BUCK MODE SWITCHING LED DRIVER

Introduction

LEDs are funny little devices. They look like light bulbs, but act like diodes. For a basic introduction to LEDs, see the excellent explanation at <insert link here>.

For the purpose of this discussion, the important point is that all LEDs require some sort of current regulation. You can't just wire one directly to a battery like you can with an incandescent bulb. This is because LEDs are diodes; when you apply a forward voltage, nothing happens until the voltage is raised to a magic value that is called, not surprisingly, the *forward voltage* of the LED. At $V_{forward\ voltage}$ the LED conducts electricity and emits light, but the conduction is like water going over a waterfall: it goes faster and faster until it hits the bottom, or in the case of the LED, the current increases more and more until the LED burns out. Some form of current regulation is needed to avoid disaster and let the LED live a long and prosperous life.

The simplest form of current regulation is just a current limiting resistor, as show in Fig 1.

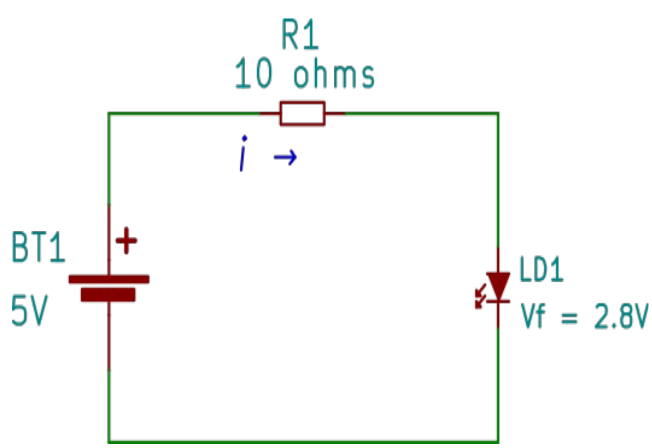

Fig 1

We can find the current in this simple example by applying Ohm's Law:

V = I*R, or I = V/R.

V is the 5 volts from the battery minus the forward voltage of the LED:

V = 5V – 2.8V = 2.2V.

Plugging these into $I = V/R$ we get:

$I = 2.2V/10\Omega$ or 0.22 Amps. This could also be written as 220 milliamps or 220mA.

Our 10 ohm resistor has limited the current to 220mA. We can get a larger current by using a smaller value of R, or a smaller current by using a larger value of R.

It seems that we have our current limiting problem solved, but let's look at the disadvantage of using a simple current limiting resistor. In Fig 1, the 220 mA is flowing through both the LED and R. This means that part of the power is used to light the LED, but part of the power is also used to heat the resistor. Since we wanted a light and not a heater, the power dissipated in the resistor is wasted. But how much power is wasted? We can figure it out.

Power, P, is current times voltage. In the case of a known resistance R and current I, the equation for power is: $P = I^2 * R$, or Power = Current Squared multiplied by the resistance. So the power dissipated by the resistor R is:

$P = 220mA^2 * 10\Omega = 0.48W$

So the wasted power is just a tad under a half Watt. This means we will need at least a ½ W resistor unless we want a puff of smoke.

The power at the LED is:

$P = V * I = 2.8V * 220mA = 0.62W$

This is workable, although 44% of the power is wasted at the resistor. For small LEDs driven by a fixed voltage, we might not care. Current limiting resistors are commonly used for indicator LEDs and other low power applications.

But what happens if the voltage changes? Let's try changing the battery to 12V and see what happens. Our new current calculation:

$I = (12V – 2.8V)/10\Omega = 920mA$

If our LED is designed to run on 220mA, we have now fried our LED. It looks like we need a bigger resistor. We can calculate the resistance needed to get back to 220mA:

$R = V/I = (12V – 2.8V)/220mA = 41.82$ ohms. Rounded to whole numbers, 42 ohms.

So a 42 ohm resistor will let us power this same LED from 12 volts at the same brightness. The LED has the same voltage and current, so it is still dissipating the same amount of power and, as you would expect, at the same brightness, but how about the resistor power? Let's calculate it:

$P = I^2 * R = 220ma^2 * 42\Omega = 2.03W$.

Our resistor is now dissipating more than 2 Watts of power. That means that for the same 0.62W of LED power, our circuit is consuming 2.65W of power. Our resistor is dissipating 77% of the power. A big, hot resistor is a major problem with current limiting resistors for high power LEDs. The second problem is that the resistor must change value with varying input voltages.

We could solve the varying input voltage problem by designing a transistor circuit that varied its effective resistance with changes to the input voltage, but we still would have the problem of dissipating waste power as heat. Getting rid of the heat is a problem, and the waste of power is a problem for energy conscious applications such as battery operated illumination. Enter Switching LED Drivers, which help solve both of these problems.

Switching LED Drivers

Let us look at our LED driver problem with a plumbing analogy: We have a high pressure source of water, let's say a garden hose, and we want to trickle a small steady stream of water into the garden to irrigate a plant. We crack open the valve just a tiny bit and let a little trickle of water run. The valve, in this example, is acting like the resistor in our LED driver; it takes the higher voltage, or pressure, and reduces it by dissipating power. The high pressure water, squeezing though the valve, converts much of its energy to heat.

Most of us, given this example, wouldn't care that a small amount of energy is wasted in the valve. Essentially, this is the low power LED example, where a current limiting resistor works just fine. But let's say that we own a plant nursery, and this watering system will be used for thousands of plants at a time. We need to find a way to more efficiently use the water pressure.

The problem comes from running the valve in a nearly closed position, restricting the flow of water. If we open the valve completely, we eliminate the wasted power, but then we have too much water. What if we pulsed it on and off quickly? Fully open and fully closed, the valve wastes no power. If we keep the valve closed most of the time, and only pulse it open briefly, the average flow of water can be reduced to a trickle, all without wasting power.

