 By David Marsh, Contributing Technical Editor

BRAVO

UNLESS YOU'RE VERY FORTUNATE, the onset of winter most often includes an unwelcome reminder that incandescent bulbs virtually always fail at the worst moment. When you fumble in the darkness on a dismal December night, it's easy to forget that it's a miracle that bulbs work as well as they do, especially outside of relatively benign domestic applications. Inside that fragile glass envelope is a coil of wire that's as much as a metre long and less than 50 microns thick. Shrouded in an atmosphere of inert and sometimes expensive gases, this tungsten wire literally burns away the whole time it's turned on. Worse, every time you turn it on, the filament rises from ambient to more than 2500°C in a few milliseconds—or to as much as 3300°C for halogen types. This temperature boost inevitably creates a thermal shock that's often sufficient to rupture depleted wire. Another favourite failure mode occurs as the element cools and passes from a semifluid state back to solid form. In either case, when you reapply power, the weakened element sometimes snaps or molten metal flows, creating a short circuit that blows fuses or trips circuit breakers.

Incandescent bulbs have other issues, too. Because the filament emits only about 12% of the energy input as visible light, while converting the entire energy input into heat, conventional bulbs are roughly 10% efficient. The inrush current as the filament heats to working temperature is some seven to 15 times the steady-state current, creating circuit-protection issues. And good quality bulbs aren't cheap, because the necessary material science is relatively expensive. For these reasons and more, designers favour LEDs for general-purpose indicator duties. But until recently, the technology's limited light output has severely constrained application potential. It's still true that the solid-state domestic bulb or automotive headlight has yet to appear, but vendors are currently working on multichip modules to tackle power output levels of as much as 35W. Such modules will require exotic process and substrate technologies to assure viable lifetimes. But today, there's a rapidly

LEDs MAKE THE SPOT

AVAILABLE IN A MULTITUDE OF COLOURS AND OFFERING SUPERIOR RUGGEDNESS WITH EXTENDED LIFE, LEDs ARE THE NATURAL SUCCESSORS TO INCANDESCENT BULBS. TODAY'S HIGH-POWER VERSIONS OVERCOME LIMITED BRIGHTNESS TO ENABLE NEW AUTOMOTIVE, CONSUMER, AND INDUSTRIAL APPLICATIONS.

LIGHT

growing generation of high-power emitters that vastly increases illumination choices for lower power levels. Such devices especially suit high-reliability aerospace, automotive, and industrial applications. Elsewhere, lifestyle-accessory consumer applications, such as cameraphone flash units, torches, and spotlights, look set to consume huge volumes.

Currently, lifetime issues restrict most high-power LED applications to a power level of about 1W. For representative devices, this amount of power produces a white-light output of some 25 lumens that rivals high-temperature tungsten-halogen sources for luminous efficiency. Better yet, the lifetime of a 1W device extends to 50,000 to 100,000 hours, which is more than 100 times longer than an average bulb. In 1999, Agilent/Philips subsidiary and high-power LED pioneer Luxeon announced its Lumiled brand, with available colours including cyan, blue, green, orange, red, and white. As recently as last year, Osram announced its Golden Dragon family, which offers amber, blue, green, verde green (cyan), yellow, and white devices that are now available in production volumes. This year, both vendors face new competition from Cree's Xlamp and Nichia's Jupiter families, both of which also offer key-product white LEDs. Some lower specification devices are available from Taiwanese vendor Super Bright Optoelectronics, and Japanese optoelectronics specialists Rohm and Stanley Electric clearly have access to appropriate technology; Stanley, for instance, last year preannounced a high-power white LED but has yet to openly show a product.

Although white LEDs boast the broadest application potential, they are also the most difficult to construct economically. The obvious way to derive white light using LEDs is to combine RGB (red, green, and blue) sources. Sharp, for example, uses this approach in its 400-mW GM5WA06260A device. Advantages include a wide colour gamut and independent colour tuning, but the necessity of separate devices increases assembly costs.

Alternatively, it's possible to convert UV (ultraviolet) or blue emissions to white-light wavelengths by coating the base device with suitable phosphors. Osram holds the original patent that describes such luminous conversion techniques, and Nichia holds and licenses many of the key patents for blue- and UV-device processes. As a result, the companies enjoy several cross-licensing arrangements, which extend to vendors including Cree. The UV approach involves laying red, green, and blue phosphors over a UV source—a theoretically simple but relatively inefficient technique. It also holds the danger of unwanted and potentially damaging UV leakage. Taking the other route, Nichia in 1996 patented its YAG (yttrium-aluminium-garnet) process, using blue devices built in an InGaN (indium-gallium-nitride) technology and a yellowish phosphor that converts the base wavelength into a colour temperature value that's typically around 6500°K. Luxeon, too, uses a similar proprietary process with a unique gel compound for its white Lumiled devices that the company says results in more consistent, "warmer" white-light characteristics than those of its competitors (**Reference 1**).

