CHARACTER CONTRAST

by Edward F. Kelley[*]

Measuring the contrast of a dark character on a light display surface is one of the hardest measurement results to obtain with accuracy. It is even more difficult when the display is subjected to ambient lighting conditions. With some display technologies, the full-screen contrast (sequential contrast) can be indicative of the contrast obtained when measuring small groups of pixels or even single-pixel character strokes. In general, however, that will not be the case because of light scattering within the display surface and any electronic irregularities that reduce contrast on a pixel scale. We are concerned here with an accurate measurement of the contrast regardless of how well the eye can see the contrast that we measure. Vision models can be applied after we attempt an accurate measurement.[1] (Of course, it is assumed that the development of the appropriate vision models dealing with such detail contrasts properly accounted for scattered light in the detection systems employed should that have been necessary.)

Veiling glare is the problem. Light from the bright areas can scatter within the detector and contaminate the dark areas. The scattering can occur between the lens elements or off their edges, off other objects such as apertures and shutters, and off any interior surfaces. The reason for the adjective “veiling” is that this type of glare tends to be uniformly present, and often the person using the detector cannot see the contamination. When such scattering in the detector is extreme it manifests itself as patterns of rings, spikes, and disks that are often called lens flare, which can be a useful artistic artifact to indicate a bright source of light in photography and videography.

However, for making accurate measurements, veiling glare is a serious problem. Often it is much more of a factor than people would think. For example, using an array camera with a charge-coupled-device (CCD) or compensated-metal-oxide-semiconductor (CMOS) detection array, contaminations of a small black area on a white screen can be well over 1000 %![2] Is this an indication that the expensive scientific-grade camera that we just purchased is inadequate? No. We must be aware of the limitations of the instrumentation we use so that we don’t expect the impossible. The secret to good metrology is to be aware of the limitations and know how to work around them in order to obtain an accurate measurement result. A good metrologist can often make a good measurement with a piece of junk whereas an inexperienced person can foul up a measurement with the best equipment available. Much depends upon attitude and awareness.

Errors in a small-area black of 1000 % or more can be alarming. But it really isn’t surprising when we think about what we are confronting. Consider an emissive display with a full-white-screen luminance of 250 cd/m^2. If the contrast of a single-pixel-wide character on that white screen is actually 250:1, then the black-character-stroke would have an actual luminance of 1 cd/m^2. If the veiling glare in the camera contaminated that black luminance measurement by only 10 cd/m^2 then that black measurement would be 11 cd/m^2 instead of 1 cd/m^2 and would be in error by 1000 %. The contrast would be incorrectly measured at 260:11 or about 24:1 instead of 250:1, and this would represent a 91 % error in the contrast. On the other hand, if we are viewing a relatively low character contrast display such as a display in a high-ambient-light environment, the character-stroke black for a display with an ambient white level of 250 cd/m^2 may have an actual ambient contrast of 5:1, or the black would actually have a luminance of 50 cd/m^2 under ambient conditions. The contamination of 10 cd/m^2 would only amount to 20 % of the black luminance in this case, and the ambient contrast with glare would be measured at 260:60 or 4.3:1, an error of only 14%. For a display with an ambient contrast of 3:1 and an ambient white luminance of 250 cd/m^2, a 10 cd/m^2 veiling-glare contamination would amount to a 12%

---

[*] * Optoelectronics Division, Electronics and Electrical Engineering Laboratory, U.S. Department of Commerce.

error in the black luminance of 83 cd/m^2 and give an ambient contrast of 260:93 or 2.8:1 that is in error only 7 %. Thus, the higher the contrast, the more serious it becomes to ignore the veiling-glare contribution to the measurement. (As a very rough rule of thumb, in a nontrivial scene many complicated camera lenses introduce approximately 4 % or 5 % of the average scene luminance as veiling glare, which is 10 cd/m^2 using our hypothetical 250 cd/m^2 display.)