We try pulsing the water on and off and end up blowing dirt everywhere. The *average* flow of water is a trickle, but the pulses are like water cannons, blasting away each time the valve is opened. We need some method of smoothing out the water flow; some sort of accumulator. Adding a tank, as in Figure 2, does the trick. The valve is pulsed open as needed to keep the water in the tank at a more or less constant level. From there, gravity drains the water in a steady trickle, just as we wanted. The tank smooths the pulses into a slow, steady flow. We see this technique employed in systems all around us: pressure tanks , flywheels, and mufflers, just to name a few.

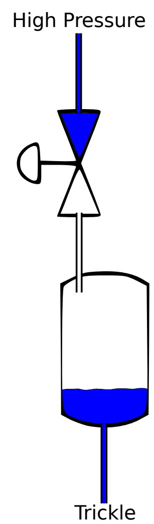

Fig 2

Now let's take our plumbing example and see how it might look in an electrical circuit. In Figure 3 we have replaced resistor R1 with inductor L1, and we have added a switch to

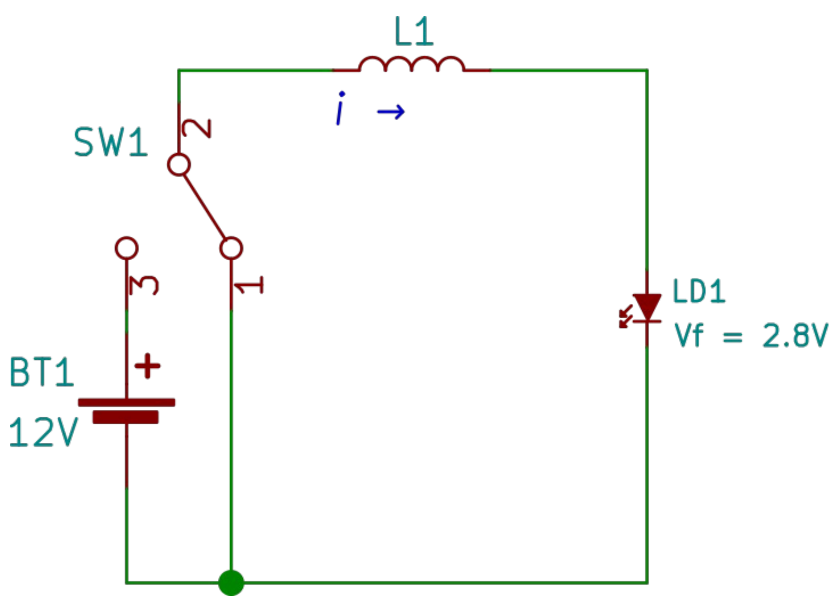

Fig 3

bypass the battery. While resistors simply resist the flow of electricity, inductors are tiny electromagnets. Inductors are made from a coil of wire wound around a ferris core.

When a voltage is applied to an inductor, a magnetic field starts to form inside the inductor. This magnetic field stores energy. When the voltage is removed, the magnetic field collapses, returning the stored energy as a voltage and current. The inductor works like an electrical flywheel, storing and releasing energy. With a flywheel, energy is converted to rotation; with the inductor, energy is turned into a magnetic field. The important point is that the inductor stores and releases energy, which we will be able to use to our benefit.

Looking at Figure 3, at the start of our experiment switch SW1 is thrown to connect the battery (contact 2 to 3). This causes 12V to be applied to the series circuit of L1 and LD1. You might expect the LED to light, but it will not immediately turn on. The magnetic field building inside the inductor impedes the flow of electrical current. Initially, all current is blocked by the inductor, so the full 12V appears across the inductor, leaving nothing for the LED. As the magnetic field grows, the inductor's opposition to current decreases until eventually the LED switches on. The inductor's opposition to current drops until it is essentially zero, which is bad news for the LED. With nothing left to regulate the flow of current, the LED will fry.

To make this work without killing the LED, we need to turn off the battery before the current reaches dangerous levels. Let's say we have some method of monitoring the current, which we want to keep around 220mA, as we did with the resistor regulated circuit. We throw the switch, connecting the 12V battery to the circuit. The inductor's magnetic field starts to build, and the LED comes on. We watch our current meter until the current hits 220mA. At this point we throw the switch the other way (contact 2 to 1) to bypass the battery. What happens then? It might seem that the LED would turn off, but remember that magnetic energy stored in the inductor? That magnetic field will collapse, releasing energy. Just as a flywheel will continue to spin at the same speed, the current through the inductor will continue at 220mA; the magnetic field is driving the LED. The switch bypassing the battery is essential to complete the circuit; recall that electricity only flows in complete circuits, going around and around like a race track.

Energy has to come from somewhere, so the inductor current powering the LED is slowly draining energy from the magnetic field. As the field reduces, so does the current. We watch the current drop until it gets to some minimum acceptable level; let's say 200mA. At 200mA, we throw the switch, once again connecting the battery. Current from the battery continues to light the LED, but it also rebuilds the magnetic field in the inductor. If we keep throwing the switch back and forth at the high and low limits, the LED will light continuously near our target current level (from 200mA to 220mA in this example). But here is the clever part: unlike the resistor in the first circuit, the inductor in this circuit does not burn up energy as heat. The inductor stores and releases energy without loss (ideally – more about losses later). We have found a much more efficient way to power our LED from a 12V battery.

How fast will we have to throw the switch? That depends on the size of the inductor and the power consumed by the LED. For a typical small LED lighting project, the switching speed will be somewhere between 50 thousand to 1 million times per second; we are going to need really fast hands and a very durable switch.

Fortunately, we don't need to do this manually. Transistors are excellent switches, and they can operate at very high speeds.

Figure 4 shows a simplified diagram of a transistor switching LED driver.