Theoretically, white light is colourless, but this ideal isn't the case for white LEDs. In practice, their colour temperature varies hugely, from warm white to distinctly bluish hues across a range that spans about 4500 to 8500°K. It's therefore important to choose a characteristic that suits your application. It soon becomes obvious that making this choice is highly subjective; hence, it's essential to have a variety of samples available. First, it's useful to examine how the vendors grade their products. Because LED characteristics vary considerably among batches, manufacturers bin according to light output and wavelength. For example, Osram's datasheet for its LW-W5SG has four brightness bins that represent average luminosity values of 22.5 to 36 lumens, or—in units familiar from regular LED data sheets—7500 to 12,000 mcd (millicandelas). Interpreting brightness metrics immediately becomes a key issue that some confuse with the "1W" label that vendors apply to this product sector. Immediately separate this power rating from luminous intensity. It simply describes the forward-voltage-versus-current profile, which typically centres around 3.5V dc and 350 mA. Luminous efficiency—or the light output power per watt—is a useful relative measure that appears in every data sheet, but, unless you're a photographer, you may need some help in appreciating chrominance and luminance specifications (see **sidebar** "Demystifying the data sheet"). But whatever you do, *never* look di-

AT A GLANCE

- ▷ LEDs of 1W rival tungsten-halogen bulbs for luminous efficiency.
- ▷ Heat dissipation limits light output and demands efficient thermal design.
- ▷ Dedicated low-cost ICs simplify constant-current drive.
- ▷ Tests prove that FR4 is a viable alternative to metal-core substrates.
- ▷ Switch-mode drivers increase efficiency but can raise EMC issues.

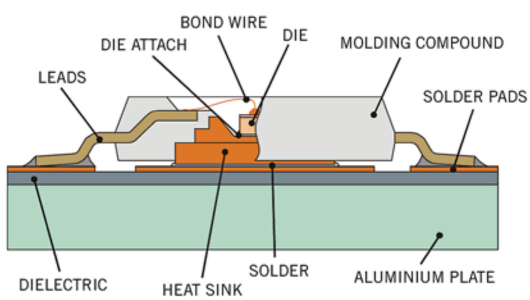

Figure 1 The copper slug beneath Osram's Golden Dragon die flows heat into the pc board, here using an aluminium substrate.

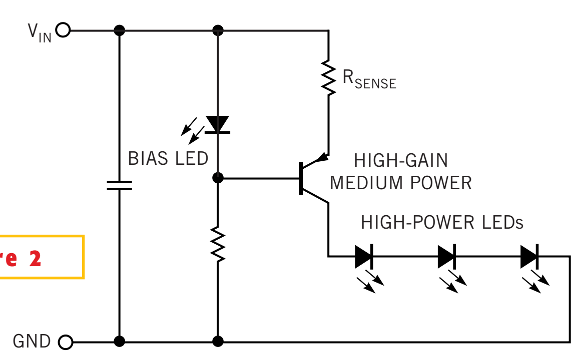

Figure 2 A low-power LED's negative-temperature coefficient stabilises a transistor-based current source.

rectly into a high-power LED while it's on; the intensity could cause discomfort and even damage to your eyes.

POWER DISSIPATION CONSTRAINS OUTPUT

Ultimately, power dissipation and its effect on device lifetime constrain light-output levels. For example, Luxeon offers several 5W devices, including a white version for portable applications, such as torches (flashlamps), that pumps out some 120 lumens at a drive current of 700 mA. The company derates lifetime to just 500 hours, which is nevertheless an approximately 50-fold improvement over a standard torch bulb. But, as with lower power devices, the key lies in getting the heat out of the device to limit its junction temperature to 135°C under worst-case conditions. Today, the widest range of available products comprises 1W devices, in which a spread of characteristics exists between the competitors' data sheets. For example, the crucial forward voltage that determines the device's dissipation at a given current level ranges from 2.75 to 4.3V dc at a junction temperature of 25°C. Maximum junction-temperature ratings span 105 to 135°C, and temperature-coefficient values span −2 mV/°C to as much as −4.5 mV/°C, depending on device construction.

Several considerations arise from these points. Because today's high-power white LEDs derive white light by pumping yellowish phosphors with blue light, the colour temperature changes a little with varying drive current and temperature. That is—and unlike a conventional LED—varying current levels change the

TABLE 1—REPRESENTATIVE RELATIVE LED INTENSITY VERSUS DRIVE CURRENT

Drive current (mA)	Brightness (lux at 250 mm)	Drive current (relative, %)	Brightness (relative, %)	Forward voltage (V)	Power dissipation (W)
350	149	100	100	3.2797	1.148
300	133	86	89	3.2224	0.967
250	116	71	78	3.1650	0.791
200	95	57	64	3.1039	0.621
150	73	43	49	3.0374	0.456
100	50	29	34	2.9606	0.296
50	25	14	17	2.8582	0.143

DEMYSTIFYING THE DATA SHEET

Like incandescent-bulb makers, high-power-LED vendors typically express light output in lumens, rather than the candela units that appear in a low-power LED's data sheet. Although a conventional bulb radiates in a relatively uncontrolled pattern, LEDs conventionally have lenses that direct light toward the target. Luxeon, for instance, offers its devices with choices that include "batwing" (low-dome) or Lambertian (high-dome) lenses with viewing angles of 110 and 140°, respectively. Cree's Xlamp 7090 similarly includes an integral Lambertian lens assembly with a 100° viewing angle. By contrast, Osram's Golden Dragon 120° viewing package is naked, with lenses available as optional accessories.