Note that the veiling-glare contribution adds to both the black and the white measurement when we are trying to measure both white and black areas on the same screen at the same time. When we measure the full-screen contrast, known also as the sequential contrast, the effects of veiling glare cancels out in the division when we make the contrast calculation because the contamination for each full-screen is directly proportional to the luminance of each screen. However, this is only for full-screen contrasts. Whenever different luminances are on the same screen, the veiling-glare contributions to those luminances do not cancel out when considering contrasts.

So, how do we use our cameras or small-area spot detectors to make a measurement of a small dark area on a white screen? One way is to employ a replica mask.[3] A replica mask is an object placed near the small dark area we want to measure that is the same size as the area to be measured and for which we know its luminance L_M. We measure the white area luminance L_h, measure the black-pixel-area luminance L_d, and then measure the luminance of the replica L_R. If we know what the luminance of the replica should be (L_M), then we subtract that luminance form the measured replica luminance to obtain the veiling-glare luminance $L_G = L_R - L_M$. That veiling-glare luminance can then be subtracted from the measured white and the black luminances to obtain a better measurement result that accounts for the glare, $L_W = L_h - L_G$, $L_K = L_d - L_G$. The ratio of the veiling-glare-subtracted white and black is the true contrast of the small area of black, $C = L_W/L_K$.

In the case of an emissive display in a darkroom, we arrange for a black-opaque replica mask with its surface parallel to the display surface—see Figure 1. A white display screen can light up a normal room; so the darkroom must be of sufficient quality so that there is no contribution to the measured luminances from scattered light in the room or any objects in the room—this is essential. Additionally, the light-measuring device (LMD) or detector, either an array camera or a spot luminance meter, must be sufficiently far back from the screen so that reflections off it don't contribute to the measured luminances. For some apparatus that problem may require making measurements from a slightly off-normal direction to avoid the influence of the detector in the measurement results. Particularly for high-contrast-capable displays, stray light contributions from the apparatus or the room must be carefully controlled so that they are below the noise level of the measurements of the black areas. If the darkroom's stray light is properly controlled, then the actual luminance L_M of the replica will be zero. For such a case the measured luminance L_R of the replica is used as the glare luminance: $L_G = L_R$ with $L_M = 0$.

In order to make small-area measurements in a uniform ambient environment, the reflectance of the replica material will render it with a finite intrinsic luminance L_M that is not zero. In Figure 2 we show a top view of a display in an integrating sphere. This is for illustration purposes only; the integrating sphere dimensions shown would not be large enough to provide a uniform illumination of the display. Generally as a rule-of-thumb, the integrating-sphere diameter should be approximately seven times the size of the object being measured. A large sample of the replica material must be placed within the uniform illumination region near the display. It must be large enough so that the detector will not be exposed to any bright areas within the sphere when it measures the luminance L_M of the replica material (otherwise veiling-glare in the detector will provide us with the wrong luminance of the replica material). When using an integrating sphere, the display must be turned away from the measurement port so that its normal is $\theta = 6°$ to $10°$ from the center of the measurement port. If it would be useful to know the hemispherical diffuse reflectance ρ_M of the replica material, then a white reflectance standard may also be placed near the display to measure the illuminance E. Knowing the reflectance might be useful for using the replica material in other uniform-illumination situations as with a sampling sphere rather than a large integrating sphere.

Making the proper replica is not always easy. For a large character like a 48 point sans-serif capital “I” or other multi-pixel shape of that size it may be possible to cut some opaque black-plastic material to the exact same size (within 5% or so). But for a character stroke that might be only a single pixel wide, it may be essentially impossible to cut such a shape successfully. Noting that most of the glare that contributes to the contamination of a narrow line or straight character stroke comes from the immediate white area next to the stroke, we can cut a very narrow triangle of black material using a razor blade and place it a little distance away from the character of interest where the thickness of the triangle is the same as the character stroke—see Figure 3. That should provide an adequate replica to determine the veiling glare contribution. Whereas we would determine the white level by measuring full pixels, that may not be the right thing to do for the black pixels or the replica. Generally, there is a strong glare contribution at the boundary of a lit pixel and a dark area so that the luminance profile is not sharp but rounded as depicted in Figure 3. To avoid that rounded boundary, we would use a smaller area within the dark pixels or replica to estimate the luminances encountered. While this may not be a perfect solution to the problem, the contrast obtained usually gets us much closer to the actual contrast than if we didn’t attempt to account for veiling glare.

