The LED luminaire based on an Altman 360 used a simple power supply to deliver DC power from an autotransformer.

Low-end dimming: the exception that proves the rule

BY KARL G. RULING

"IT'S THE EXCEPTION that proves the rule," is a common expression, but what does it mean? Look it up on the Internet, and you will find lots about what people mean when they say it—much of it contradictory. However, in this case, I mean it's the odd situation that sees if the rule works. In this case, the rule is the testing procedure outlined in the draft standard *BSR E1.69, Reporting the Low-End Dimming Performance of Entertainment Luminaires Using LED Sources*, and the oddities are two home-brew LED luminaires with mains dimming. While working at home, I tried the testing and data presentation procedures outlined in the draft *BSR E1.69* to see if they are practical: Can a person without access to a professional photometric testing lab do them?

> . . . running the tests on . . . these [two home-brewed units] certainly . . . would prove the rule.

Standards are written by committees. They are never the work of one person, but I myself aim to help write photometric standards that can be implemented by high school or college faculty and students. The test results might be more credible if done by a professional testing laboratory, but it's good if a competent amateur can get reasonable results with easily-affordable equipment. A testing procedure that can't be done by a lay person puts the test and what it tells us outside the secular world, into the realm of experts. That's fine if a test really takes such advanced knowledge

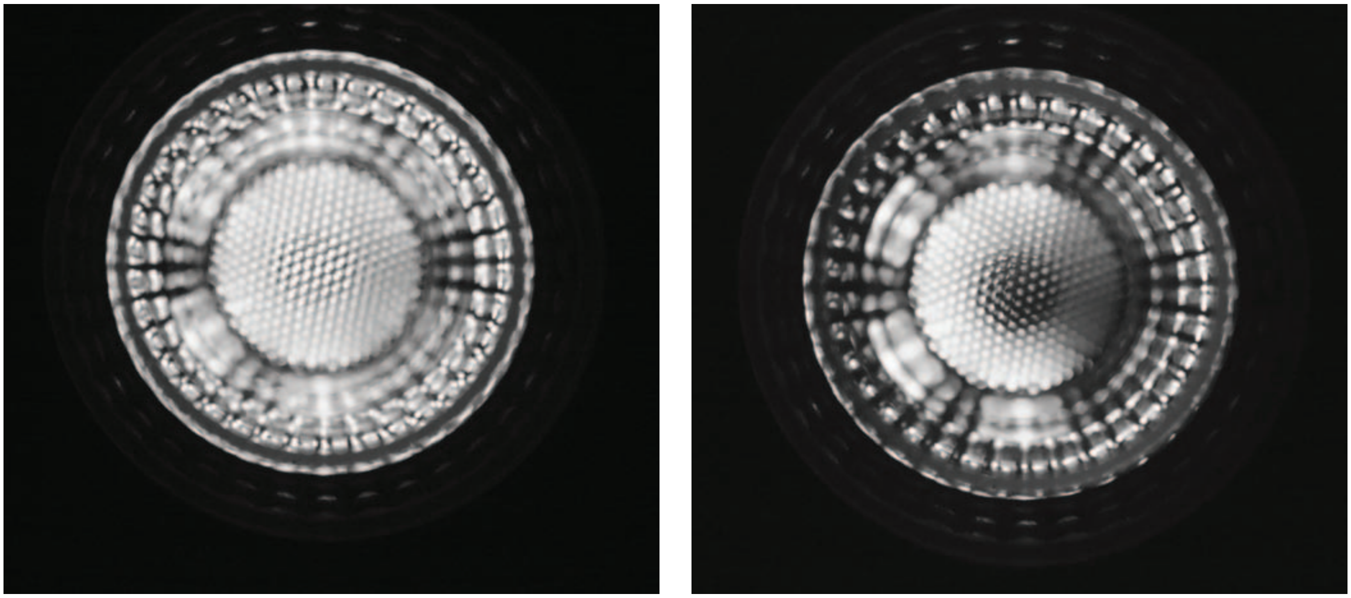

The Xerox copier lens could project an image of the Sleek LED PAR 20 face. It was fully illuminated when run at full, but parts of it dimmed out as I reduced power to it.

or sophisticated equipment that only a professional should do it—I will not be doing my own CT scans—but simpler is better, if possible.