The term "Lambertian" refers to an ideal, perfectly diffuse surface, where the brightness remains constant from any viewing angle. Osram's Colin Beale expands: "The luminous intensity of a Lambertian emitter varies according to the cosine of the angle of view. A bare die or a device such as the Golden Dragon without a lens is virtually a Lambertian emitter. Devices with lenses tend to deviate from this model." He says that a Lambertian reflection from a surface appears to have the same luminance in all directions, as the area the viewer sees varies according to the cosine law. For instance, at an incident angle of 60°, the area of view is twice that at 0°: "As the luminance is the same at 60°, the actual luminous intensity is half."

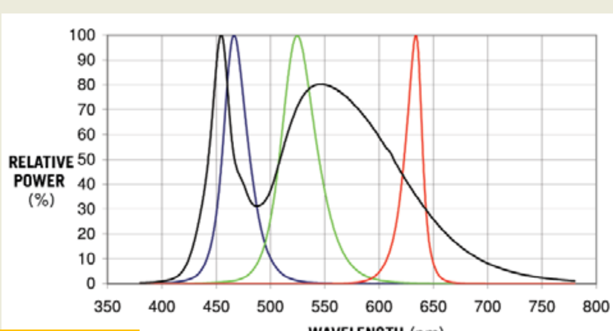

Figure A

Spectral output graphs plot luminous output versus wavelength.

Examining data-sheet terms, a candela—or one candlepower in the original Imperial units—represents "the luminous intensity of a source that emits 555-nm monochromatic radiation, with a radiant intensity of 1/683W per steradian in a given direction." The steradian is the SI (standard-international) unit of solid angular measure, in which one sphere comprises 4π steradians. In this context, it relates to a cone of light emanating from the point source, enclosing the angle equivalent to the LED's field of view.

Crucially, the raw candela value is independent of distance from the source and is thus a measure of luminous intensity. By contrast, the lumen is a measure of luminous flux and embraces the total light output in all directions. It gives rise to the lux metric, which is a measure of illuminance, or the source's ability to illuminate an area. The lux measures the energy content that strikes 1m^2 of a globe with a radius of 1m—that is, 1 steradian. (One lux equals 1 lumen/m^2 or 1 candela-steradian/m^2.) The intention of these concepts is to allow you to calculate the amount of light emanating from a point source that falls on an object at an arbitrary distance, as an inverse-square-law relationship applies; if a source is 1m away from an object at its 1-steradian point, the power falling on the object is one-fourth its previous value, if you move it 1m farther away

Chromaticity and colour-temperature concepts may be even less familiar. The visible spectrum ranges between wavelengths of around 400 nm (blue) and 720 nm (red), and coloured LEDs have spectral characteristics that typically centre on 470 nm (blue), 530 nm (green), and 625 nm (red) (**Figure A**). This spectral-output graph plots relative power output versus wavelength for Cree's white and coloured LEDs, with white emitters appearing as the black line that peaks at around 450 and 550 nm. In the case of white LEDs, chromaticity diagrams show *Cx* and *Cy* coordinates according to the CIE-1931 colour-space convention. Adopted in 1931 by the CIE (Commission International d'Eclairage), these chromaticity graphs describe hues within a standard envelope that represents the eye's perceptual ability within a measurement syschromaticity as well as the device's brightness. In practice, your eye is unlikely to notice the chromaticity change. But assuming that you drive each device at its 350-mA rating, an average device dissipates about 1.25W. This power dissipation in a surface-mount package that measures only a few square millimetres demands careful management to avoid exceeding its maximum junction temperature. For this reason, device manufacturers include on the underside of the package a solid copper slug that carries heat away through the pc board. Depending on the power dissipation of the assembly, conventional FR4 pc board might suffice, or a metallic pc-board substrate may be necessary (**Figure 1**).

The large variation in forward-voltage specification compounds the necessity not to overdrive the device. This requirement demands a constant-current driver, preferably with PWM (pulse-width-modulation) capability that you can use for dimming or to minimise light output variations via software corrections. As a practical example, assume that you're designing a three-LED light cluster for an automotive application. Connect these devices in series, as their varying forward voltage variations makes it impossible to accurately share current between parallel connections. Also, a parallel connection would require approximately 350 mA per device, requiring a more-than-1A supply for this three-LED example. The automotive environment infers a nominally 13.5V rail that's typically 12.6 to 14.4V while the alternator is maintaining battery charge. This voltage dips massively during cranking and may transmit high voltage or even negative-voltage transients during normal running. Be-

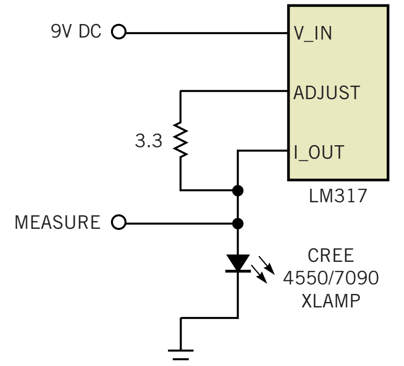

Figure 3 National Semiconductor's LM317 builds a precision current source for Cree's Xlamp demo board.

tem that relies on additive colour mixing. The additive principle combines reference wavelengths of red, green, and blue light, with "white" conventionally lying within the *Cx*=0.32, *Cy*=0.31 area.