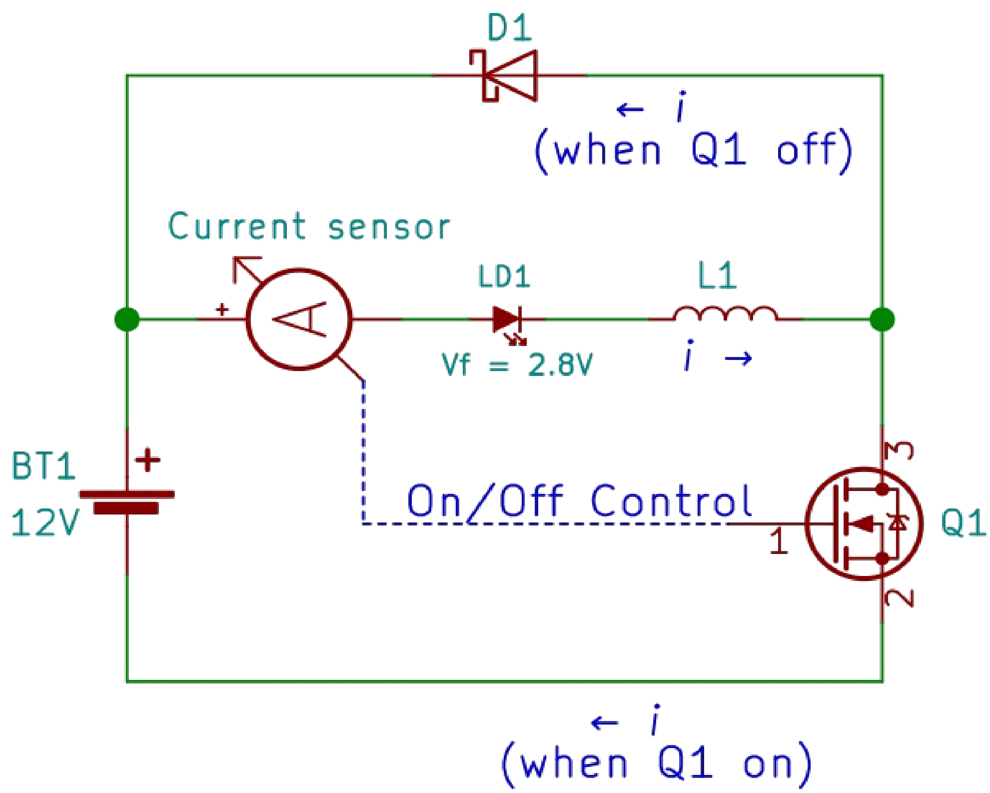

Fig 4

Transistor Q1 replaces the switch from our previous example. When the gate, pin 1, of transistor Q1 is turned on, the transistor becomes a closed switch, allowing current to flow from pin 3 to pin 2. When the gate is turned off, the transistor acts like an open switch, blocking the flow of current.

The current sensor monitors current *i* flowing through the LED and inductor. If this current is below our low limit, 200mA, the current sensor turns on transistor Q1. When the current reaches our high limit, 220mA, the current sensor turns off transistor Q1.

At first power-up, there is no current flowing. Our current sensor recognizes that zero amps is less than our lower limit of 200mA, so the current sensor switches on transistor Q1. This applies 12V across the series circuit of LD1 and L1. The magnetic field inside of inductor L1 starts to rise, current *i* increases, and LED LD1 lights. Current *i* continues to increases until reaching the high limit of 220mA. At this point, the current sensor switches off transistor Q1.

With Q1 switched off, the 220mA of current though inductor L1 is maintained by the collapsing magnetic field in L1, but it can no longer flow through Q1 since the transistor is switched off. This is the purpose of diode D1, which provides an alternate path to complete the circuit. Diodes act like one-way valves, allowing current to flow in only one direction. When transistor Q1 is switched on, 12V is applied to the diode in the reverse direction; no current can flow through the diode. When Q1 switches off, diode D1 is no longer subjected to the reverse 12V. Current *i* can no longer flow through Q1, but it can now freely flow through diode D1. With the 12V battery removed from the circuit, the LED is coasting on the power stored in inductor L1. Because of this coasting analogy, a diode used in this fashion is often called a *freewheel diode;* the current freewheels through the diode.

As the current freewheels it powers the LED, robbing energy from the inductor. This causes the current level to drop. When it reaches the low limit of 200mA, the current sensor once again switches on transistor Q1 and the cycle repeats.

This cycle is continuous, alternating between adding power and coasting. The powered portion of the cycle finds the current flowing in a clockwise loop through the lower half of the circuit diagram. In the coasting portion, current flows in a counter-clockwise direction through the upper half of the circuit.

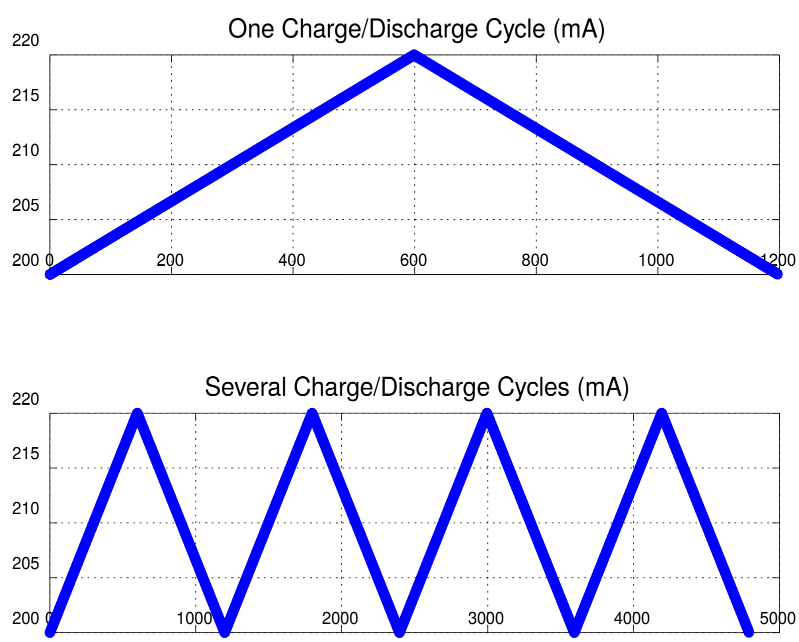

Fig 5

The current flowing through inductor L1 and LED LD1 is graphed in Figure 5. With transistor Q1 switched on, the current ramps up from 200mA to 220mA. When Q1 switches off, the current coasts down from 220mA to 200mA. This cycle continues indefinitely, for the entire time that the LED is lit.

