Why can LEDs benefit from optical feedback?

Why can LEDs benefit from optical feedback?

The steady improvements in light-emitting diode (LED) technology have made these devices increasingly suitable for both illumination and fluorescence excitation in microscopy – and indeed for macroscopical applications too. Compared with incandescent and arc lamp sources, they run much cooler and are inherently more stable. Although what we refer to in Cairn as their “point intensity” (by which we mean their radiant intensity per unit area of the source) still may not be as good as for some arc lamps, it’s still way higher than for incandescent sources. That’s important for efficient illumination in microscopy. And although a given LED has only a limited spectral range, there is a wide enough choice to cover the entire optical spectrum and the near infrared and ultraviolet too, and their outputs an easily be combined by use of appropriate chromatic reflectors, or dichroic mirrors as they are usually misnamed! Furthermore, unlike those other sources, they can be switched on and off on timescales in the nanosecond range. So what is there not to like?

The potential performance shortcoming becomes apparent when you begin to exploit their fast switching ability. Different LED wavelengths are produced by slightly different technologies, which means that their propensity to this particular shortcoming is greater for some wavelengths than others. The accompanying oscilloscope screen shots should make this clear.

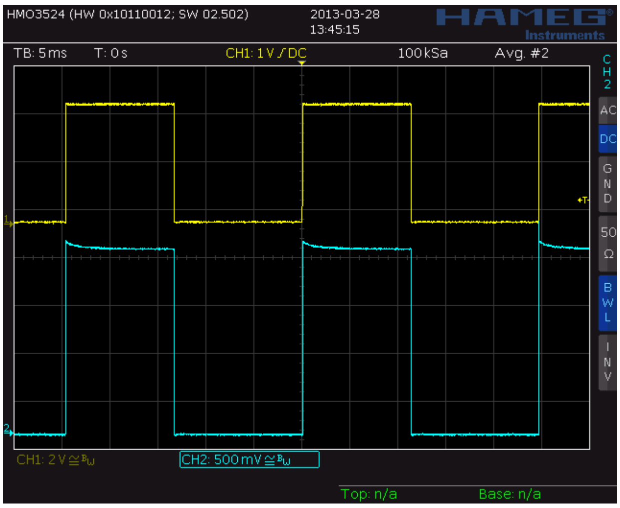

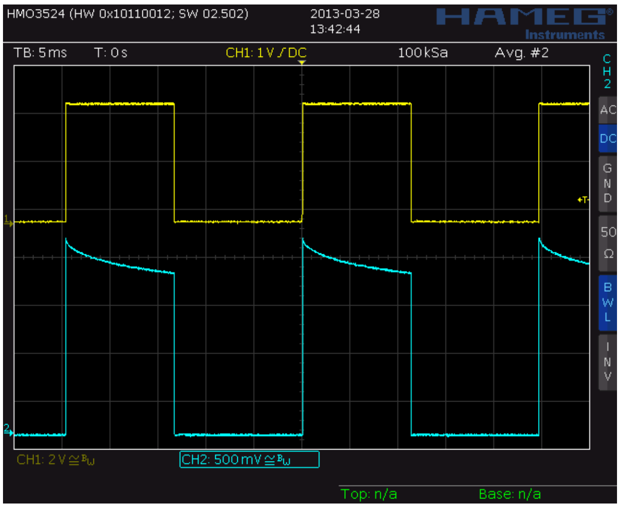

The yellow trace shows a pulsed current going through the LED, and the blue trace is the optical output, for a 505nm and 590nm LED respectively. The timescale is 5msec per division. One can see that in both cases the optical output declines during the “on” period. The effect isn’t too bad for the 505nm LED, but for the 590nm one it’s quite serious. So what is going on here?

This is actually a temperature effect, caused by the tendency for LEDs to become less efficient as they get hotter, but it’s worth explaining in much more detail than that, because the physics here is frequently misunderstood. Although the light from an LED is "cold", in that it is not thermally emitted, it's not a completely efficient process, so inevitably some heat is generated as well. It's the generation of this heat, and how efficiently it can be removed, that limits the amount of light that an LED can be safely made to produce. Like most other semiconducting devices, the maximum safe junction temperature (which in this case is the light-emitting one of course) is limited to about 150 degrees C before it wings its way to semiconductor heaven. Clearly the more heat we can conduct away, the more current we can pass, and hence the more light we can generate. Therefore these devices – and indeed any other "power" semiconductors – have a metal surface to which a heat sink of some form or other can be attached.

The source of possible misunderstanding is that the temperature of this metal surface is not the LED junction temperature! It would be great if that were so, but there is inherently a thermal resistance between the junction and the case, so the junction temperature will always be higher, and perhaps substantially so. Clearly the LED designers strive to reduce this resistance, but they have their own laws of physics to contend with here. As a general rule, the maximum rating of a power semiconductor is quoted for a case temperature of 25 degrees, so if the maximum safe junction temperature is 150 degrees, then the temperature difference between junction and case is a massive 125 degrees. That difference and the thermal resistance between junction and case sets the rate at which heat can be removed, so if we want to double the power handling by having a 250 degree temperature gradient instead, we'd need to get the case temperature down to -100 degrees! Therefore, although it makes sense to keep the case temperature as low as we sensibly can, this doesn't help nearly as much as you might think. Put the other way, a case temperature of 50 degrees (getting quite warm) would reduce the power rating by only 100/125, or 20%.