Vendors most often specify white LEDs in terms of colour temperature, which relates to the chromaticity graph to represent the colour that an ideal black-body radiator emits at a given temperature. In practice, few white-light sources produce the same colour as a black-body radiator, so the data-sheet colour temperature of a light source is typically an approximation that's known as the CCT (correlated-colour-temperature). The Osram LW-W5SG's typical colour temperature is 6500°K, where 3200°K is a typical for a low-voltage halogen bulb; 5400°K is a photographer's "mean noon sunlight" value; 6000°K is a metal-halide discharge lamp, such as illuminates motorway construction areas; and a clear blue sky lies between 12,000 and 18,000°K. Osram bins its white Golden Dragon devices into one of eight groups, each of which has subtly different *Cx*/*Cy* balance within the CIE-1931 "tongue" envelope that appears on the left in **Figure B**. The expanded graph on the right shows groups in transition from blue to red as the plot's x axis moves from left to right, together with the 5K-to-8L bins that the company employs.

Although these concepts are academically interesting, as the main text of this article reveals, a more practical approach evaluates samples by making relative measurements under representative conditions. This situation is especially true when it's difficult to equate vendors' specifications like-for-like. Although illumination must please the eye, a luxmeter soon reveals that the eye can be insensitive to relatively large intensity changes. Because white LEDs produce light with strong blue and yellow/green peaks, some luxmeters may give readings that don't accurately compare with light from conventional sources. They nevertheless provide a useful relative measure among white LEDs.

The scientific way to measure colour temperature employs a colorimeter, but because perception is so highly subjective, it's more productive to invite decision makers to view a side-by-side demonstration. *But on no account should you directly view a high-power LED;* its output can hurt your eyes. Project its output onto a pure-white card, which makes it easier to perform subjective assessments of both colour and intensity.

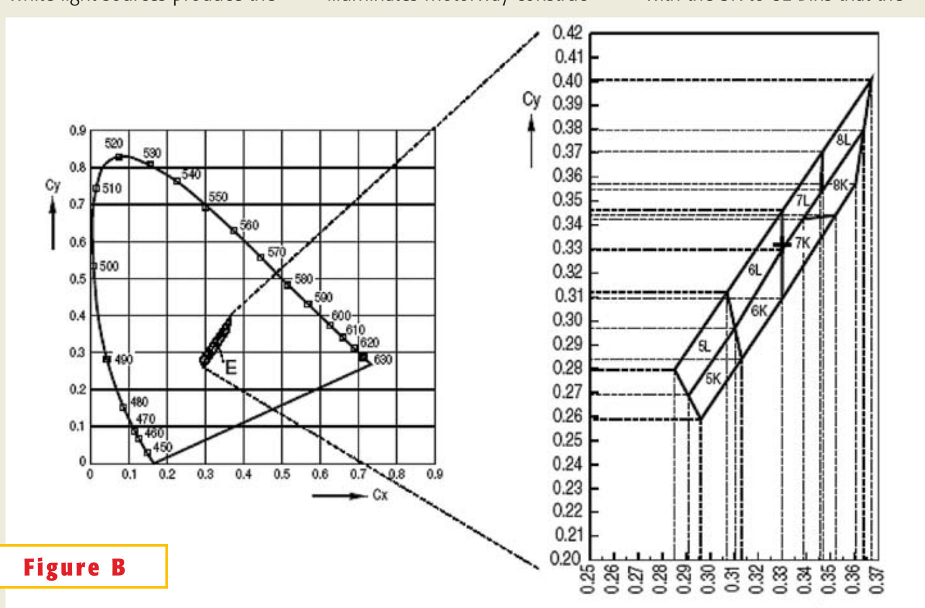

Figure B

Chromaticity graphs for white LEDs plot correlated-colour-temperature within the CIE-1932 window.

cause it's unlikely that lighting equipment will need to operate to full specification while the starter-motor engages, the low voltage issue is less critical than transient protection. But a few calculations quickly determine that the standard approach that powers normal LEDs—deriving a current via a series resistor from a regulated power rail—is unsatisfactory at current levels of more than 10 times a low-power device's operating point.

You may think of using a pnp pass transistor, a reference resistor, and a dual-diode bias network to derive the constant current. With so few connections, this configuration is especially attractive for single-sided boards. It's also cheap and capable of good performance within a small area using contemporary SOT23-6 packaged transistors, such as the ZX5T2E6 from Zetex. This 20V-rated device has a minimum gain of 300 and features a typical saturation voltage of just 50 mV at 350 mA, assuming 3.5-mA base current. You can improve the circuit's temperature stability by substituting a low-power LED in place of the double-diode, being careful to match the LED's negative temperature-coefficient against that of the transistor (**Figure 2**). This is also a cheap way to get a separate power-on indicator. A more compact and possibly even cheaper alternative employs National Semiconductor's venerable LM317 voltage regulator in current-regulation mode. Now available in SOT23-style packaging, this 1.5A-rated chip has a maximum input-output differential of 40V and features internal power limitation. Cree's demo board uses it for its 4550 and 7090 Xlamp LEDs (**Figure 3**).