This circuit is a type of oscillator with several interesting characteristics. While it oscillates, the oscillation is a small current ripple impressed on a much larger steady current. Recall that our original intention was to have a constant current driver for our LED. The steady current is that constant current. The ripple is an unwanted but mostly harmless side effect of this sort of control system. The ripple causes a slight variation in the brightness of the LED, but the ripple frequency is far too fast to be seen by the human eye.

The frequency of oscillation is natural function of the circuit; in other words, there is no clock driving the system, but instead it oscillates on its own. Oscillation frequency is largely a function of the inductance of inductor L1 and the voltage difference between the supply voltage and the forward voltage of the LED, which in this example is 12V – 2.8V, or 9.2V.

The voltage difference, 9.2V, provides the potential to drive current *i* though inductor L1 and LED LD1. The larger the voltage difference, the faster the current rises to the high limit, and the faster the circuit oscillates. A larger voltage difference increases the oscillation frequency; a lower voltage difference lowers the oscillation frequency.

The size, or inductance, of inductor L1 determines how much energy it can store, and consequentially, how fast the magnetic field will build for a give voltage. A larger inductor takes longer to charge up magnetically, so the larger the inductor the lower the oscillation frequency. Using a larger inductor will lower the frequency; using a smaller inductor will raise the frequency.

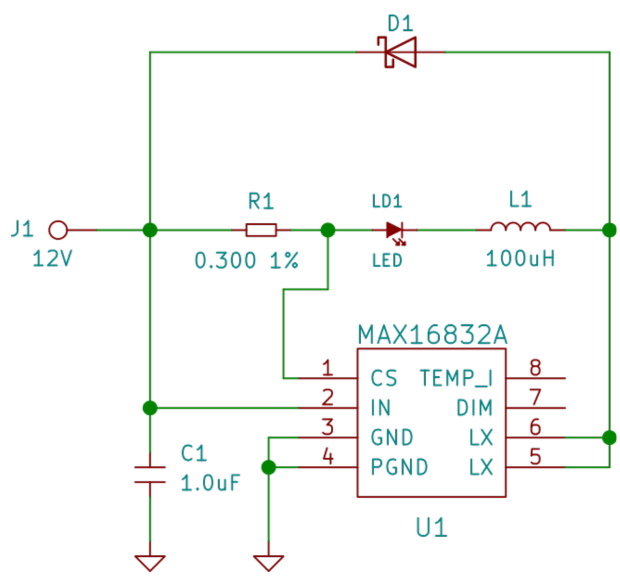

Fig 6

A Practical Example

In practice, the current sensor in our example is generally a specialized Integrated Circuit (IC). For smaller drivers, this IC may also contain transistor Q1. Figure 6 shows an example using a MAX16832A buck-mode switching LED driver IC.

A quick inspection reveals that the circuit in Figure 6 is functionally identical to that in Figure 4. The current sensor has been replaced by resistor R1; transistor Q1 is now incorporated into integrated circuit U1; and a capacitor, C1, has been added. LED LD1, inductor L1, and diode D1 are identical, and the general layout remains the same.

The current sensor is now replaced by R1 and circuitry inside of U1. The LED current, which we will still call *i*, flows through R1. This current produces a voltage across R1 which follows Ohm's Law:

$V_{R1} = i * R1$

In other words, the voltage across R1 is proportional to the current driving the LED. This type of voltage is often called a *current sense* voltage, since it is used to measure, or sense, current. The current sense voltage is measured across pins 1 and 2 of U1. When the current sense voltage drops below the low limit, U1 switches on its internal transistor, allowing current to flow into pins 5 and 6 and exit to ground through pins 3 and 4. This is exactly the same as transistor Q1 in our previous example.

When the current sense voltage reaches the high limit, U1 switches off its internal transistor, stopping the flow of current to ground. This forces the current to freewheel through diode D1. Once again, this action is identical to the previous example.

The resistance of R1 sets the level of the current through the LED. U1 is made with fixed high and low limit settings. $V_{current\_sense\_high}$ and $V_{current\_sense\_low}$ are fixed voltages specified on U1's data sheet, but since $V_{R1} = V_{current\_sense} = i * R1$, we can select R1 for any desired current, within the capabilities of U1. Integrated circuits like U1 have detailed data sheets to help with this sort of selection.

Sensing current as a voltage drop across a resistor is a very common technique. There is a trade off between precision and power loss: a larger resistor provides more voltage for the same current, which is easier to sense, but also burns up more power heating itself. The solution is usually a very small but high precision resistor. In this example we use a 0.300 ohm resistor with 1% tolerance, which provides enough precision for our lighting application without being ridiculously expensive.

The current switching process used in a system like this produces electrical spikes that we call noise. It's not noise in the audible sense, normally, but electrical noise that is impressed on power lines and control signals. Capacitor C1 is used to filter out some of this noise, helping to keep it out of IC U1. Capacitor C1 acts like a pressure tank, smoothing the voltage. The technical section that follows deals with noise in greater detail.

Dimming and Temperature Control

A look at U1 shows two pins that we haven't used: DIM and TEMP_I. Let's start with TEMP_I.

TEMP_I is an analog dimming and thermal foldback input. TEMP_I has an internal current source of 25uA which can be used to drive a resistor to generate a voltage, much like the current sense resistor previously mentioned. When the voltage at TEMP_I is above 2V, the current driving the LED is at its full value, as set by R1. As the TEMP_I voltage begins to drop below 2V, the LED current is reduced. A variable resistor can be connected between TEMP_I and ground, allowing for dimming of the LED. This is called *analog dimming* since the LED current is controlled by an analog (not digital) input.

More commonly, TEMP_I is used for thermal protection. LEDs are more efficient than incandescent lamps, but LEDs still make heat. And more troubling, the heat from an LED is all right there in the tiny LED; very little is radiated away. This is why high-power LEDs are always connected to heat sinks.