When using an array camera for our detector we should always try to use sufficient magnification to get 10 to 20 (preferably) detector pixels covering a single display pixel because of this rounding of the luminance profile into the dark areas—we want to clearly see the boundary region we need to avoid. Unless the display is very low contrast, a 16-bit array camera is often needed to span the luminance range encountered. If that is not the case or if the contrast exceeds the 16-bit camera’s capability, then two exposures will be needed, one for the black reading, and one for the white reading. In using array cameras, it is also important that they be photopic; that is, their spectral sensitivity must be very similar to the spectral luminous efficiency for photopic vision—the $V(\lambda)$ curve.

In making small-area measurements we have only discussed direct-view displays here. Front projection displays may also have small-area measurements made using replica masks as well as replicas that produce shadows. The idea is the same. We subtract from the white and black luminance measurements the luminance of some region of the same size—a replica—that represents the amount of stray light that we encounter.[4]

Using replica masks to measure small-area and character contrasts can be one of the most difficult measurements to make. However, not using some technique to eliminate the effects of veiling glare can produce measurement results for small-area black levels and associated contrasts that are very inaccurate—even absurd. Ultimately, it would be very helpful to have deconvolution techniques that would allow us to fully account for veiling glare, but such techniques would have to be tested to agree with actual measurements such as made with replica masks.

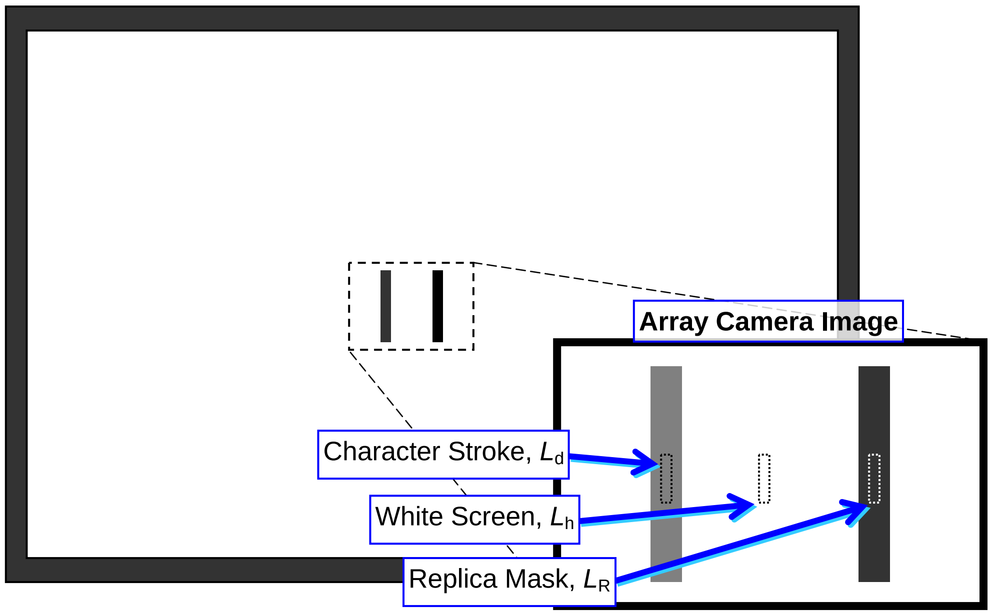

Figure 1. Small-area measurement for an emissive display in a darkroom using an array camera. The luminance measurement areas are noted by small dashed rectangles in the array-camera image inset.