Figure 4

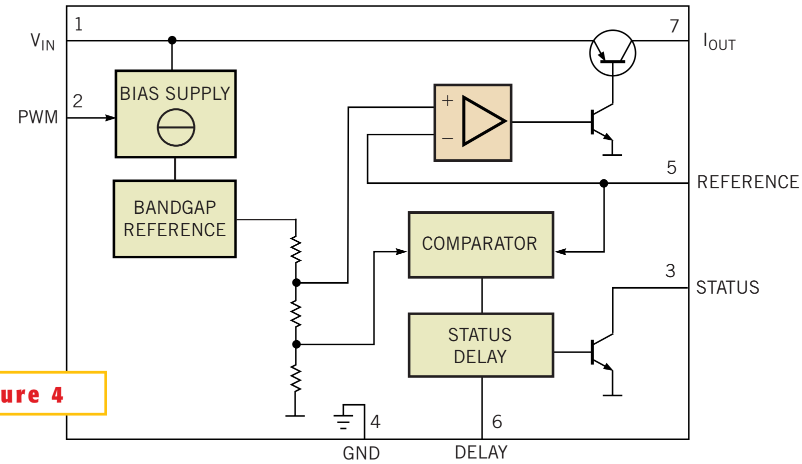

Comprehensive protection circuitry within Infineon's TLE4242G withstands automotive extremes.

DRIVER ICs SIMPLIFY REGULATION

With a few additions, it's possible to protect the simple transistor circuit from the faults and transients that accompany automotive and industrial environments. But these steps increase cost, consume additional pc-board real estate, and can't compensate for low-voltage dips. A growing breed of high-power LED-driver chips eases these issues. If you can ignore undervoltage-regulation concerns, Infineon's TLE4242G linear IC assures simple operation with protection that's tuned for automotive applications. With a minimum input overhead requirement of around 0.7V dc, it can drive three series-connected LEDs at nominal automotive input voltages. Conceptually, this chip comprises the pnp pass-transistor circuit, together with integral overtemperature, overvoltage, overcurrent, and reverse-polarity protection (**Figure 4**). Available in a seven-pin TO-263 surface-mount package, the device also features a high-impedance PWM input that also facilitates sleep mode and an open-collector status output with a programmable delay facility.

You program LED current using a single resistor that dissipates less than 100 mW, because only some 180 mV appears across it. But, being a linear circuit, the TLE4242G dissipates power that approximates the product of differential input-output voltage and LED-drive current. The chip also consumes some bias current, but it is negligible compared with the power through the device. With the vehicle's power rail at 14.4V and three LEDs with minimal forward voltages of around 2.75V, the chip thus dissipates just over 2W. The package has a junction-to-case thermal impedance of 3°C/W, and maximum junction temperature is 150°C. Interestingly, and despite the device's 500-mA maximum rating, the TLE4242G's data sheet makes several references to a 300-mA drive-current datum. As Infineon works closely with its former stablemate Osram, this observation suggests examining the implications of running at a reduced current level. Expect a slight light-output penalty but a substantial reduction in power dissipation that directly impacts lifetime. According to Colin Beale, technical sales manager for Osram's opto-semiconductors, "We estimate that due to 50% lumen depreciation, the lifetimes and relevant junction temperatures for this die technology and package are in the region of more than 10,000 hours at 125°C; over 30,000 hours at 100°C; over 50,000 hours at 75°C; and in excess of 75,000 hours at 40°C."

If the TLE4242G's input overhead is too narrow or its power dissipation is too great, consider a switch-mode approach. For example, the new HV9910 from Supertex boasts more than 90% conversion efficiency and a huge 8 to 450V-dc input range. It's also possible to run the chip from rectified ac-line inputs of 85 to 265V ac, with just enough smoothing to hold the rectified input voltage at a level twice the value of the LED string. Higher input voltages work better here, because the area under the rectified sine wave spends proportionately less time close to zero. In conventional low-voltage operation, the HV9910 works in down-converter mode, driving an external FET that charges and discharges an inductor (**Figure 5a**). One resistor sets the switching frequency, which is adjustable between about 25 and 300 kHz. The current that passes through the LED is then an average of the charge/discharge cycle—a principle that similarly applies to all switch-mode drivers. To avoid instability in buck-converter operation, it's also essential to keep the input voltage above twice the total LED forward voltage. Alternatively, the circuit can operate as a buck-boost converter when the HV9910 drives as many as six devices from a 12V rail (**Figure 5b**). As well as the normal 0 to 100% PWM dimming mechanism, the HV9910 unusually includes an analogue linear-dimming input that works with