The testing and data procedures outlined in *BSR E1.69* are fairly simple. You set up the test per *ANSI E1.25*, which is ESTA's standard setting the basic conditions for photometric testing that involves shining a luminaire at a wall and measuring the illumination levels produced. *E1.69* then requires you to dim the luminaire from 10% output to 0%, and then back up again to 10%, recording the brightest illumination produced at each dimming step. The number of dimming steps for digitally controlled luminaires is set by the control method: for eight-bit control, there would be 26 steps from 10% to 0%; for 16-bit, there would be 6,554 steps. (That's a lot of steps, but it might be changed. This is a first draft, a starting place for figuring out what is reasonable.) For analog voltage control, the dimming control signal steps are arbitrarily set at 0.5% increments—10%, 9.5%, 9.0%, and so on—for a total of 21 steps. For mains dimming, the control steps are much the same, with a total of 21 steps over the RMS voltage range of the mains

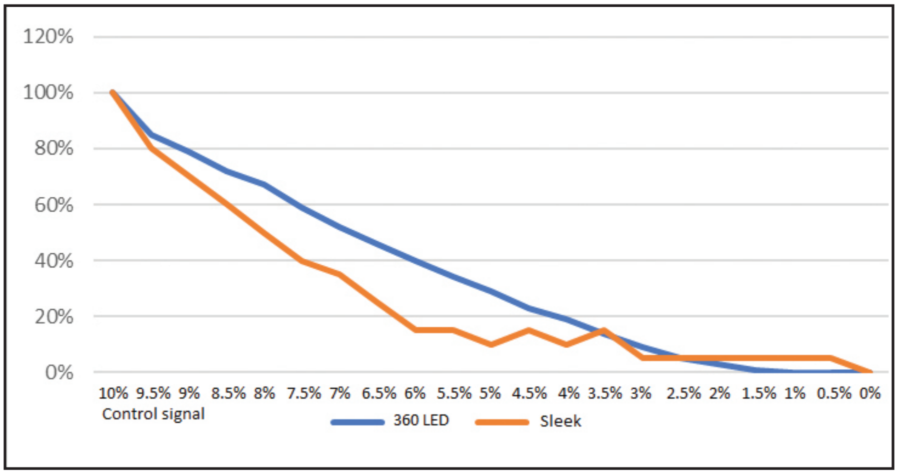

The low-end dimming outputs of the two luminaires, when plotted on the same graph, show significant differences. Both appeared to dim smoothly; the bobbles in the Sleek curve are reading differences of one or two lux, down at the level of noise. Not shown on this chart is that 0% for the Sleek unit was reached at about 44 V, while the 360 LED unit did not go black until below 24 V.

power that produces 10% to 0% dimming. (With mains dimming, the line voltage is both the electrical power and the control signal.) Then the results are reported with illumination level as a function of control signal level in tabular form, graph, or both. The illumination levels are normalized so that the illumination level produced at the 10% setting is the maximum level, 100%, and all the other levels are expressed as a percentage of that 10% control signal level, rather than as absolute illumination levels.

My working-at-home experiments were done with two home-brew LED luminaires dimmed with an autotransformer. This equipment was what I had access to, but also mains dimming is the most controversial method. It works and is particularly handy for retrofits in systems that lack a distributed control network, but mains-dimmed luminaires are the ones that generate the most discussion of what constitutes "theatrical quality" dimming. With money and time I could run tests on LED luminaires using a variety of control schemes, but running the tests on my two home-brew units was work enough for now—and these certainly were exceptional units that would prove the rule.

Why RMS voltage?

