in your environment, this is a wonderful thing. For flashlights trying to get ever brighter, this is a terrible thing because it takes over twice as much power to generate twice as much light while the battery life drops to less than half - due to the escalating inefficiencies discussed earlier - for a single incremental increase in brightness. However, if you want to maximize battery life by using lower brightness settings this is a great thing. Why? For every incremental decrease in brightness, you more than double the battery life and after a few minutes, your eyes adapt to the new brightness level so the apparent drop in brightness actually decreases. By reducing the amount of light to match your task, you can dramatically improve the battery life.

The eye's logarithmic response to light suggests that brightness levels should be spaced in a logarithmic way. That is, the steps between brightness settings should be a constant factor - such as 2x or 1.4x - instead of using a more linear approach. This provides the appearance of the steps being equally spaced. This same method is used in the audio industry for volume controls. Because of the eye's adaptive characteristics, a continuous control is inappropriate for a flashlight. The reason is that you are quick to turn up the light but slow to turn it back down. In fact, if you are asked to turn up the

light for a minute and then turn it back down to where it started, the resulting brightness is usually brighter than before starting the exercise. Human factors should never be forgotten when designing control systems. Thus, a stepped approach is the preferred method.

Taking all of the above together suggests that we can use lower levels of light for many tasks and gain benefits in the process. Not only do you increase battery life with lower output levels, you will be able to see further when you really need to and you reduce your visibility footprint. Here is the scenario. Imagine you are using a low light level to navigate - such as following a trail in the mountains at night - which you can easily do on less than 1 lumen. Also imagine you can get maximum output with a simple press of the button. After hiking for a while you come to a place where you need to look far ahead. You hold your light above your head, point your light to where you want to see and then press the button for maximum light output. Because your eyes are dark adapted, you will be able to see much further with the maximum amount of light than if you had been following the trail with a brighter light. Depending on circumstances, you may be able to see twice as far on the same amount of light. That is a huge increase in distance without buying a bigger flashlight.

For the tactical crowd, select 0.1 lumens to sneak up on your enemy. The light footprint is so low that they will not see you coming. Upon arrival, a quick press of the button not only blinds your enemy, but it gives you the tactical advantage because the enemy is also brightly lit. Releasing the button allows you to disappear back into the darkness with enough light to move undetected.

Putting all the light in a narrow beam is great for looking at an object at a great distance. However, the contrast from the bright beam center to the dark area just outside the beam is so great that you will have difficulty seeing objects just outside the beam - to the extent that trying to walk using such a beam can be difficult or dangerous. A general purpose light should provide enough light outside the main beam so your eyes can move comfortably between the beam center and the outer areas and take in a large viewing area. Better still is to provide a smooth transition between the two areas. The broader the transition zone, the lower the contrast and the safer you will be when using the light to navigate rough terrain. You will find that you can use less total light with a smooth broad beam. In fact, you might want to carry two lights - one with a broad beam for walking and general use and another with a narrow beam for spotting distant objects.

User Interface

As you may have gathered from the previous discussion, there is a lot to be gained by providing a flashlight with more than an on/off switch. Adding a brightness control extends the utility. Having the flashlight automatically detect and compensate for different battery systems allows operational flexibility. And there are always a few extra features that would be nice to have, such as automatic emergency signaling, a way to find the flashlight in the dark or a way to prevent accidental turn-on.

As soon as you start adding features to the flashlight, you increase the complexity of the user interface - not to mention the design and cost. Let's take adding a brightness control as an example. Will there be a single button (click codes), two buttons (brighter / dimmer), n buttons (one for each brightness), linear slider or a rotary switch (volume control)? Is the brightness control integrated with the on/off switch or is it separate. Do you have to worry about making the controls water proof or making them work after being abused? How many brightness levels will there be? Will the user have access to all brightness levels all the time? Will the brightness control be continuous or a step function? How will brightness levels be spaced. Are there other features that will need controls and how will they interact with the brightness control?

And how will the flashlight's electronics detect the user's input? On/off switch? Change in resistance? Change in capacitance? Change in inductance? Mechanical coupling? Optical coupling? Magnetic coupling? Digital or analog? The possibilities seem endless.

While you are thinking about all this, you should never forget that the main purpose of a flashlight is to turn on and generate light. So this function should remain as simple to use as any other flashlight. Features that are used the most often should be the simplest to use - a user should never have to scroll through a list of brightness settings to get to his favorite brightness setting.

A microcomputer can be used to help implement the user interface, whether it involves a single control or multiple controls. The microcomputer is also handy in implementing automatic functions and performing other tasks. Having the intelligence in software instead of in hardware provides us with the opportunity to tailor various characteristics and provide automatic controls that are not easily done in hardware. Further, it allows us to add new features with relative ease.

Summary

Designing a flashlight around an LED allows us to design a very rugged and sophisticated light. A typical laundry list of desired features includes small size, high maximum brightness, multiple brightness presets, rotary brightness control, constant brightness regulation, thermal regulation, a smoothly tapered beam pattern, support for multiple battery chemistries and rechargeable batteries, long battery life, battery end-of-life warning, simple controls, automatic emergency signal, switch lock-out, automatic turn off after some time, a find-me-in-the-dark capability and the ability to customize these things. These features can now be found in a single compact flashlight.

HDS Systems designs advanced flashlights and incorporates the superior capabilities discussed in this paper. The latest flashlights can be seen at www.HDSsystems.com .