Figure 5

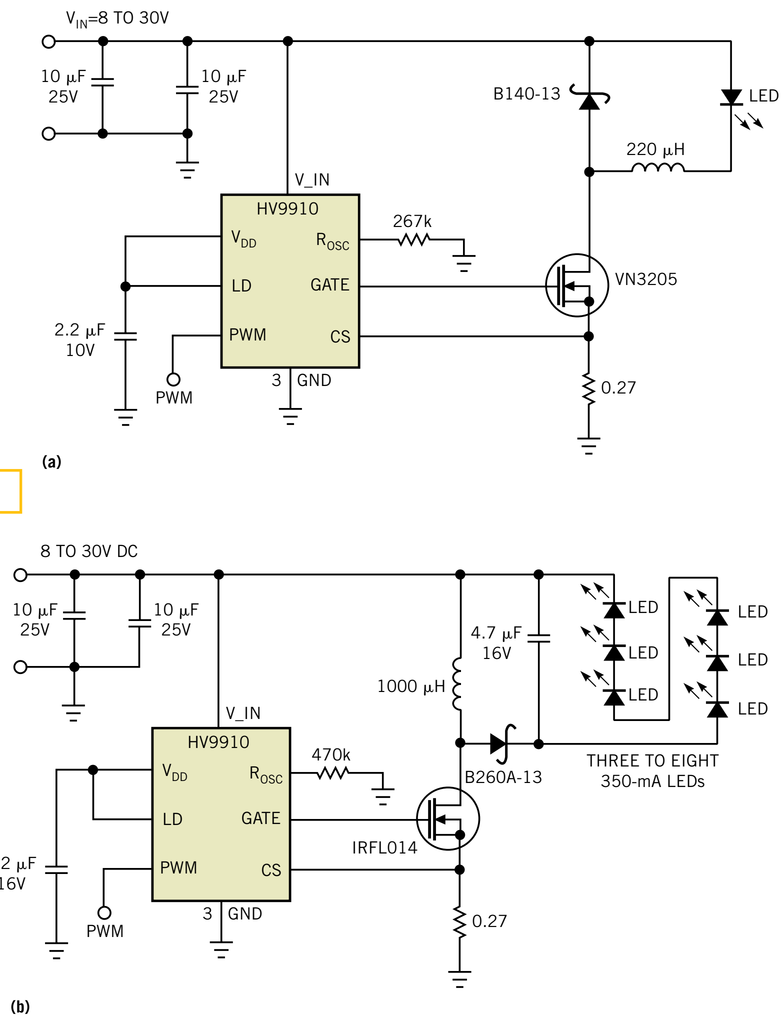

As a downconverter, the HV9910 from Supertex drives one LED from 8 to 30V dc supplies (a). Rearranging the HV9910 as a buck-boost converter drives multiple LEDs from low-voltage supplies (b).

PWM or on its own. Applying a 0- to 250-mV control voltage to the linear-dimming pin overrides the internal 250-mV threshold for the current-sense pin, reducing the output current level.

Other vendors with appropriate switch-mode drivers include Melexis, Microsemi, and Zetex. For example, the MLX10801 from Melexis is an eight-pin surface-mount chip that's primarily intended as a voltage-to-current downconverter. Targeting automotive use, it requires nominally 13.8V dc to drive one high-power LED at 350 mA in the downconverter configuration. The device contains a programmable EEPROM that holds calibration constants, allowing you to reprogram default parameters to suit alternative applications. To drive multiple series-connected LEDs, rearrange the circuit as a step-up converter by replacing the commutation diode with series-connected LEDs in much the same way as the Supertex example. In this configuration, the circuit drives three series-connected LEDs at voltages down to 5V dc and operates safely to about 38V dc.

Microsemi offers a range of drivers, such as its LX1992, that you can use to drive high-power LEDs by adding external power MOSFETs. Zetex similarly offers devices, such as its ZXSC310 and ZXSC400. With an allowable input voltage range of 0.8 to 8V dc, these devices are principally intended to operate as upconverters that can drive high-power LEDs from primary-cell sources. Rearranging the circuitry similarly but a little differently from the Supertex example allows the ZXSC310 to drive as many as three LEDs in a 12V environment (**Reference 2**). Because Zetex designed these devices for battery-powered applications, their efficiency is high even at low input-voltage values and reaches about 94% in the upconverter configuration. The circuit uses the pulse-frequency-modulation principle at frequencies up to 200 kHz to drive an external power transistor. Although the upconverter-reference example shows a MOSFET driving the inductor, Alan Buxton, marketing manager at Zetex, is keen about stressing the usefulness of the company's range of high-efficiency bipolar switching transistors: "In low-voltage applications, bipolar transistors require only about 0.7V to control conduction, rather than 1.5V minimum for a low-gate-threshold MOSFET. This consideration is essential for single-cell and even dual-cell designs."

TESTS RESOLVE SUBJECTIVE ISSUES

With a characteristic as subjective as lighting quality, the only option is to assess competitors' products on a board that properly represents your end-user application. This approach allows you to evaluate less-than-obvious points—such as optimal LED spacing and lens selection for even illumination—as well as brightness, colour temperature, and electrical and thermal characteristics. In their race to boast the highest possible luminance, vendors encourage you to drive their LEDs at the maximum current rating. For this reason, device vendors favour metal-core boards, such as Bergquist's T-Clad and Sericut's aluminium material, although other possibilities include the direct-bond-copper substrate from vendors such as Curamik that traditionally appears in chip-on-board and power-module designs. Although these materials are highly efficient, they're at best several times the price of FR4. Metal-core boards become even more expensive if you need more than one signal layer, further restricting application potential, creating pc-board-routing issues, or both. Also, don't overlook manufacturing implications; relatively few pc-board and assembly houses can work with metal-core products.