Even with a good heat sink, hot weather or reduced ventilation could cause an LED lighting fixture to overheat. Rather than damage the LED, TEMP_I can be used to limit the maximum temperature of the LED. Instead of connecting a variable resistor to TEMP_I, we connect a negative temperature coefficient (NTC) thermistor. The resistance of the NTC thermistor drops as it gets hotter. If we select the proper NTC thermistor, the voltage at TEMP_I will just reach the 2V threshold at our target temperature. Any hotter and the voltage drops below 2V, lowering the current to the LED, which keeps the LED from getting too hot. It's an automatic temperature control system. If the LED gets too hot, the controller reduces the power and brightness to keep the temperature under control. For this to work, the NTC thermistor needs to be in thermal contact with the LED. Generally this done by soldering the thermistor onto the printed circuit board in close proximity to the LED.

DIM is a digital dimming input for the LED driver. Before we discuss the details, let's talk a bit about analog verses digital dimming. With analog, the output is a continuously varying representation of the input; in other words, the output is an *analog* of the input. The current driving the LED smoothly varies in relation to an analog dimming input.

Digital, on the other hand, is either on or off, with nothing smooth about it. We could digitally dim our office lights by leaving them off for half the day. On the average, they would be half as bright, but it wouldn't be very useful since half the time we would be squinting and the other half sitting in the dark. Digital dimming, to be useful, must turn on and off fast enough that human eyes cannot see the on/off cycles. When visible, we call this flicker. Some people are sensitive to flicker, and are even bothered by the 120Hz flicker in old-school fluorescent lights. Because of this, digital dimming is usually operated at rates of a couple of hundred Hertz or higher.

Given the problem with flicker, why would we use digital dimming? Why not just use analog? The answer is dimming range. Analog dimming works well for brightness levels above 40%, but LEDs begin to behave poorly at current levels lower than that. Color shift is one of the bigger problems. For full-range dimming, digital is the method of choice. We use a method called ‘Pulse Width Modulation’, which is usually abbreviated to PWM. PWM varies the on/off time, or duty cycle, while keeping a fixed frequency. Figure 7 illustrates this concept.

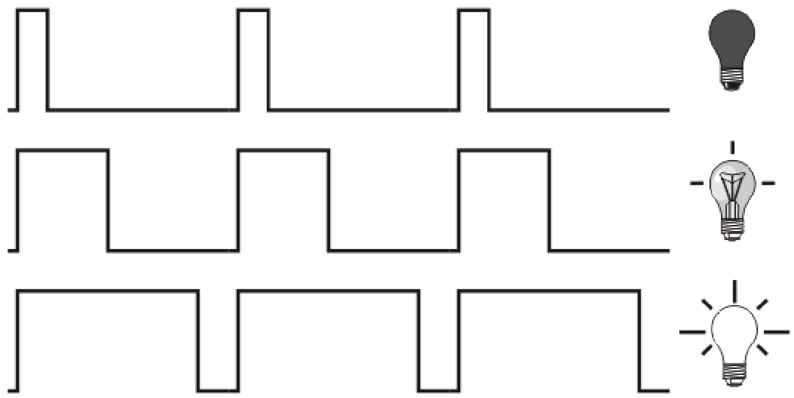

Fig 7

DIM is a logic level digital dimming input for the LED driver. Applying a logic level 0 to DIM turns off the LED; a logic level 1 turns on the LED. This is typically done a couple of hundred times a second (200Hz) or more by an external microcontroller. A compatible microcontroller dimmer circuit is described here: <insert link>.

Electronic Noise

Switching power supplies, by their very nature, deal with rapidly switching electrical currents. An unfortunate side effect of rapid switching is the generation of electromagnetic interference, normally abbreviated as EMI. EMI can be directly radiated from printed circuit board traces, which act like tiny antennas. EMI can also be carried by wires connected to the switching power supply; wires such as power leads and control signals.

Minimizing radiated EMI is accomplished by keeping the current loops small, using lower switching frequencies, and, in some cases, with physical shielding.

Conducted EMI is minimized by adding low-pass filters to power lines and control signal lines. The low-pass filters allow power or low frequency signals to pass, while absorbing or shunting high-frequency noise. Often these low-pass filters are composed of a ferrite bead which absorbs high frequencies, and a small capacitor that shunts high frequencies to ground.

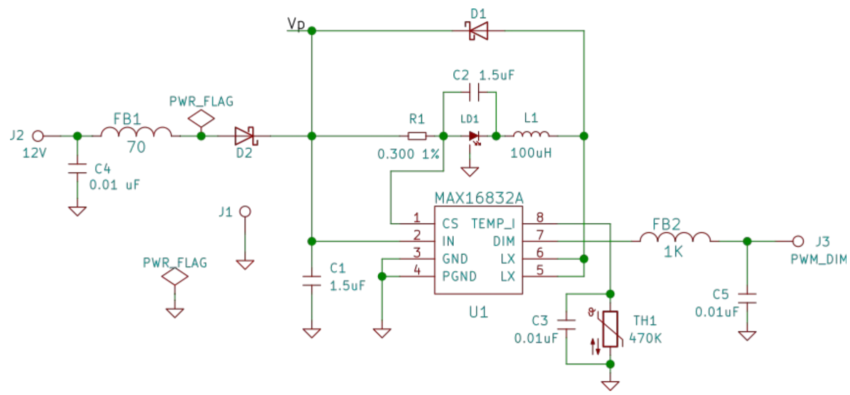

Fig 8

A Real-World Example

Building on our MAX16832A example in Figure 6, Figure 8 shows an actual production version of a MAX16832A LED driver. The production version adds a PWM input for digital dimming, a temperature control thermistor, and low-pass filtering for EMI suppression.

Figure 8 adds a few new pieces to our LED driver. First let's deal with the superficial changes: The PWR_FLAG symbols are not real components; they are virtual components that tell the PCB design software that the attached networks are power networks. They are used by automatic design checking tools that help keep us from making mistakes. The ground on LD1 is for a thermal pad, soldered to the PCB board to conduct heat away from the LED. The thermal pad doesn't need to be grounded, but generally is to help suppress EMI.