One thing that needs to be clarified in the standard is what the mains power measure is to be used to described control level with a mains-dimmed luminaire. I used "RMS voltage" because it is a simple way to describe how a dimmer controls the power to a dimmed load. I used an autotransformer for these experiments, which outputs a smooth sine wave of varying amplitude depending on the autotransformer setting. However, most modern, conventional theatrical lighting system dimmers use some version of phase control to dim a load, with power being delivered only for part of each AC cycle. If the conduction angle is 180° for each half-cycle, the dimmer is on full; reduce the conduction angle to 120°, and the dimmer output is a bit lower. Take the conduction angle down to 0°, and the dimmer is off. The output waveform at any setting other than full is a chopped sine wave, something that looks a bit like Yosemite's Half Dome on an oscilloscope.

RMS voltages are how we usually describe AC voltages; we do it so often that we assume that the RMS voltage of an AC circuit is what the voltage is, as though the voltage were a fixed value. It's not. On a 120 V RMS circuit the voltage varies through a full cycle from 0 to about 170 V peak. Common lamp cord has a maximum voltage rating of 300 V, which looks like a lot if we think of the line voltage as 120 V, but the line voltage actually is 170 V at its peak. A 300 V rating is not so high if we keep 170 V in mind.

We can describe the output of a phase control dimmer in terms of its RMS voltage, its root mean square voltage. It's a convenient way to describe the average voltage on the output of a dimmer using phase control. If the conduction angle is set at 120°, turning on before the AC wave hits its peak of 170 V, the RMS voltage would be just under 108 V RMS. For 60° of each half cycle the voltage is nothing; 108 V is the average.

RMS voltage is a good way to describe the output of a dimmer, whether it is a phase-control dimmer, an autotransformer dimmer, resistance dimmer, or something really unusual, such as a saturable core reactor or magnetic amplifier. It is a reasonable way to describe a dimmer's output—but we have to recognize that it may not be perfectly accurate for predicting how a mains-dimmed LED luminaire will behave on a dimmer. Given the variety of dimmers in the field, a mains-dimmed luminaire with active control electronics might sample the power waveform to determine the RMS voltage and take that as the output level signal. That's reasonable, but a manufacturer might assume that any modern dimmer will be a phase-control dimmer and design the control circuit to define the level from the conduction angle. Then the luminaire won't dim if powered through an autotransformer or resistance dimmer. RMS voltage is not a perfect parameter for characterizing dimming, but it's about all we can use.

I used an autotransformer dimmer because I have one available and because a phase-control dimmer with the 360 LED power supply would have been problematic. The input to that power supply is a transformer, an inductive load. Cheap phase-control dimmers, which are what I have, tend not to have the same conduction angles for each half-cycle when connected to inductive loads. The result is a DC offset in the output and a stinking, hot transformer.

Furthermore, the simple power supply I built will charge the capacitors up to the peak AC voltage. Never mind frying the transformer—I could dim a phase-control dimmer feeding the power supply down to a 120° conduction angle, 108 V RMS, but the AC peak would still be 170 V. The power supply output to the LED would still be a full 17 V. I would get no dimming at all until the conduction angle was less than 90°, half the dimming range. Plotting the bottom 10% of the dimming range was fussy enough without compressing it further.

The luminaires

One of the luminaires was the unit based on an Altman 360 I used for "Wasting Watts Revisited" published in the Spring 2017 issue of *Protocol*. For the tests in that article I used a variable, low-voltage, analog power supply; it wasn't mains-dimmed. This time I powered it through a simple unregulated power supply consisting of a 10:1 step-down transformer, a bridge rectifier, and two 82,000 µF capacitors in parallel for a total filter capacitance of 164,000 µF. That offered extremely smooth dimming. With the luminaire at full, I could switch off the power, and it would take 20 seconds to fade to black. Varying the input voltage to the power supply varied the output voltage and thus the drive current to the LED, which was limited by a ballast resistor.

The other luminaire was a Sleek 7.2 W PAR 20 mounted in a socket at the end of a coffee-can lens tube. The objective lens was from an old Xerox copier. With this I was able to project an image of the PAR 20's face, or put it slightly out of focus for a smoother field. The lamp dimmed as I varied the voltage supplied to it, but some parts of the lamp face dimmed faster than others. I used the part that stayed illuminated the longest as the "brightest spot" to be measured per *BSR E1.69*. I have no idea how the Sleek lamp dimmed, whether by decreasing the lamp current, pulse-width modulation, or some combination of both, but it dimmed.