Assembly practices also vary among devices. Specifically, Luxeon's emitters require a thermally conductive epoxy to glue the heat-sink slug beneath the device to the pc board. Luxeon recommends a hot-bar soldering process to secure the leads to the pad lands, followed by a functional test. The working assembly then receives a heat cycle to cure the glue (**Reference 3**). The company's competitors are keen about stressing their products' compatibility with conventional infrared reflow processes, but be aware that it takes substantial heat to reliably flow the metal slug to a board—and it's impossible to hand-solder prototypes this way. Accordingly, Luxeon's recommended thermal epoxy—Emerson & Cuming's Amicon E3503 material—or alternatives from suppliers such as Loctite are highly appropriate during development.

If you plan to use low-cost FR4 and prefer conservative design practice, it's interesting to assess brightness versus drive current and the power dissipation that results. **Table 1** plots these values for a single Cree XL7090WHT-L100-W2-J-0001 device that has spectral qualities in the centre of the white emission region. This test, which is hardly rigorous, uses the company's PCB112B sample assembly built over a 1.2-mm-thick aluminium substrate that measures some 25 mm square. Because this mass is inadequate for full-power tests, mounting holes allow you to add an external heat sink—in this case, a 5.6°C/W part from Aavid Thermalloy (FarnellInOne part number 707-510). Here, the test assembly uses a layer of Warth's double-sided self-adhesive Kool-Pad thermally conductive sheet, which glues the device to the heat sink. In a TO-3's area, this 0.14-mm-thick aluminium-based sheet adds about 0.49°C/W, constraining the device's temperature rise to around 18°C at full power. Notice that the choice of a 3.3Ω sense resistor sets the LM317's output current to around 380 mA, so it's preferable to use **Figure 3**'s "measure" terminal for the direct-drive current-source purpose that Cree intends. This connection is essential for PWM drive.

Starting with 350 mA for one hour, the test's forward-voltage readings allow 15 minutes at each successive power level so the device temperature can stabilise. The relative brightness measurements come from arranging the assembly 250 mm away from, and directly on-axis with, a luxmeter in a darkened room. Such close proximity alleviates slight off-axis and stray background-illumination effects. If you try this test or a similar one, notice that there's relatively little subjective brightness difference between 300 and 350 mA, but power dissipation drops by about 16% at the lower level. Within the test's limits of measurement error, the results confirm that device brightness approximates a straight-line function of its drive current. (In fact, the characteristic appears slightly parabolic.)

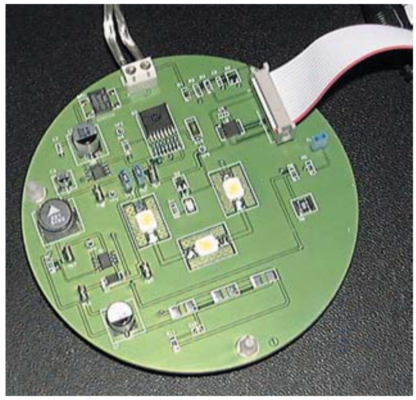

Figure 6 Careful design allows low-cost FR4 to dissipate several watts from multiple LEDs.

To assess FR4's suitability for heat dispersal, the 110-mm-diameter prototype in **Figure 6** accommodates three LEDs and two onboard power supplies. It also has various jumpers to enable options, such as external power. With careful pad-land design and a bit of inventiveness, it's straightforward to prototype LEDs from Cree, Luxeon, and Osram using the same pc board. The onboard power supplies comprise a linear circuit using Infineon's TLE4242G, and a home-brewed switch-mode circuit built on a microcontroller. (See **Reference 4** for more on this approach.) Here, a 68HC908-series microcontroller from Freescale principally provides a PWM-dimming facility. The pc board employs four layers of 35-micron copper with a layout that optimises each layer for heat removal and dispersion. Designed using the custom-pads facility in Version 3.1 of the Pulsonix suite, the LED pad-land design includes 47 30-thou-diameter (0.75-mm) thermal vias with 16-thou (0.4-mm) plated-through-holes that carry heat from the top layer to the bottom; from there, heat disperses via solid copper planes on what is the top side in the end-user application (**Figure 7**). Because the linear regulator dissipates around 1.8W in normal operation, another set of thermal vias connects its tab to both the ground plane and the opposite side of the board.