Now for the real additions: J1, J2, and J3 are solder pads for ground, power, and PWM dimming. FB1 and FB2 are ferrite beads used for EMI suppression. Each ferrite bead is accompanied by a 0.01uF capacitor, forming low-pass filters. D2 is a diode on the power input. D2's entire purpose is to protect against reverse polarity; if someone accidentally wires the power backward, D2 blocks it. TH1 is an NTC thermistor used for temperature control. C3 is a bypass capacitor, which is needed to stabilize the TEMP_I input. C2 is a parallel capacitor for the LED, which is the subject of our next section.

The working current for this production design is around 700mA, as set by R1.

To Cap or Not To Cap

Figure 8 adds a 1.5uF capacitor in parallel with LED LD1. Why? What does this do? It wasn't used on our previous circuits, so why add it now? An electrical engineer will tell you that the parallel capacitor reduces the effective impedance of the LED, but what does this really mean. To understand, we need to look at what the capacitor is doing.

Capacitors act like electrical pressure tanks, storing charge as a voltage (pressure) and releasing it later. When you apply a voltage to a capacitor, electrons pour in and charge the capacitor. When the voltage is removed, the capacitor retains these electrons, acting like a tiny battery.

To see how we use this with an LED, we need to think back to how our switching LED driver works. Recall that the transistor switches on and current begins to flow through our current sense resistor, the LED, and the inductor. The LED lights, maintaining its forward voltage, as LEDs operate at a nearly constant voltage. The current increases until it reaches the high limit setting, at which point the transistor switches off and the current freewheels through the freewheel diode until the current drops to the low limit and the cycle continues, creating an oscillation. The frequency of this oscillation is largely dependent on the applied voltage, the forward voltage of the LED, and the inductance of the inductor.

With this in mind, let's look at what happens when we add capacitor C2 in parallel with LED LD1. During the transistor-on part of the cycle, when inductor L1 is building its magnetic field, capacitor C2 is absorbing charge, building an internal electric field. Capacitor C2 releases this charge during the transistor-off part of the cycle, when the current is freewheeling. This is analogous to the pressure tank on a water pump. When the pump turns on, it supplies water and pressurizes the tank. When the pump turns off, water is supplied by the pressure tank until the pressure drops too low, at which point the pump comes on again.

The net effect is that it takes longer for the LED current to rise and fall, lowering the operating frequency. We can do the same thing by using a larger inductor, but inductors are much larger and more expensive than capacitors. By using a capacitor in parallel with the LED, we save money and space. In addition, the capacitor smooths the current flowing through the LED, which improves efficiency.

Buck vs Boost

Switching LED drivers fall into two camps: buck drivers and boost drivers.

Buck mode starts with a higher supply voltage and reduces it to something compatible with the LED; in other words, it bucks the voltage.

Boost mode works the other way around. It starts with a lower supply voltage and boosts it to something high enough to operate the LED.

Switching drivers can also be Buck-Boost drivers, where the circuitry is able to seamlessly shift between buck and boost modes. Buck-boost drivers can operate with lower and higher supply voltages. Obviously, buck-boost operation involves a more complicated, more expensive circuit.

Buck mode drivers are commonly used for plug-in devices, where higher voltages are available, and for mobile applications where 12V power is standard.

Boost drivers are common for battery operated portable devices where the weight and size of multiple batteries is a problem. The boost driver also compensates for drooping battery voltage that occurs with some battery chemistries as the battery wears out.

Buck-boost drivers are reserved for special applications with wide voltage swings or where the voltage is uncertain.

What's In a Name?

The MAX16832A integrated circuit in our example is a Buck Mode Switching High-Brightness LED Driver with Integrated MOSFET, High-Side Current Sense, and Hysteretic Control. What does all of that mean?

Buck mode, as we previously defined, means the driver *bucks* the input voltage to step it down to a useful level.

Switching means digital; the power device (transistor) in this circuit is either on or off, not somewhere in between as with analog.

High-Brightness means it has enough current to power high-brightness LEDs.

Integrated MOSFET means the transistor, in this case a MOSFET, is contained inside the integrated circuit. More powerful LED drivers will often use an external transistor driven by the integrated circuit. The external transistor can be sized as needed.

High-Side Current Sense refers to the location of the current sense resistor R1. It can be located above the LED, high side, or below the LED, low side, as referenced to ground. Both work just fine, but they require specific circuitry, which is why the location is specified in the name. Likewise the transistor can be high side or low side. Since the transistor is integrated into this device, its location was left out of the name.

Hysteretic Control refers to the control scheme. Hysteresis is the dependence of the state of a system on its history, and so it is with this method of control. Recall the triangle wave going from 200mA to 220mA; the low and high current limits in our example. When the current reaches the low limit the transistor switches on. When the current reaches the high limit the transistor switches off. For all the current values in between the low and high limits, how does the controller know whether to have the transistor on or off? The answer is that the controller remembers its history: if the previous limit was the low limit, the transistor stays on. If the previous limit was the high limit, the transistor stays off. This dependence on history is why this is called Hysteretic Control.

Technical Details

Switch mode LED drivers are tiny control systems, subject to all of the stability concerns of any feedback controller. Control system design is an engineering specialty of its own, often involving graduate level classes. Fortunately for us, most of the complex work has been done for us by the IC manufacturers, who strive to make their ICs easy to use. Always carefully review the data sheet and application notes for any IC that you elect to use. The manufacturer will generally offer guidelines and design notes. Following these will avoid a lot of pain and suffering.

Figure 9 shows the printed circuit layout of our production LED driver.