I do not claim that either of these luminaires is representative of LED luminaires sold on the entertainment lighting market. They probably are not. However, either one could be used on a show. The Altman 360-based luminaire is one I had made for a puppet theatre company needing a silent pin-spot generating almost no heat.

The laundry room lab

My working-at-home laboratory was my laundry room. It is the one room in the house other than the bathroom where I can control the light level. It's big enough to offer the required throw, the washer and dryer are convenient, stable surfaces for the luminaire and autotransformer, and I could hang my photometer on the wine rack at the far end. However, the air conditioning there is minimal, so the ambient room temperature was 30° C—right at the upper limit allowed per *ANSI E1.25*. It was unpleasant. I did not want to do any more testing than necessary for this article. The work would have been more pleasant and faster if I had an assistant or had reconfigured the equipment so my voltmeter and lux meter's displays were next to each other—or if I did this testing in November instead of July. These changes are so obvious and simple, I don't think my sweat should be held as an argument against the *BSR E1.69* procedure.

Tests and results

Because the room was hot, I wanted to be there as short a time as possible, and didn't do the full set of testing required by *BSR E1.69*. It requires measuring the light levels produced as the control signal is stepped up from 0% to 10% as well as down from 10% to 0% to reveal hysteresis, but after doing the measurements in one direction in a dim, hot room, I figured I had enough data. The procedure fundamentally works, and shows that these two luminaires do indeed dim smoothly at the low end, but they will not track each other if they are run on the same mains dimmer. The Sleek unit will have dimmed to black well before the 360 LED reaches its 10% level.

I also didn't follow the procedure perfectly. You are supposed to measure the peak output with full line voltage, find the mains voltage when the output is 10% of that, then find the mains voltage at the bottom when the output is nothing. From those two voltage points—the 10% light output level and the 0% light output level—you set all the control signal steps in between. Getting the 0% light level was a tedious job of turning the autotransformer up a bit, then down, then up . . . . In frustration, on the 360 LED, the first luminaire I tested, I looked into the lens and turned the knob until I could just barely see the LED glow. I used that autotransformer output voltage as the bottom of the scale and set the steps in between. However, when I actually did the tests with my photometer set on its most sensitive scale, the two steps above the set 0% level at the low end showed no output at the meter hanging on the wine rack—but I could see the glow if I looked into the lens. Members of the Photometrics Working Group have discussed how a luminaire's output looks different if we stare into the light rather than look at what it's illuminating. I did not make this mistake with the Sleek unit.

If I plot the outputs of the two luminaires on the same graph, they show different curves—even if we allow for my mistake setting the bottom end of the 360 LED measurements. What's not on the graph, but is required by *BSR E1.69*, is that the mains voltage levels for 10% and 0% output be reported, and these two luminaires are quite different. The 10% and 0% control levels I set for the 360 LED were 33.90 V and 22.68 V mains voltages. For the Sleek unit, those two values were 52.10 V and 43.70 V—nowhere near close.

Where we go from here

BSR E1.69 was offered for public review from 19 June through 3 August 2020. Six people responded: three "Yes with comments," and three "No with reasons." Many of the comments were of the "it would be better if" and word-smithing sort. This really is to be expected. ***Very*** rarely does a standard get simple "Yes" responses on its first review. Unless people simply don't care about the standard, there almost is always someone offering advice on how to make it better. The Photometrics Working Group will consider all the comments and objections, and will make changes to the draft. However, we at least now have some practical experience on how this standard might actually work.

Look for another public review, probably in early 2021. ■

Thank you to Mike Wood of Mike Wood Consulting for calculating the RMS output voltage of a phase-control dimmer set at a 120º conduction angle.

Karl G. Ruling is ESTA's Technical Standards Manager. He also serves as *Protocol*'s Technical Editor. Karl can be reached at karl.ruling@esta. org.