Tests with Osram's Golden Dragon parts again confirm that the eye finds difficulty in discriminating between current-level shifts of less than 50 mA, especially toward full-current operation. The prototype shows that choosing a conservative drive level of nominally 275 mA constrains the worst-case temperature rise to about 35°C. Unsurprisingly, this rise occurs in the region of the TLE4242G current regulator. With a centre-to-centre spacing of just 22.5 mm, the temperature around the LEDs on both sides of the test board rises by less than 30°C. A practically negligible thermal gradient between top and bottom sides demonstrates the efficiency of the thermal via system. Scanning the surface with an infrared thermometer also shows that the heat quickly dissipates to values about 10°C lower within around 2 cm of the devices, suggesting board-area reductions; if you're trying this test for production, get a proper image from a thermal camera. In the case of Osram's Golden Dragon parts, estimate the crucial LED-die temperature by measuring the spot temperature at the cathode lead, which has a thermal resistance of about 9°C/W. Remembering that the temperature coefficient is around −4 mV/°C provides a simple sanity check by

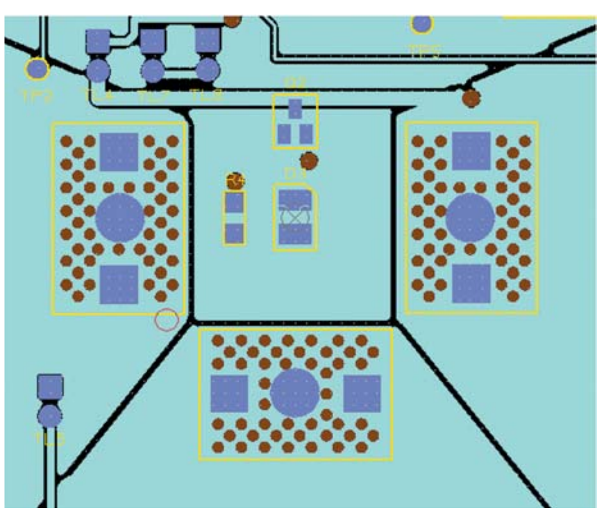

Figure 7 Thermal vias flow heat away from the LEDs and into surrounding solid copper areas.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN Europe*.

Aavid Thermalloy
www.aavid.com

Bergquist
www.bergquistcompany.com

Commission International d'Eclairage
www.cie.co.at/cie/

Cree
www.cree.com

Curamik
www.curamik.com

Emerson & Cuming
www.emersoncuming.com

FarnellInOne
www.farnellinone.com

Freescale Semiconductor
www.freescale.com

Infineon Technologies
www.infineon.com

Loctite
www.loctite.com

Luxeon
www.lumileds.com

Melexis
www.melexis.com

Microsemi
www.microsemi.com

National Semiconductor
www.national.com

Nichia
www.nichia.co.jp

Osram
www.osram.com

Pulsonix
www.pulsonix.com

Rohm
www.rohm.co.jp

Sericut
www.sericut.com

Sharp
www.sharpsma.com

Stanley Electric
www.stanley.co.jp

Super Bright Optoelectronics
www.sboopto.com

Supertex
www.supertex.com

Warth International
www.warth.co.uk

Zetex
www.zetex.com

measuring the cold-to-warm forward-voltage change across the LED string; here, results suggest a die temperature of approximately 55°C in a 20°C room.

To trade the lack of EMC concerns that a linear regulator provides when it's delivering 100% output—and thus without PWM dimming being active—for a more efficient switch-mode design, the HV-9910DB4 demo board from Supertex makes a representative test tool. At 12V-dc input, measurements using a normal DMM (digital multimeter) indicate that the sample demo board delivers 332 mA but consumes less than this figure. Applying a scope quickly explains this discrepancy, which is due to the DMM's integrating fast transients at the output that also appear at reduced amplitude at the power-supply input. In fact, the Supertex demo board is about 88% efficient, but this experience serves as a reminder that switch-mode operation can have hidden costs.

Ultimately, extrapolating results from such tests suggests that a few simple precautions allow safe design with FR4—and most likely without needing a four-layer stack, enabling further cost reductions. Be prepared to haggle with the LED vendors, because pricing is currently volatile. This situation should stabilise as more players enter the market and the early leaders adjust their profit margins. In the meantime, prices range from as much as 10€ for 1W Luxeon emitters from catalogue distributors, such as FarnellInOne, to less than 3€ (10,000) for competitive parts. Prices for dedicated driver chips centre around 1€ (10,000), but don't neglect the extra costs that switch-mode circuits incur for external MOSFETs and inductors.

Also, be sure to regularly check for new developments as this fast-moving marketplace matures. For example, at press time, Osram announced its new "Ostar" 120-lumen device, which uses multiple RGB thin-film emitters. Referring to November's Electronica show in Munich, Germany, Cree Lighting's business and sales development manager, Bob Henry, confirms: "We are announcing a 50-lumen device at 350 mA." This development represents a times-two improvement in luminous efficiency over first-generation devices.□

You can reach Contributing Editor David Marsh at forncett@btinternet.com.

References

1. Mueller-Mach, R, Gerd O Mueller, Michael R Krames, and Troy Trottier, "High-power phosphor-converted light-emitting diodes based on III-nitrides," *IEEE Journal on Selected Topics in Quantum Electronics*, Volume 8, No. 2, March/April 2002. www.luxeon.com.

2. "A high-power LED driver for low-voltage halogen replacement," Application Note AN-44, Zetex, 2004, www.zetex.com.

3. "Luxeon emitter assembly information," Application Brief AB-10, www.luxeon.com.

4. Fosler, Ross, "Use a microcontroller to design a boost converter," *EDN*, March 4, 2004, pg 74.