Oscilloscope traces

Now let's get back to those oscilloscope traces. Since the sag is a temperature effect, we can actually use it as a measure of the LED junction temperature, so this is telling us that the junction temperature changes on a timescale of some milliseconds when the current is changed. Two important things follow. First, it should be clear that no way can we can we get rid of this effect by regulating the case temperature – the required temperature fluctuations would be enormous and the timescales are all wrong anyway. Second, and perhaps rather more usefully, it's giving us some useful information on the thermal capacity of the LED junction and its immediate environs. Since the junction temperature doesn't change immediately with the current, this means that we can pass higher currents for shorter periods without sending the device heavenwards, on a timescale revealed (in this case milliseconds) by the temperature effect. That can be very handy!

Even though this is all very useful information, the effect is nevertheless a potential pain, so that's where optical feedback comes in. Take a look at these next two traces.

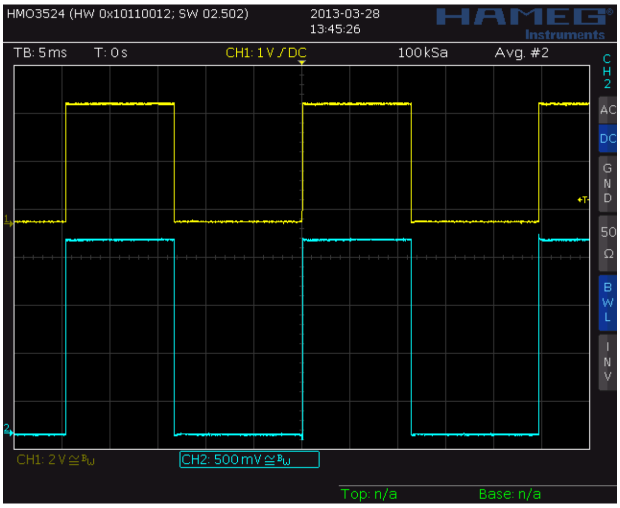

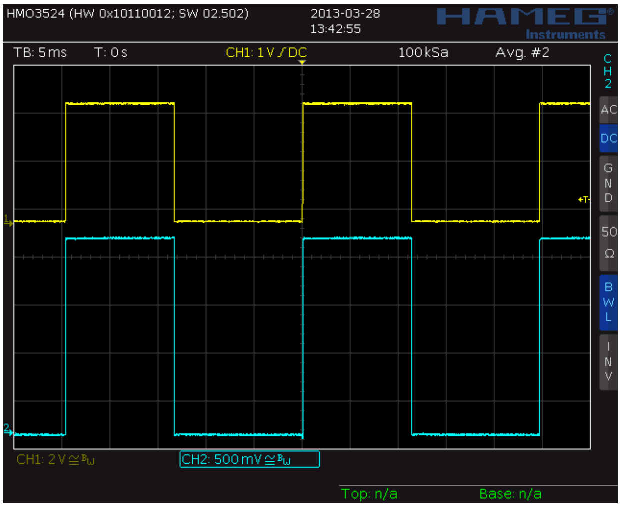

These are for the same two 505 and 594nm LEDs, both with optical feedback this time, and the traces are now perfectly square. The thermal effect has a timecourse on the order of milliseconds, whereas the optical feedback can be applied on a microsecond timescale, which is sufficiently faster to deal with it completely. All we need to do this is a photodiode and amplifier that looks at just a fraction of the light from the LED, which it can do by being "out of the way" of the main optical pathway, so there is no loss of useable light. This is actually so easy to do that we incorporated it in our original OptoLED design way back in 2003, and it's been a key feature of that product ever since.

However, there are two potential traps for the unwary! First, in order to "square up" the optical output, the current during a light pulse is going to increase somewhat, so one must be careful not to cause this to overdrive the LED. Some sort of protection circuit is therefore required, although in practice we need something like that anyway.

Spectral output in relation to temperature

The second trap is a rather more insidious one. This is that the spectral output of an LED can change with temperature, or possibly with the current itself. It's always advisable to use an LED with cleanup filters, to block any emissions that are outside the required waveband, so now one must ensure that the photodiode is seeing the same spectral range as the filtered output, and that means siting the photodiode downstream of the filtering. We found this to be particularly important when using a 365nm LED with a filter to select out the spectral range around 340nm for excitation of the calcium indicator fura2. But since we had actually anticipated the possible need for this, it was rather nice to find an application for which it was actually necessary!

But finally, optical feedback isn't useful just on short timescales. If you're not pulsing an LED, you might think that optical feedback isn't really going to be useful, but we have had people needing stable optical outputs over periods of days or even longer, so for them it's been very reassuring to know that optical feedback will guarantee that too. It's not just for the short events in life….