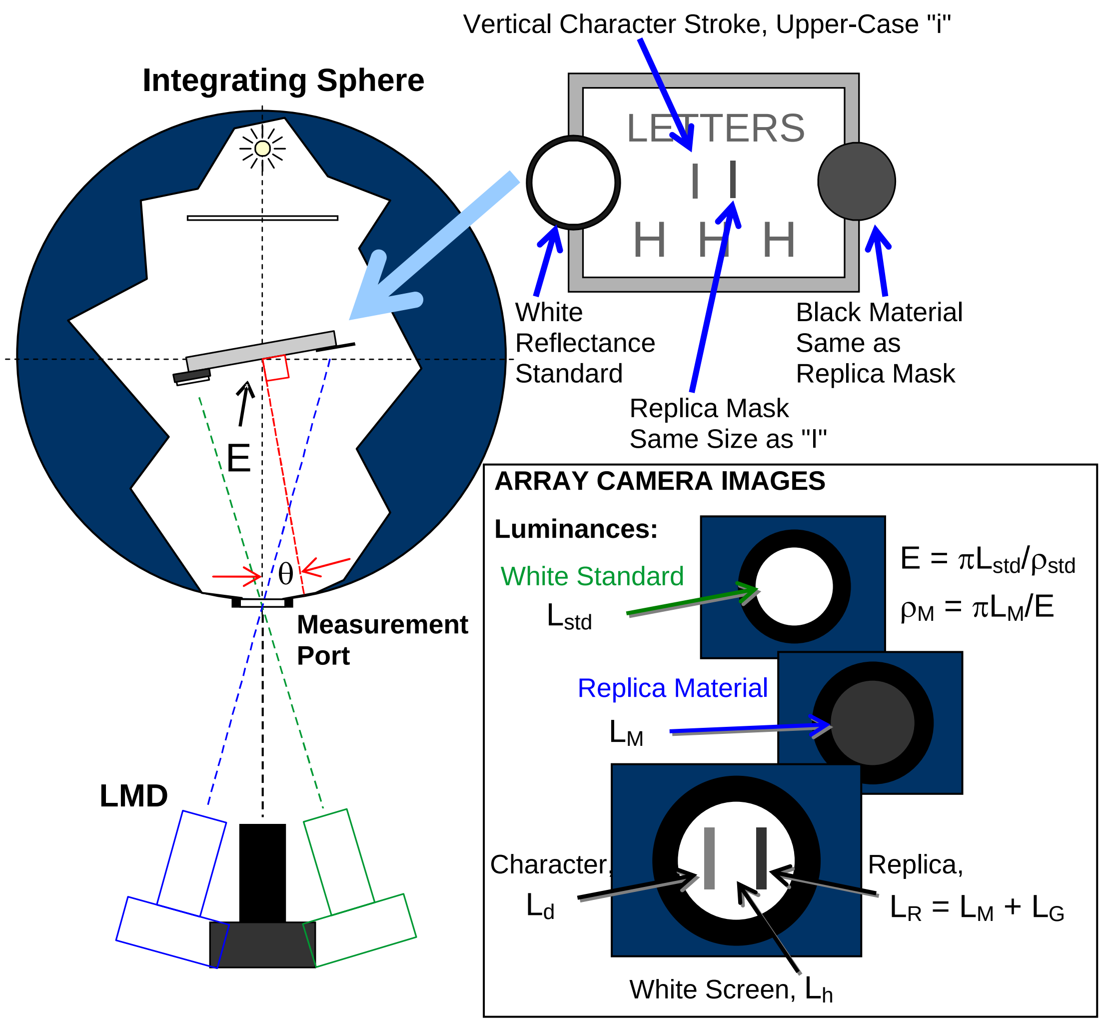

Figure 2. Small-area measurements of a display in a uniform-ambient environment using a replica mask and an array camera for the light-measuring device (LMD). The replica-mask luminance L_M is measured assuring no bright areas are visible from the camera. The integrating sphere should be larger than shown here for illustration purposes.