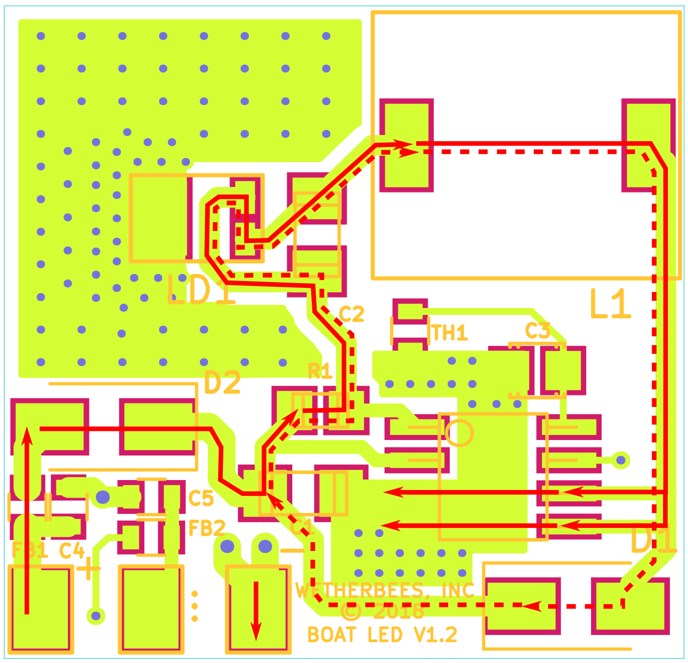

Fig 9

The two current flow loops are indicated by the red arrows. The solid red arrows show the current flow when the internal transistor is switched on, flowing from the plus solder pad, through FB1, through reverse protection diode D2, through current sense resistor R1, through LED LD1, through inductor L1, and finally through the internal transistor of U1 back out the negative solder pad.

The dashed red arrows show the current flow when the internal transistor is switched off. The current flows from inductor L1, through freewheel diode D1, through sense resistor R1, and through LED LD1 back to inductor L1.

These two current loops act as antennas, radiating EMI from the board. The amount of radiated energy depends on the area defined by these loops. Keeping the area small is important for stability and noise reduction.

The design guidelines on the data sheet call for the current sense resistor, R1, to be as close as possible to the IC U1. They also call for the freewheel diode D1 to be physically on the opposite side of U1. Both of these measures are to reduce the amount of noise coupled into the current sense resistor.

Note the copper areas stitched with numerous vias. The vias are plated holes that connect the upper copper of the board to the lower copper layer. These vias conduct heat from the LED and IC to the lower copper layer, which is normally attached to a heat sink.

The vias also conduct electricity, but only a few are needed for this purpose. The large number is necessary for good heat transfer. On the board, most of the bottom layer is used for a heat sink and ground plane, which is common for this type of circuit.

In the real world, inductor L1 is not an ideal inductor. The inductor is composed of coils of copper wire wrapped around a ferrite core. This means that the inductor has some series resistance, and capacitance between the coils of wire. By selecting a high-quality inductor we keep the series resistance low, which improves efficiency. But there isn't much we can do about the capacitance. The effect of this capacitance is that the inductor forms an L-C network that will oscillate at a characteristic frequency. Figure 10 shows this. Instead of the ideal triangle waveform shown in Figure 5, we see high-frequency oscillations on the back side of each triangle wave. These occur when the switching transistor turns off, allowing L1 to self-resonate, or ring.

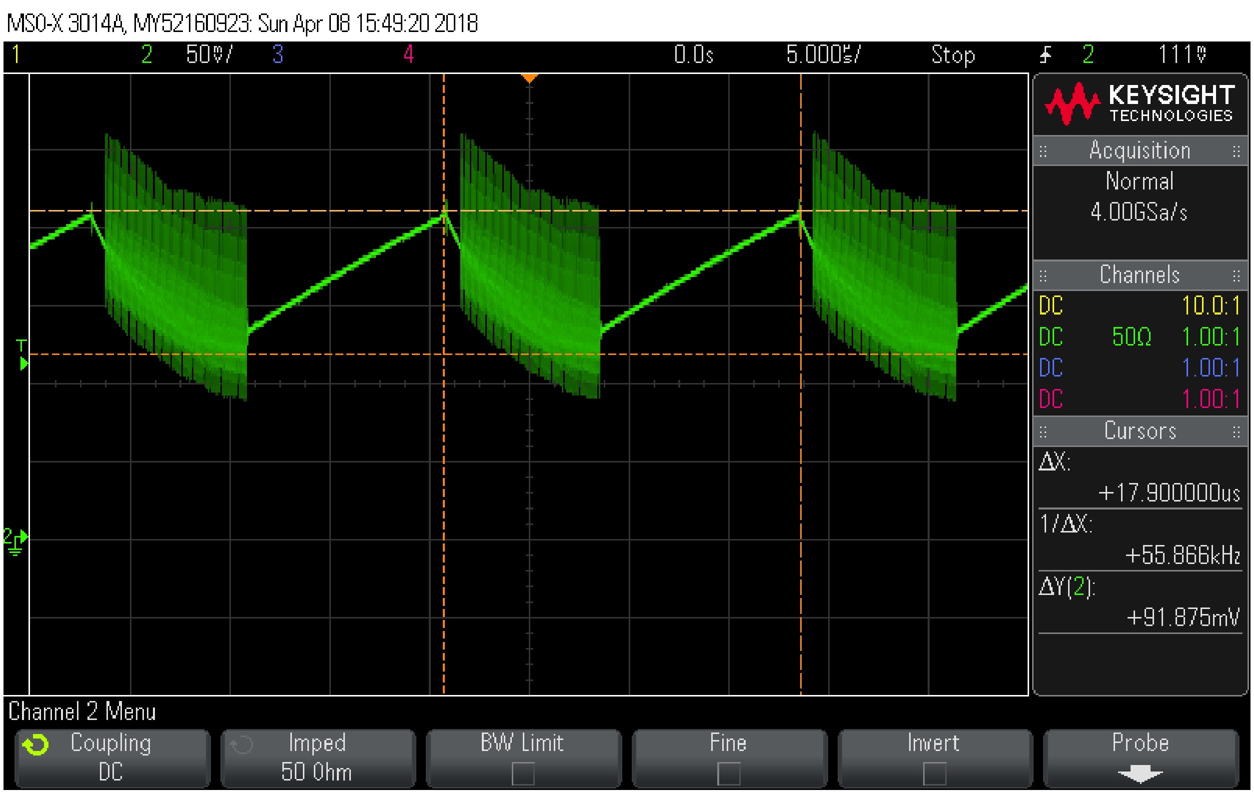

Fig 10

Getting Fancy

This same driver design is extended in Figure 11, which is another production design. Figure 11 uses three white LEDs, tripling the brightness, and one red LED for night use.

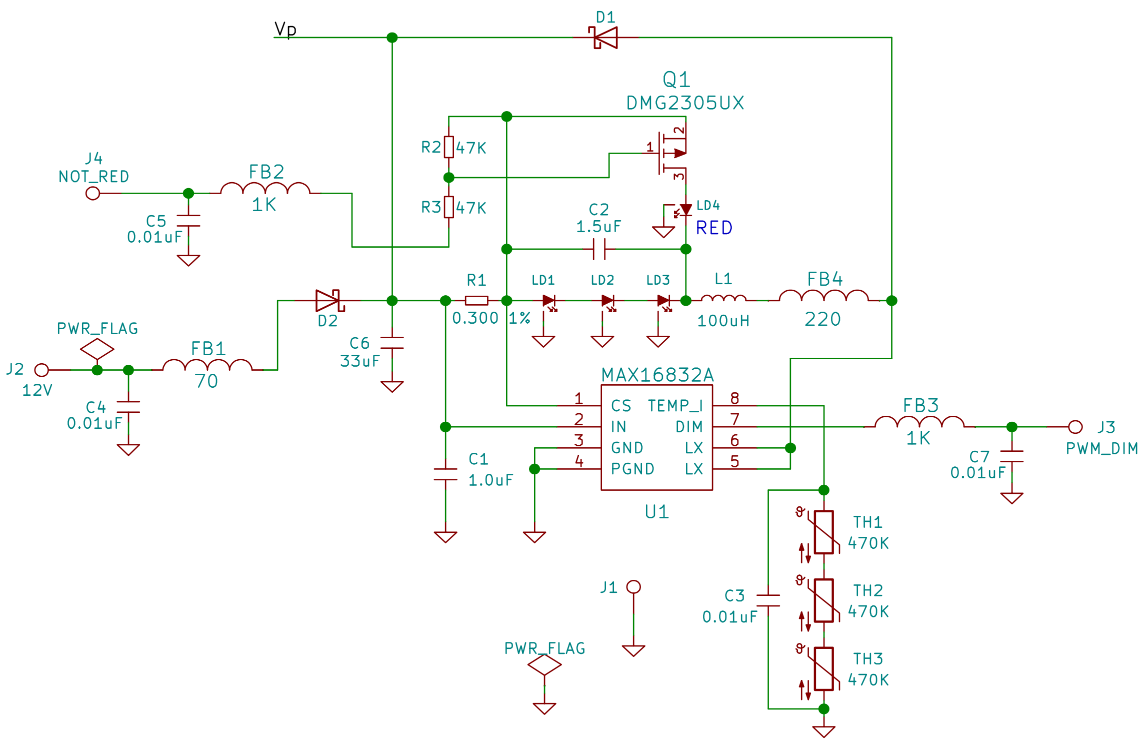

Fig 11

The three white LEDs, LD1, LD2, and LD3, are connected in series. The same 700mA current flows through all three white LEDs. This triples the light output without increasing losses in the rest of the circuit, greatly increasing efficiency. Any number of LEDs can be wired in series like this, limited only by their combined forward voltages, which must be lower than the supply voltage and within the limits of the integrated circuit.

The red LED is for night use, allowing for light without destroying night vision. Switching between white and red is done in an interesting fashion. The red LED is wired in parallel with the three white LEDs. The forward voltage of the red LED is lower than the combined forward voltages of the three white LEDs. Because of this, the red LED would light and the three white LEDs would not receive sufficient voltage to light. However, the red LED is wired in series with transistor Q1. The gate of Q1 is normally biased off by resistor R2. With Q1 switched off, no current can flow through the red LED, which causes the three white LEDs to light. When the NOT_RED input is grounded, R3 pulls the gate of Q1 low which causes Q1 to switch on. With Q1 switched on, the current goes through the red LED, turning off the three white LEDs.

This design is used for marine berth lights which have round brass shades. A round printed circuit board is used to fit this form factor, with wiring routed through the central lamp 'nipple', as shown below in Figure 12. The entire back side is a ground plane heat sink.

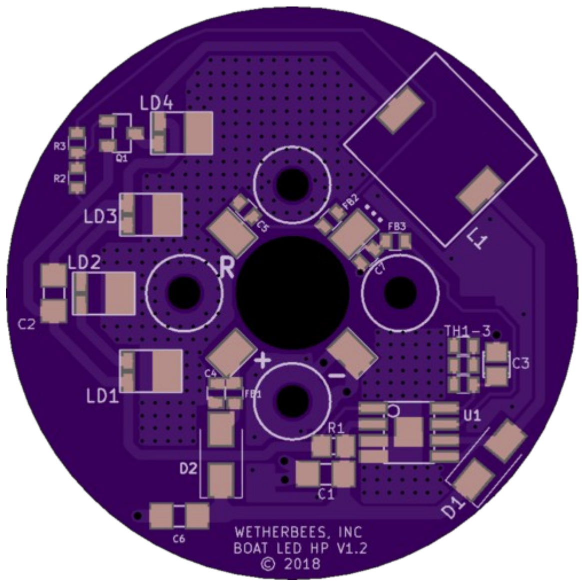

Fig 12

Conclusion

Buck Mode Switching LED drivers with hysteretic control are an inexpensive and relatively easy to implement form of high-efficiency LED driver. Controller ICs are offered in a variety of forms from several manufacturers. Pay attention to the manufacturer's recommendations and physical board layout to ensure stable operation.