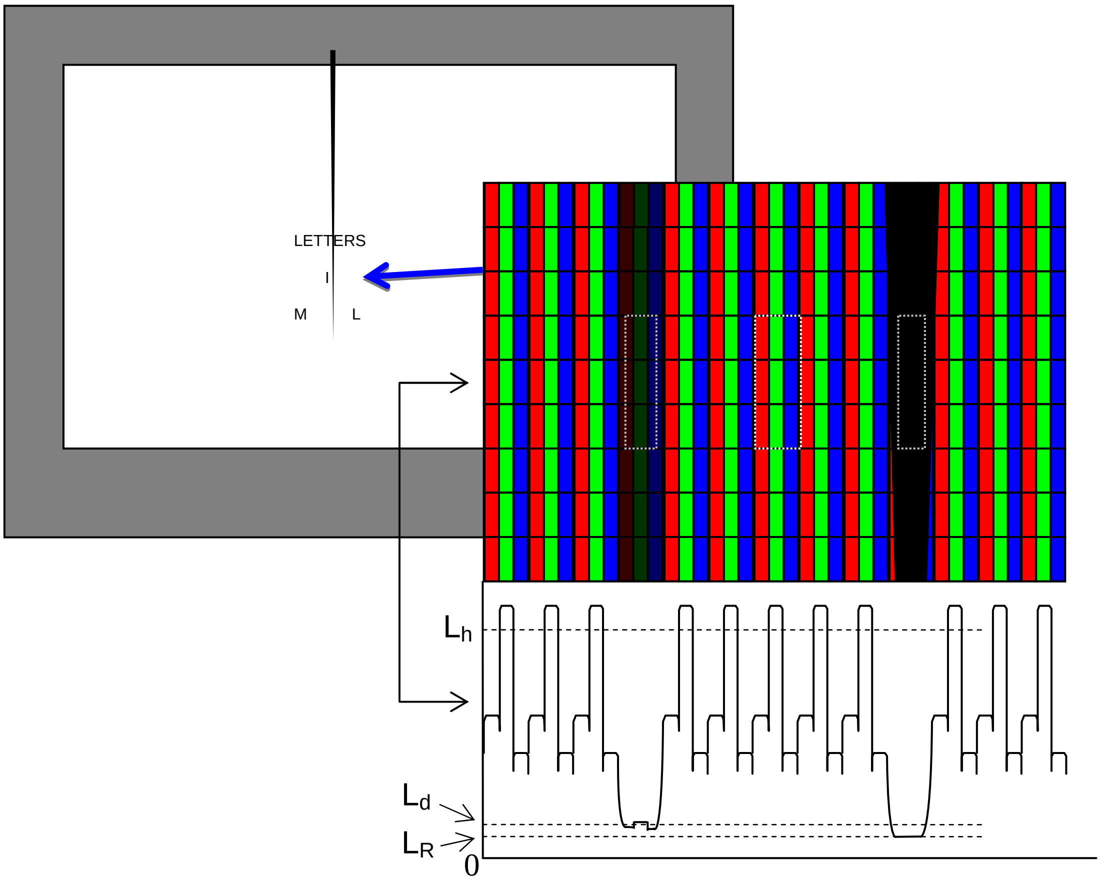

Figure 3. Narrow triangular replica mask used to match the thickness of the character stroke of the sans-serif letter "I." The dark regions are measured with areas thinner than a pixel width to avoid the rounding of the luminance profile that occurs from glare from the immediate proximity of the bright subpixels. (The luminance profile shown is for illustration purposes and is not necessarily scaled to actual relative levels encountered in practice.)

---

[1] T. G. Fiske and L. D. Silverstein, "Estimating Display Modulation by 2-D Fourier Transform: A Preferred Method," Journal of the Society of Information Display, Vol. 14, No. 1, pp. 101-105, January 2006.
[2] J. W. Roberts and E. F. Kelley, "Measurements of Static Noise in Display Images," Proceedings of the SPIE V4295B-27, Electronic Imaging Symposium, San Jose, CA, pp. 211-218, January 23, 2001.
[3] P. A. Boynton and E. F. Kelley, "Small-Area Black Luminance Measurements on a White Screen Using Replica Masks," 1998-SID International Symposium Digest of Technical Papers, Society for Information Display, Vol. 29, pp. 941-944, Anaheim, CA, May 17-22, 1998.
[4] P. A. Boynton and E. F. Kelley, "Stray Light Compensation in Small Area Contrast Measurements of Projection Displays," Projection Displays VIII, Proceedings of the SPIE, Vol. 4657, pp. 122-130 (January 2002).